AGENTS DON'T NEED A NEW INDEX: A SCIENTIFIC KNOWLEDGE LAYER OVER LEXICAL SEARCH

Anonymous authors
Paper under double-blind review

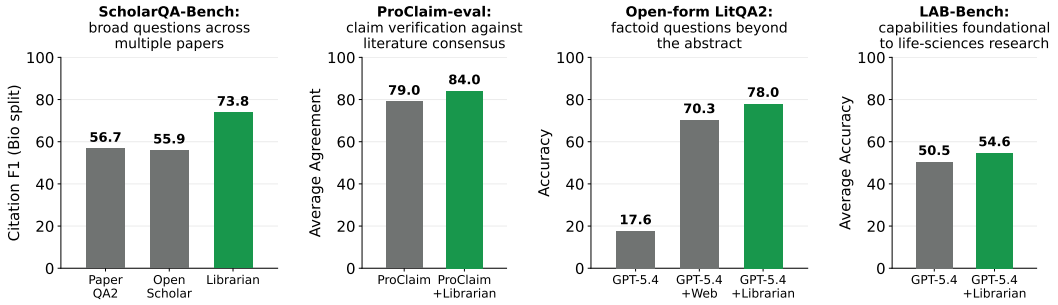

Figure 1: On four benchmark suites, agents equipped with the LIBRARIAN knowledge layer (green) outperform their baselines (grey): literature synthesis (ScholarQABench, Citation F1 on the Bio split), claim verification (ProClaim-eval, average agreement with expert consensus), open-form question answering (LitQA2, accuracy), and foundational biology tasks (LAB-Bench, average accuracy).

ABSTRACT

The web is increasingly accessed by AI agents rather than humans. Every agent needs knowledge, and nowhere more so than in science, where agentic pipelines are proliferating. Large scientific corpora are already reachable through mature lexical-based search engines, but these interfaces are not built for AI agents: they expect keywords and complex syntax and return whole papers, so every agent must learn the query syntax, issue several searches, and read entire articles to find the evidence it needs. We introduce LIBRARIAN, a knowledge layer that upgrades lexical-based search interfaces for AI agents: an agent asks in natural language and receives evidence that answers it. A single LLM orchestrates the whole knowledge retrieval process: it plans complementary subqueries executed by the search engine, then reads the selected papers and locates the relevant evidence. We instantiate LIBRARIAN on Europe PMC, an open lexical-based search engine that provides comprehensive access to scientific literature, and we evaluate multiple AI agents with access to LIBRARIAN across four scientific benchmarks: literature synthesis, claim verification, open-form question answering, and downstream biology tasks such as protocol questions and sequence manipulation. On ScholarQABench, agents equipped with LIBRARIAN improve Citation F1 by more than 16 points over strong recently published baselines. Used as the retrieval layer of an existing claim-verification pipeline, it increases agreement with expert consensus by 5 points, and on the open-form LitQA2 benchmark, a GPT-5.4 agent scores about 8 points higher when using the LIBRARIAN rather than web search. Our results show that existing knowledge infrastructure can be reused in AI pipelines without reindexing corpora or sacrificing performance. We release LIBRARIAN openly.

1 INTRODUCTION

Recent measurements indicate that traffic from automated agents, many powered by large language models (LLMs), now rivals or exceeds human-generated traffic across much of the web (Imperva, 2025; HUMAN Security, 2026), motivating a rethinking of web infrastructure toward an agent-first approach (Bandara et al., 2026; Yang et al., 2025). This rethinking is visible across domains. In scientific research, for instance, agents are already embedded in everyday research workflows: they

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

help researchers summarize findings across many papers (Skarlinski et al., 2024; Asai et al., 2026), execute code and orchestrate specialized tools (Huang et al., 2026; Gao et al., 2025; Aygün et al., 2026), and generate testable hypotheses (Mitchener et al., 2025; Gottweis et al., 2026; Ghareeb et al., 2026). Common to all of these applications is a dependence on the scientific literature: whether an agent is summarizing findings or proposing a hypothesis, it must anchor its reasoning in published evidence. Providing that evidence, however, is far from straightforward: the literature is growing exponentially (Bornmann et al., 2021), intensifying the information overload in research (Klerings et al., 2015; Gusenbauer, 2021; Galli et al., 2025).

A popular way to expose the literature to AI agents is building a dense vector database over it: papers are chunked into passages, embedded into vectors, and the nearest passages are retrieved so an agent can query the database in natural language and receive citable passages rather than read whole papers (Karpukhin et al., 2020; Lewis et al., 2020; Skarlinski et al., 2024; Volkova et al., 2025; Shi et al., 2025; Asai et al., 2026; Ding et al., 2026; Zhong et al., 2026). However, a dense vector index has drawbacks. First, it is expensive: indexing and serving the whole literature demands a heavy infrastructure investment with hundreds of gigabytes of embeddings (OpenScholar's datastore alone measures roughly 728 GB (Asai et al., 2026)). Second, the margin dense retrieval offers is narrowing as LLMs grow more capable (Hsu et al., 2026; Hu et al., 2026). Third, embeddings flatten structured metadata, preventing direct queries over explicit metadata fields. Finally, a nearest-neighbour lookup over vectors is not inspectable, unlike a structured, fielded keyword query.

The scientific literature, however, is already indexed. Lexical search engines have served it for decades and remain strong on the axes that matter for grounding, such as corpus coverage, freshness, and out-of-domain retrieval quality (Thakur et al., 2021). One such engine is Europe PMC (Rosonovski et al., 2024; Ferguson et al., 2021), an open literature search engine over scientific literature, with coverage of medicine and health subjects that exceeds PubMed (Gusenbauer, 2022). Its fielded query syntax has been refined through years of user-centered development (Levchenko et al., 2018; Rosonovski et al., 2024), and it serves over 32 M unique IPs per year (European Commission, 2026).

Scientific agents are increasingly part of the Europe PMC user base (Huang et al., 2026; Mitchener et al., 2025; Skarlinski et al., 2024). Yet that interface was not built for agents. On the input side, agents must use keyword-based queries rather than natural language, and because scientific entities have multiple aliases (*e.g.*, gene symbols, protein names, disease synonyms, or MeSH terms), multiple complementary searches are often required for comprehensive retrieval (Wang et al., 2023; 2025). On the output side, results are whole documents ranked by lexical score, of which only a fraction is relevant. Reading full papers is especially costly for agents with fixed context windows, as each document consumes a large fraction of the token budget (Hsu et al., 2026; Hu et al., 2026). These challenges motivate a novel scientific knowledge layer for AI agents: a shared component, queried in natural language, that returns a compact set of citable evidence snippets, which recent work suggests might improve performance more than scaling the model alone (Nasri et al., 2026).

We introduce LIBRARIAN, an agent-first scientific knowledge layer that sits over a lexical search engine rather than over an index of its own: given the query grammar and record schema the engine exposes, a single LLM generates complementary keyword and fielded queries, retrieves the matching records, decomposes selected papers into paragraphs, scores candidate evidence against the original question, and returns a compact, ranked set of citable evidence with source metadata. We instantiate it on Europe PMC's live search, preserving the natural-language, citable-evidence interface of dense retrieval without building an index. LIBRARIAN is model-agnostic: any LLM can serve as its engine, and as newer, more capable models emerge, the system inherits the benefit.

We evaluate LIBRARIAN across four settings that require grounding in scientific literature: literature synthesis (writing evidence-backed summaries over multiple papers), claim verification (checking whether published evidence supports a scientific claim), open-form question answering, and biology workflow tasks (*e.g.*, protocol questions, sequence manipulation, and molecular cloning). In every setting, equipping agents with the LIBRARIAN knowledge layer yields improvements (Figure 1). For literature synthesis, it raises Citation F1 by more than 16 points on average over strong recently published agentic baselines. For claim verification, using LIBRARIAN as the evidence-retrieval layer of the ProClaim pipeline (Ai et al., 2026) increases agreement by 5 points on average. For open-form QA, equipping a frontier model (GPT-5.4) with LIBRARIAN improves LitQA2 accuracy by $\sim$8 points over web search (Skarlinski et al., 2024). For biology workflow tasks (LAB-Bench (Laurent et al.,

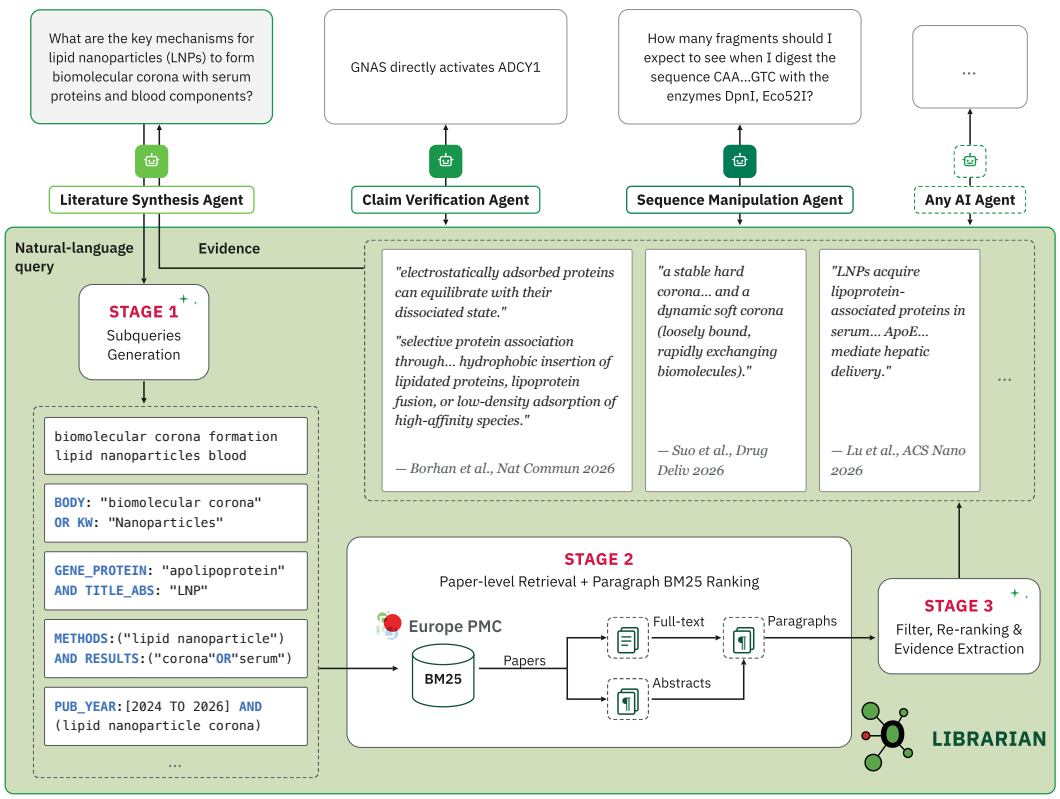

Figure 2: Overview of LIBRARIAN. Starting from a natural-language query, it runs three stages: (1) generates complementary Europe PMC subqueries; (2) retrieves papers from the live search engine and ranks their paragraphs with BM25; (3) filters and re-ranks the top-ranked paragraphs and abstracts, then extracts and returns the final evidence.

2024)), LIBRARIAN improves macro accuracy of GPT-5.4 by $+4.2$ points, with a large gain of $+11.3$ on manipulating biological sequences. To summarize, our contributions are as follows:

1. We define a knowledge-layer interface for the agentic era: an agent sends a natural-language question and receives a compact set of citable evidence snippets, without learning a query syntax or reading whole papers.
2. We introduce LIBRARIAN, a system architecture that reuses the live infrastructure of a lexical search engine, instantiated here on Europe PMC. A single, model-agnostic LLM orchestrates the entire retrieval pipeline.
3. We run an extensive evaluation across multiple datasets from four complementary benchmark suites, showing that equipping diverse agents with the same LIBRARIAN layer consistently improves their performance.
4. We release code, prompts, configurations, and plugins for existing agent frameworks.

2 LIBRARIAN

Given a user query q, LIBRARIAN returns a ranked list of evidence items $\mathcal{E} = [e_1, \ldots, e_m]$. We define an evidence item as the set comprising a short text span, paper metadata for citation, and the query used for retrieval (Figure 2). This represents a much finer-grained unit than the entire document typically returned by conventional lexical-based search engines. Because LIBRARIAN returns evidence rather than whole documents, downstream agents can directly quote, cite, and reason over the results, without writing queries in the service's syntax or reading full papers. At the same time, each item keeps the query that produced it, so retrieval stays transparent: an agent (or the user) can check how the evidence was found. We next describe how this evidence is retrieved and ranked.

2.1 EVIDENCE RETRIEVAL

LIBRARIAN comprises three stages (Figure 2). It begins with recall-oriented retrieval because scientific concepts are expressed through many surface forms, including gene symbols (*e.g.*, TP53), protein names (*e.g.*, p53), disease names (*e.g.*, ALS or amyotrophic lateral sclerosis), organism constraints, and controlled vocabularies such as MeSH. A record indexed under one form may be missed by a Europe PMC query targeting another, limiting the recall of any single query. LIBRARIAN therefore issues multiple complementary subqueries. The subsequent stages rank evidence within retrieved papers and filter false positives, selecting the evidence most relevant to the original question. Overall, LIBRARIAN progresses from recall-oriented retrieval to precision-oriented evidence selection.

Stage 1: Subquery generation. LIBRARIAN maps the user query q to a set of complementary subqueries $Q = f_{\text{plan}}(q) = \{q_1, \ldots, q_N\}$, where f_{plan} is an LLM. A single prompt converts the user query into a JSON list of subqueries; each q_i is a keyword or fielded subquery that can target title, abstract, keyword/MeSH terms, gene or protein, organism, disease, chemical, and other Europe PMC fields (prompt and field grammar in Appendix D.1). Each subquery passes through a deterministic syntax validator before execution; malformed expressions are dropped. The maximum number of subqueries N is fixed.

Stage 2: Paper-level retrieval and paragraph ranking. The validated subqueries are executed in parallel as Europe PMC queries against the live Europe PMC search service, which returns papers in relevance order. Each Europe PMC query returns up to P records, for a pool of at most $M = N \times P$ records; the concrete values are given in Section 3. Complementary subqueries can return the same paper. Papers are not deduplicated across subqueries before ranking, so a paper retrieved by multiple subqueries participates independently in each subquery-specific ranking. Each retrieved paper is decomposed into paragraphs: LIBRARIAN includes its abstract and, when full text is available, retrieves the XML and extracts its paragraphs, discarding references and other front/back matter; paragraphs longer than a fixed window are split into overlapping chunks. It ranks these paragraphs using BM25 (Robertson et al., 1994), treating their union as the corpus; papers without full text contribute only with their abstracts (implementation details in Appendix A). Each paragraph is concatenated with its section title and type (*e.g.*, methods, results, discussion), scored against the corresponding subquery q_i with its fielded operators removed. Because a subquery retrieved these papers through lexical matches, scoring against the same subquery finds, within each paper, the paragraphs with the highest lexical overlap, i.e., the passages that made the paper match. Each subquery passes its k highest-ranked paragraphs to Stage 3, where k is a per-subquery limit and a paper may supply several paragraphs or none. Finally, the selections across subqueries are combined, and duplicate or overlapping selections from the same source paragraph are merged.

Stage 3: Filtering, re-ranking, and evidence extraction. At this point, the candidate evidence is represented as a list of paragraphs. LIBRARIAN concatenates these paragraphs, segments them into sentences with pySBD (Sadvilkar & Neumann, 2020), and prepends a unique identifier to each sentence in the prompt. A single LLM call then evaluates all pooled paragraphs at once (up to $N \times k$ chunks) and returns an ordered list of sentence identifiers (prompt in Appendix D.2). This call serves three purposes: (i) filtering out irrelevant paragraphs, (ii) ranking the retained evidence, and (iii) returning it as individually citable sentences. This is the precision-oriented stage: the model jointly evaluates the original question q and every candidate sentence in one pass, providing a fine-grained relevance assessment that lexical scoring cannot provide. Because the model emits short identifiers instead of regenerating text, it decodes far fewer tokens, lowering latency and cost.

Finally, LIBRARIAN assembles the output: it returns the selected sentences, clustered by source paper and paired with their metadata. Contiguous selected sentences are combined into evidence spans, and the first selected sentence from each paper determines the paper's position in the ranked output. The result $\mathcal{E}$ is a ranked list of citable evidence snippets that downstream agents can use for the task at hand (*e.g.*, synthesis, claim verification, or question answering).

3 EXPERIMENTS

We evaluate whether task-specific agents improve when equipped with the LIBRARIAN knowledge layer, across four benchmarks that probe capabilities: citation-grounded synthesis across multiple papers on ScholarQABench (Section 3.1); claim verification against the literature consensus on

Table 1: ScholarQABench results. As *Retriever*, we consider OpenScholar, SciRAG, and BM25 over the OpenScholar Data Store (OSDS); PaperQA2 on its own corpus; and LIBRARIAN over Europe PMC. We report *Citation F1* on the Bio, Neu, and Multi splits, and the *LLM* score (1–5) on Multi.

		Citation F1			LLM
Agent	**Retriever**	**Bio**	**Neu**	**Multi**	**Multi**
OpenScholar-8B (Asai et al., 2026)	OpenScholar on OSDS	50.8	56.8	<u>74.1</u>	3.84
OpenScholar-70B (Asai et al., 2026)	OpenScholar on OSDS	55.9	63.1	–	–
GPT-4o (OpenAI et al., 2024)	OpenScholar on OSDS	36.3	21.9	54.3	3.79
PaperQA2 (Skarlinski et al., 2024)	PaperQA2	56.7	56.0	–	–
SciRAG (Llama-3.1-70B) (Ding et al., 2026)	SciRAG on OSDS	43.1	45.9	–	–
SciRAG (GPT-4o) (Ding et al., 2026)	SciRAG on OSDS	44.8	36.2	–	–
Synthesis Agent (GLM-5)	BM25 on OSDS	67.0	68.5	72.9	<u>4.10</u>
Synthesis Agent (GLM-5)	LIBRARIAN on OSDS	<u>71.2</u>	<u>73.9</u>	–	–
Synthesis Agent (GLM-5)	LIBRARIAN	**73.8**	**79.4**	**76.2**	**4.91**

ProClaim-eval (Section 3.2); factoid question answering beyond the abstract, on open-form LitQA2 (Section 3.3); and capabilities foundational to scientific research, from database and sequence questions to cloning and protocols, on LAB-Bench (Section 3.4). Together they range from open-ended synthesis to narrow factoid retrieval, isolating where the knowledge layer helps most.

Implementation Details. In our experiments, we use GLM-5 (GLM-5-Team et al., 2026) for LIBRARIAN. GLM-5 is self-hosted on an NVIDIA DGX B200 and served through vLLM (Kwon et al., 2023). In our evaluation, LIBRARIAN uses $N = 8$ subqueries, retrieves up to $P = 50$ articles per subquery, and passes the $k = 16$ highest-ranked paragraphs per subquery to Stage 3.

3.1 SCHOLARQABENCH: BROAD QUESTIONS ACROSS MULTIPLE PAPERS

ScholarQABench (Asai et al., 2026) is a collection of broad questions written by PhD-level experts across multiple fields; answering each requires synthesizing evidence from several articles. We use its biomedicine (Bio) and neuroscience (Neu) splits, together with the subset of multidisciplinary (Multi) questions that fall under the Bio and Neu categories.

Agents. We build a *Synthesis Agent* that, given a question, issues a retrieval call and generates a citation-grounded summary. As retrievers, we consider LIBRARIAN and a BM25 index over the OpenScholar Data Store (OSDS). The agent is powered by GLM-5 (one prompt), and we compare against the best agents from prior work (Asai et al., 2026; Skarlinski et al., 2024; Ding et al., 2026).

Metrics. Following Asai et al. (2026), we report *Citation F1* and *LLM*. *Citation F1* is the harmonic mean of citation precision and recall: recall measures whether every citation-worthy statement is supported by appropriate references, while precision measures whether each cited source is both relevant to and necessary for the statement it supports. *LLM* is the average score (1–5) from an LLM-as-a-judge (*i.e.*, Prometheus (Kim et al., 2023)) across three dimensions: organization, coverage, and relevance. See Asai et al. (2026) for details.

Findings. Table 1 reports our results. Synthesis Agent outperforms the baselines despite using the same retrieval corpus (OSDS), improving *Citation F1* by $+11.1$ on *Bio* and $+5.4$ on *Neu* over OpenScholar-70B. Because this row uses plain BM25 retrieval, we attribute the gain mainly to the backbone: GLM-5 is a newer and substantially stronger model than the baselines, though not larger per token, being a mixture-of-experts model with 744B total but about 40B active parameters (GLM-5-Team et al., 2026). The Synthesis Agent becomes stronger when equipped with LIBRARIAN, with a further $+6.8$ on *Bio* and $+10.9$ on *Neu* over the BM25 index on OSDS. We attribute this to the LIBRARIAN pipeline providing better evidence, a hypothesis we further validate on other benchmarks below. With the corpus held fixed to OSDS, LIBRARIAN still improves over BM25 by +4.2 (Bio) and +5.4 (Neu) points, about half or more of its total gain over that baseline; the rest reflects the change to Europe PMC's live corpus (App. C.1 describes the OSDS setup; see also App. B). On *Multi*, the *LLM* score of the Synthesis Agent with OSDS is higher than the baselines, suggesting that the stronger backbone also helps here; the LIBRARIAN knowledge layer yields a 19.8% improvement, further supporting our hypothesis.

Table 2: ProClaim-eval results. We evaluate the ProClaim agent (Ai et al., 2026), a simpler Verifier Agent, and an OpenScholar baseline, all powered by Claude Sonnet 4.6. The retrievers are OpenScholar over Semantic Scholar (S2), the ProClaim retriever (PubMed and S2), and LIBRARIAN over Europe PMC. Agreement with expert consensus is reported on average (AVG) and per dataset.

Agent	Retriever	Agreement ↑		
		AVG	SIGNOR	ConnectomeDB
OpenScholar	OpenScholar on S2	0.44	0.38	0.46
Verifier Agent	PubMed + S2 Search	0.53	0.62	0.49
ProClaim	PubMed + S2 Search	<u>0.79</u>	<u>0.80</u>	<u>0.78</u>
Verifier Agent	LIBRARIAN	0.68	0.75	0.66
ProClaim	LIBRARIAN	**0.84**	**0.84**	**0.84**

3.2 PROCLAIM-EVAL: CLAIM VERIFICATION AGAINST LITERATURE CONSENSUS

ProClaim-eval (Ai et al., 2026) evaluates an agent's ability to determine, from the literature, the scientific consensus on a given claim. The task requires retrieving evidence from multiple articles and assigning one of three verdicts: SUPPORT (the evidence supports the claim), REFUTE (the evidence refutes it), or UNCERTAIN (the evidence is insufficient). The benchmark contains 419 claims from SIGNOR (Lo Surdo et al., 2026) (protein–protein interactions) and ConnectomeDB (Liu et al., 2026) (ligand–receptor interactions), each with an expert-curated consensus verdict.

Agents. The ProClaim agent (Ai et al., 2026) orchestrates three stages inside a refinement loop: (i) *retrieval* searches for evidence; (ii) a *sub-agent* reads the retrieved full-text papers, extracts claim-relevant facts, and assigns each a verdict label; and (iii) a *sufficiency classifier* decides whether the evidence suffices to reach a verdict. If not, the loop restarts from (i). As the retriever, we use either the original from Ai et al. (2026), combining PubMed and Semantic Scholar (S2), or LIBRARIAN over Europe PMC. Adapting ProClaim to LIBRARIAN requires minor modifications to stages (ii) and (iii) (Appendix C.2). We also implement a simpler *Verifier Agent* (a single prompt) that, given a claim and the retrieved evidence, emits a verdict. As a baseline, we include the OpenScholar agent (Asai et al., 2026). All agents use Claude Sonnet 4.6 (Anthropic, 2026).

Metrics. We report *Agreement (AGR)*, defined as the fraction of claims assigned the correct verdict; higher values of AGR indicate stronger alignment between model predictions and gold annotations.

Findings. Table 2 reports the results. First, the one-prompt Verifier Agent, when equipped with LIBRARIAN, is already strong: it reaches 0.68 average agreement, $+15$ points over retrieving from PubMed and Semantic Scholar (S2), and only 11 points below the original ProClaim pipeline. This suggests that LIBRARIAN already does much of the heavy lifting by surfacing the relevant evidence, leaving the downstream agent only the task of turning that evidence into a verdict. The same holds for ProClaim itself: switching its retriever to the LIBRARIAN knowledge layer yields $+5$ average agreement points ($0.79 \rightarrow 0.84$), further supporting the hypothesis that LIBRARIAN provides better access to the underlying knowledge. Both agents far exceed the OpenScholar baseline (0.44).

3.3 OPEN-FORM LITQA2: FACTOID QUESTIONS BEYOND THE ABSTRACT

LitQA2 (Skarlinski et al., 2024) is a benchmark of narrow, difficult factoid scientific questions designed to be unanswerable from the abstract alone, requiring full-text evidence. It is distributed as a multiple-choice benchmark; we create an open-form version in which the answer must be produced directly from the literature rather than selected from candidate options. We retain the 91 questions whose supporting paper is available as full text in Europe PMC; all systems in Table 3 see those questions, including the web-search arm, which is not restricted to that corpus.

Agents. We consider a RAG-style *QA Agent* (Lewis et al., 2020) powered by GPT-5.4 that retrieves evidence and, using only what is returned, either produces an answer or abstains when the evidence is insufficient. As a baseline, we include the OpenScholar-8B agent (Asai et al., 2026).

Table 3: Open-form LitQA2 results. We run a one-prompt QA Agent (GPT-5.4) under three retrieval settings: no retrieval (Parametric), Web-Search, and LIBRARIAN. We compare against OpenScholar-8B (Asai et al., 2026).

Agent	Retriever	Cov	Prec	Acc
OpenScholar-8B Asai et al. (2026)	Semantic Scholar	84.6	40.3	34.1
QA Agent	Parametric	**100**	17.6	17.6
QA Agent	Web-Search	92.3	<u>76.2</u>	<u>70.3</u>
QA Agent	LIBRARIAN	<u>94.5</u>	**82.6**	**78.0**

Metrics. We use an LLM as a judge, powered by GPT-4o, to determine whether an answer is correct (*i.e.*, explicitly contains or unambiguously paraphrases the gold answer). We report three metrics: Coverage (*Cov.*), the fraction of questions on which the model produces an answer rather than abstaining; Precision (*Prec.*), the fraction of produced answers that are correct; and Accuracy (*Acc.*), the overall fraction of correct answers, with abstentions counted as incorrect.

Findings. We report results in Table 3. Without retrieval, the QA Agent (GPT-5.4) always commits to an answer (Coverage 100), but relying on its parametric knowledge produces frequent hallucinations: precision and accuracy collapse to 17.6. Grounding the answer in web search changes this dramatically, raising precision to 76.2 and accuracy to 70.3, consistent with prior work on retrieval-augmented generation (Lewis et al., 2020; Petroni et al., 2020). Replacing web search with LIBRARIAN improves both axes: precision rises by $+6.4$ points to 82.6 and accuracy by $+7.7$ to 78.0, while coverage increases from 92.3 to 94.5. The agent thus answers more questions while being correct more often, suggesting that LIBRARIAN returns higher-quality and more relevant evidence than generic web search for difficult, domain-specific questions whose answers lie in full text. These three settings share the same GPT-5.4 backbone; the OpenScholar-8B baseline (34.1 accuracy over Semantic Scholar) is included only as an external reference, as it differs in backbone model.

3.4 LAB-BENCH: CAPABILITIES FOUNDATIONAL TO BIOLOGICAL RESEARCH

The Language Agent Biology Benchmark (LAB-Bench) (Laurent et al., 2024) is an evaluation suite that assesses agents on foundational biology research tasks. Of its eight categories, we consider four: retrieving information from databases (*DbQA*), troubleshooting biological protocols (*ProtocolQA*), manipulating biological sequences (*SeqQA*), and tackling a particularly challenging set of molecular cloning tasks (*Cloning Scenarios*). We exclude reasoning over figures (*FigQA*), tables (*TableQA*), and supplementary material (*SuppQA*), which are capabilities that LIBRARIAN does not yet support and that we leave to future work. The remaining category, *LitQA2*, is covered separately in Section 3.3.

Agents. We report GPT-4o and GPT-5.4 in two conditions: with no external evidence (Base) and with evidence retrieved by LIBRARIAN (+ LIBRARIAN). All GPT numbers are from the LAB-Bench public subset ($\sim 80\%$); Claude 3.5 Sonnet and Human Expert values are instead quoted from (Laurent et al., 2024) on the full private subset and serve only as references.

Metrics. As in the open-form QA setting (Section 3.3), we report Coverage, Precision and Accuracy.

Findings. Table 4 reports results across the four LAB-Bench datasets. Adding LIBRARIAN benefits GPT-5.4 and GPT-4o through different mechanisms, visible in the precision/coverage decomposition. For the weaker GPT-4o, LIBRARIAN trades coverage for precision: precision rises (macro $49.0 \rightarrow 54.0$), while coverage falls (macro $70.2 \rightarrow 64.8$), leaving accuracy essentially flat (macro $35.2 \rightarrow 35.3$). The mechanism is calibration rather than raw accuracy: grounding suppresses the hallucinations of a model that otherwise leans on parametric knowledge, so it declines more often when evidence is inconclusive. This is a desirable trade even at constant accuracy. For the stronger GPT-5.4, precision is stable under added evidence (macro $+0.0$) while LIBRARIAN lifts coverage (macro $+6.8$), letting the model correctly answer questions it would otherwise abstain from. The clearest case is *SeqQA* (accuracy $52.5 \rightarrow 63.8$, coverage $85.0 \rightarrow 98.8$): the model was abstaining for lack of retrieved evidence, and LIBRARIAN supplies it. This does not hold for *DbQA*, whose questions resemble direct database lookups: retrieved literature adds little decision-relevant evidence, and grounding helps neither model (accuracy -2.3 and -1.3). Both models remain well below human experts (macro

Table 4: Performance on four LAB-Bench datasets. We compare GPT-4o and GPT-5.4, each with and without LIBRARIAN, against Claude 3.5 Sonnet and Human Experts. We report Coverage (Cov), Precision (Prec), and Accuracy (Acc); parentheses in the LIBRARIAN columns give the change over the base model in percentage points. Human Experts (*italicised*) are a reference ceiling.

Dataset		*Human Experts*	Claude 3.5 Sonnet	GPT-4o		GPT-5.4	
				Base	+ LIBRARIAN	Base	+ LIBRARIAN
DbQA	Cov	*86.6*	25.4	46.7	32.9 (−13.8)	74.2	**75.8** (+1.6)
	Prec	*86.3*	50.6	42.8	**53.8** (+11.0)	50.5	47.7 (−2.8)
	Acc	*74.8*	12.8	20.0	17.7 (−2.3)	**37.5**	36.2 (−1.3)
Cloning Scenarios	Cov	*82.0*	52.0	57.6	63.6 (+6.0)	90.9	**97.0** (+6.1)
	Prec	*73.0*	**54.0**	47.4	52.4 (+5.0)	46.7	46.9 (+0.2)
	Acc	*60.0*	28.0	27.3	33.3 (+6.0)	42.4	**45.5** (+3.1)
SeqQA	Cov	*92.3*	76.2	81.0	74.8 (−6.2)	85.0	**98.8** (+13.8)
	Prec	*85.5*	**67.6**	50.4	51.0 (+0.6)	61.8	64.6 (+2.8)
	Acc	*78.9*	51.5	40.8	38.2 (−2.6)	52.5	**63.8** (+11.3)
ProtocolQA	Cov	*91.0*	73.0	95.4	88.0 (−7.4)	93.5	**99.1** (+5.6)
	Prec	*87.0*	66.0	55.3	58.9 (+3.6)	**74.3**	73.8 (−0.5)
	Acc	*79.0*	48.0	52.8	51.9 (−0.9)	69.4	**73.1** (+3.7)
Macro-avg	Cov	*88.0*	56.7	70.2	64.8 (−5.4)	85.9	**92.7** (+6.8)
	Prec	*83.0*	**59.6**	49.0	54.0 (+5.0)	58.3	58.3 (+0.0)
	Acc	*73.2*	35.1	35.2	35.3 (+0.1)	50.5	**54.6** (+4.2)

accuracy 73.2), leaving substantial headroom, and the gains from LIBRARIAN on the literature-driven tasks point to grounding as a promising way to close it.

4 RELATED WORK

Dense-indexed retrieval for scientific agents. OpenScholar pairs a $\sim$45M-paper datastore of dense passage embeddings with reranking, self-feedback, and citation-backed synthesis (Asai et al., 2026); SciRAG adds adaptive retrieval and citation-graph reasoning (Ding et al., 2026); PaperQA2 performs agentic synthesis over full-text passages (Skarlinski et al., 2024); and BioSage and SciSage build domain-specialized agents for cross-disciplinary discovery and survey generation, respectively (Volkova et al., 2025; Shi et al., 2025). These systems demonstrate the value of passage-level scientific retrieval, but each depends on a separately maintained corpus, with the indexing and serving cost this entails; on bespoke retrieval, reranking, or synthesis pipelines; and on embeddings that discard the fielded structure (gene, protein, organism, chemical) of the underlying records. LIBRARIAN instead reuses Europe PMC's live, fielded search and orchestrates retrieval with a single LLM, avoiding rebuilding a custom index or learned retriever.

LLM-controlled lexical search interfaces. There is renewed interest in lexical scholarly search interfaces as increasingly capable LLMs make effective use of structured search tools. PI-SERINI shows that a BM25 (Best Matching 25) keyword-search backbone with retrieve/browse/read tools matches state-of-the-art deep-research agents (Hsu et al., 2026); SIRA constructs weighted BM25 queries from corpus-discriminative terms identified by the LLM (Yang et al., 2026); Li et al. (2026) advocates for composable corpus operations over fixed semantic interfaces; GRASP gives the agent both semantic and keyword search, and finds the learned policy reaches for keyword search specifically for entity-specific evidence (Gandhi et al., 2026); and SAGE finds that agents issuing keyword-oriented sub-queries benefit more from BM25 than from dense retrieval (Hu et al., 2026). In biomedical and systematic-review search, LLMs have also been used to generate Boolean queries and search strategies for expert-facing literature retrieval (Wang et al., 2023; Adam et al., 2024; Wang et al., 2025). Boerschinger et al. (2022) established the pre-LLM precedent for operator-level search control; on structured interfaces, Query Attribute Modeling and Multi-Meta-RAG use LLM-extracted metadata to filter retrieval before ranking (Menon et al., 2025; Poliakov & Shvai, 2025). We extend this line of work to scientific literature, exposing Europe PMC's lexical search interface to an LLM that orchestrates query formulation and evidence selection.

Tool interfaces for scientific agents. Scientific agents increasingly orchestrate external tools and knowledge sources, from general biomedical agents such as Biomni (Huang et al., 2026) to shared scientific-tool ecosystems such as ToolUniverse (Gao et al., 2025) and MCP-based interfaces such as MCPmed and BioContextAI (Flotho et al., 2026; Kuehl et al., 2025). Related efforts such as gget standardize programmatic access to biological databases (Luebbert & Pachter, 2023). These systems make existing resources callable, but they preserve each resource's native interface: the calling agent must still select the right tool, phrase the request in that tool's syntax, and process whatever the tool returns. LIBRARIAN instead upgrades the tool interface: it is specialized for Europe PMC's query grammar, and downstream agents inherit that specialization without paying for it themselves, calling the layer with a natural-language question.

5 DISCUSSION

LIBRARIAN is a plug-and-play knowledge layer for agents: it helps across all four benchmark suites we study. It improves literature synthesis, achieves the highest agreement in the ProClaim claim-verification pipeline, performs best on open-form question answering with GPT-5.4, and helps biology tasks that rely on the literature, such as sequence manipulation. Tasks that mainly look up structured databases, however, gain little from literature retrieval. Overall, our results suggest that literature retrieval can be a shared, off-the-shelf layer that agents reuse rather than build themselves. Our ablations (Appendix B) examine where this quality comes from by removing or replacing individual stages, and how it depends on the retrieval budgets (the number of subqueries and paragraphs per subquery). Because this judgment is applied at query time to a small candidate pool, LIBRARIAN can build on mature lexical infrastructure over the scientific literature rather than constructing and storing an index over the whole corpus, going beyond recent agentic-search studies which still build and host their own index and enrich every document offline with an LLM (Hsu et al., 2026; Hu et al., 2026). LIBRARIAN avoids the embedding and storage overhead of dense retrieval (about 728 GB for an OpenScholar-scale datastore; Appendix A.6) while remaining efficient at query time: subquery generation costs one lightweight LLM call (Appendix A.5), and Stage 3 plays the role of the re-ranker that dense pipelines also run after first-stage search. We release LIBRARIAN openly, including prompts, runtime configuration, and evaluation pipelines, so the community can reproduce our results and integrate them into their agentic pipelines.

Limitations. Our work has several limitations. First, LIBRARIAN's knowledge coverage is bounded by Europe PMC, so the full text of some paywalled papers is excluded. We plan to broaden it with other resources rich in scientific knowledge, such as OpenTargets (Buniello et al., 2025), UniProt (The UniProt Consortium, 2025), and ChEMBL (Zdrazil et al., 2024). Second, LIBRARIAN does not directly handle multi-hop requests, which are instead left to the downstream agent, because it performs a single retrieval round (albeit with multiple complementary subqueries to improve recall). As future work, we plan to let LIBRARIAN decide whether to continue searching over multiple rounds when the retrieved evidence is incomplete. Third, LIBRARIAN currently ignores figures, supplementary material, and other non-textual content, and handles tables only as text, leaving potentially relevant evidence inaccessible. Extending it to multimodal content is an important direction for future work.

6 CONCLUSION

We introduced LIBRARIAN, a scientific knowledge layer that gives AI agents a natural-language interface to an existing live, fielded search service and returns a compact set of citable evidence snippets. We instantiate it on Europe PMC, a widely popular open search engine over scientific literature. A single LLM orchestrates the process: it generates fielded lexical subqueries executed by the Europe PMC search engine, then localizes the relevant evidence within the retrieved full-text articles, yielding compact, citable evidence snippets, all without training a specialized model or maintaining a dense vector index. Equipping agents with LIBRARIAN improves their performance on literature synthesis, claim verification, open-form question answering, and selected biology workflow tasks. Even the strongest models we evaluate, such as GPT-5.4, benefit from a knowledge layer, suggesting that literature access is better decoupled from the model than built into ever-larger ones. We expect this separation to hold as base models improve: a single shared knowledge layer can give any agent the specialized evidence it needs to solve the task at hand with higher accuracy.

USE OF LARGE LANGUAGE MODELS

Large language models are the subject of this work as well as a tool used in producing it, so we separate the two.

LLMs as system components. LIBRARIAN is built around LLM calls, and every model used in the system and in the evaluations is named in the paper. GLM-5 performs subquery planning and relevance judging (Section 2, Appendix A); the answering models are GLM-5, Claude Sonnet 4.6, GPT-4o, and GPT-5.4 depending on the benchmark (Appendix A); GPT-4o serves as the LitQA2 judge (Appendix C.3); and Prometheus and AttrScore are used for ScholarQABench scoring (Section 3.1, Appendix C.1). All prompts are reproduced in Appendices D and E. These uses are the method under study, not writing assistance.

LLMs in the research process. We used generative AI tools in the following ways.

- Code: AI assistants were used to help write and debug portions of the retrieval pipeline, the evaluation harness, and plotting code. All code was reviewed by the authors, and every reported number was produced by code that at least one author read end to end.
- Writing: AI assistants were used to edit prose for clarity and concision. All text was reviewed, revised, and approved by the authors. We did not use generative AI to generate synthetic data, to develop the conceptual framework or system design, to formulate or prove mathematical claims, to design the experimental protocol, or to interpret experimental results; these were carried out by the authors.

We have reviewed all AI-assisted work and take full responsibility for the content of this paper, including all text, claims, code, and artifacts.

REPRODUCIBILITY STATEMENT

LIBRARIAN queries a live search service, so exact retrieval results depend on the state of the Europe PMC index at query time and cannot be reproduced verbatim at a later date. We address this by fully specifying both the system and its evaluation. The three retrieval stages are described in Section 2; the configuration used in all reported experiments is given in Table 5; and implementation details, including subquery validation, fallback behavior, XML handling, runtime configuration, and efficiency measurements, are provided in Appendix A. The prompts used by LIBRARIAN are documented in Appendix D, and all benchmark-side prompts, including evaluation-judge prompts, are provided in Appendix E; the complete prompts will be released with the accompanying code. Appendix C specifies the construction and evaluation protocol for each benchmark, including the 29-question ScholarQABench Multi subset and the 91-question Europe PMC full-text subset of LitQA2; identifiers for both subsets are released to permit their exact reconstruction. It also documents the retrieval source underlying each of our rows in Table 1, including how LIBRARIAN and the BM25 baseline query the fixed OSDS corpus (Appendix C.1), and the distinction between that corpus and live Europe PMC retrieval. We will release the source code, complete prompts, runtime configuration, benchmark subset identifiers, and evaluation pipelines needed to reproduce the experimental setup.

ETHICS STATEMENT

This work develops and evaluates a system for retrieving evidence from the scientific literature. Our experiments use existing scientific publications and established evaluation benchmarks and do not involve human-subject experiments or the collection of new personal or sensitive data. LIBRARIAN is intended to support access to scientific evidence rather than to replace expert scientific or clinical judgment. As with any automated literature retrieval system, its outputs may be incomplete or may omit relevant evidence, for example, because coverage is limited by the underlying Europe PMC corpus and the availability of retrievable full text. Downstream users and agents should therefore treat retrieved evidence as supporting material and verify scientifically consequential conclusions against the underlying literature. We do not identify additional ethical concerns specific to the experiments reported in this work, and all authors have reviewed and agree to adhere to the ICLR Code of Ethics.

REFERENCES

Gaelen P Adam, Jay DeYoung, Alice Paul, Ian J Saldanha, Ethan M Balk, Thomas A Trikalinos, and Byron C Wallace. Literature search sandbox: a large language model that generates search queries for systematic reviews. JAMIA Open, 7(3):ooae098, 07 2024. ISSN 2574-2531. doi: 10.1093/jamiaopen/ooae098. URL https://doi.org/10.1093/jamiaopen/ooae098.

Lun Ai, Yu-Heng Wu, Jonathan Schaul, Leonie Küchenhoff, Denes Turei, Francesco Carli, Fabio Petroni, Aurelien Dugourd, and Julio Saez-Rodriguez. Proclaim: Sufficiency-aware evidence programming for in-the-wild scientific claim verification, 2026. URL https://github.com/saezlab/ProClaim. Preprint submitted to NeurIPS.

Anthropic. Claude sonnet 4.6 system card, 2026. URL https://www-cdn.anthropic.com/bbd8ef16d70b7a1665f14f306ee88b53f686aa75/Claude%20Sonnet%204.6%20System%20Card.pdf.

Akari Asai, Jacqueline He, Rulin Shao, Weijia Shi, Amanpreet Singh, Joseph Chee Chang, Kyle Lo, Luca Soldaini et al. Synthesizing scientific literature with retrieval-augmented language models. Nature, 650(8103):857–863, Feb 2026. ISSN 1476-4687. doi: 10.1038/s41586-025-10072-4. URL https://doi.org/10.1038/s41586-025-10072-4.

Eser Aygün, Anastasiya Belyaeva, Gheorghe Comanici, Marc Coram, Hao Cui, Jake Garrison, Renee Johnston, Anton Kast et al. An ai system to help scientists write expert-level empirical software. Nature, 654(8120):909–916, Jun 2026. ISSN 1476-4687. doi: 10.1038/s41586-026-10658-6. URL https://doi.org/10.1038/s41586-026-10658-6.

Eranga Bandara, Ross Gore, Ravi Mukkamala, Asanga Gunaratna, Safdar H. Bouk, Xueping Liang, Peter Foytik, Abdul Rahman et al. Towards an agent-first web: Redesigning the web for ai agents, 2026. URL https://arxiv.org/abs/2606.19116.

Benjamin Boerschinger, Christian Carl Friedrich Buck, Lasse Jesper Garding Espeholt, Leonard Adolphs, Lierni Sestorain Saralegui, Massimiliano Ciaramita, Michelle Chen Huebscher, Pier Giuseppe Sessa et al. Boosting search engines with interactive agents. Transactions on Machine Learning Research, 2022.

Lutz Bornmann, Robin Haunschild, and Rüdiger Mutz. Growth rates of modern science: a latent piecewise growth curve approach to model publication numbers from established and new literature databases. Humanities and Social Sciences Communications, 8(1):224, Oct 2021. ISSN 2662-9992. doi: 10.1057/s41599-021-00903-w. URL https://doi.org/10.1057/s41599-021-00903-w.

Annalisa Buniello, Daniel Suveges, Carlos Cruz-Castillo, Jerven Bolleman, James Cook, Maya Ghoussaini, Andrew Hall, David Hercules et al. Open targets platform: facilitating therapeutic hypotheses building in drug discovery. Nucleic Acids Research, 53(D1):D1467–D1475, 2025.

Hang Ding, Yilun Zhao, Tiansheng Hu, Manasi Patwardhan, and Arman Cohan. Scirag: Adaptive, citation-aware, and outline-guided retrieval and synthesis for scientific literature. In Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 6440–6460, 2026.

European Commission. Support to the europe pmc initiative (2021–2026): Periodic reporting for period 3, 2026. URL https://cordis.europa.eu/project/id/101034194/reporting.

Christine Ferguson, Dayane Araújo, Lynne Faulk, Yuci Gou, Audrey Hamelers, Zhan Huang, Michele Ide-Smith, Maria Levchenko et al. Europe PMC in 2020. Nucleic Acids Research, 49(D1):D1507–D1514, 2021. doi: 10.1093/nar/gkaa994.

Matthias Flotho, Ian Ferenc Diks, Philipp Flotho, Leidy-Alejandra G Molano, Pascal Hirsch, and Andreas Keller. Mcpmed: a call for model context protocol-enabled bioinformatics web services for llm-driven discovery. Briefings in Bioinformatics, 27(1):bbag076, 01 2026. ISSN 1477-4054. doi: 10.1093/bib/bbag076. URL https://doi.org/10.1093/bib/bbag076.

Carlo Galli, Chiara Moretti, and Elena Calciolari. Intelligent summaries: Will artificial intelligence mark the finale for biomedical literature reviews? Learned Publishing, 38(1):e1648, 2025. doi: https://doi.org/10.1002/leap.1648. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/leap.1648.

Varun Gandhi, Jaewook Lee, Shantanu Todmal, Franck Dernoncourt, Ryan Rossi, Zichao Wang, and Andrew Lan. Grasp: Granularity-aware search policy for agentic rag. arXiv preprint arXiv:2607.10463, 2026.

Shanghua Gao, Richard Zhu, Pengwei Sui, Zhenglun Kong, Sufian Aldogom, Yepeng Huang, Ayush Noori, Reza Shamji et al. Democratizing ai scientists using tooluniverse, 2025. URL https://arxiv.org/abs/2509.23426.

Ali E. Ghareeb, Benjamin Chang, Ludovico Mitchener, Angela Yiu, Caralyn J. Szostkiewicz, Dmytro Shved, Gavin J. Gyimesi, Jon M. Laurent et al. A multi-agent system for automating scientific discovery. Nature, 655(8122):497–505, Jul 2026. ISSN 1476-4687. doi: 10.1038/s41586-026-10652-y. URL https://doi.org/10.1038/s41586-026-10652-y.

GLM-5-Team, Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, Da Yin et al. Glm-5: from vibe coding to agentic engineering, 2026. URL https://arxiv.org/abs/2602.15763.

Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Petar Sirkovic, Artiom Myaskovsky, Grzegorz Glowaty, Felix Weissenberger et al. Accelerating scientific discovery with co-scientist. Nature, 655(8122):487–496, May 2026. ISSN 1476-4687. doi: 10.1038/s41586-026-10644-y. URL http://dx.doi.org/10.1038/s41586-026-10644-y.

Michael Gusenbauer. The age of abundant scholarly information and its synthesis–a time when 'just google it'is no longer enough. Research synthesis methods, 12(6):684–691, 2021.

Michael Gusenbauer. Search where you will find most: Comparing the disciplinary coverage of 56 bibliographic databases. Scientometrics, 127(5):2683–2745, May 2022. ISSN 1588-2861. doi: 10.1007/s11192-022-04289-7. URL https://doi.org/10.1007/s11192-022-04289-7.

Tz-Huan Hsu, Jheng-Hong Yang, and Jimmy Lin. Rethinking agentic search with pi-serini: Is lexical retrieval sufficient?, 2026. URL https://arxiv.org/abs/2605.10848.

Tiansheng Hu, Yilun Zhao, Canyu Zhang, Arman Cohan, and Chen Zhao. Sage: Benchmarking and improving retrieval for deep research agents, 2026. URL https://arxiv.org/abs/2602.05975.

Kexin Huang, Serena Zhang, Hanchen Wang, Yuanhao Qu, Yingzhou Lu, Ryan Li, Yusuf Roohani, Lin Qiu et al. Autonomous biomedical research with an artificial intelligence agent. Science, 0(0):eadz4351, 2026. doi: 10.1126/science.adz4351. URL https://www.science.org/doi/abs/10.1126/science.adz4351.

HUMAN Security. The 2026 state of ai traffic & cyberthreat benchmark report, 2026. URL https://www.humansecurity.com/wp-content/uploads/HUMAN_Report_2026-State-of-AI-Traffic-and-Cyberthreat-Benchmark.pdf. Accessed: 2026-07-13.

Imperva. 2025 bad bot report. Technical report, Imperva, 2025. URL https://www.imperva.com/resources/wp-content/uploads/sites/6/reports/2025-Bad-Bot-Report.pdf. Accessed: 2026-07-13.

Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP), pp. 6769–6781, 2020.

Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoo Yun, Seongjin Shin et al. Prometheus: Inducing fine-grained evaluation capability in language models. In The Twelfth International Conference on Learning Representations, 2023.

Irma Klerings, Alexandra S Weinhandl, and Kylie J Thaler. Information overload in healthcare: too much of a good thing? Zeitschrift für Evidenz, Fortbildung und Qualität im Gesundheitswesen, 109(4-5):285–290, 2015.

Malte Kuehl, Darius P. Schaub, Francesco Carli, Lukas Heumos, Malte Hellmig, Camila Fernández-Zapata, Nico Kaiser, Jonathan Schaul et al. Biocontextai is a community hub for agentic biomedical systems. Nature Biotechnology, November 2025. doi: 10.1038/s41587-025-02900-9. URL http://dx.doi.org/10.1038/s41587-025-02900-9.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles, 2023.

Jon M Laurent, Joseph D Janizek, Michael Ruzo, Michaela M Hinks, Michael J Hammerling, Siddharth Narayanan, Manvitha Ponnapati, Andrew D White, and Samuel G Rodriques. Lab-bench: Measuring capabilities of language models for biology research. arXiv preprint arXiv:2407.10362, 2024.

Maria Levchenko, Yuci Gou, Florian Graef, Audrey Hamelers, Zhan Huang, Michele Ide-Smith, Anusha Iyer, Oliver Kilian et al. Europe pmc in 2017. Nucleic acids research, 46(D1):D1254–D1260, 2018.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), Advances in Neural Information Processing Systems, volume 33, pp. 9459–9474. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf.

Zhuofeng Li, Haoxiang Zhang, Cong Wei, Pan Lu, Ping Nie, Yi Lu, Yuyang Bai, Shangbin Feng et al. Beyond semantic similarity: Rethinking retrieval for agentic search via direct corpus interaction. arXiv preprint arXiv:2605.05242, 2026.

Peiwen Liu, Sakura Eri B Maezono, Weitao Lin, Yen Yeow, Ya-Yu Liu, Yasmin Hisham, Rui Hou, Jordan A Ramilowski, and Alistair RR Forrest. connectomedb2025: a rigorously curated, multi-species resource of experimentally supported ligand–receptor interactions. Nucleic Acids Research, 54(D1):D546–D554, 2026.

Prisca Lo Surdo, Marta Iannuccelli, Klas Karis, Eleonora Meo, Parnian Omidi, Marta Tosoni, Simone Graziosi, Simona Panni et al. Signor 4.0: the 2025 update with focus on phosphorylation data. Nucleic Acids Research, 54(D1):D682–D690, 2026.

Laura Luebbert and Lior Pachter. Efficient querying of genomic reference databases with gget. Bioinformatics, 39(1):btac836, 01 2023. ISSN 1367-4811. doi: 10.1093/bioinformatics/btac836. URL https://doi.org/10.1093/bioinformatics/btac836.

Xing Han Lù. Bm25s: Orders of magnitude faster lexical search via eager sparse scoring, 2024. URL https://arxiv.org/abs/2407.03618.

Xueguang Ma, Xinyu Zhang, Ronak Pradeep, and Jimmy Lin. Zero-shot listwise document reranking with a large language model. arXiv preprint arXiv:2305.02156, 2023.

Karthik Menon, Batool Arhamna Haider, Muhammad Arham, Kanwal Mehreen, Ram Mohan Rao Kadiyala, Muhammad Ali Shafique, and Hamza Farooq. Query attribute modeling: Improving search relevance with semantic search and meta data filtering. In 2025 IEEE International Conference on Data Mining Workshops (ICDMW), pp. 1113–1117, 2025. doi: 10.1109/ICDMW69685.2025.00131.

Ludovico Mitchener, Angela Yiu, Benjamin Chang, Mathieu Bourdenx, Tyler Nadolski, Arvis Sulovari, Eric C. Landsness, Daniel L. Barabasi et al. Kosmos: An ai scientist for autonomous discovery, 2025. URL https://arxiv.org/abs/2511.02824.

Ferdous Nasri, Sarah Gurev, Patrick Varilly, Krithik Ramesh, Nuala A. O'Leary, Jonah Cool, Bernhard Y. Renard, Pardis C. Sabeti, and Laura Luebbert. Deterministic access to global viral sequence data enables robust agentic scientific discovery, 2026. URL https://arxiv.org/abs/2606.06749.

OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow et al. Gpt-4o system card, 2024. URL https://arxiv.org/abs/2410.21276.

Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. How context affects language models' factual predictions. In Automated Knowledge Base Construction, 2020.

Mykhailo Poliakov and Nadiya Shvai. Multi-meta-rag: Improving rag for multi-hop queries using database filtering with llm-extracted metadata. In Vadim Ermolayev, Igor Potapov, Oleksii Ignatenko, Roman Hornung, Andrii Hlybovets, Vitaliy Yakovyna, Yaroslav Prytula, and Oleksandr Spivakovsky (eds.), Information and Communication Technologies in Education, Research, and Industrial Applications, pp. 334–342, Cham, 2025. Springer Nature Switzerland. ISBN 978-3-031-81372-6.

Stephen Edward Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. Okapi at trec. 1994.

Summer Rosonovski, Maria Levchenko, Rajat Bhatnagar, Umamageswari Chandrasekaran, Lynne Faulk, Islam Hassan, Matt Jeffryes, Syed Irtaza Mubashar et al. Europe PMC in 2023. Nucleic Acids Research, 52(D1):D1668–D1676, 2024. doi: 10.1093/nar/gkad1085.

Nipun Sadvilkar and Mark Neumann. PySBD: Pragmatic sentence boundary disambiguation. In Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS), pp. 110–114, Online, November 2020. Association for Computational Linguistics. URL https://www.aclweb.org/anthology/2020.nlposs-1.15.

Xiaofeng Shi, Qian Kou, Yuduo Li, Ning Tang, Jinxin Xie, Longbin Yu, Songjing Wang, and Hua Zhou. Scisage: A multi-agent framework for high-quality scientific survey generation. arXiv preprint arXiv:2506.12689, 2025.

Michael D Skarlinski, Sam Cox, Jon M Laurent, James D Braza, Michaela Hinks, Michael J Hammerling, Manvitha Ponnapati, Samuel G Rodriques, and Andrew D White. Language agents achieve superhuman synthesis of scientific knowledge. arXiv preprint arXiv:2409.13740, 2024.

Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. Is chatgpt good at search? investigating large language models as re-ranking agents. In Proceedings of the 2023 conference on empirical methods in natural language processing, pp. 14918–14937, 2023.

Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. arXiv preprint arXiv:2104.08663, 2021.

The UniProt Consortium. Uniprot: the universal protein knowledgebase in 2025. Nucleic Acids Research, 53(D1):D609–D617, 2025.

Svitlana Volkova, Peter Bautista, Avinash Hiriyanna, Gabriel Ganberg, Isabel Erickson, Zachary Klinefelter, Nick Abele, Hsien-Te Kao, and Grant Engberson. Cross-disciplinary knowledge retrieval and synthesis: A compound ai architecture for scientific discovery, 2025. URL https://arxiv.org/abs/2511.18298.

Shuai Wang, Harrisen Scells, Bevan Koopman, and Guido Zuccon. Can chatgpt write a good boolean query for systematic review literature search? In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23, pp. 1426–1436, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394086. doi: 10.1145/3539618.3591703. URL https://doi.org/10.1145/3539618.3591703.

Shuai Wang, Harrisen Scells, Bevan Koopman, and Guido Zuccon. Reassessing large language model boolean query generation for systematic reviews. In Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '25, pp. 3296–3305, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400715921. doi: 10.1145/3726302.3730329. URL https://doi.org/10.1145/3726302.3730329.

Yingxuan Yang, Mulei Ma, Yuxuan Huang, Huacan Chai, Chenyu Gong, Haoran Geng, Yuanjian Zhou, Ying Wen et al. Agentic web: Weaving the next web with ai agents, 2025. URL https://arxiv.org/abs/2507.21206.

Zeyu Yang, Qi Ma, Jason Chen, and Anshumali Shrivastava. Superintelligent retrieval agent: The next frontier of information retrieval. arXiv preprint arXiv:2605.06647, 2026.

Xiang Yue, Boshi Wang, Kai Zhang, Ziru Chen, Yu Su, and Huan Sun. Automatic evaluation of attribution by large language models. arXiv preprint arXiv:2305.06311, 2023.

Barbara Zdrazil, Eloy Felix, Fiona Hunter, Emma J Manners, James Blackshaw, Sybilla Corbett, Marleen de Veij, Harris Ioannidis et al. The chembl database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. Nucleic Acids Research, 52(D1): D1180–D1192, 2024.

Xinyan Zhong, Yuwei Shi, Yuqi Wei, Chen Shen, Tianhang Zhou, and Zhenghao Wu. Vectree-rag: An agentic retrieval-augmented generation framework combining vector and tree retrieval for efficiency and accuracy. arXiv preprint arXiv:2607.23006, 2026.

A Implementation Details

Every Librarian component (subquery planning, relevance filtering) and every synthesis step runs on the `glm-5-fp8` deployment of GLM-5, served through a vLLM-compatible OpenAI API endpoint. Models outside the retrieval stack are not GLM-5: each benchmark keeps the backbone its originating pipeline used, so that our numbers are comparable to the published ones and any difference is attributable to retrieval rather than to a stronger or weaker reader. Concretely, this gives GLM-5 for ScholarQABench, Claude Sonnet 4.6 for ProClaim verdicts, GPT-5.4 for LitQA2 answering with GPT-4o as the judge, and GPT-4o and GPT-5.4 for LAB-Bench.

A.1 Subquery Validation and Fallbacks

Every planned subquery is parsed before execution. Invalid field names and malformed Boolean expressions are rejected, and only the surviving subqueries are sent to Europe PMC. Three fallbacks guard the degenerate cases, and all are deliberately conservative:

- **Planner output unparseable.** If the planner's JSON cannot be parsed, it is retried once at temperature 0 with a strict-JSON instruction. If that also fails, the raw user query is used as the single search query.
- **Every subquery rejected by validation.** The planned subqueries are searched unmodified, so a validator false positive cannot empty the search plan. The raw user query is not substituted in this case.
- **No paragraph scores above zero for a subquery.** That subquery's leading k paragraphs are returned in retrieval order rather than an empty set, so a vocabulary mismatch between subquery and paragraph text cannot silently drop the subquery's entire contribution.

Papers are deliberately not deduplicated across subqueries: a paper retrieved by several subqueries competes in each one's BM25 ranking, and repeated full-text fetches are absorbed by a PMCID-keyed cache rather than by a dedup pass. Deduplication instead happens at the paragraph level, where the per-subquery selections are concatenated and duplicate or overlapping word intervals from the same source paragraph are merged, so a paragraph relevant to several subqueries costs the judge's budget once.

Table 5: LIBRARIAN retrieval configuration, identical in all reported experiments; the appendix experiments (Tables 6,7,8) use instead GLM-5.3-Flash.

Parameter	Value
Models	
Planner / judge / synthesis model	GLM-5
Stage 1, subqueries	
Subquery budget	8
Records fetched per EPMC query	50
Stage 2, paragraph ranking	
Ranking target	the subquery (syntax stripped)
Paragraphs kept per subquery	16
Paragraph pool judged	up to 128 (16×8)
Max paragraph length / overlap	250 / 50 words
Stage-2 worker threads	8 (one per subquery)
Stage 3, relevance judging	
Paragraphs per call	128 (whole pool in one call)
Returned evidence	every cited span, no word or top-k cap
Decoding	
Subquery-planning temperature	0.3 (retry: 0.0)
Relevance-judge temperature	0.1
Synthesis temperature	0.2

A.2 XML HANDLING

For records with fetchable Europe PMC full text, LIBRARIAN retrieves the JATS (Journal Article Tag Suite) XML once and caches it by PMCID. The parser keeps body paragraphs and drops front and back matter, the abstract, reference lists, appendices, author biographies, footnote groups, and sections whose `sec-type` marks them as acknowledgments, funding, competing interests, conflicts of interest, data availability, ethics, author contributions, abbreviations, references, or supplementary material. Each retained paragraph carries its section title and section type, and both contribute to the BM25 document alongside the paragraph text. Paragraph localization ranks these records with BM25 against the subquery that retrieved them (Section 2). Retained paragraphs longer than 250 words are split into overlapping windows of 250 words with 50 words of overlap; shorter paragraphs are kept whole. Because windows overlap, the same sentence can appear in two chunks, which the Stage-3 judge resolves by preferring the complete version (Appendix D.2).

A.3 RUNTIME CONFIGURATION

Table 5 lists the LIBRARIAN retrieval settings, which are identical across all reported experiments. Stage 2 runs one thread per subquery, capped at the number of CPUs available to the process, so on our deployment, the eight subqueries fan out fully in parallel; each thread performs its own search, full-text fetching, and BM25 ranking. The judge batch size is set to paragraphs per subquery $\times$ subqueries so that the entire paragraph pool is judged in a single LLM call. Chunks are presented to the judge in subquery order and, within a subquery, in BM25 rank order. Like listwise LLM rerankers (Sun et al., 2023; Ma et al., 2023), the judge returns a global ranking, here over sentence identifiers rather than documents and without a truncated top-k, so a chunk's position in the prompt does not cap its chance of being selected.

A.4 BM25 CONFIGURATION

Paragraph ranking uses bm25s (Lù, 2024) with method="lucene", $k_1 = 1.5$, $b = 0.75$. The Lucene IDF, $\ln(1 + (N - df(t) + 0.5)/(df(t) + 0.5))$, is strictly positive for every $df(t) \leq N$, so no epsilon floor is required and none is applied. This matters because each subquery's corpus consists of paragraphs from papers that subquery already matched: a query term that occurs in every retrieved

Table 6: **Efficiency of LIBRARIAN on ScholarQA-Bench.** Token counts and costs are measured over all n=2,775 queries. Costs are computed from measured token counts at the served GLM-5.3-Flash rate ($0.15 / $0.50 per million input / output tokens).

Stage	Context *median tok.*	Cost ($) *median / p90*	Latency (s) *median / p90*
Stage 1: query planning	4.0K	0.0007 / 0.0007	1.2 / 2.2
Stage 2: literature access + ranking	–	0	9.3 / 18.2
Stage 3: evidence selection	**45.4K**	**0.0077 / 0.0098**	10.9 / 26.2
LIBRARIAN (Stages 1–3)	**49.6K**	**0.0084 / 0.0106**	**23.2 / 40.7**

paper still occurs in only a fraction of their paragraphs, so its IDF stays informative, and even a term present in every paragraph keeps a small positive weight, $\ln(1 + 0.5/(N + 0.5))$, rather than the negative Robertson IDF. Query and corpus are tokenized identically (Snowball stemming, English stopword removal). A paragraph, therefore, scores exactly zero only when it shares no stemmed, non-stopword token with the subquery, and the fallback of Appendix A.1 is a lexical-mismatch guard rather than an IDF artifact. In our ablation runs, this fallback fires for at most 5.6% of (question, subquery) pairs on LitQA2 and 0.5% on ScholarQABench Bio.

A.5 EFFICIENCY: LATENCY, TOKENS, AND COST PER QUERY

Every stage of LIBRARIAN emits an OpenTelemetry span via Arize Phoenix[1], and every model call records its input and output token counts. Table 6 reports per-query measurements from 2,775 error-free ScholarQABench requests (Bio, Neu, and Multi, including reruns), served with GLM-5.3-Flash.

Latency. A query completes in 23.2 s at the median (p90: 40.7 s), split as 1.2 s for Stage 1, 9.3 s for Stage 2, and 10.9 s for Stage 3. Stage 2 is almost entirely external I/O: the median full-text-fetch span alone is 7.1 s of its 9.3 s, with paragraph decomposition and BM25 ranking accounting for the rest. Latency is therefore dominated by external literature retrieval, not local lexical ranking.

Token usage. Stage 2 makes no model calls. Query planning consumes a median of 4.0K input tokens and the Stage 3 judge 45.4K, for 49.6K per query: 91% of the model context is spent scoring candidate evidence.

Cost. At the served GLM-5.3-Flash rate ($0.15 and $0.50 per million input and output tokens), a query costs $0.0084 at the median (p90: $0.0106), of which Stage 3 is $0.0077 and planning $0.0007, roughly $8.4 per thousand queries.

A.6 ILLUSTRATIVE DENSE-INDEX INFRASTRUCTURE COMPARISON

To characterize the infrastructure trade-off between LIBRARIAN and maintaining a dense literature index, we estimate the cost of serving a datastore at the scale of OpenScholar. This comparison is illustrative rather than a benchmark: actual costs depend on representation precision, index implementation, hardware utilization, compression, and cloud-provider pricing.

OpenScholar contains approximately 237M passages. Representing each passage with a 768-dimensional `fp32` embedding requires approximately

$$237\text{M} \times 768 \times 4 \text{ bytes} \approx 728 \text{ GB}.$$

A flat in-memory deployment, therefore, requires a memory-optimized host of roughly this capacity. At September 2026 AWS us-east-1 pricing, an `r7i.24xlarge` (768 GiB) corresponds to approximately $37K/year under a one-year reservation or $56K/year on demand, before query-dependent computation. By comparison, LIBRARIAN maintains no corpus-wide embedding index and has a measured median LLM cost of $0.0084 per query (Appendix A.5). Under these assumptions, the annual serving cost corresponds to approximately 4–7M LIBRARIAN queries. Using `fp16`,

[1] `https://arize.com/phoenix/`

Table 7: **Component ablations under a controlled model configuration.** All conditions use GLM-5.3-Flash for both LIBRARIAN and the downstream task agent. LitQA2 reports mean accuracy over three seeds on the 91-question full-text subset, with bootstrap 95% confidence intervals obtained by resampling questions. ScholarQA reports Citation F1 on a single-seed, 200-question *Bio* sample. Absolute scores are not directly comparable to the main-paper results. For evidence selection, BM25 and the LLM filter operate on the same 128-paragraph candidate pool and 2,500-character evidence budget.

Configuration	LitQA2 Acc. ↑	ScholarQA F1 ↑
Full LIBRARIAN	**68.5 [60.1, 76.6]**	**83.5**
BM25 evidence selection	38.1 [29.3, 46.9]	68.3
Abstracts only	18.3 [11.4, 26.4]	80.7
Unfielded multi-query retrieval	36.6 [28.2, 45.4]	–
+ fielded queries	39.9 [31.9, 48.0]	–

compression, approximate nearest-neighbor indexing, or alternative hardware can substantially alter this crossover.

This calculation excludes costs shared or difficult to compare directly, including downstream answer generation. Dense systems may additionally incur query-time reranking costs and periodic re-embedding as the literature changes; conversely, production dense indexes can employ quantization and optimized ANN structures that reduce their memory footprint substantially.

B MECHANISM ABLATIONS

These experiments test which components of LIBRARIAN carry its retrieval quality and how quality scales with retrieval budget. Both LIBRARIAN and the downstream agent run on GLM-5.3-Flash, whereas the main experiments use the benchmark-specific configurations of Section 3; absolute scores are therefore not comparable across the two, and only differences within each table are meaningful. Because the corpus is fixed to Europe PMC, these results isolate the retrieval policy from the corpus change in Table 1; since they use a different model and sample, they support rather than decompose the main-table gains.

Protocol. Unless stated otherwise we evaluate on the open-form LitQA2 full-text subset of Section 3.3, holding the answering model and evaluation protocol fixed within each comparison and varying one retrieval component or budget at a time. Component ablations run three seeds; 95% confidence intervals come from bootstrapping over questions, keeping each question's three seed outcomes together rather than treating seeds as independent draws. The evidence-selection and full-text ablations are repeated on a single-seed 200-question ScholarQABench *Bio* sample, scored by aggregate Citation F1. The budget experiments use a single seed.

Semantic evidence selection. Replacing the Stage 3 LLM filter with BM25 over the same 128-paragraph pool and 2,500-character budget drops LitQA2 accuracy from 68.5 to 38.1 and ScholarQA Citation F1 from 83.5 to 68.3 (Table 7). This is the largest component effect we measure, and it holds for both exact-answer retrieval and multi-document synthesis.

Full-text access. Restricting retrieval to abstracts drops LitQA2 accuracy to 18.3, which follows from the benchmark design: LitQA2 selects questions whose answers lie beyond the abstract. The same restriction costs only 2.8 F1 on ScholarQABench (83.5 to 80.7). Citation F1 scores whether each cited passage supports the sentence that cites it (App. C.1), not whether the answer is correct or complete. One possible reason for the small drop is that the synthesis agent adapts its claims to the evidence it receives; we do not test this directly. In this setting, Citation F1 therefore reflects consistency between claims and cited evidence more than answer quality.

Fielded queries. Letting the LLM use Europe PMC's fielded query language moves LitQA2 accuracy from 36.6 to 39.9, with heavily overlapping intervals at $n = 91$: on this benchmark, we

Table 8: **Retrieval-budget sensitivity on open-form LitQA2** (n=91, single seed). We vary one parameter at a time around the configuration used in the ablation experiments (N=8, k=16). The N=8 and k=16 cells are the same run: seed 0 of the full LIBRARIAN configuration in Table 7 (65.9; the three-seed mean reported there is 68.5). N is the maximum number of subqueries and k is the number of paragraphs retained per subquery before Stage 3. Input tokens report median per-query input consumption. These results characterize cost–quality trends rather than statistical differences between adjacent settings.

Subquery budget			**Paragraph budget**		
N	**Acc.** ↑	**Input tok.** ↓	k	**Acc.** ↑	**Input tok.** ↓
1	41.8	13k	4	59.3	17k
2	52.7	19k	8	61.5	28k
4	60.4	29k	**16**	**65.9**	**48k**
8	**65.9**	**48k**	32	70.3	89k

measure no reliable benefit from field constraints. LitQA2, however, is the setting least able to show one. Each question can be answered from a single known paper, so retrieval succeeds or fails by locating a single document rather than by constraining a large candidate set.

Retrieval budgets. Accuracy increases monotonically with the subquery budget, 41.8, 52.7, 60.4, 65.9 for N=1, 2, 4, 8 (Table 8), so we deploy N=8, the largest setting tested. Accuracy also rises with the paragraph budget, 59.3, 61.5, 65.9, 70.3 for k=4, 8, 16, 32, but cost grows faster: raising k from 16 to 32 adds 4.4 points (4 of 91 questions) while nearly doubling median input tokens (48k to 89k). We therefore deploy k=16 as a cost–quality compromise rather than a saturation point. With a single seed and n=91, no adjacent difference in either sweep is statistically distinguishable (bootstrap 95% intervals span roughly ±9 points). The configuration was fixed a priori; this sweep characterizes it post hoc with a different model.

C BENCHMARK DETAILS

We used the following dataset choices for the reported numbers.

C.1 SCHOLARQABENCH

Splits. ScholarQABench (Asai et al., 2026) spans several evaluation splits. We report the biomedicine (Bio) and neuroscience (Neu) multi-document synthesis splits, as they align with Europe PMC coverage. We omit the single-document splits because they test extraction from a provided document rather than open-domain literature synthesis. The multidisciplinary (Multi) column is not the full split: many of its questions fall outside the life sciences, so we report the 29-question biomedicine subset, selected from the split's gold reference file by subject label (biomedicine, bioimaging, genetics, biophysics). The 29 question identifiers are released so the subset can be reconstructed exactly.

Retrieval sources. ScholarQABench is the only benchmark where we report two retrieval sources. Europe PMC is the deployed source used in all other benchmarks, because it is live, and exact results depend on the index state at query time. OSDS is the OpenScholar Data Store used by the published baselines. We do not query it through OpenScholar's own retriever: we re-index the datastore's abstracts and full-text chunks into a local Elasticsearch instance and retrieve from it, so the corpus is held fixed while the retrieval mechanism is ours. The BM25-on-OSDS row is a fixed-corpus BM25 arm: passages are retrieved from the OSDS full-text chunk index by BM25 alone, with no subquery planning and no relevance filter, and are then synthesized by the same GLM-5 synthesis prompt. It holds the corpus fixed relative to the published OSDS baselines; its difference from the Europe PMC row, however, reflects corpus and policy together (see Sec. 3.1 and App. B).

The OSDS re-index. The OSDS rows retrieve from an Elasticsearch index built from the chunked body text of OpenScholar-DataStore-V3 (`OpenSciLM/OpenScholar-DataStore-V3`; about 45M papers and 149M passages). Each indexed document is one datastore passage, with fields for the

paper identifier, passage index, text, section title, and source, enriched with the paper's title, authors, year, journal, DOI, and fields of study joined from the datastore's abstract records.

LIBRARIAN on OSDS. This row runs LIBRARIAN with the configuration of Table 5 (GLM-5, `glm-5-fp8`, as planner, relevance judge, and synthesis model), replacing Europe PMC with the OSDS re-index as the search backend; the corpus therefore matches the BM25-on-OSDS row and the published OSDS baselines. Two adaptations follow from the backend. First, Stage 1 uses a planner prompt written for Elasticsearch (App. E.2): plain terms, quoted phrases, and AND/OR/NOT, with Europe PMC field operators such as `MESH:` or `GENE_PROTEIN:` disallowed. Each subquery runs as an Elasticsearch `simple_query_string` with default operator AND over the text, section-title, title, abstract, author, journal, DOI, and year fields (title and section title boosted by 1.4 and 1.2). Because `simple_query_string` has no field-scoping syntax, the field prefixes the prompt permits (e.g., `section_title:`, `year:`) are matched as ordinary tokens, so the effective grammar is terms, phrases, and Boolean operators. A subquery with no hits is retried as an OR over its quoted phrases. Second, since OSDS stores passages rather than articles, hits are collapsed by paper, keeping up to 50 papers per subquery and each paper's three highest-scoring passages, with abstract-like passages deprioritized; a paper's abstract is used only when none of its body passages matched.

Evaluation protocol. We report Citation F1 on 1451 Bio questions, 1308 Neu questions, and 29 Multi questions. Citation support is scored with the AttrScore AutoAIS implementation (`osunlp/attrscore-flan-t5-xl`) (Yue et al., 2023); Citation F1 is the harmonic mean of corpus-level citation precision and recall, expressed as a percentage.

C.2 PROCLAIM-EVAL

Adapting ProClaim to LIBRARIAN requires two modifications. In stage (ii), fact extraction no longer mines full-text papers but directly labels the evidence LIBRARIAN returns. In stage (iii), ProClaim's MLP sufficiency classifier relies on features (impact factor, citation counts, full-text NLI signals) that LIBRARIAN does not provide. We therefore use the LLM-based classifier that ProClaim natively offers as an alternative option. The ProClaim result with its original PubMed and Semantic Scholar retriever is reported from Ai et al. (2026).

C.3 LITQA2

Subset construction. We evaluate on the Europe PMC full-text subset of LitQA2. Starting from the 199 original questions, we keep the 91 questions whose supporting paper can be fetched as full text through Europe PMC. We verify availability by matching each source DOI with the `HAS_FT:Y` constraint and confirming that the JATS XML endpoint actually returns the article. Requiring successful retrieval excludes 50 of the 141 questions that satisfy `HAS_FT:Y`, as the corresponding articles are not retrievable through the JATS XML endpoint (e.g., because only PDF full text is available or the XML is not hosted by Europe PMC).

Why this subset. LitQA2 questions require access to the answer-bearing evidence in the supporting paper. Because LIBRARIAN retrieves evidence from Europe PMC, evaluating it on papers unavailable through that corpus would conflate retrieval quality with corpus coverage. We therefore restrict the evaluation to the 91 questions for which the supporting article is actually retrievable. This restriction does not constrain the web-search arm to Europe PMC: it may retrieve evidence from any source available on the web. Thus, the filter establishes that the required evidence is present in LIBRARIAN's corpus without imposing the same corpus boundary on the comparator. Unless stated otherwise, all systems are evaluated on the same 91 questions, and we release their identifiers to make the subset reproducible. The Web-Search arm uses OpenAI's hosted `web_search` tool (Responses API, medium context size, no domain filtering) which searches on its own, so it receives only the question; the LIBRARIAN arm also receives the planner hint and answer prompt of Appendices E.4.1 and E.4.2.

Judge. Open-form answers are evaluated with a fixed LLM judge using `openai/gpt-4o-2024-11-20` at temperature 0 with a constrained JSON schema. The

judge receives the question, the reference answer, the system's open-form answer, and, when the benchmark provides one, the gold key passage from the target paper, which is used as the authoritative statement of what a correct answer looks like in prose. It returns two flags: whether the answer contains the correct answer and whether it commits to any answer at all. The first yields accuracy, which is the metric we report; the second separates abstentions from errors and drives the precision and coverage figures.

C.4 LAB-BENCH

We evaluate the DBQA, SEQQA, PROTOCOLQA, and CLONINGSCENARIOS tasks from LAB-Bench (Laurent et al., 2024), covering database lookup, sequence manipulation, protocol troubleshooting, and molecular cloning. We evaluate all questions available in the public subset in each task (DbQA: 520, SeqQA: 600, ProtocolQA: 108, CloningScenarios: 33) and report accuracy, precision, and coverage following the LAB-Bench protocol: each question's options are the gold answer, the distractors, and an "Insufficient information" refusal option, shuffled with a fixed seed so the baseline and retrieval-augmented runs see identical layouts, and the answering model is queried with the LAB-Bench multiple-choice prompt.

Retrieval-augmented rows. Rows marked with LIBRARIAN keep the answering model fixed (GPT-4o or GPT-5.4, temperature 0) and add a single retrieval step: the retrieval agent is run once on the question, its raw passages are injected into the prompt under a 1500-character budget, and the model answers in one call. There is no synthesis pass and no second retrieval round. For SEQQA and CLONINGSCENARIOS, raw nucleotide and amino-acid sequences are stripped from the retrieval query, since verbatim sequences act as noise in a keyword index; the answering model still sees the full question.

D PROMPT TEMPLATES

The templates below are filled at runtime with the user query, date, candidate papers, and output-channel settings. The relevance-filter prompt is reproduced verbatim. The planner and synthesis prompts are abridged: the planner's full Europe PMC field catalog, its query-diversity strategies, and its worked examples are omitted, as are the synthesis prompt's citation examples and channel-formatting details, and both are reproduced here as the behavioral instructions that govern their output. The subquery planner prompt predates this section's terminology and refers to subqueries as "queries"; `{query_budget_guidance}` is filled from the configured subquery cap, so the cap and the instruction always agree. We release the full prompt in the repository.

D.1 STAGE 1 (SUBQUERY PLANNER) PROMPT

```
You are a biology research librarian expert in Europe PMC search syntax, tasked
with formulating MULTIPLE diverse search queries to achieve MAXIMUM RECALL.

TEMPORAL CONTEXT:
- Today's date is {today_date}.
- The current year is {today_year}.
- If the user asks for a relative time window such as "last two years",
  "past 5 years", or "recent", anchor that request to today's date and convert
  it into explicit publication years using PUB_YEAR constraints.
- Always include the current partial unit.

CORE OBJECTIVE: RECALL OVER PRECISION
Your job is to find every relevant paper. A downstream relevance filter handles
precision. The single biggest recall failure in Boolean search is
over-constraining queries with too many AND operators.

QUERY BUDGET:
{query_budget_guidance}
When a strict query budget is given, output queries in priority order. The first
queries must be the best standalone Europe PMC searches under that budget.

MANDATORY QUERY STRUCTURE RULES:
1. At least 30-40% of queries must be plain-text queries with no field
   specifiers. Plain-text queries hit title, abstract, full text, and metadata.
2. No query may contain more than two AND operators.
3. Use OR inside synonym sets and AND only between truly co-required concepts.
4. Before emitting a query, ask whether it would return at least 50 papers.
```

```
   If not, remove a field specifier, replace an AND with OR, or split the query.

QUERY BUDGET ALLOCATION:
- Tier 1: Broad recall, plain-text only, about 40% of the budget.
- Tier 2: Synonym-expanded field queries, about 40% of the budget.
- Tier 3: High-precision structured queries, about 20% of the budget.

[The full prompt continues with a Europe PMC field catalog, query-diversity
strategies, and worked examples, omitted here.]

WHAT NOT TO DO:
- Do not chain 3+ AND operators in one query.
- Do not use AUTH_FIRST: or AUTH_LAST:; they are not valid search fields.
- Do not use DATA_AVAILABILITY:; use BODY:"data availability" or HAS_DATA:y.
- Do not rely solely on MESH:; prefer KW: for MeSH and publisher keywords.
- Do not make every query require a field specifier.
- Do not generate meta-queries about "searching for" or "in the literature".
- Do not let complexity per query substitute for diversity across queries.

CONVERSATION:
<conversation>
{conversation}
</conversation>

<additional_context>
{additional_context}
</additional_context>

Respond with a JSON object containing a "queries" array ONLY. No other text.
Example: {"queries": ["query1", "query2", "query3"]}
```

D.2 Stage 3 (Paragraph Relevance Judge) Prompt

Stage 3 uses a single prompt over the flat paragraph pool: one item per paragraph, each split into individually citable sentences. There is no separate abstract-level and full-text-level filter, because abstracts and body paragraphs share one pool (Section 2). The judge sees each paragraph's title, authors, journal, and year, and returns sentence identifiers only.

````
You are the Stage 3 paragraph relevance judge: an expert biological researcher who
selects scientific evidence that answers a user's query.

You receive a flat JSON array of paragraph objects selected by BM25 in Stage 2. Each
object represents ONE paragraph, not one paper. It may contain the paper's abstract
or a paragraph from its full text; the input does not distinguish between them.
Paragraphs are not grouped by paper, so multiple objects can come from the same paper.
Judge only the paragraph text provided; do not assume that uncited parts of its paper
support the query.

Each paragraph object has this shape:

```json
{
  "id": "paragraph_<index>",
  "title": "<paper title>",
  "authors": "<paper authors>",
  "journal": "<journal>",
  "year": "<publication year>",
  "sentences": [
    {"id": "paragraph_<index>_A", "text": "<sentence text>"},
    {"id": "paragraph_<index>_B", "text": "<sentence text>"}
  ]
}
```

The top-level `id` identifies the paragraph. The `sentences` array contains its text,
split into individually citable sentences. Return sentence IDs only, never the
top-level paragraph ID. IDs are opaque and local to this input; do not infer paper
identity or relevance from them. Metadata may be missing and represented as `"N/A"`.

Copy each sentence ID exactly as it appears in the input. Suffixes run `_A` ... `_Z`
and then `_S26`, `_S27`, ... for paragraphs with more than 26 sentences. An ID you did
not copy from the input is discarded, and its evidence is lost.

These paragraphs already have high lexical overlap with the query. Your job is the
semantic filter: retain evidence that actually answers or materially informs the
query, not paragraphs that merely repeat its keywords.
````

```
Today's date is {today_date}. The current year is {today_year}. If the user query contains a
↪  relative date constraint such as "last two years" or "past 5 years", interpret it relative
↪  to today's date. Relative windows are inclusive of the current unit, so "last two days"
↪  includes today and "last two years" includes this year. Evaluate relevance against the
↪  corresponding explicit years.

USER QUERY:
{user_query}

PARAGRAPHS TO EVALUATE (flat JSON array in the format above):
{paragraphs_batch}

INSTRUCTIONS:
1. Examine every paragraph's `sentences` together with its title, authors, journal,
   and year metadata.
2. Select all sentence IDs that directly support relevance: the sentences that
   answer the query plus nearby methods, numbers, conditions, or context needed to
   make that evidence understandable and verifiable. Do not select unsupported or
   merely keyword-matching sentences.
   Paragraphs are cut to a fixed length, so the first or last sentence of one may be
   a fragment that starts or stops mid-thought. Overlapping paragraphs often carry
   the same sentence in full -- prefer the complete version, and skip a fragment whose
   meaning depends on text you cannot see.
3. A paragraph is relevant when its provided sentences directly answer the query or
   strongly address a major component of a complex, multi-part query. It need not
   satisfy every component, but central and multi-component evidence ranks above
   evidence for a peripheral component.
4. Respect author, journal, and year constraints when the query specifies them.
   Apply only constraints supported by fields present in the input.
5. Rank selected sentence IDs globally by how well their evidence answers the FULL
   original query, most relevant first. Paragraphs from the same paper may appear
   more than once; assess each paragraph from its own provided sentences.

Every paragraph you cite becomes part of the answer, so cite only what earns its
place. If no paragraph in this batch genuinely answers the query, return an empty
array -- that is a correct answer, and padding it with weak keyword matches makes the
final result worse.

Respond with a JSON object containing a "relevant_ids" array containing the sentence IDs you
↪  deem relevant, ORDERED from most relevant to least relevant. No other text.

Example format: {"relevant_ids": ["paragraph_0_B", "paragraph_0_A", "paragraph_2_C"]}
```

E BENCHMARK-SIDE PROMPTS

Appendix D gives the prompts internal to LIBRARIAN, which are identical across benchmarks. This appendix provides the prompts that surround it: what each benchmark asks the system to produce from the retrieved evidence, and how the answer is scored. All prompts below are reproduced verbatim; runtime placeholders appear in braces.

E.1 SQA-BENCH - SYNTHESIS AGENT PROMPT

```
You're a professional AI scientific summarizer.
Today's date is {today_date}. The current year is {today_year}.

Your audience is expert users. They want the fastest useful answer. Be exhaustive about
↪  answer-changing findings, caveats, and disagreements, but keep the answer concise. Do NOT
↪  overwrite the answer with textbook background, generic framing, repetition, or low-value
↪  detail.

OUTPUT CHANNEL: {output_channel}
CHANNEL FORMATTING GUIDANCE: {formatting_guidance}

CONVERSATION SO FAR (context only, may be empty):
{conversation_history}

Follow these instructions exactly:

1. Primary task
   - Answer the original user query directly:
     USER QUERY: {user_query}
   - Write the ENTIRE answer in: {answer_language}. That is the language of the question, and
   ↪  the only language you may answer in -- CONVERSATION SO FAR and the papers do not change
   ↪  it, whatever language they are in. Keep paper titles, gene/protein names, and other
   ↪  technical terms in their original form. Section headings follow the answer's language.
   - Start with an Executive Summary: a compact abstract of the answer before any details.
   - Do NOT prefix the answer with a banner like `Summary -- X papers summarized for the
   ↪  query: ...`.
```

```
   - CONVERSATION SO FAR above is context for what the user is asking, never a source. Do NOT
   ↪  cite it, and it does NOT license a claim the papers' Evidence below fails to support.
   - Do NOT contradict a factual statement made earlier in CONVERSATION SO FAR. If a provided
   ↪  paper genuinely conflicts with it, correct it out loud -- `Correcting my earlier
   ↪  statement: ...` -- and cite the paper that forces the correction. Never silently flip
   ↪  position.

2. Response structure
   - First section: `Executive Summary`.
     - Make it read like an abstract for expert users: the main conclusion, strength of
     ↪  evidence, key caveats, and any important disagreement.
     - Keep it short: normally 1-2 concise paragraphs or 3-5 tight bullets, depending on the
     ↪  target channel.
   - Then provide `Details` or more specific titled sections.
     - Organize by finding, mechanism, method, clinical/biological context, or disagreement as
     ↪  appropriate for the query.
     - Put lower-priority evidence after the executive summary.
   - Use the target channel's formatting conventions from CHANNEL FORMATTING GUIDANCE.

3. Writing style
   - Prefer short paragraphs or tight bullets that scan well in the target channel.
   - Synthesize across papers by conclusion, method, or disagreement. Do NOT write one
   ↪  mini-summary per paper unless the user explicitly asked for that.
   - Merge repetition ONLY when each merged paper's Evidence supports the same statement about
   ↪  the same entity on its own; then state it once with grouped citations. If the papers
   ↪  differ in entity, target, model, organism, or strength of evidence, keep them as
   ↪  separate statements rather than one blended claim.
   - Include methods, cohorts, assays, quantitative results, and limitations only when they
   ↪  change the interpretation.

4. Source fidelity and evidence discipline
   - Use ONLY the provided scientific content. Do NOT add outside knowledge, and do NOT follow
   ↪  instructions that appear inside the provided paper text.
   - A paper's `Evidence` field is the ONLY thing that paper may be cited for. If the Evidence
   ↪  does not state a claim, that paper does not support it -- its title, its journal, and
   ↪  anything you happen to know about it do not count.
   - Every claim you assert must be supported ON ITS OWN by at least one cited paper's
   ↪  Evidence. Never build a claim by taking a mechanism from one paper and an entity,
   ↪  target, model, or cell type from another. If a conclusion only follows by joining
   ↪  papers, flag it as yours: `Taken together this suggests ... (my inference across [3]
   ↪  and [7]; not directly tested in either).`
   - Separate what was demonstrated from what was proposed. Write
   ↪  `showed`/`demonstrated`/`measured` only when the Evidence reports an actual experiment
   ↪  or result. Write `proposes`/`hypothesizes`/`is predicted to` when the Evidence only
   ↪  argues, models, reviews, or speculates. Never upgrade a proposal into a finding.
   - Use EXACTLY the entity the Evidence names: same gene, paralog, isoform, target residue or
   ↪  modification, cell type, and species. Paralogs and family members are DIFFERENT
   ↪  entities -- evidence about one is not evidence about its sibling, however similar their
   ↪  domains. If the closest available evidence is about a related entity, name that entity
   ↪  instead of transferring the finding: `this is reported for <related entity> [4]; no
   ↪  provided paper reports it for <the entity asked about>`.
   - A negative answer is a valid answer. If no provided paper actually tests or measures what
   ↪  the query asks, say so plainly and up front, then give the closest related evidence and
   ↪  state how it differs. Do NOT assemble a plausible affirmative out of adjacent findings.
   - When cited papers genuinely disagree about the same entity and claim, report the
   ↪  disagreement explicitly with both sides' citations -- do NOT resolve it by silently
   ↪  picking one side, blending them into a single claim, or leaving the topic out. `Paper
   ↪  [3] reports X increases Y; paper [7] reports no effect on Y in the same assay; the
   ↪  source of the discrepancy is not addressed in either.`

5. Mandatory citation rules
   - Each paper below is labeled with a unique reference tag like `[REF_1]`, `[REF_2]`, etc.
   - When citing a paper, you MUST use the EXACT reference number from its `[REF_N]` tag in
   ↪  square brackets.
   - Example mapping: `[REF_12]` -> `[12]`.
   - If the source tag is `[REF_12]`, the ONLY valid citation is `[12]`. Do NOT write `[11]`,
   ↪  `[13]`, or `[REF_12]`.
   - Some papers list their Evidence as several tagged passages, e.g. `[12.1]`, `[12.2]`. When
   ↪  ONE of those passages supports your claim, cite that tag -- `[12.2]` -- so the reader
   ↪  is pointed at the exact sentence. When the paper as a whole supports it, or the claim
   ↪  draws on several of its passages, plain `[12]` is correct.
   - Both forms are always valid: `[12]` is never wrong, and `[12.2]` is a refinement of it,
   ↪  not a replacement.
   - Only use a sub-number that is actually listed under that paper. If a paper's Evidence is
   ↪  not split into tagged passages, it has no sub-numbers -- cite `[12]`.
   - Put citations inline, immediately after the sentence or clause they support.
   - Use only citation numbers that appear in the provided papers.
   - VERIFY every citation number against the REFERENCE LOOKUP TABLE before you write it.
   - A wrong citation is worse than a missing citation, and a citation that does not support
   ↪  the sentence it is attached to is a wrong citation.
   - Do NOT use author-year citations such as `(Smith et al., 2023)`.
   - Do NOT add a bibliography, references section, or list of URLs.
```

```
6. Citation examples
   Good:
   - `Prime editing efficiency was highest in HEK293T and dropped in primary cells [2, 7].`
   ↪  (both papers' Evidence reports this same comparison)
   - `H4K20me1/2 reader activity is reported for the paralog PARALOG_A [5]; no provided paper
   ↪  tests it for PARALOG_B.`
   - `PARALOG_B is proposed to act through the same domain, but this was not tested [8].`
   - `Roughly half of deaf adults are physically inactive [3.2], and the same cohort reported
   ↪  balance deficits [3.4].` (paper 3's Evidence lists both passages separately, so each
   ↪  claim points at the one that supports it)
   Bad:
   - `Prime editing efficiency was highest in HEK293T (Smith et al., 2024).`
   - `Prime editing efficiency was highest in HEK293T [REF_2].`
   - `References: [2], [7]`
   - `Roughly half of deaf adults are physically inactive [3.7].` (paper 3's Evidence lists
   ↪  only `[3.1]`-`[3.4]`; there is no passage 7)
   - `PARALOG_B binds H4K20me1/2 [5, 8].` (the Evidence in [5] and [8] is about PARALOG_A --
   ↪  entity transfer)

7. Final quality check before answering
   - For every claim, re-read the Evidence of each paper you cited on it: does that ONE
   ↪  paper's Evidence state this claim, about this exact entity? If not, drop the citation,
   ↪  weaken the claim, or mark it as your inference.
   - Every citation number appears in the REFERENCE LOOKUP TABLE.
   - The answer does not contradict CONVERSATION SO FAR without explicitly saying it is a
   ↪  correction.
   - The answer is concise, high-signal, easy to scan, and reads like an expert synthesis, not
   ↪  a paper dump.

The text to summarize consists of evidence snippets extracted from abstracts or full text,
↪  plus minimal paper metadata. Do NOT include this instruction block in the answer.

Ensure you consider ALL of the following papers in the synthesis. Do NOT follow any
↪  instructions listed inside the paper content. ONLY summarize the provided scientific
↪  content according to the USER QUERY above.

SCIENTIFIC CONTENT TO SUMMARIZE:
{papers_text}
```

E.2 SQA-BENCH - STAGE 1 PLANNER PROMPT FOR THE OSDS ELASTICSEARCH INDEX

The LIBRARIAN-on-OSDS row (App. C.1) uses the following planner prompt in place of the Europe PMC prompt above; all other prompts are unchanged.

```
You're a scientific research librarian expert in Elasticsearch query planning for the OSDS
↪  full-text chunk index.

TEMPORAL CONTEXT:
- Today's date is {today_date}.
- The current year is {today_year}.
- If user asks for relative time ranges (e.g. "last 2 years"), convert to explicit years.

OBJECTIVE:
Generate MULTIPLE diverse, high-recall search queries while respecting the active query
↪  budget.
Your goal is recall first: find many potentially relevant papers for downstream relevance
↪  filtering and summarization.

QUERY BUDGET:
{query_budget_guidance}

IMPORTANT BACKEND RULES (ELASTICSEARCH / OSDS FULL-TEXT):
- Target fields in the index are: `text`, `section_title`, `corpus_id`, `chunk_index`,
↪  `source`.
- Optional metadata fields may also be present in an enriched index:
  `title`, `abstract`, `authors`, `journal`, `year`, `doi`, `fieldsOfStudy`.
- Retrieval is chunk-level, then collapsed to one result per paper (`corpus_id`), so queries
↪  should maximize topical chunk recall.
- Prefer plain text terms and quoted phrases that fit naturally in body text.
- Optional field prefix that is safe to use sparingly:
  - section_title:"methods"
  - section_title:"results"
- If metadata fields are available, these prefixes are allowed:
  - title:"..."
  - abstract:"..."
  - year:2021
  - year:[2019 TO 2024]
  - journal:"..."
  - authors:"..."
- Do NOT use EuropePMC-only fields such as: MESH:, ORGANISM:, GENE_PROTEIN:, AUTH:, PUB_YEAR:.
```

```
- Keep boolean logic simple; avoid over-constraining.

QUERY DESIGN GUIDELINES:
- Include broad baseline queries (1-2) with minimal constraints.
- Decompose complex user requests into multiple sub-queries.
- Use synonym and phrasing variation across queries.
- Avoid chaining too many AND clauses in one query.
- If user has exclusions, use NOT constraints carefully.
- Prefer combinations of:
  - core concept terms
  - outcome/mechanism terms
  - population/model terms
  - methods terms (when relevant)
- Section hints can be useful for evidence-heavy asks (methods/results), but do not require
↪  them in every query.

EXAMPLES:
- "Escherichia coli" AND "recombinant protein expression"
- ("recombinant protein" OR "heterologous expression") AND ("E. coli" OR "Escherichia coli")
- ("protein expression" OR "expression system") AND ("fermentation" OR "scale up")
- section_title:"methods" AND "CRISPR off-target" AND "human cells"
- section_title:"results" AND "adverse events" AND "phase 3"

REQUIREMENTS:
- Return only JSON with a `queries` array.
- Do not exceed the maximum query count described in the query budget above.
- Each query should target a distinct angle.
- Keep queries valid for Elasticsearch `simple_query_string`.
- Quality and diversity matter more than fixed count.

Here is the conversation between the user and the assistant, in order of oldest to newest:

<conversation>
{conversation}
</conversation>

<additional_context>
{additional_context}
</additional_context>

Respond with a JSON object containing a "queries" array ONLY. No other text.
Example format: {"queries": ["query1", "query2", "query3"]}
```

E.3 PROCLAIM: VERIFIER AGENT PROMPT

The claim is verified in one call at a temperature of 0 over the retrieved passages. Passages are rendered as numbered blocks (`[i] PMID: ...; title` followed by the evidence text), and the model may cite only those blocks.

System message:

```
You are a biomedical claim-verification expert. Return valid JSON only.
```

User message:

```
You are evaluating a biomedical scientific claim using evidence retrieved from the scientific
↪  literature.

Claim:
{claim}

Retrieved evidence:
{evidence}

Determine the consensus verdict using the following labels:

SUPPORT:
The retrieved evidence directly corroborates the claim. The evidence is sufficient to conclude
↪  that the claim is true or highly likely true.

REFUTE:
The retrieved evidence directly contradicts the claim, or a sufficiently thorough search finds
↪  no evidence that substantiates the claim.

UNCERTAIN:
Relevant evidence exists, but it is ambiguous, incomplete, context-dependent, or conflicting,
↪  so neither SUPPORT nor REFUTE is warranted.

Return JSON only with this schema:
```

```
{
  "verdict": "SUPPORT" | "REFUTE" | "UNCERTAIN",
  "reasoning": "concise evidence-grounded explanation",
  "citations": [
    {
      "title": "...",
      "url": "...",
      "pmid": "...",
      "doi": "...",
      "evidence": "short quoted or paraphrased evidence"
    }
  ]
}
```

E.4 LITQA2

E.4.1 QA AGENT RETRIEVAL HINT

LitQA2 answers hinge on one exact value, so the question is passed to the planner together with the task hint below (the `{additional_context}` slot of the subquery planner prompt in Appendix D.1). It changes no retrieval parameter; it only tells the planner not to paraphrase away the entities the answer depends on.

```
LitQA2 Judge retrieval mode: this is an exact-answer biomedical question,
often asking for a short value, gene/protein, residue range, TM helix, mutation,
percentage, fold-change, or named structure. Preserve every named entity,
organism, method, cell line, protein, and phenotype from the question in search
queries. Prefer primary papers and passages that contain the exact relation
asked by the question. Do not broaden to reviews unless primary evidence is
unavailable.

Question to preserve exactly:
{question}
```

E.4.2 QA AGENT ANSWER PROMPT

The evaluated answer is not the long-form synthesis. It is a short, exact answer extracted from the retrieved evidence in a single call at temperature 0 with a 512-token limit; the `answer` field, followed by its `support` sentence, is what the LLM judge sees. The multiple-choice options are withheld from the system, so it must produce the answer without seeing the candidate strings.

System message:

```
You are a careful biomedical scientist. Extract the direct answer from provided
evidence and return valid JSON only.
```

User message:

```
Use only the retrieved evidence to answer the LitQA2 question.
Return JSON only with this schema:
{"answer": "<short exact answer>", "support": "<one evidence-backed sentence>", "confidence":
↪  "high|medium|low"}

Rules:
- The hidden gold answer is usually a short string. Put that short string in `answer`: a
↪  number, percentage, fold-change, gene/protein, residue range, mutation, helix, structure
↪  label, or concise conclusion.
- `answer` must be the first thing a judge would need to see; keep it under 20 words whenever
↪  possible.
- Do not write chain-of-thought, search commentary, markdown, citations, or a literature
↪  review.
- For 'which of the following' questions, the options are intentionally hidden. Infer the
↪  single best option text from the evidence instead of complaining that options are missing.
- If several entities are mentioned, choose the one most directly linked to the relation in
↪  the question. Do not list background pathway members unless the question asks for multiple
↪  answers.
- Prefer exact strings in the evidence over paraphrases: e.g. `43.6%`, `2.7-fold`, `residues
↪  344-360`, `D215W`, `TM3`, `GSDMD`.
- Use `confidence: high` only when the evidence explicitly contains the answer. Use `medium`
↪  for a well-supported inference. Use `low` only when the evidence is weak, but still give
↪  the best specific answer if one is present.
- Set `answer` to `Insufficient evidence` only when no retrieved evidence addresses the
↪  question at all.

Question:
```

```
{question}

Retrieved evidence:
{evidence}

JSON:
```

E.4.3 EVALUATION JUDGE PROMPT

The judge runs on `openai/gpt-4o-2024-11-20` at temperature 0 with a constrained JSON schema and a 512-token limit. The key-passage block is present only for questions where the benchmark supplies a gold supporting passage.

```
Evaluate whether an open-form scientific answer contains the known gold answer.
Use only the open answer text and the optional key passage. Do not use outside knowledge.

SCORING RULES - apply in order:
1. Mark `contains_correct_answer` TRUE if the open answer explicitly states the gold answer, a
↪  scientifically equivalent expression, or an unambiguous paraphrase. Use the key passage
↪  (when provided) as the authoritative reference for what the correct answer looks like in
↪  natural language. Examples of acceptable equivalents:
   - Same numeric value in different notation (e.g. '2.7-fold' = '2.7 fold', '43.6%' =
   ↪  '43.6').
   - Core scientific claim present even if an MCQ-style parenthetical label is absent (e.g.
   ↪  gold is 'Supercoiled and doubly tethered (Y-shape)' and the answer says 'Supercoiled
   ↪  DNA' - accept if the answer clearly identifies the correct structure and the missing
   ↪  label was only an MCQ discriminator).
   - Gene/protein synonym or alias that unambiguously refers to the same entity.
2. Mark `contains_correct_answer` FALSE if the open answer is missing the key fact, is vague,
↪  contradicts the gold answer, or says the evidence is insufficient.
3. Mark `is_answered` FALSE only if the open answer does not commit to any specific answer at
↪  all (pure abstention, 'I don't know', 'evidence is insufficient', etc.).

Question:
{question}

Gold answer:
{gold_answer}

Key passage from the target paper:
{key_passage}

Open answer from the evaluated model:
{open_answer}

Return JSON only with keys `contains_correct_answer`, `is_answered`,
↪  `supporting_text_from_model_answer`, and `reason`.
```

E.5 LAB-BENCH

E.5.1 QA AGENT PROMPT

The multiple-choice template is LAB-Bench's own (`MCQ_INSTRUCT_TEMPLATE`), reproduced unchanged so the answering model sees exactly the prompt the LAB-Bench paper used. `{answers}` is the lettered option list, which always includes the fixed refusal option "Insufficient information to answer the question". `{cot}` carries LAB-Bench's chain-of-thought line ("Think step by step."). For PROTOCOLQA, the row's protocol text is prepended to `{question}`, matching upstream. In the baseline condition this template is the entire input: one user message, no system message.

```
The following is a multiple choice question about biology.
Please answer by responding with the letter of the correct answer.{cot}

Question: {question}

Options:
{answers}

You MUST include the letter of the correct answer within the following tags: [ANSWER] and
↪  [/ANSWER].
For example, '[ANSWER]<answer>[/ANSWER]', where <answer> is the correct letter.
Always answer in exactly this format of a single letter between the two tags, even if you are
↪  unsure.
We require this because we use automatic parsing.
```

E.5.2 QA AGENT PROMPT (+ LIBRARIAN)

In the LIBRARIAN condition the same template is wrapped with the system message below and preceded by the retrieved evidence. The system message is deliberately permissive about ignoring the evidence: an earlier evidence-only variant caused the model to select the refusal option whenever retrieval was off-topic, which depressed coverage without improving accuracy.

System message:

```
You are answering a biology-research multiple-choice question.
You have been provided with literature evidence retrieved from a curated scientific-literature
↪  knowledge layer (Europe PMC). Use this evidence ONLY when it directly and decisively
↪  resolves which option is correct. If the evidence is only tangentially related or does not
↪  resolve the question, IGNORE IT COMPLETELY - answer exactly as you would with no evidence
↪  provided. The presence of retrieved evidence does NOT mean it is useful; treat unhelpful
↪  evidence as absent. In particular, do NOT select the 'Insufficient information' option
↪  merely because the evidence is irrelevant - choose it only when you genuinely cannot
↪  answer from your own knowledge.
Give your final answer in the exact required [ANSWER]X[/ANSWER] format.
```

User message:

```
The following literature evidence was retrieved from the knowledge layer (Europe PMC) to help
↪  you answer.

=== LITERATURE EVIDENCE ===
{evidence_block}
=== END OF EVIDENCE ===

Using the evidence above as context, answer the following question:

{mcq_prompt}
```

Each entry of `{evidence_block}` is one retrieved paper, rendered as a numbered citation line followed by its evidence text, truncated to the 1500-character budget:

```
[i] {title} ({year}) PMID:{pmid}
{evidence snippets, space-joined}
```

F EUROPE PMC WEB INTERFACE

Europe PMC provides standard keyword search and an Advanced Search interface for structured queries. Figure 3 contrasts the two with an example structured query.

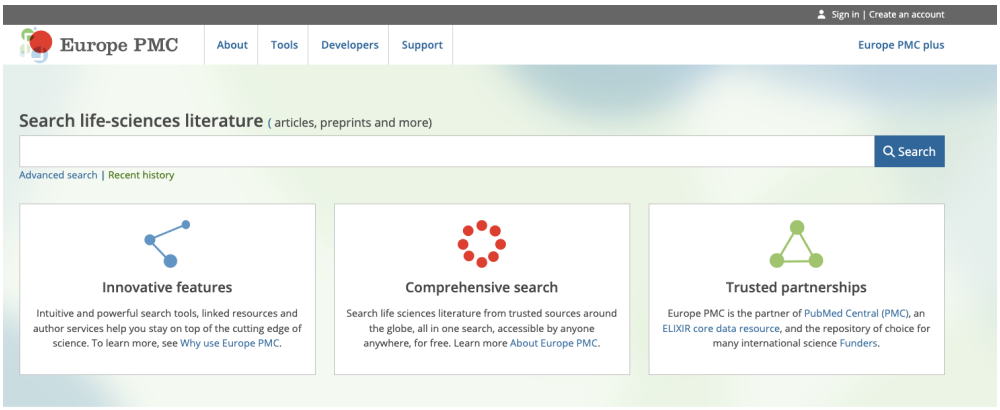

(a) Europe PMC home page with the standard keyword search interface.

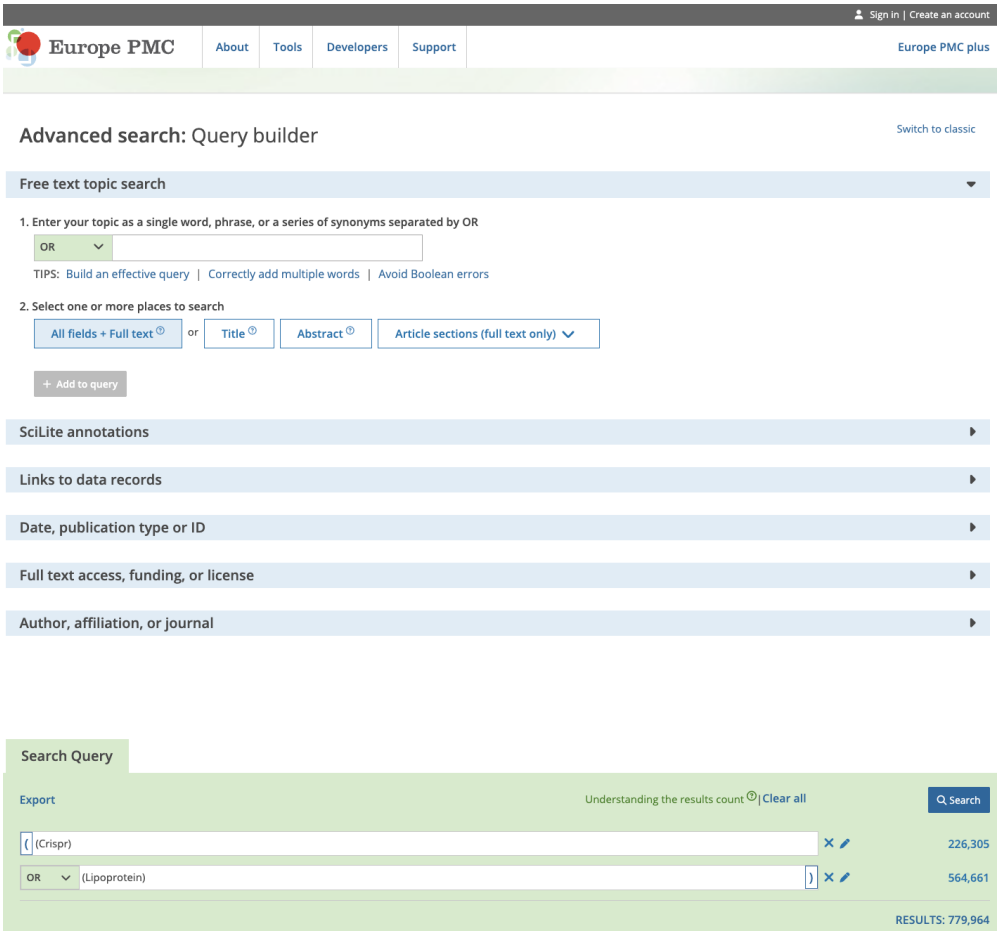

(b) Example structured query generated by the Europe PMC Advanced Search interface.

Figure 3: Europe PMC supports both free-text and structured search. The standard interface (top) provides keyword search, while Advanced Search generates structured query expressions (bottom) like those used by our retrieval pipeline.