

BRIEF-PRO: Universal Context Compression with Short-to-Long Synthesis for Fast and Accurate Multi-Hop Reasoning

Jia-Chen Gu*, Junyi Zhang*, Di Wu, Yuankai Li, Kai-Wei Chang, Nanyun Peng

University of California, Los Angeles

{gujc, junyizhang2002}@ucla.edu, {diwu, kwchang, violetpeng}@cs.ucla.edu

Abstract

As retrieval-augmented generation (RAG) tackles complex tasks, increasingly expanded contexts offer richer information, but at the cost of higher latency and increased cognitive load on the model. To mitigate this bottleneck, especially for intricate multi-hop questions, we introduce BRIEF-PRO. It is a universal, lightweight compressor that distills relevant evidence for a given query from retrieved documents into a concise summary for seamless integration into in-context RAG. Using seed data consisting of relatively short contexts (fewer than 1k words), BRIEF-PRO is trained to perform abstractive compression of extended contexts exceeding 10k words across a wide range of scenarios. Furthermore, BRIEF-PRO offers flexible user control over summary length by allowing users to specify the desired number of sentences. Experiments on four open-domain multi-hop question-answering datasets show that BRIEF-PRO generates more concise and relevant summaries, enhancing performance across small, large, and proprietary language models. With the 70B reader model, 32× compression by BRIEF-PRO improves QA performance by 4.67% on average over LongLLMLingua’s 9×, while requiring only 23% of its computational overhead¹.

1 Introduction

Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) has emerged as a powerful paradigm for enhancing the factual grounding and knowledge breadth of large language models (LLMs) (Meta AI, 2024a; Anthropic, 2024; OpenAI, 2025; Google, 2025). By dynamically retrieving relevant information from a vast corpus and incorporating it into the LLM’s context, RAG systems aim to mitigate issues like hallucination and outdated knowledge (Ji et al., 2023). However, a significant bottleneck

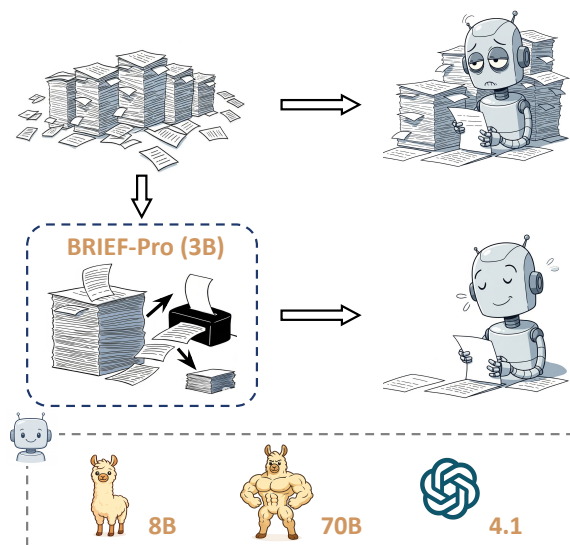


Figure 1: A comparison of the inference process with and without the lightweight BRIEF-PRO. Retrieved documents are compressed into a highly dense textual summary relevant to the query before being prepended, thereby reducing the cognitive load caused by the extended context on a range of larger models, including 8B, 70B, and proprietary models.

in scaling RAG (Shao et al., 2024) to real-world applications with numerous retrieved documents is the proportional increase in latency (Jiang et al., 2023b; Xu et al., 2024a). Feeding such an extensive context not only dramatically slows down inference but also increases the cognitive load imposed on the model (Shi et al., 2023; Mallen et al., 2023; Xu et al., 2024b; Liu et al., 2024). This challenge is amplified for intricate *multi-hop* questions that demand reasoning across multiple disparate sources (Yang et al., 2018; Ho et al., 2020; Trivedi et al., 2022). The large volume of information can overwhelm LLMs, hindering evidence integration and diluting key signals, leading to suboptimal performance even with relevant content.

To address these challenges, context compression has emerged as a crucial technique for improving the efficiency and effectiveness of the RAG sys-

* Equal contribution.

¹Code and data: <https://github.com/JasonForJoy/BRIEF>

tem. However, existing context compression methods often struggle to retain all critical information when the input context becomes extensive (Jiang et al., 2024) as illustrated in Appendix B.3. This limitation arises from their inability to capture interdependencies across large or multiple documents. Moreover, these methods usually do not allow users to set the compression budget (Xu et al., 2024a; Li et al., 2025), making it difficult to balance compression rate and information preservation for specific applications.

To mitigate this critical bottleneck and unlock the full potential of RAG for complex reasoning, we introduce improved BRIEF-PRO (Figure 1), a universal, lightweight compressor for Bridging Retrieval and Inference through Evidence Fusion. BRIEF-PRO is designed to distill relevant evidence for a given query from retrieved documents into a concise summary. Unlike previous studies on short-context compression (Xu et al., 2024a; Yoon et al., 2024; Li et al., 2025), our curated training data enables BRIEF-PRO to generalize to much longer inputs, as required in real-world RAG scenarios. The key innovation of BRIEF-PRO is its long-context synthetic data pipeline, developed from short-context seed data (fewer than 1k words), enabling abstractive compression of 10k+ word contexts across diverse scenarios. Furthermore, BRIEF-PRO allows users to control compression length by designing an instruction-conditioned paradigm, specifying the desired number of sentences to balance conciseness and informativeness for specific application requirements.

We evaluated BRIEF-PRO on four open-domain multi-hop question-answering (QA) datasets: the extended MuSiQue (Trivedi et al., 2022), HotpotQA (Yang et al., 2018), and 2WikiMultiHopQA (Ho et al., 2020) from LongBench (Bai et al., 2024b), and LongSeal (Pham et al., 2025), with context lengths ranging from 4.9k to 14.8k words. The extensive contexts in these datasets challenge traditional compression methods. Experimental results consistently demonstrate that BRIEF-PRO generates more concise and relevant summaries. This significantly improves the QA accuracy and inference latency across a wide range of small (Llama-3.1-8B-Instruct), large (Llama-3.1-70B-Instruct) (Meta AI, 2024a), and proprietary (GPT-4.1-nano) (OpenAI, 2025) language models. Specifically, with the 70B reader model, 32× compression by BRIEF-PRO

improves QA performance by 4.67% on average over LongLLMLingua’s 9× (Jiang et al., 2024), while requiring only 23% of its computational overhead. These findings highlight BRIEF-PRO’s potential to advance the scalability and effectiveness of RAG for complex retrieval and generation tasks.

In summary, our contributions in this paper are three-fold: (1) This study pioneers the exploration of multi-hop reasoning and compression of RAG for *long contexts of 10k+ words* across diverse scenarios. (2) A synthetic data pipeline, built on short-context seed data, is designed to synthesize *long-context* training data for compression learning. (3) BRIEF-PRO, trained on the curated dataset, generates concise summaries that *accelerate the inference and enhance the accuracy* of a wide range of small, large, and proprietary language models.

2 Related Work

The processing and understanding of long contexts presents several challenges, including increased inference costs, longer latency, and decreased performance due to redundant and distracting information. Many efforts have been made to compress long contexts. One line of research proposes compressing long contexts into soft prompts that can be used by LMs, such as GIST (Mu et al., 2023), AutoCompressors (Chevalier et al., 2023). However, these soft prompts are usually tailored to particular tasks and require fine-tuning to align with the representation space of LMs, which severely limits their compatibility. Another line of work proposes compressing contexts into textual summaries, such as RECOMP (Xu et al., 2024a), CompAct (Yoon et al., 2024), EXIT (Hwang et al., 2025), and BRIEF (Li et al., 2025), and our method belongs to this category. RECOMP distills the summarization ability of extreme-scale proprietary LLMs into an in-house abstractive compressor. CompAct employs an active strategy to recurrently acquire new information from documents and compress it into a compacted context. BRIEF curates a synthetic data pipeline built by open-source models to enhance the awareness of multi-hop reasoning and scalability. Compared to soft prompts, this approach yields more interpretable textual summary that can transfer across different LMs, and can be applied to black-box LMs without requiring gradient updates. However, these compression methods are limited to scenarios where

the context is relatively short and provide limited user control over compression. The LLMingua family (Jiang et al., 2023b, 2024) is most relevant to this work, proposing demonstration- and token-level compression that leverages a small LM to calculate perplexity and prune redundancy. But their budget allocation may inadvertently discard crucial details in an attempt to meet an assigned compression budget, resulting in a loss of fidelity. While CoLoR (Seo et al., 2025) also employs a synthetic data method for training a compressor, its data synthesis process, supervision strategy, pipeline design, and target context length are all substantially different from ours. Despite these efforts, our work addresses the underexplored challenge of compressing significantly longer contexts while offering flexible, user-controllable compression, advancing beyond prior methods limited to shorter inputs and rigid budgets.

3 BRIEF-PRO

3.1 Problem Formulation

The proposed architecture consists of two modules: a compressor \mathcal{C} and an LM \mathcal{M} . For every input query \mathbf{x} , an off-the-shelf passage retriever (Karpukhin et al., 2020; Izacard et al., 2022) returns a query-related long context \mathbf{D} , which may consist of multiple short documents or a single long document. Then, the compressor \mathcal{C} takes as input the concatenation of query \mathbf{x} , query-related documents \mathbf{D} , and *an optional user-specified compression instruction \mathbf{i}* , and outputs a summary \mathbf{s} . The summary \mathbf{s} captures the core information with respect to \mathbf{x} with significantly fewer words. Finally, the input query \mathbf{x} and the compressed summary \mathbf{s} are fed into an off-the-shelf LM \mathcal{M} . The compressor \mathcal{C} is trained on the corpora we curated in this work, while the LM \mathcal{M} remains frozen and can be any off-the-shelf LM. In this work, we train a 3B autoregressive model to serve as an *abstractive* compressor \mathcal{C} and adopt a wide range of 8B, 70B, and proprietary models as the LM \mathcal{M} . The compressor \mathcal{C} is intentionally designed to be substantially smaller than the LM \mathcal{M} , as we aim to reduce the computational cost of encoding the lengthy query-related documents.

3.2 BRIEF-PRO Inference

Compared to previous work, we propose a novel context compression paradigm featuring both *user-controllable* and *automated* budget allocation. For

user-controllable compression, we explicitly append a user-specified compression instruction \mathbf{i} that indicates the expected number of sentences in the summary after the context. This direct instruction guides the model to produce a summary of the specified length. In contrast, for automated context compression, the model implicitly determines the optimal number of sentences for the compressed summary based on the inherent content and its internal learned representations. This allows for a more hands-off approach, where the system intelligently decides the summary length, making it ideal for scenarios where a pre-defined sentence count is not feasible or desired. Appendix A presents the detailed inference prompts for the compressor and reader models, respectively.

3.3 Data Collection

Traditional context compression methods can only handle contexts of limited length and often lack flexibility (Xu et al., 2024a; Yoon et al., 2024; Li et al., 2025), offering limited control over the semantic granularity of the preserved information. In contrast, to develop a more capable and user-controllable context compression model, it is essential to curate a high-quality dataset that captures the nuances of information relevance across various query types. This can help the model learn to prioritize essential context as instructed by the user while eliminating redundancy. This work presents a cost-efficient training recipe by designing a synthetic data pipeline that leverages short-context seed data for long-context compression learning, as illustrated in Figure 2. Specifically, our investigation delves into three key practices in terms of expanding the input context, curating the target summary, and designing a user-controllable compression mechanism.

3.3.1 Short-to-Long Context Expansion

Building models capable of compressing very long contexts requires genuinely long, coherent training examples. Since most available data consists of shorter documents (Yang et al., 2018; Trivedi et al., 2022), our goal is to expand these into extended-context data for robust training.

Source Location For each document, we first identify its topic and use this information to locate the corresponding Wikipedia page from which the document originates. Specifically, the structured Wikipedia corpus released by Izacard et al. (2022)

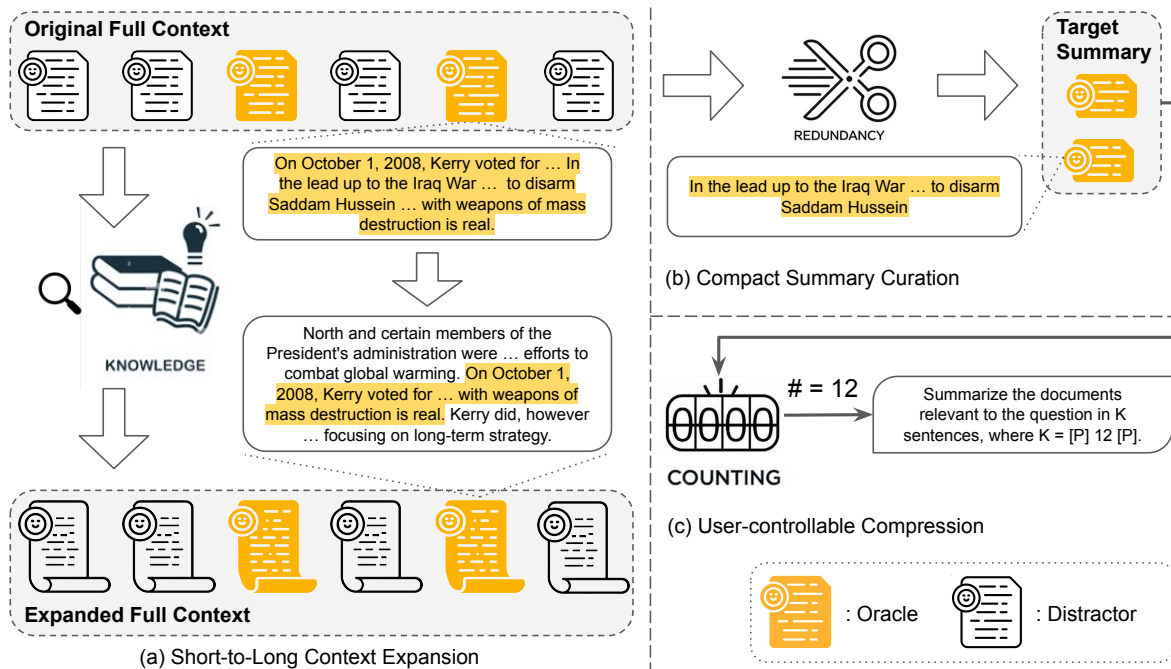


Figure 2: An overview of the synthetic data pipeline for training BRIEF-PRO. Starting with a mixture of oracle and distractor documents for a given query, the pipeline: (a) expands each document by looking up an external knowledge corpus, (b) curates a compact summary by further reducing redundancy in the oracle documents, and (c) generates a user-controllable compression instruction by counting the number of sentences in the compact summary.

is leveraged to efficiently retrieve the most relevant page. For documents that cannot be directly matched to an existing Wikipedia page in this corpus, we further turn to the Wikipedia website to search for content on the same topic².

Context Expansion Once the appropriate Wikipedia source is located, the precise position of the document within that page is pinpointed. Using this position as a reference, the document can be expanded by including a specified number of sentences before and after its original location, thereby naturally enriching the context and extending the content. To control the extent of this expansion, each document is assigned a specific expansion ratio, defined as the number of sentences in the expanded document relative to the original. To diversify context lengths in the training data, the expansion ratio for each document is randomly sampled from a predefined normal distribution with a mean of 20. This approach ensures that the expanded context lengths are, on average, substantially longer, while also maintaining a broad range of context lengths to better support model generalization.

²Querying other verifiable, authoritative web sources holds promise for further expanding the scope and ensuring the availability of information beyond fixed, Wikipedia-style text.

3.3.2 Compact Summary Curation

Given a set of retrieved documents, a pressing research challenge is to identify which text segments within them are the most helpful and can effectively support answering a specific question. Existing QA datasets offer oracle document annotations (Yang et al., 2018; Trivedi et al., 2022), and our pipeline operates under the assumption of their availability. However, these annotations often exhibit redundancy, which requires additional processing to distill truly essential information. Therefore, to improve the compression rate and extract more precise context, our work extends beyond existing annotations by actively pruning the oracle documents to eliminate this redundancy.

Helpfulness Definition For a question x , the helpfulness of a sentence p_{ij} in an oracle document d_i is determined by the LM’s end-task performance when that sentence is removed. Formally, we compare the log likelihood assigned to the target output y by an LM \mathcal{M} before removing the sentence and after. A sentence is considered unhelpful if its removal increases the likelihood of correctly answering that question.

Head-Tail Iterative Pruning In this work, we explore a practical strategy to achieve this by pruning

the head and tail sentences of each oracle document, hypothesizing that critical information is often centrally located. For each oracle document, we iteratively check whether each head sentence is unhelpful and remove, continuing this process until a helpful sentence is identified. Similarly, we apply the same iterative pruning process to the tail sentences of the document. This approach allows us to identify a more compact, continuous, and helpful text segment within the oracle document. Finally, the concatenation of these pruned oracle documents is considered as the target summary.

3.3.3 User-controllable Compression

User-specified Compression Instruction User-controllability is implemented by strategically inserting a clear instruction \mathbf{i} immediately following the provided context. This instruction, phrased as "*Summarize the documents relevant to the question in K sentences, where $K = [P] \#\# [N]$* " where $\#\#$ is an integer placeholder, empowers the user to directly specify the desired length of the compressed output. By allowing users to assign the number of sentences, we provide a flexible and intuitive mechanism to control the granularity of the summary, ensuring the resulting output aligns precisely with their information needs.

Instruction Data Creation To create an instruction-tuning data pair containing: 1) an instruction specifying the number of sentences, and 2) its corresponding summary, we leverage the already constructed summaries in Section 3.3.2. For each summary, we count the number of sentences it contains, let's say this count is k . Then, we formulate the instruction part of the pair as "*Summarize the documents relevant to the question in K sentences, where $K = [P] k [N]$* ." The corresponding summary for this instruction is simply the summary we initially constructed. This direct pairing ensures that the model learns the precise relationship between a numerical sentence constraint in the instruction and the actual length of the summary. By generating a diverse set of such pairs across various contexts and their pre-computed summaries, we build a robust dataset for instruction tuning, enabling the model to generalize and produce summaries of user-specified lengths.

3.4 BRIEF-PRO Training

The curated dataset \mathcal{D}_{comp} is utilized to fine-tune the compressor \mathcal{C} , a Llama-3.2-3B-Instruct

Training Set Statistic	Number
Sample Number	45.2k
Context Words	
- Average	6.0k
- Standard Deviation	3.5k
Summary Words	
- Average	0.2k
- Standard Deviation	0.3k

Table 1: The statistics of the curated training data.

model (Meta AI, 2024b). This model strikes an effective balance between computational efficiency and performance. Its relatively compact size, compared to larger LLMs, makes it well-suited for context compression without demanding excessive computational resources. The fine-tuning process follows the standard next token objective, formulated as:

$$\max_{\mathcal{C}} \mathbb{E}_{(\mathbf{x}, \mathbf{D}, \mathbf{i}, \mathbf{s}) \sim \mathcal{D}_{comp}} \log p_{\mathcal{C}}(\mathbf{s} | \mathbf{x}, \mathbf{D}, \mathbf{i}). \quad (1)$$

Discussion We respectfully note that the novelty of our method lies not merely in combining existing techniques but in the careful design and training of a lightweight compression model that achieves strong long-context summarization at low inference cost, preserves downstream task quality, and generalizes across diverse datasets. Achieving this balance is non-trivial and requires tailored training strategies, aspect-based supervision, and length-controlled outputs, which are not provided by off-the-shelf summarization or LLM-based pipelines.

4 Experiments

4.1 Experimental Settings

BRIEF-PRO Series BRIEF-PRO was initialized from Llama-3.2-3B-Instruct (Meta AI, 2024b). The original training sets of MuSiQue (Trivedi et al., 2022), HotpotQA (Yang et al., 2018), and LongAlign (Bai et al., 2024a) were used as the seed data of our synthesis pipeline. Table 1 presents the statistics of the curated training data. We evaluated a series of BRIEF-PRO models that feature diverse user-controllable and automated compression scenarios. In the user-controllable setting, we defined three compression levels of HIGH, MEDIUM, and LOW, which were empirically set to compress to 5, 10, and 20 sentences, respectively. The AUTO mode denotes dynamic thinking, where the model

	# samples	# context words
MuSiQue	200	11.2k
HotpotQA	200	9.2k
2WikiMultiHopQA	200	4.9k
LongSeal	254	14.8k

Table 2: The statistics of the evaluation data.

decides the number of sentences based on the complexity of the task.

Reader Series To demonstrate that BRIEF-PRO can benefit a wide range of models, small (Llama-3.1-8B-Instruct), large (Llama-3.1-70B-Instruct), and proprietary (GPT-4.1-nano)³ language models were used as the reader \mathcal{M} . Readers may refer to Appendix B.2 for more implementation details.

Datasets We evaluated the BRIEF-PRO series on the following four open-domain multi-hop QA datasets: MuSiQue (Trivedi et al., 2022), HotpotQA (Yang et al., 2018), and 2WikiMultiHopQA (Ho et al., 2020) which are extended versions from LongBench (Bai et al., 2024b), and LongSeal (Pham et al., 2025). Table 2 presents the statistics of the evaluation data. Additionally, to further demonstrate the superior cross-domain robustness of BRIEF-PRO over previous compression methods, we refer readers to Table 6 in Appendix B.3 for results on three non-Wikipedia-based datasets.

Metrics **Exact match (EM)** and **F1** of answer strings were reported for QA performance. **Compression rate** was defined as the ratio of the number of words in the retrieved documents before compression to the number of words in the compressed summary after compression. A higher compression rate indicates a shorter summary.

Baselines We compared the BRIEF-PRO series with four main categories. (1) **Long-context LLMs** including: • **FILM-7B** (An et al., 2024). • **ProLong-8B** (Gao et al., 2025). They present training recipes to enhance the long-context capabilities. (2) **Non-compression** denotes prepending the full retrieved documents without compression. (3) **Extractive compression** methods including: • **RECOMP (Extractive)** (Xu et al., 2024a) ranks sentences based on whether it is useful as input

³gpt-4.1-nano-2025-04-14 for its favorable cost.

for LM. • **EXIT** (Hwang et al., 2025) classifies sentence-level relevance with lightweight single-token predictions, and reassembles only the high-relevance sentences in their original order. • **Rerank Top- k** denotes reranking the set of retrieved documents and keeping only the top- k documents. (4) **Abstractive compression** methods including: • **RECOMP (Abstractive)** (Xu et al., 2024a) distills the summarization knowledge of proprietary LLMs (gpt-3.5-turbo) into an abstractive compressor T5-large. • **BRIEF** (Li et al., 2025) enhances compression necessitating multi-hop reasoning by curating synthetic data. • **Llama-3.1-3B-Instruct** denotes the off-the-shelf official release without further fine-tuning. • **LongLLMLingua** (Jiang et al., 2024) performs both demonstration-level and token-level compression, leveraging their perplexity calculated by a causal LM Llama-2-7B-Chat. • **GPT-4.1-nano** was prompted to summarize the documents with respect to the question. • **BRIEF-PRO-AUTO_{L7C}** denotes that BRIEF-PRO-AUTO is initialized from **Llama2-7B-Chat**, allowing a fair head-to-head comparison with LongLLMLingua under the same base model. Note that the context length of Llama-2-7B-Chat is 4K tokens, which is shorter than the average length of our training data. Consequently, we truncate the distractor documents in our training data, which substantially degrades the performance of BRIEF-PRO-AUTO_{L7C} relative to BRIEF-PRO-AUTO, while still outperforming LongLLMLingua. Readers may refer to Appendix B.1 for more baseline details.

4.2 Experimental Results

Table 3 presents the evaluation results on four multi-hop QA tasks with three different reader models. BRIEF-PRO demonstrates promising multi-hop performance in both QA and document compression. Compared to non-compression, BRIEF-PRO-AUTO achieves an average compression rate of 32x, while still outperforming it by 6.70%, 0.60%, and 7.27% across three different reader models, respectively. These results indicate that existing LLMs still struggle with the demands of heavy long-context understanding, while a lightweight, highly specialized long-context compressor can help identify relevant information and alleviate the cognitive burden on reader models, enabling fast and accurate multi-hop reasoning. Although the curated training data for FILM-7B and ProLong-8B improves their long-context capabilities, their performance still underperforms BRIEF-PRO sig-

Category	Method	MuSiqQue			HotpotQA			2WikiMultiHopQA			LongSeal			Average	
		EM	F1	Rate	EM	F1	Rate	EM	F1	Rate	EM	F1	Rate	QA	Rate
Long-context LLMs	ProLong-8B	19.50	27.42	1x	42.50	54.43	1x	36.00	42.94	1x	7.09	11.71	1x	30.20	1x
	FILM-7B	26.50	36.45	1x	47.50	61.43	1x	38.50	46.14	1x	9.84	15.10	1x	35.18	1x
Llama-3.1-8B-Instruct															
Non-compression		20.50	28.88	1x	40.00	54.29	1x	37.00	46.63	1x	12.60	16.81	1x	32.09	1x
Extractive	RECOMP (Extractive)	16.00	23.39	41x	35.00	48.48	36x	27.50	35.40	20x	4.72	8.18	35x	24.83	33x
	EXIT	11.00	17.19	35x	36.50	48.62	32x	36.00	41.83	16x	7.09	10.22	20x	26.06	26x
	Rerank Top-1	9.50	14.58	7x	23.50	35.66	8x	13.00	20.75	11x	6.30	9.08	15x	16.55	10x
	Rerank Top-3	15.50	23.40	2.5x	37.50	50.07	2.2x	36.50	46.14	3.0x	7.87	12.19	4.2x	28.65	3x
	Rerank Top-5	20.50	28.03	1.7x	42.00	55.16	1.5x	43.00	52.74	1.8x	9.06	12.42	2.7x	32.86	2x
Abstractive	BRIEF	5.00	12.39	22x	18.00	27.70	59x	21.00	26.38	28x	5.91	10.02	19x	15.80	32x
	RECOMP (Abstractive)	17.50	25.92	10x	28.50	40.59	10x	19.50	26.60	9x	5.51	9.88	9x	21.75	10x
	Llama-3.2-3B-Instruct	13.00	18.48	48x	30.00	41.67	41x	34.50	43.29	39x	7.48	12.03	58x	25.06	47x
	LongLLMLingua	20.50	28.75	10x	38.50	53.63	8x	43.00	50.24	5x	8.27	13.26	14x	32.02	9x
	GPT-4.1-nano	28.50	40.33	128x	38.00	53.67	90x	47.50	60.34	108x	<u>10.63</u>	<u>15.95</u>	116x	<u>36.87</u>	110x
	BRIEF-PRO-AUTO _{L7C}	22.50	32.44	13x	<u>42.50</u>	<u>56.17</u>	13x	43.00	51.55	14x	6.12	10.59	11x	33.11	12x
	BRIEF-PRO-AUTO	<u>27.50</u>	<u>36.64</u>	35x	49.00	63.05	34x	47.50	<u>56.00</u>	23x	12.99	17.62	36x	38.79	32x
	BRIEF-PRO-HIGH	24.00	31.70	84x	47.00	58.85	72x	40.00	47.79	45x	9.06	13.73	72x	34.02	68x
	BRIEF-PRO-MEDIUM	27.50	35.55	47x	51.00	63.41	42x	48.50	57.21	26x	10.24	15.67	52x	38.64	42x
	BRIEF-PRO-LOW	30.50	39.68	26x	50.50	64.36	26x	49.00	58.39	15x	11.42	16.63	31x	40.06	25x
Llama-3.1-70B-Instruct															
Non-compression		36.00	45.47	1x	51.50	65.14	1x	59.50	66.85	1x	15.35	20.05	1x	44.98	1x
Extractive	RECOMP (Extractive)	29.50	37.65	41x	40.00	53.27	36x	40.50	47.46	20x	5.91	9.43	35x	32.97	33x
	EXIT	23.00	30.85	35x	46.00	59.80	32x	52.00	58.93	16x	7.87	11.19	20x	36.21	26x
	Rerank Top-1	19.00	26.50	7x	32.00	44.42	8x	21.00	27.14	11x	6.69	9.23	15x	23.25	10x
	Rerank Top-3	26.00	35.32	2.5x	43.00	55.78	2.2x	49.50	58.75	3.0x	10.63	13.21	4.2x	36.52	3x
	Rerank Top-5	30.00	39.06	1.7x	50.00	64.45	1.5x	57.50	65.78	1.8x	11.02	14.57	2.7x	41.55	2x
Abstractive	BRIEF	11.50	18.74	22x	25.00	36.03	59x	31.00	36.53	28x	7.09	11.13	19x	22.13	32x
	RECOMP (Abstractive)	25.00	33.80	10x	37.50	50.47	10x	34.50	41.56	9x	7.09	11.10	9x	30.13	9.5x
	Llama-3.2-3B-Instruct	24.50	30.66	48x	34.50	47.26	41x	42.00	50.12	39x	9.45	12.48	58x	31.37	47x
	GPT-4.1-nano	32.00	41.46	128x	42.00	56.59	90x	51.00	64.03	108x	<u>10.63</u>	<u>16.55</u>	116x	39.28	110x
	LongLLMLingua	32.00	42.35	10x	46.00	<u>61.32</u>	8x	<u>55.50</u>	<u>64.96</u>	5x	10.24	14.87	14x	40.91	9x
	BRIEF-PRO-AUTO _{L7C}	<u>36.00</u>	<u>46.40</u>	13x	<u>47.50</u>	60.97	13x	54.50	62.68	14x	9.06	12.68	11x	<u>41.22</u>	12x
	BRIEF-PRO-AUTO	38.50	48.84	35x	53.00	67.29	34x	59.50	66.94	23x	12.99	17.59	36x	45.58	32x
	BRIEF-PRO-HIGH	38.00	45.55	84x	55.00	67.42	72x	54.00	60.50	45x	9.45	14.18	72x	43.01	68x
	BRIEF-PRO-MEDIUM	38.00	49.17	47x	54.50	68.02	42x	58.50	67.38	26x	10.63	16.41	52x	45.33	42x
	BRIEF-PRO-LOW	40.00	51.95	26x	54.50	68.35	26x	59.50	68.31	15x	12.60	16.72	31x	46.49	25x
GPT-4.1-nano															
Non-compression		24.00	33.08	1x	42.50	56.56	1x	39.00	46.04	1x	11.02	16.02	1x	33.53	1x
Extractive	RECOMP (Extractive)	19.50	27.32	41x	41.00	52.44	36x	38.00	44.81	20x	3.54	7.06	35x	29.21	33x
	EXIT	16.50	24.33	35x	42.00	54.41	32x	43.00	47.58	16x	6.69	10.65	20x	30.65	26x
	Rerank Top-1	11.00	17.99	7x	30.00	40.85	8x	23.50	27.73	11x	3.94	6.55	15x	20.20	10x
	Rerank Top-3	22.50	30.73	2.5x	40.00	52.85	2.2x	45.50	52.83	3.0x	7.87	11.83	4.2x	33.01	3x
	Rerank Top-5	21.00	29.02	1.7x	41.00	54.87	1.5x	46.00	54.71	1.8x	9.45	13.99	2.7x	33.76	2x
Abstractive	RECOMP (Abstractive)	13.00	19.13	10x	27.00	39.21	10x	25.50	30.26	9x	3.54	8.11	9x	20.72	9.5x
	BRIEF	6.00	13.55	22x	25.50	34.75	59x	30.00	35.92	28x	8.27	12.30	19x	20.79	32x
	Llama-3.2-3B-Instruct	17.50	23.60	48x	33.00	44.72	41x	38.50	47.10	39x	8.27	11.76	58x	28.06	47x
	LongLLMLingua	20.50	28.55	10x	40.50	54.90	8x	48.00	54.52	5x	5.91	11.35	14x	33.03	9x
	GPT-4.1-nano	31.50	41.24	128x	40.50	55.87	90x	51.50	63.79	108x	<u>11.02</u>	<u>16.51</u>	116x	<u>38.99</u>	110x
	BRIEF-PRO-AUTO _{L7C}	20.50	31.58	13x	<u>46.00</u>	<u>56.06</u>	13x	46.50	52.81	14x	7.87	11.35	11x	34.08	12x
	BRIEF-PRO-AUTO	<u>29.50</u>	<u>41.10</u>	35x	51.50	65.52	34x	<u>51.00</u>	<u>58.68</u>	23x	12.20	16.87	36x	40.80	32x
	BRIEF-PRO-HIGH	31.50	38.70	84x	51.00	64.13	72x	49.00	56.48	45x	11.02	14.91	72x	39.59	68x
	BRIEF-PRO-MEDIUM	30.50	41.15	47x	52.00	65.38	42x	53.50	62.08	26x	11.02	15.31	52x	41.37	42x
	BRIEF-PRO-LOW	33.00	42.96	26x	51.00	64.86	26x	52.00	60.22	15x	11.42	15.76	31x	41.40	25x

Table 3: Evaluation results on four multi-hop QA tasks with small, large, and proprietary LMs as the \mathcal{M} , respectively. **Bold** and underscore denote the best and second-best QA performance, respectively, under the model’s self-determined compression setting (i.e., the AUTO mode).

nificantly while also requiring more computational resources. Compared to LongLLMLingua, BRIEF-PRO-AUTO compresses by higher 32x than its 9x on average, while still outperforming it by 6.77%, 4.67%, and 7.77%, respectively. Compared to using the proprietary GPT-4.1-nano as the compressor, it tends to significantly compress

long contexts which results in suboptimal QA performance. For the BRIEF-PRO series, the three compression levels offer varying granularities in preserving key semantic information, as guided by the user-specified instruction. This balance of efficiency and accuracy demonstrates its robustness and versatility across diverse scenarios.

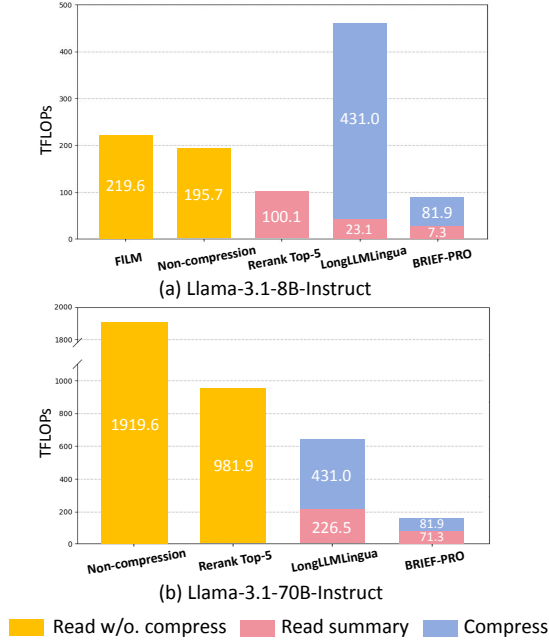


Figure 3: The comparison of TFLOPs consumption using (a) Llama-3.1-8B-Instruct and (b) Llama-3.1-70B-Instruct as the reader model.

4.3 Analysis

The improvement of latency in terms of overall computational overhead Figure 3 presents a comparison of TFLOPs consumption for processing long contexts. The profiler provided by Accelerate to count flops was adopted⁴. Specifically, when BRIEF-PRO is adopted for compression, the overall required TFLOPs are significantly lower than those needed for the original, uncompressed long contexts. The total amount of computation is reduced to 45% and 8% of what it was before compression using Llama-3.1-8B-Instruct and Llama-3.1-70B-Instruct as the LM \mathcal{M} , respectively. Compared to LongLLMLingua which uses Llama-2-7B-Chat to compute sentence and token perplexity, BRIEF-PRO consumes less than 20% and 24% of LongLLMLingua’s resources, respectively, while delivering better performance. LongLLMLingua incurs higher computational costs with the 8B reader model, because it has to divide intermediate results into segments and apply token-level compression iteratively, where each token’s perplexity based on preceding compressed segments. This substantial reduction in TFLOPs highlights BRIEF-PRO’s potential to optimize inference, especially for large-scale long context and with larger reader models, by enabling reader

⁴https://huggingface.co/docs/accelerate/usage_guides/profiler

Training Data	Length	Average QA Rate	
Llama-3.1-8B-Instruct			
Oracle++ & Distractor++	6.0k	38.79	32x
Oracle+ & Distractor+	3.6k	36.02	34x
Oracle+++	3.6k	33.76	35x
Llama-3.1-70B-Instruct			
Oracle++ & Distractor++	6.0k	45.58	32x
Oracle+ & Distractor+	3.6k	41.74	34x
Oracle+++	3.6k	41.68	35x
GPT-4.1-nano			
Oracle++ & Distractor++	6.0k	40.80	32x
Oracle+ & Distractor+	3.6k	39.11	34x
Oracle+++	3.6k	37.03	35x

Table 4: Evaluation results, averaged over four test sets, of comparing with exhaustively expanding only oracle documents. The number of + denotes the extent of expansion. Appendix B.3 presents the full results.

models to focus on compressed, more relevant information without sacrificing accuracy. We also analyzed end-to-end latency by summing the execution times of all pipeline components. Notably, with the 70B model, the overall latency was reduced to 14%, 32%, and 7% of that of Non-Compression, Rerank Top-5, and LongLLMLingua, respectively (see Appendix B.3 for details).

The comparison with exhaustively expanding only oracle documents In real-world scenarios, detailed knowledge about a fact may already be available. This raises the question: what would be the impact of *directly using this ready-made knowledge* to construct a long context? To demonstrate the effectiveness of expanding both oracle and distractor documents to form the input long context, our approach is compared against a strategy that performs exhaustive expansion solely on oracle documents, utilizing complete Wikipedia pages. The experimental results shown in Table 4 demonstrate a significant performance degradation when only oracle documents are expanded. This suggests that expanding only oracle documents might lead to a somewhat artificially "clean" context, potentially overestimating the model’s ability to handle complex, noisy inputs. In contrast, incorporating the expansion of distractor documents provides essential contextual diversity

Compression Mode	Expected	Average
HIGH	5	6.2
MEDIUM	10	10.4
LOW	20	18.0

Table 5: Average sentence counts in the generated summary in various compression modes.

and better reflects realistic long-context scenarios, where relevant and irrelevant information is often interspersed. On the other hand, our method is statistically able to synthesize significantly longer contexts (Avg. 6.0k vs. 3.6k words). The ability to process and learn from these substantially longer, noisy contexts directly contributes to the observed performance gains.

The accuracy of user-controllable instructions in terms of target sentence count The accuracy of the instruction hinges on how precisely the system can adhere to the user-specified sentence count. While the intent is to provide a flexible and intuitive mechanism for controlling summary granularity, the actual accuracy of this control can vary. Table 5 presents the average sentence counts in the generated summary in various compression modes across all four test sets. Appendix B.3 presents the detailed distribution. Although fitting the summary perfectly within the specified sentence limit is challenging, the results show that BRIEF-PRO performs well in following the HIGH and MEDIUM compression instructions. The reason lies in the sufficient training data for the target summaries within this length range. Meanwhile, achieving precise control while maintaining high summarization quality remains a technical challenge and should continue to be an active area of research.

5 Conclusion

This work introduces BRIEF-PRO, a universal, lightweight context compressor tailored for long contexts of 10k+ words across diverse scenarios, enabling fast and accurate multi-hop reasoning with RAG. BRIEF-PRO is trained on synthetic data generated through a pipeline built from short-context seed data, which synthesizes long-context training examples for compression learning. This pipeline provides a data-centric approach that supports user-controllable compression over output summary length. Experimental results demonstrate that the denoised summaries produced by BRIEF-

PRO substantially reduce computational overhead while improving QA accuracy across a wide range of small, large, and proprietary language models compared to previous compression methods.

Limitations

While BRIEF-PRO demonstrates promising results in long-context compression for RAG, several potential limitations warrant consideration. First, the model’s ability to effectively abstract and compress information from significantly longer (e.g., more than 20k words) and more complex inputs than those seen during training could be constrained. It could potentially lead to a performance degradation on contexts vastly exceeding the training data’s length or complexity. Second, although BRIEF-PRO is highly effective within the RAG framework, its performance in other long-context applications, such as few-shot learning, code completion, or long-dialogue history understanding, remains untested and could be suboptimal. Take code completion as an example, the task demands strict adherence to syntax, accurate tracking of variable definitions, precise function signatures, and coherent logical structure. A compressor optimized for natural language evidence may struggle to maintain the exactness and completeness required for reliable code generation. Without a comprehensive evaluation across these diverse long-context tasks, the generalizability and efficacy of BRIEF-PRO beyond its RAG-specific domain remain an open question.

Acknowledgement

We would like to express gratitude to the UCLANLP group members for their valuable feedback.

References

- Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, Jian-Guang Lou, and Weizhu Chen. 2024. [Make your LLM fully utilize the context](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Anthropic. 2024. [Introducing Claude 3.5 Sonnet](#).
- Yushi Bai, Xin Lv, Jiajie Zhang, Yuze He, Ji Qi, Lei Hou, Jie Tang, Yuxiao Dong, and Juanzi Li. 2024a. [Longalign: A recipe for long context alignment of large language models](#). In *Findings of the*

- Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, pages 1376–1395. Association for Computational Linguistics.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024b. [Longbench: A bilingual, multitask benchmark for long context understanding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 3119–3137. Association for Computational Linguistics.
- Alexis Chevalier, Alexander Wettig, Anirudh Ajith, and Danqi Chen. 2023. [Adapting language models to compress contexts](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 3829–3846. Association for Computational Linguistics.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. [A dataset of information-seeking questions and answers anchored in research papers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610, Online. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. 2025. [How to train long-context language models \(effectively\)](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Google. 2025. [Start Building with Gemini 2.5 Flash](#).
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing A multi-hop QA dataset for comprehensive evaluation of reasoning steps](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 6609–6625. International Committee on Computational Linguistics.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Taeho Hwang, Sukmin Cho, Soyeong Jeong, Hoyun Song, SeungYoon Han, and Jong C. Park. 2025. [EXIT: context-aware extractive compression for enhancing retrieval-augmented generation](#). In *Findings of the Association for Computational Linguistics, ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 4895–4924. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Trans. Mach. Learn. Res.*, 2022.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Comput. Surv.*, 55(12):248:1–248:38.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023a. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023b. [Llmllingua: Compressing prompts for accelerated inference of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 13358–13376. Association for Computational Linguistics.
- Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2024. [Longllmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 1658–1677. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6769–6781. Association for Computational Linguistics.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in*

- Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.*
- Yuankai Li, Jia-Chen Gu, Di Wu, Kai-Wei Chang, and Nanyun Peng. 2025. [BRIEF: bridging retrieval and inference for multi-hop reasoning via compression](#). In *Findings of the Association for Computational Linguistics: NAACL 2025, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, pages 5449–5470. Association for Computational Linguistics.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Trans. Assoc. Comput. Linguistics*, 12:157–173.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 9802–9822. Association for Computational Linguistics.
- Meta AI. 2024a. [Introducing Llama-3.1: Our Most Capable Models to Date](#). Blog post.
- Meta AI. 2024b. [Llama 3.2 community license release](#). Meta AI.
- Jesse Mu, Xiang Li, and Noah D. Goodman. 2023. [Learning to compress prompts with gist tokens](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- OpenAI. 2025. [Introducing GPT-4.1 in the API](#).
- Thin Pham, Nguyen Nguyen, Pratibha Zunjare, Weiyuan Chen, Yu-Min Tseng, and Tu Vu. 2025. [Sealqa: Raising the bar for reasoning in search-augmented language models](#). *CoRR*, abs/2506.01062.
- Minju Seo, Jinheon Baek, Seongyun Lee, and Sung Ju Hwang. 2025. [Efficient long context language model retrieval with compression](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15251–15268, Vienna, Austria. Association for Computational Linguistics.
- Rulin Shao, Jacqueline He, Akari Asai, Weijia Shi, Tim Dettmers, Sewon Min, Luke Zettlemoyer, and Pang Wei Koh. 2024. [Scaling retrieval-based language models with a trillion-token datastore](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, and Denny Zhou. 2023. [Large language models can be easily distracted by irrelevant context](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 31210–31227. PMLR.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [MuSiQue: Multihop questions via single-hop question composition](#). *Trans. Assoc. Comput. Linguistics*, 10:539–554.
- Eric Winglian and Axolotl Contributors. 2025. [Axolotl: Open-source llm fine-tuning library](#). Accessed: 2025-07-21.
- Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2024a. [RECOMP: improving retrieval-augmented lms with context compression and selective augmentation](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2024b. [Retrieval meets long context large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380. Association for Computational Linguistics.
- Chanwoong Yoon, Taewhoo Lee, Hyeon Hwang, Minbyul Jeong, and Jaewoo Kang. 2024. [Compact: Compressing retrieved documents actively for question answering](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 21424–21439. Association for Computational Linguistics.

A Prompts

The following are the prompts for the compression model and the reader model. For the reader model, we follow the prompt used in LongBench (Bai et al., 2024b).

Prompt for Compression Model (Auto)

Write a high-quality summary of the provided documents with respect to the question.
This is the question: {QUESTION}
These are the documents:
{DOCUMENTS}
This is the summary:

Prompt for Compression Model (User-controllable)

Write a high-quality summary of the provided documents with respect to the question.
This is the question: {QUESTION}
These are the documents:
{DOCUMENTS}
This is the summary:
Summarize the documents relevant to the question in K sentences, where K = [P] {LENGTH} [NP]

Prompt for Reader Model

Answer the question based on the given passages. Only give me the answer and do not output any other words.
The following are given passages.
{DOCUMENTS}
Answer the question based on the given passages. Only give me the answer and do not output any other words.
Question: {QUESTION}
Answer:

B Experimental Details

B.1 Baselines

We compared the BRIEF-PRO series with four main categories:

(1) **Long-context LLMs** including: • FILM-7B (An et al., 2024) initialized from Mistral-7B (Jiang et al., 2023a). • ProLong-8B (Gao et al., 2025) initialized from Llama-3-8B (Dubey et al., 2024). They present training recipes to enhance long-context capabilities. FILM leverages synthesized, information-intensive long-context QA data, while ProLong combines long-context code repositories and books with high-quality short-context data.

(2) **Non-compression** denotes prepending the full retrieved documents without compression.

(3) **Extractive compression** methods including: • RECOMP (Extractive) (Xu et al., 2024a) formulates extractive compression as a sentence ranking problem, and the sentence is evaluated based on whether it is useful as input for the LM. • EXIT (Hwang et al., 2025) splits a document into sentences, classifies sentence-level relevance with lightweight single-token predictions, and reassembles only the high-relevance sentences in their original order to preserve coherence and key information. • Rerank Top- k denotes reranking the set of retrieved documents using Contriever (Izacard et al., 2022) trained on MS MARCO dataset and keeping only the top- k documents.

(4) **Abstractive compression** methods including: • RECOMP (Abstractive) (Xu et al., 2024a) distills the summarization knowledge of proprietary LLMs (gpt-3.5-turbo) into an abstractive compressor T5-large. • BRIEF (Li et al., 2025) is a T5-based context compressor that can only handle a maximum sequence length of 512 tokens at a time. To accommodate this limitation, long contexts were uniformly divided into chunks of up to 512 tokens. Each chunk was then compressed using the trained compressor, and the compressed results of each chunk were concatenated to form the overall compressed summary. • Llama-3.1-3B-Instruct denotes the off-the-shelf official release without further fine-tuning. • LongLLMLingua (Jiang et al., 2024) uses LLMLingua (Jiang et al., 2023b) as the backbone and further improves its perception of key information pertinent to the question. We followed their 2,000 tokens compression constraint. It performs both coarse-grained, demonstration-level compression and fine-grained, token-level compression, leveraging the perplexity of each demonstration or token calculated by a causal LM Llama-2-7B-Chat. • GPT-4.1-nano was prompted to summarize the documents with respect to the question.

B.2 Implementation

The LoRA technique (Hu et al., 2022) was adopted for efficient compressor training on our curated data for three epochs. AdamW (Loshchilov and Hutter, 2019) was used as the optimizer, and the batch size was set to 64. We utilize the Axolotl library (Winglian and Contributors, 2025) for efficient parallel training. The entire training process takes about two days on $2 \times$ NVIDIA A100 80GB GPUs. The Llama models were accessed via HuggingFace^{5 6} and GPT-4.1 was accessed via OpenAI⁷.

⁵<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

⁶<https://huggingface.co/meta-llama/Llama-3.1-70B-Instruct>

⁷<https://platform.openai.com>

B.3 More Experimental Results and Analysis

Category	Method	NarrativeQA			Qasper			MultiFieldQA-en			Average	
		EM	F1	Rate	EM	F1	Rate	EM	F1	Rate	QA	Rate
Long-context LLMs	ProLong-8B	14	31.63	1x	18.5	28.78	1x	22	51.44	1x	27.73	1x
	FILM-7B	10	28.47	1x	21	45.25	1x	22	54.34	1x	30.18	1x
Llama-3.1-8B-Instruct												
Non-compression		7.50	29.43	1x	20.50	47.62	1x	21.33	52.77	1x	29.86	1x
Extractive	RECOMP (Extractive)	4.50	16.19	94x	13.50	33.14	16x	<u>17.33</u>	43.82	16x	21.41	42x
	EXIT	<u>5.50</u>	20.11	40x	13.00	24.28	15x	10.67	25.33	14x	16.48	23x
	Rerank Top-3	5.00	20.06	3x	14.00	39.31	3x	16.50	43.10	3x	23.00	3x
Abstractive	RECOMP (Abstractive)	3.50	13.69	17x	12.00	20.52	17x	8.00	30.14	12x	14.64	15x
	Llama-3.2-3B-Instruct	<u>5.50</u>	19.81	45x	11.50	29.28	12x	14.67	42.30	16x	20.51	24x
	LongLLMLingua	4.50	21.56	17x	<u>14.50</u>	<u>41.18</u>	3x	<u>17.33</u>	<u>44.62</u>	4x	<u>23.95</u>	8x
	GPT-4.1-nano	6.50	24.46	91x	13.50	35.89	19x	14.67	41.95	29x	22.83	46x
	BRIEF-PRO-AUTO	<u>5.50</u>	<u>21.73</u>	32x	16.50	42.70	7x	24.67	54.21	11x	27.55	17x
	BRIEF-PRO-HIGH	6.50	19.10	97x	14.50	36.77	22x	24.00	52.43	29x	25.55	49x
	BRIEF-PRO-MEDIUM	6.50	20.40	66x	14.50	40.95	12x	24.00	52.23	16x	26.43	31x
	BRIEF-PRO-LOW	7.00	21.84	40x	15.00	40.50	7x	24.00	53.88	10x	27.04	19x
Llama-3.1-70B-Instruct												
Non-compression		15.00	35.20	1x	21.00	48.28	1x	21.33	54.38	1x	32.53	1x
Extractive	RECOMP (Extractive)	6.00	16.62	94x	<u>18.50</u>	38.53	16x	<u>18.00</u>	<u>46.43</u>	16x	24.01	42x
	EXIT	8.00	19.70	40x	17.50	27.96	15x	10.67	25.75	14x	18.26	23x
	Rerank Top-3	<u>7.50</u>	21.62	3x	<u>18.50</u>	37.58	3x	17.33	45.10	3x	24.61	3x
Abstractive	RECOMP (Abstractive)	5.50	17.16	17x	12.00	22.86	17x	11.33	33.78	12x	17.11	15x
	Llama-3.2-3B-Instruct	7.00	19.95	45x	16.50	35.23	12x	12.80	39.31	16x	21.80	24x
	LongLLMLingua	8.00	21.09	17x	18.00	<u>45.26</u>	3x	16.67	45.87	4x	<u>25.82</u>	8x
	GPT-4.1-nano	<u>7.50</u>	22.91	91x	14.00	37.41	19x	13.33	41.44	29x	22.77	46x
	BRIEF-PRO-AUTO	8.00	<u>22.83</u>	32x	22.00	47.40	7x	20.67	51.16	11x	28.68	17x
	BRIEF-PRO-HIGH	7.00	18.65	97x	20.50	42.82	22x	21.33	50.70	29x	26.83	49x
	BRIEF-PRO-MEDIUM	6.50	20.02	66x	19.50	46.16	12x	20.67	51.93	16x	27.46	31x
	BRIEF-PRO-LOW	7.50	20.71	40x	19.00	46.37	7x	20.67	50.63	10x	27.48	19x
GPT-4.1-nano												
Non-compression		10.50	28.25	1x	20.50	44.47	1x	16.67	50.09	1x	28.41	1x
Extractive	RECOMP (Extractive)	5.00	15.62	94x	<u>17.50</u>	37.46	16x	12.67	43.72	16x	22.00	42x
	EXIT	<u>7.50</u>	18.45	40x	11.50	25.40	15x	9.33	27.12	14x	16.55	23x
	Rerank Top-3	<u>7.50</u>	19.10	3x	16.00	38.04	3x	13.33	42.67	3x	22.77	3x
Abstractive	RECOMP (Abstractive)	5.50	14.01	17x	8.00	18.32	17x	7.33	28.83	12x	13.67	15x
	Llama-3.2-3B-Instruct	5.50	17.15	45x	12.00	32.39	12x	12.67	41.55	16x	20.21	24x
	LongLLMLingua	4.00	18.24	17x	15.50	<u>39.53</u>	3x	<u>14.00</u>	<u>44.88</u>	4x	22.69	8x
	GPT-4.1-nano	<u>7.50</u>	24.50	91x	16.50	38.64	19x	12.67	41.04	29x	<u>23.48</u>	46x
	BRIEF-PRO-AUTO	8.50	<u>19.53</u>	32x	19.00	42.21	7x	16.67	48.17	11x	25.68	17x
	BRIEF-PRO-HIGH	7.50	18.88	97x	16.50	38.78	22x	18.00	48.53	29x	24.70	49x
	BRIEF-PRO-MEDIUM	7.50	19.25	66x	18.00	42.54	12x	16.67	49.79	16x	25.63	31x
	BRIEF-PRO-LOW	8.50	20.85	40x	17.00	41.87	7x	17.33	49.75	10x	25.88	19x

Table 6: Evaluation results on three non-Wikipedia-based QA datasets across diverse domains, using small, large, and proprietary LMs as the \mathcal{M} , respectively. **Bold** and underscore denote the best and second-best QA performance, respectively, among different compression methods under the model’s self-determined compression setting (i.e., the AUTO mode).

Table 6 presents the evaluation results on three non-Wikipedia-based QA datasets across diverse domains

using three different reader models. The selected datasets are the extended NarrativeQA (literature and film) (Kočiský et al., 2018), Qasper (science) (Dasigi et al., 2021), and MultiFieldQA-en (multiple fields) (Bai et al., 2024b) from LongBench (Bai et al., 2024b), with context lengths of 18.4k, 3.6k, and 4.5k words, respectively. BRIEF was not included because it was not designed for such diverse evaluations.

BRIEF-PRO still demonstrates promising performance in this evaluation. Compared to other compression methods, BRIEF-PRO-AUTO achieves the best QA performance. Notably, BRIEF-PRO-HIGH attains the highest compression rate while still outperforming other compression baselines. Nevertheless, all compression methods fall short of the non-compression baseline in QA performance on these datasets, highlighting the need for further advances in robust compression across diverse domains.

Training Data	Avg. Length	MuSiqQue			HotpotQA			2WikiMultiHopQA			LongSeal			Average	
		EM	F1	Rate	EM	F1	Rate	EM	F1	Rate	EM	F1	Rate	QA	Rate
Llama-3.1-8B-Instruct															
Oracle+ & Distractor+	6.0k	27.50	36.64	35x	49.00	63.05	34x	47.50	56.00	23x	12.99	17.62	36x	38.79	32x
Oracle+ & Distractor+	3.6k	27.00	33.84	36x	43.00	56.93	39x	43.50	50.94	23x	13.78	19.18	39x	36.02	34x
Oracle++	3.6k	22.50	29.71	42x	42.00	55.37	37x	39.50	48.01	23x	14.17	18.81	39x	33.76	35x
Llama-3.1-70B-Instruct															
Oracle+ & Distractor+	6.0k	38.50	48.84	35x	53.00	67.29	34x	59.50	66.94	23x	12.99	17.59	36x	45.58	32x
Oracle+ & Distractor+	3.6k	36.50	45.07	36x	47.00	60.97	39x	51.50	60.23	23x	13.78	18.90	39x	41.74	34x
Oracle++	3.6k	32.50	41.64	42x	48.50	61.67	37x	52.50	62.33	23x	14.57	19.73	39x	41.68	35x
GPT-4.1-nano															
Oracle+ & Distractor+	6.0k	29.50	41.10	35x	51.50	65.52	34x	51.00	58.68	23x	12.20	16.87	36x	40.80	32x
Oracle+ & Distractor+	3.6k	32.00	39.00	36x	47.50	60.79	39x	50.00	57.24	23x	11.02	15.35	39x	39.11	34x
Oracle++	3.6k	28.00	34.54	42x	46.00	58.95	37x	47.50	55.46	23x	10.24	15.55	39x	37.03	35x

Table 7: Evaluation results of comparing with exhaustively expanding only oracle documents.




Figure 4: The distribution of sentence counts in the generated summary over four test sets in various compression modes.




Figure 5: The performance of compressors under different context length. We expand the scope of retrieved documents from the top 20 to the top 100 based on the validation set in the Musique dataset (Trivedi et al., 2022). For BRIEF-PRO, we follow the AUTO setting. For LongLLMLingua (Jiang et al., 2024), we follow their 2,000-token compression constraint. The reported score is the average of the F1 and EM scores.




Figure 6: The comparison of end-to-end latency using (a) Llama-3.1-8B-Instruct and (b) Llama-3.1-70B-Instruct as the reader model. For each method, we computed the total running time by summing the execution times of all components in the pipeline. Experimental results demonstrate that BRIEF-PRO consistently reduces the overall end-to-end latency. Notably, on the 70B model, latency is reduced to only 14%, 32%, and 7% of that of No-Compression, Rerank Top-5, and LongLLMLingua, respectively.

B.4 Qualitative Examples

We provide qualitative examples of correct and incorrect compression cases. Sentences highlighted in blue indicate key information for answering the question.

In the correct case, this is a two-hop question: the key information is distributed across two passages in an extremely long context, requiring the compression model to identify sub-questions and precisely locate the relevant sentences. BRIEF-PRO performs well at locating key information in such long contexts, thanks to the high-quality training data synthesized by our method, which is highly representative of this setting.

In the incorrect case, BRIEF-PRO shows a typical mistake in context compression for multi-hop questions: it fails to capture relevant information in one of the hops. This problem is common across existing methods and becomes more severe with long contexts, where locating the relevant information is much harder. It reflects a broader challenge for the field and warrants further investigation.

Correct Case (10343 words -> 282 words | Compression rate: 36x)

- **Question:** Robbie Tucker plays in what series that follows a group of friends who run an Irish bar?

- **Context Before Compression (10343 words):**

Passage 1: John Franks (judge) Sir John Franks (1770–1852), was an Indian judge. Franks was the second son of Thomas Franks (1729–1787), of Ballymagooly, County Cork, by Catherine, daughter of Rev. John Day. He was born in 1770, and graduated at Trinity College, Dublin, B.A. 1788, LL.B. 1791. He was called to the Irish Bar 1792. He went on the Munster circuit, and had a good practice as chamber counsel.

...

Robbie Tucker (born April 5, 2001) is an American actor. His best known role to date is that of Fenmore Baldwin on the CBS soap opera The Young and the Restless. **Tucker has also starred on other series, such as Criminal Minds, FlashForward and It's Always Sunny in Philadelphia.** He has also appeared in the films Prom and Little Fockers. In 2012, Tucker was nominated at the 33rd Young Artist Awards for his performance in Prom and won for his role in The Young and the Restless. He is also the brother of actress Jillian Rose Reed.

Filmography

...

It's Always Sunny in Philadelphia is an American sitcom created by Rob McElhenney and developed with Glenn Howerton for FX. It premiered on August 4, 2005, and was moved to FXX beginning with the ninth season in 2013. It stars Charlie Day, Howerton, McElhenney, Kaitlin Olson, and Danny DeVito. The series follows the exploits of "The Gang", **a group of narcissistic and sociopathic friends who run the Irish dive bar Paddy's Pub in South Philadelphia,** Pennsylvania, but spend most of their free time drinking, scheming, arguing amongst themselves, and plotting elaborate cons against others, and at times each other, for personal benefit, financial gain, revenge, or simply out of boredom. With its 15th season, the series became the longest-running live-action American comedy series, replacing The Adventures of Ozzie and Harriet, which ran for 14 seasons (435 episodes) between 1952 and 1966.

...

Podcast On November 9, 2021, Howerton, Day, and McElhenney started The Always Sunny Podcast, an episode-by-episode-recap of the show, with Megan Ganz as producer. Occasional guest stars include Kaitlin Olson, David Hornsby, Mary Elizabeth Ellis, Michael Naughton, Lin-Manuel Miranda, Cormac Bluestone and Danny DeVito.

• **Context After Compression (282 words):**

Robbie Tucker (born April 5, 2001) is an American actor. His best known role to date is that of Fenmore Baldwin on the CBS soap opera *The Young and the Restless*. **Tucker has also starred on other series, such as *Criminal Minds*, *FlashForward* and *It's Always Sunny in Philadelphia*.** He has also appeared in the films *Prom* and *Little Fockers*. In 2012, Tucker was nominated at the 33rd Young Artist Awards for his performance in *Prom* and won for his role in *The Young and the Restless*. He is also the brother of actress Jillian Rose Reed.

***It's Always Sunny in Philadelphia* is an American sitcom** created by Rob McElhenney and developed with Glenn Howerton for FX. It premiered on August 4, 2005, and was moved to FXX beginning with the ninth season in 2013. It stars Charlie Day, Howerton, McElhenney, Kaitlin Olson, and Danny DeVito. The series follows the exploits of "The Gang", **a group of narcissistic and sociopathic friends who run the Irish dive bar Paddy's Pub in South Philadelphia**, Pennsylvania, but spend most of their free time drinking, scheming, arguing amongst themselves, and plotting elaborate cons against others, and at times each other, for personal benefit, financial gain, revenge, or simply out of boredom. With its 15th season, the series became the longest-running live-action American comedy series, replacing *The Adventures of Ozzie and Harriet*, which ran for 14 seasons (435 episodes) between 1952 and 1966. In December 2020, the series was renewed for a total of four additional seasons, bringing it to 18 seasons. The 16th season premiered on June 7, 2023. The show has received critical acclaim, with many lauding the cast performances and dark humor. It has amassed a large cult following.

- **Answer:** It's Always Sunny in Philadelphia (✓)
- **Ground Truth Answer:** It's Always Sunny in Philadelphia

Incorrect Case (11179 words -> 140 words | Compression rate: 80x)

• **Question:** Roger Stuart Woolhouse is a biographer of a philosopher commonly known as what?

• **Context Before Compression (11179 words):**

Passage 1: Philip Carlo Philip Carlo (April 18, 1949 – November 8, 2010) was an American journalist and best selling biographer of Thomas Pitera, Richard Kuklinski, Anthony Casso, and Richard Ramirez. Carlo had amyotrophic lateral sclerosis (ALS), commonly known as "Lou Gehrig's Disease".

...

John Locke (; 29 August 1632 – 28 October 1704) was an English philosopher and physician, widely regarded as one of the most influential of Enlightenment thinkers and **commonly known as the "father of liberalism"**. Considered one of the first of the British empiricists, following the tradition of Francis Bacon, Locke is equally important to social contract theory. His work greatly affected the development of epistemology and political philosophy. His writings influenced Voltaire and Jean-Jacques Rousseau, and many Scottish Enlightenment thinkers, as well as the American Revolutionaries. His contributions to classical republicanism and liberal theory are reflected in the United States Declaration of Independence. Internationally, Locke's political-legal principles continue to have a profound influence on the theory and practice of limited representative government and the protection of basic rights and freedoms under the rule of law. Locke's theory of mind is often cited as the origin of modern conceptions of identity and the self, figuring prominently in the work of later philosophers such as Jean-Jacques Rousseau, David Hume, and Immanuel Kant.

...

Roger Stuart Woolhouse (1940–2011) was an English philosopher, an expert on empiricism and rationalism and a biographer of John Locke. He was born in Wath-upon-Dearne and educated at Saltburn Primary School, Sir William Turner’s Grammar School, London University (Philosophy) and then Selwyn College, Cambridge for his Doctorate. From 1969 until his retirement in 2001, Woolhouse worked in the Department of Philosophy at the University of York. Cambridge University Press requested Woolhouse write a biography of Locke, the last major biography being Maurice Cranston’s 1957 work. Woolhouse’s biography appeared in 2007. After his death, York’s Department of Philosophy founded the Roger Woolhouse Prize, an annual £500 prize awarded to MA Philosophy students.

...

Descartes, Spinoza, Leibniz: The Concept of Substance in Seventeenth Century Metaphysics (Routledge, 1993). Locke: A Biography (Cambridge University Press, 2007). Starting with Leibniz (Continuum, 2010).

• **Context After Compression (140 words):**

Roger Stuart Woolhouse (1940–2011) was an English philosopher, an expert on empiricism and rationalism and a biographer of John Locke. He was born in Wath-upon-Dearne and educated at Saltburn Primary School, Sir William Turner’s Grammar School, London University (Philosophy) and then Selwyn College, Cambridge for his Doctorate. From 1969 until his retirement in 2001, Woolhouse worked in the Department of Philosophy at the University of York. Cambridge University Press requested Woolhouse write a biography of Locke, the last major biography being Maurice Cranston’s 1957 work. Woolhouse’s biography appeared in 2007. After his death, York’s Department of Philosophy founded the Roger Woolhouse Prize, an annual £500 prize awarded to MA Philosophy students.

Works Locke (Prentice Hall / Harvester Wheatsheaf, 1984). The Empiricists (Oxford University Press, 1988). Descartes, Spinoza, Leibniz: The Concept of Substance in Seventeenth Century Metaphysics (Routledge, 1993). Locke: A Biography (Cambridge University Press, 2007). Starting with Leibniz (Continuum, 2010).

- **Answer:** John Locke (✗)
- **Ground Truth Answer:** Father of Liberalism

C Discussions

C.1 Hallucination Risk

We would like to raise the discussion about the risk of hallucinations in abstractive and extractive compression. While extractive methods avoid generating new tokens, they are not immune to hallucinations. In particular, by removing essential context, extractive compressors may yield misleading interpretations; by elevating marginal sentences, they may distort relevance; and by retaining unresolved references (e.g., pronouns, temporal markers), they may create spurious associations. Thus, both extractive and abstractive approaches carry inherent risks, though manifested in different ways. We acknowledge and share the ambition that mitigating hallucinations should be a key objective, regardless of whether compressors are abstractive or extractive.

C.2 Importance of Training a Long-context Compressor

The goal of long-context compression is to reduce inference cost without degrading downstream performance. While modern LLMs can perform aspect-based and length-controlled summarization, using them directly for compression increases inference cost and undermines this objective. Our focus

extends beyond absolute context length to include efficiency, cost, and accessibility. Specifically, our contribution is to train a lightweight compression model that achieves strong long-context compression at low inference cost while preserving task quality, which is a non-trivial objective that requires careful training design rather than simply invoking modern LLMs. Long-context compression can be positioned as a complementary approach to future long-context models, enhancing efficiency and reasoning robustness rather than merely extending the maximum token length.

Furthermore, our experimental results show that fine-tuning a compression model for these tasks is also non-trivial. The core challenge lies in constructing effective training data. Our preliminary experiments show that the definition of the target summary strongly influences both training difficulty and model performance. For example, we currently adopt the Head-Tail Iterative Pruning method because removing unhelpful sentences in the middle produces multiple discrete text segments rather than a coherent and continuous summary. While this approach can yield more concise outputs, it substantially increases training difficulty and results in suboptimal model performance. These findings reinforce our belief that the training method proposed in this work makes an important contribution to efficient compression model training.

C.3 Explanation of Head-Tail Assumption

We agree that critical information can be distributed across multiple sentences, and this does not contradict our implementation. For example, if important content appears at the beginning or end of a document, our proposed Head-Tail Iterative Pruning method will halt pruning and retain those segments. Our design aims to identify a more compact, continuous, and informative text span within each oracle document. By retaining a continuous central segment, our approach captures the majority of salient information while maintaining textual coherence. This design is motivated by our preliminary experiments showing that removing middle sentences often produces fragmented and disconnected spans, which significantly increases the difficulty of compression training. Pruning head and tail sentences thus offers a practical balance, preserving continuity while focusing on the most informative portion of the document.

C.4 Head-to-Head Comparison against GPT-4.1-nano

While GPT-4.1-nano shows strong performance across benchmarks, our focus with BRIEF-Pro is on democratizing LLM training and inference. BRIEF-Pro enables high-quality compression using smaller, open-weight models, making it accessible to researchers and practitioners who do not have access to proprietary or extremely large LLMs. Training a compressor, as required by BRIEF-Pro, is a one-time and extremely low cost that allows repeated efficient inference on downstream tasks without relying on expensive API calls or large-scale models. In this sense, BRIEF-Pro prioritizes accessibility, reproducibility, and computational efficiency, complementing rather than directly competing with very large, closed-weight LLMs like GPT-4.1.

C.5 Scale across Multiple Retrievals

BRIEF-Pro primarily focuses on single-retrieval compression, with key contributions including user-controlled compression and support for longer contexts. While our current experiments are limited to single-retrieval settings, BRIEF-Pro is modular and scalable, allowing each retrieval to be compressed individually before aggregation in multi-retrieval scenarios. We believe this approach can be seamlessly integrated into iterative retrieval pipelines, and exploring such multi-retrieval applications is an important direction for future work.

C.6 Clarification on the Definition of Sentence Usefulness and Model Dependence

In our definition of helpfulness, we used Llama-3-8B-Instruct to compute log-likelihoods. This model was not cherry-picked to maximize our final results. Instead, we simply adopted a commonly used open-source model without tuning this choice.

For each sentence, we recompute the log-likelihood of the correct answer after removing that sentence, and we label the sentence as unhelpful and remove it if and only if this removal makes the correct answer

more likely. In other words, the threshold is simply “remove if the log-likelihood increases,” and no extra hyperparameters are introduced or tuned.

We acknowledge that the likelihood results and the final constructed data might be model-specific, but the model is used only once offline to construct the training supervision. Our experiments show that using data constructed with this single, fixed model already yields consistent gains across multiple, diverse reader models, including black-box ones, suggesting that the induced notion of usefulness generalizes beyond the specific LM used for data synthesis.

C.7 Data Contamination Analysis

MuSiqQue	HotpotQA	2WikiMultiHopQA	LongSeal	NarrativeQA	Qasper	MultiFieldQA-en
0.095	0.086	0.029	0.008	0.0007	0.002	0.062

Table 8: Average per-sample maximum 5-gram overlap between each test set and the training corpus.

Since our training pipeline uses Wikipedia for data expansion, one potential concern is the possibility of overlap between the expanded training data and the test sets, which could lead to data leakage. We would like to clarify that no data leakage occurs during compressor training. The compressor is query-aware, and each training example consists of a document context, a query, and an oracle text (i.e., the compressor output). The queries and oracle texts in the training set do not overlap with those in the test set, preventing the compressor from observing identical query–oracle compression patterns during training.

To provide a more rigorous contamination analysis, we further quantify the potential overlap between document contexts in the training and test sets. For each test example, we compute its 5-gram overlap with every training example, take the maximum overlap score, and then report the average of these maxima for each dataset. The 5-gram overlap is computed as follows:

$$\text{Overlap}(\text{context}_{test}, \text{context}_{train}) = \frac{|\mathcal{G}_5(\text{context}_{test}) \cap \mathcal{G}_5(\text{context}_{train})|}{|\mathcal{G}_5(\text{context}_{test})|}, \quad (2)$$

where $\mathcal{G}_5(x)$ denotes the set of unique 5-grams (contiguous 5-word shingles) extracted from text x after normalization and tokenization.

Table 8 presents the 5-gram overlap results. Across all datasets, the average maximum 5-gram overlap is below 0.1, and for some datasets it is below 0.01, indicating only slight (or nearly no) contextual overlap between the training and test sets. Therefore, the training-test contamination in our expanded training data is negligible and unlikely to affect our results.