
Accelerating Inference of Retrieval-Augmented Generation via Sparse Context Selection

Yun Zhu¹, Jia-Chen Gu³, Caitlin Sikora², Ho Ko², Yinxiao Liu¹, Chu-Cheng Lin²,
Lei Shu¹, Liangchen Luo¹, Lei Meng², Bang Liu⁴, Jindong Chen¹

¹Google DeepMind

²Google

³University of California, Los Angeles

⁴Université de Montréal & Mila

{yunzhu, csikora, hoko, canoee, kitsing, leishu, luolc, leimeng, jdchen}@google.com

gujc@ucla.edu, bang.liu@umontreal.ca

Abstract

Large language models (LLMs) augmented with retrieval exhibit robust performance and extensive versatility by incorporating external contexts. However, the input length grows linearly in the number of retrieved documents, causing a dramatic increase in latency. In this paper, we propose a novel paradigm named Sparse RAG, which seeks to cut computation costs through sparsity. Specifically, Sparse RAG encodes retrieved documents in parallel, which eliminates latency introduced by long-range attention of retrieved documents. Then, LLMs selectively decode the output by only attending to highly relevant caches auto-regressively, which are chosen via prompting LLMs with special control tokens. It is notable that Sparse RAG combines the assessment of each individual document and the generation of the response into a single process. The designed sparse mechanism in a RAG system can facilitate the reduction of the number of documents loaded during decoding for accelerating the inference of the RAG system. Additionally, filtering out undesirable contexts enhances the model’s focus on relevant context, inherently improving its generation quality. Evaluation results of two datasets show that Sparse RAG can strike an optimal balance between generation quality and computational efficiency, demonstrating its generalizability across both short- and long-form generation tasks.

1 Introduction

Large language models (LLMs) have attracted increasing attention and exhibited impressive abilities to understand instructions and generate fluent outputs in natural language [4, 27, 35, 34]. Nevertheless, LLMs inevitably manifest hallucinations [13] due to their struggle with factual errors and inability to secure the accuracy of generated text solely by the parametric knowledge they encapsulate [39, 26]. Feeding the source of truth to LLMs in the format of retrieved context segments [29] alleviates this problem. The technique is widely known as Retrieval-Augmented Generation (RAG) [20, 21, 8].

Although the RAG framework is empirically shown to be effective, it can be expensive to scale up. This is because it requires prepending relevant documents retrieved from an external knowledge corpus to the queries [8]. As a result, the input length grows linearly in the number of documents, causing a dramatic increase in latency when using a standard Transformer whose the latency scales quadratically with the input length. Some prior works such as Fusion-in-Decoder (FiD) [12] and Parallel Context Windows (PCW) [28] have proposed to alleviate this issue. Yet these methods fail to strike an optimal balance between generation quality and computational efficiency. FiD was originally

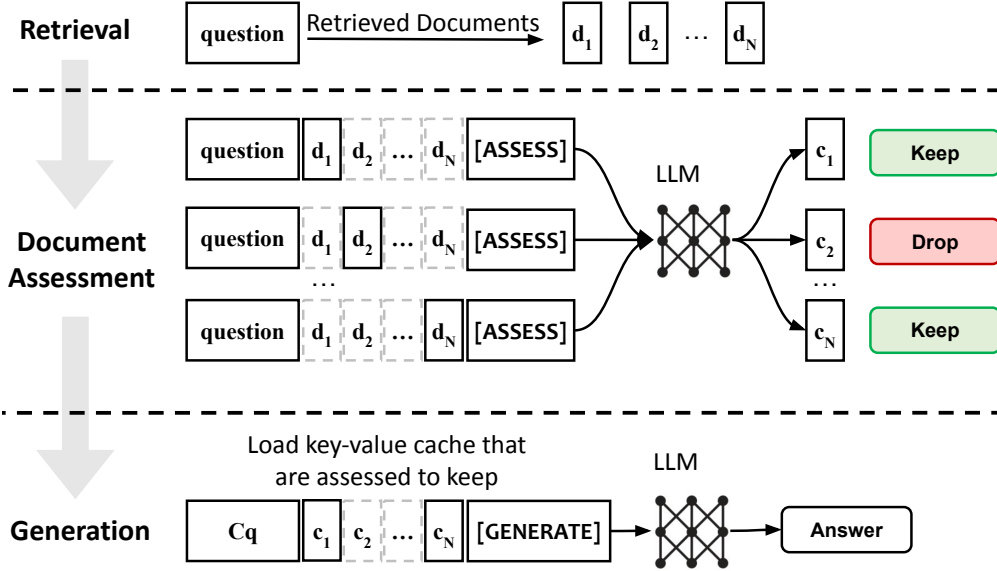


Figure 1: An overview of Sparse RAG at inference. Each of the retrieved documents will be assessed by LLMs to decide whether to keep or drop by estimating a relevance score. Then, load the documents that are considered to keep for generation.

designed for the encoder-decoder architecture, and thus is not compatible with currently prevalent decoder-only architectures without significant changes. While PCW can be applied to decoder-only LLMs, it only speeds up the model pre-filling and still incurs high latency since the whole context window cache is still being attended to when decoding each token. Moreover, the heavy reliance of generation on the retrieved knowledge raises significant concerns about the model’s behavior and performance in scenarios where retrieval may fail or return inaccurate results [33]. A typical approach for mitigating this issue is to rely on an external classifier to rank or filter the documents before prepending them to the input [36], but this process requires extra model calls which adds new complexity to inference.

In light of the issues above, we propose a novel paradigm named Sparse RAG. It generally operationalizes by massive pre-filling where the key-value cache was generated through a single forward pass of the input tokens and selective decoding where the output is generated by attending to highly relevant tokens auto-regressively. Previous works where the length of the retrieved contexts during pre-filling are equal to that during decoding are called *dense*-RAG in this paper. Different from that, Sparse RAG allows the pre-filled context to be significantly larger than the decoding context, where partial low-quality context is dynamically dropped according to their relevance to the input query. Furthermore, Sparse RAG combines the assessment of each individual context and the generation of the response into a single process. Special control tokens are used to prompt LLMs to assess each retrieved context for picking high-quality ones. Then only the key-value caches of those picked contexts are loaded into LLMs for decoding with another control tokens.

The design of such a sparse mechanism in RAG system brings two unique advantages. Firstly, by reducing the number of key-value cache loads during the decoding process, LLMs can achieve lower latency as they are typically constrained by memory usage. Besides, filtering out undesirable contexts enhances the model’s focus on relevant context, inherently improving its generation quality. To demonstrate the effectiveness and efficiency of the proposed method, we evaluate on two datasets of PopQA [24], and Biography [25]. Experimental results show that Sparse RAG can achieve the similar or better performance while keep much better latency compared with standard dense-RAG or PCW-RAG, demonstrating its generalizability across both short- and long-form generation tasks.

2 Related Work

Retrieval-Augmented Generation RAG is a family of techniques for generating output structures while using retrieved nearest-neighbor structures as a reference. It typically involves two stages:

Table 1: Comparisons with existing RAG-related works.

Approach	Corrective	No extra model	Prefill efficiency	Decode efficiency
RAG [19]	No	Yes	No	No
Corrective RAG [36]	Yes	No	No	No
PCW RAG [28]	No	Yes	Yes	No
Sparse RAG	Yes	Yes	Yes	Yes

retrieval and generation. Retrieval finds most similar contexts based on BM25 or learned embedding, where the context could be represented as token embedding [16, 37], dense embedding [5] or raw text [9, 12, 19]. Once those contexts are retrieved, different architectures are leveraged to incorporate them into the model. Popular approaches includes concatenation [12, 19] and cross-attention [3, 18]. In recent years, the LLM architecture evolves towards a single decoder only model with significant larger sizes. To this end, concatenation of raw text [19] is becoming popular for being simple and practical. Many advanced approaches have been developed on top of it in recent years. Self-RAG [2] is proposed to selectively retrieve knowledge and introduce a critic model to decide whether to retrieve. Yoran et al. [38] designed an NLI model to identify the irrelevant context and improve robustness. SAIL [23] is tuned on instructions to insert retrieved documents before instructions. Toolformer [30] is pre-trained for calling APIs such as Wikipedia. Jiang et al. [14] actively anticipate future content and decide when and what to retrieve in long-form generation. CRAG [36] makes the attempt to explore and design corrective strategies for RAG to improve its robustness of generation. Specifically, an external T5 model is trained and used to determine the usefulness of the retrieved context. Generally, these approaches target on exploiting retrieval as a useful tool to augment generation or whether retrieval is necessary.

Efficiency in RAG The efficiency of LLM inference is a widely explored research area, where different categories of approaches have been studied. Some works are towards architecture level acceleration such as efficient attention [31], Mixture of Expert [6], Transformer-alternative architectures [7], etc. While some other works are algorithm-level acceleration like quantization [22] or speculative decoding [17]. Nevertheless, most previous works are not RAG specific. Recently, RAG Cache [15] was proposed as a multilevel dynamic caching system tailored for RAG from the system perspective. Another stream of the work try to process each document individually so that no cross document attention computation is required, such as FiD [12] or PCW [28]. FiD encoded each retrieved passage independently from other passages and decoded by attending over the concatenation of the resulting representations of all the retrieved passages. PCW carved a long context into chunks (“windows”), restrict the attention mechanism to apply only within each window, and re-use the positional embeddings across the windows.

Comparison with previous works that are the most relevant to our work is illustrated in Table 1. This work aims to strike an optimal balance between generation quality and computational efficiency. It is notable that the extra classifier in CRAG requires maintaining extra model with more complex serving infrastructure. When there is N contexts retrieved. There is totally $N + 1$ model runs. Our work also relies an classification process to refine the retrieved documents. However, our approach relies on the “internal” classification process that is aligned with the generation process. The total model run in our case is 1.

3 Sparse RAG

Sparse RAG is designed for the decoder-only model architecture, which is the default case for most popular LLMs. Figure 1 presents an overview of Sparse RAG at inference, which designs the document assessment strategies to improve the robustness of generation. The key hypothesis of our approach is that the RAG task and per context assessment is a similar tasks and the model could do both tasks in one shot. We found simple and effective approaches to enable them.

3.1 Training Process

Our work is based on the assumption that a certain amount of RAG training data is accessible, which allows us to effectively tailor and adapt existing LLMs to our specific needs. In the training phase, we integrate an additional Per Context Assessment (PCA) task into the training mixture. This PCA

task is included alongside the primary RAG task to enhance the model’s ability to assess and respond accurately within various contexts. By incorporating the PCA task, we aim to improve the overall performance and contextual understanding of the LLMs, ensuring they are more adept at handling a diverse range of scenarios presented during their use.

Data Augmentation with LLMs For typical RAG data, one question-answer pair can be mapped to multiple retrieved contexts using either BM25 or existing standalone retriever. However, there are cases where no golden label regarding the quality of every retrieved context is available.

To collect these missing labels, we leveraged two off-the-shelf LLMs [1, 34]—PALM2 and Gemini—to get the assessment for each context. We observe empirically that a second round of prompting for critique, especially using a different model from the initial round, ensures the best quality labels. We provide our prompts in Table 9 in the Appendix. We compare different model combinations for labeling to human ground truth labels in Section 4.

Multitasking Data Format The LLM is trained on a mixture of two types of tasks, “Rate” assessment and answer generation. Specifically, we format the inputs and outputs of the two task types as

- Assessment: {Question}{Context}{Control_Assessment}{Rate}
- Generation: {Question}{Context_1}...{Context_N}{Control_Generation}{Answer}

where Rate and Answer are targets for the generative tasks while everything before them are inputs. Control_Assessment and Control_Generation are fixed tokens to ensure the LLM can differentiate the two tasks.

Parallel Contexts Since each context is rated independently in PCA task, we introduce this independence in the primary RAG task as well so two tasks could reuse KV cache during inference time. Thus for the generation task, we enforce no cross attention between different two contexts following the Parallel Context Windows [28] to process the input for training. Specifically, we modify two things in the standard LM training process. First, we change the attention masks to be block-wise, and restrict that context_i and context_j will not attend to each other. Question, Context_i and Control_Generation still follow the default attention mechanism, i.e. in the causal attention case, the latter could attend to the previous ones. Second, we adjust the position encoding to “mimic” the situation that Context_i is not visible to context_j that each context will have its own incremental positional ids. However, the position id of Control_Generation is designed to back to normal, where its position considers all the length of contexts.

3.2 Inference Process

Given the question and retrieved contexts, Sparse RAG handles the assessment task and generation task in one single pass.

Per Context Assessment Similar to the training process, when pre-filling the KV cache, each retrieved context is treated independently. Then these KV cache will be used to score the input, which is the concatenation of the token of Control_Assessment and the token of “Good”. The score is the probability of this tokens as the next tokens.

Generation The generation uses a filtered KV cache, where only K out of N KV cache will finally be loaded. We use a simple threshold based filtering approach: we drop the context when its score is less than σ . Once those caches are loaded, Control_Generation will be used to prompt the model to generate the answers.

Table 2: Auto-rater comparison to ground truth.

Auto-labeling method		Average F1	F1 Label 0	F1 Label 1
Rater model	Critic model			
PALM2 XL	n/a	0.729	0.765	0.694
PALM2 XL	PALM2 XL	0.781	0.820	0.741
Gemini Ultra	n/a	0.761	0.807	0.716
Gemini Ultra	Gemini Ultra	0.704	0.747	0.660
PALM2 XL	Gemini Ultra	0.728	0.776	0.680
Gemini Ultra	PALM2 XL	0.821	0.861	0.781

4 Evaluation of Per Context Assessment

4.1 New Benchmark: Natural Question Per Context Assessment

We isolated a subset of 50 questions, each with 10 retrieved contexts, from the Natural Questions dataset. We assigned 3 raters to each question-context pair from a pool of 7 raters and provided the instructions in Section B.2.

We aggregated responses for all 3 raters for each context, selecting the majority decision 0 or 1 for each context. We found that raters unanimously agreed on 351 out of 500 context, with 30% of the documents considered relevant. For questions where raters were not unanimously decided, a specialist rater was assigned to investigate more carefully and set the best label to correct mistakes of the other raters. This resulted in 6 additional documents considered relevant out of the entire dataset, boosting the portion of relevant documents to 31% and slightly increasing alignment with the auto-rater approaches (average F-score increase of 1.4% across auto-rater methods using these corrections as the ground truth).

4.2 LLM rater comparisons

We tested several different LLM-based automatic labeling methods—different combinations of models and prompts—for creating training data for the classifier in Sparse RAG. We compared several of these auto-rater approaches by creating a ground-truth relevance dataset using human labeling. The auto-rater comparison using the revised human labels as the ground-truth is shown in Table 2. We find that combining two different models in two rounds—initial prompting and critique—provides the labels that are most closely aligned with the human labels. We hypothesize that the different representations learned by two different models are able to capture the most nuance in the input sequences, leading to better relevance judgements. We also observe that Gemini Ultra appears slightly less effective at critiquing model outputs than PALM2 XL.

5 Evaluation of Sparse RAG

5.1 Benchmarks and Metrics

PopQA is a large-scale open-domain question answering (QA) dataset, consisting of 14k entity-centric QA pairs. Each question is created by converting a knowledge tuple retrieved from Wikidata using a template. We follow the setup from [36] and use Contriver [11] to retrieve the related contexts. Since PopQA is lacking the per-context assessment label, we adopted the “Gemini + PALM2” combination to get training labels. We split the dataset into training, validation and test set with 8:1:1 ratio. Since the answer is usually short, we report **Exact Match (EM)** and **F1** scores.

QMSum [40] is a human-annotated benchmark for query-based multi-domain meeting summarization task, which consists of 1,808 query-summary pairs over 232 meetings in multiple domains. To adopt it in the RAG domain, we divide each conversation into different contexts where the average context contains 300 words, and filter the data point where the conversation containing ground truth label is split into multiple contexts. Note that this dataset has human labeled per-context assessment that we could leverage during training. Eventually we get 250 training data, 70 validation and 77 test set. The target of this data is longer and we report **RougeLSum** and **F1** scores.

5.2 Baselines

RAG We evaluated the standard RAG [20] where an LM generates output given the query prepended with all the top ranked documents using the same retriever as in our system.

PCW RAG We applied the Parallel Context Windows [28] to the RAG process, where no cross attention is applied between documents.

Corrective RAG It is the same with RAG except that an external T5-XXL classifier is trained using heuristic labels [36]. This classifier is used to process all the document and decide the rank. Note to facilitate the fair comparison, we did not adopt the "web search" feature of this paper.

5.3 Experimental Configuration

The base LLMs used in the paper Gemini [34]. Although our approach could be applied to different training stage of the model, we apply LoRA tuning [10] to enforce alignment on top of the foundation LLMs due to its low resource requirement and wide usage. Note that the same LoRA tuning is applied to Sparse RAG and all baselines. In all our experiment, we apply LoRA in self-attention and use the default rank as 4. By default, we use the XXS size which could be running on-device.

During training, we use 64 Tensor Processing Units (TPU) V3 chips for PopQA while use 128 Units for . The batch size is 64. We use the Adafactor optimizer [32] with a learning rate of 0.003. The training dropout rate is 0.05. We leverage the metrics of validation set to pick the best checkpoint. During inference, the temperature is set to 0.5. Unless specifically noted, we use sampling decoding with sample number 1 for our experiments.

5.4 Inference Setup and Metrics

Evaluation of Sparse RAG was conducted on a Samsung S21 Ultra, utilizing the device’s CPU to assess real-world performance on a relatively mid-tier smartphone compared to the latest flagship models. Inference configuration consisted of fixed token lengths for queries, contexts and generated responses. This setup allows for evaluating the system’s efficiency and effectiveness under resource constraints typical of mobile devices, providing insights into its practical applicability for on-device question-answering tasks.

Specifically, the overall inference process considers two stages.

- Prefill stage: For baseline RAG Model, we measure the total time taken to process all input tokens (question and all contexts). For PCW RAG and Sparse RAG models: We take advantage of these models’ ability to cache question processing. We first measure the time to process the question alone. Then, we measure the time to process different contexts with the pre-processed kv-cached question. We use **Encoding Speed (ES)** (tokens per second (t/s)) to measure the efficiency of prefill stage.
- Decoding stage: To assess decoding speed comprehensively, we generate output sequences using the same question but varying the amount of relevant context data considered, ranging from the top 1 most relevant document to the top K . For each context size, we produce output sequences of different lengths. This systematic approach allows us to evaluate the impact of both context size and response length on decoding speed. We use **Decoding Speed (DS)** (tokens per second (t/s)) to measure the efficiency of decoding stage.

5.5 Main Results

We report the results with both short form answer (PopQA) and long form answer (QMSum) in Table 3, where both quality and latency metrics are shared. “ K ” is the number of chosen memories.

Notably, our proposed approach achieves the best quality while being the most efficient during inference compared to other “Dense RAG” approaches. It can be seen that Sparse RAG shares the same pre-filling efficiency with PCW, due to the parallel context encoding. However Sparse RAG can achieve significant better decoding efficiency than standard RAG and PCW RAG. In detail, out of 20 retrieved contexts, the Sparse RAG has average of 7.84 context for PopQA and 4.45 context

Table 3: Quality & Efficiency Tradeoff. For both short form and long form generation tasks, Sparse RAG achieves both higher quality and efficiency compared to “dense” RAG approaches.

	ES	PopQA				QMSum			
		EM	F1	K	DS	F1	RougeLSum	K	DS
RAG [19]	56.28	65.43	69.99	20	6.65	21.43	18.20	20	6.64
PCW-RAG[28]	147.58	65.04	69.54	20	6.65	20.18	16.95	20	6.64
CRAG [36]	-	66.52	70.99	8.90	-	-	-	-	-
Sparse RAG	147.58	67.17	71.16	7.84	12.28	23.96	20.10	4.45	16.05

for QMSum. This leads to almost **double** or even **triple** decoding speed. Meanwhile, Sparse RAG achieves higher quality metrics than the dense counterparts, which shows that Sparse RAG effectively filtered noisy and irrelevant contexts.

On the other hand, it is interesting to compare Corrective-RAG(CRAG) with Sparse RAG on PopQA task since Corrective-RAG trained their own external classifier with T5 XXL based on PopQA data. Results suggested that our approach, albeit using an “in-place” classifier, is outperforming the CRAG with external classifier. Note the encoding and decoding speed of CRAG is not comparable because it involves multiple model running from the classifier.

5.6 Analysis

5.6.1 Impact of Confidence Threshold

To show how the metrics changes with different thresholds, the performance of Sparse RAG corresponding to different thresholds is illustrated in Table 4. As the threshold gradually increases, the system will filter out more contexts, which would always reduce the latency during inference. In terms of generation quality, it can be seen that the performance of Sparse RAG was significantly improved as thresholds increased at the beginning, which shows the effectiveness of filtering out irrelevant contexts. Then, the performance was stable and dropped slightly. The reason might be that some truly relevant contexts are accidentally filtered out.

Table 4: Sampling a few confidence threshold on PopQA and QMSum. Higher confidence threshold means filtering more memories. Efficiency wise, filtering more always helps; While in terms of quality.

Threshold	PopQA				QMSum			
	EM	F1	K	DS	F1	RougeLSum	K	DS
0.05	66.95	70.97	9.75	10.61	22.85	19.49	7.92	11.92
0.1	66.84	70.66	8.72	11.70	23.78	19.98	6.68	12.89
0.15	67.17	71.16	7.84	12.28	23.43	19.66	5.77	13.01
0.2	66.77	70.54	7.13	12.88	23.2	19.79	5.05	14.54
0.25	65.75	69.64	6.56	13.00	23.96	20.1	4.45	16.05
0.3	63.86	68.2	5.98	13.08	23.84	19.99	3.93	16.38

5.6.2 Ablation of Prefill Documents

One question one may have is that the better performance may rely on the “massive” prefilling but in practical situation, we can retrieve smaller amount of context which have better quality. To this end, we compare different numbers of the prefill document: 10 vs 20 for PopQA, and show the results in Table 5. Even in the setting lower prefilled case, the previous conclusion of Sparse RAG remains. It achieves better quality with higher speed of decoding. It makes sense that the gap between dense RAG and sparse RAG is relative smaller at 10 prefill documents. This is because the headroom is smaller as the top 10 documents naturally has less to filter and make it less “sparse”. Yet our approach still achieves best trade-off, showing the capability to generalize with different numbers of input documents.

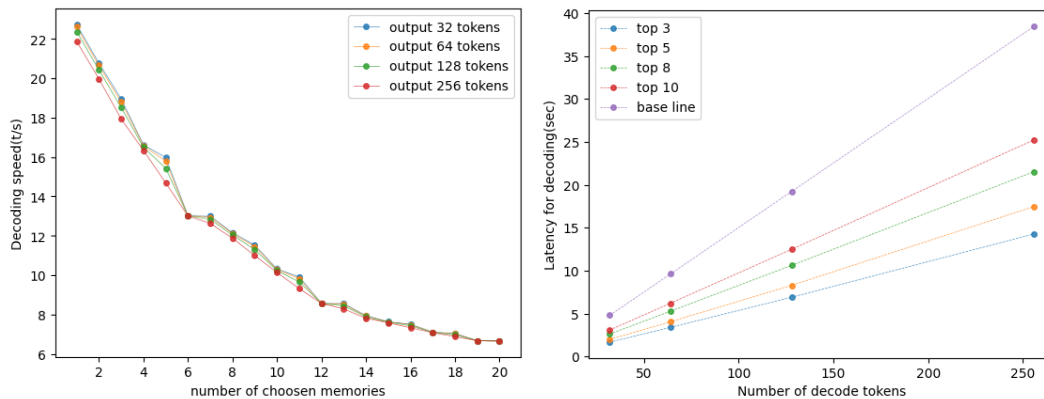
Table 5: Ablation on Different number of prefill documents for PopQA.

Approach	Prefill Documents	EM	F1	K	ES	DS
RAG	10	64.66	68.67	10	80.74	10.31
PCW RAG	10	63.9	68.58	10	147.48	10.31
Sparse RAG	10	65.86	70.2	7.79	147.48	12.33
RAG	20	65.43	69.99	20	56.28	6.65
PCW RAG	20	65.04	69.95	20	147.58	6.65
Sparse RAG	20	67.17	71.16	7.84	147.58	12.28

5.6.3 Inference Efficiency Ablations

Focusing on the decoding stage, we delve into the performance differences when varying the number of retrieved memories (i.e., top-K documents) and the length of the generated response. This analysis provides valuable insights into the computational overhead associated with incorporating additional context and generating more extensive outputs, ultimately shedding light on the practical trade-offs for real-time question-answering on resource-constrained devices.

As illustrated in Fig 2b, we observed that the baseline method requires over 50% (varies from 52% to over 55% depending on number of output tokens) more time to generate outputs of varying lengths compared to the Sparse RAG approach. This significant disparity in generation speed is primarily attributed to the increased computational burden incurred by the baseline method when considering a larger number of memories. As shown in Fig 2a, this heightened demand for computational resources results in a notable slowdown in terms of tokens per second (t/s). This observation underscores the efficiency advantages offered by Sparse RAG, especially in scenarios where a substantial amount of context needs to be considered during the decoding process.



(a) Decoding speed with different number of memories (b) E2E Latency for decoding different number of tokens

Figure 2: Inference Efficiency Comparison

5.6.4 Additional Ablations

Further ablation studies are detailed in Appendix A. These include comparisons between Silver Labels and LLM Labels (Appendix A.1), analyses based on Foundation Model Size (Appendix A.2), and evaluations using the Golden Context Label during inference (Appendix A.3).

6 Conclusion

This paper presents Sparse RAG to address the challenges of increased input length and latency. Through a novel approach of massive pre-filling and selective decoding, Sparse-RAG efficiently manages the key-value cache of retrieved documents, allowing the LLMs to focus on highly relevant tokens. This selective attention mechanism not only reduces the computational burden during

inference but also enhances the generation quality by filtering out irrelevant contexts. The evaluation on two datasets validates Sparse RAG's ability to achieve a balanced trade-off between high-quality generation and computational efficiency, proving its versatility and effectiveness for both short- and long-form content generation tasks. This innovative paradigm showcases the potential for improving LLM performance in various applications by optimizing context management and inference processes.

References

- [1] Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, et al. 2023. PaLM 2 technical report. *CoRR*, abs/2305.10403.
- [2] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. *CoRR*, abs/2310.11511.
- [3] Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*, pages 2206–2240. PMLR.
- [4] Tom B Brown, Benjamin Mann, Nick Ryder, et al. 2020. Language models are few-shot learners. In *Advances in neural information processing systems*, pages 1877–1901.
- [5] Michiel De Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Fei Sha, and William Cohen. 2021. Mention memory: incorporating textual knowledge into transformers through entity mention attention. *arXiv preprint arXiv:2110.06176*.
- [6] William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.
- [7] Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- [8] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Retrieval augmented language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR.
- [9] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR.
- [10] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- [11] 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.
- [12] Gautier Izacard and Edouard Grave. 2020. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282*.
- [13] 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.
- [14] Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active retrieval augmented generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 7969–7992. Association for Computational Linguistics.
- [15] Chao Jin, Zili Zhang, Xuanlin Jiang, Fangyue Liu, Xin Liu, Xuanzhe Liu, and Xin Jin. 2024. Ragcache: Efficient knowledge caching for retrieval-augmented generation. *arXiv preprint arXiv:2404.12457*.

- [16] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*.
- [17] Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pages 19274–19286. PMLR.
- [18] Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020. Pre-training via paraphrasing. *Advances in Neural Information Processing Systems*, 33:18470–18481.
- [19] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- [20] 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*.
- [21] Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lemao Liu. 2022. A survey on retrieval-augmented text generation. *CoRR*, abs/2202.01110.
- [22] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Xingyu Dang, and Song Han. 2023. Awq: Activation-aware weight quantization for llm compression and acceleration. *arXiv preprint arXiv:2306.00978*.
- [23] Hongyin Luo, Tianhua Zhang, Yung-Sung Chuang, Yuan Gong, Yoon Kim, Xixin Wu, Helen Meng, and James R. Glass. 2023. Search augmented instruction learning. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 3717–3729. Association for Computational Linguistics.
- [24] 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.
- [25] Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12076–12100. Association for Computational Linguistics.
- [26] Dor Muhlgay, Ori Ram, Inbal Magar, Yoav Levine, Nir Ratner, Yonatan Belinkov, Omri Abend, Kevin Leyton-Brown, Amnon Shashua, and Yoav Shoham. 2023. Generating benchmarks for factuality evaluation of language models. *CoRR*, abs/2307.06908.
- [27] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- [28] Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2022. Parallel context windows for large language models. *arXiv preprint arXiv:2212.10947*.

- [29] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- [30] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *CoRR*, abs/2302.04761.
- [31] Noam Shazeer. 2019. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*.
- [32] Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- [33] 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 *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 31210–31227. PMLR.
- [34] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- [35] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971.
- [36] Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Corrective retrieval augmented generation. *arXiv preprint arXiv:2401.15884*.
- [37] Dani Yogatama, Cyprien de Masson d’Autume, and Lingpeng Kong. 2021. Adaptive semiparametric language models. *Transactions of the Association for Computational Linguistics*, 9:362–373.
- [38] Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2023. Making retrieval-augmented language models robust to irrelevant context. *CoRR*, abs/2310.01558.
- [39] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. Siren’s song in the AI ocean: A survey on hallucination in large language models. *CoRR*, abs/2309.01219.
- [40] Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, et al. 2021. Qsum: A new benchmark for query-based multi-domain meeting summarization. *arXiv preprint arXiv:2104.05938*.

A Ablation Study

A.1 Silver Labels vs LLM Labels

In Corrective RAG, the T5 model was trained with silver labels that comes from the title matching [36]. We collect the same silver labels to replace the LLM labels and train the model with this new dataset. The comparison result is shown in Table 6.

From these results, we can observe that the quality of the labels generated by LLMs is higher than that of the silver labels from Yan et al. [36]. We hypothesize that the superior quality of the LLM-generated labels is attributable to our methodology, which involved a two-round process of soliciting responses from two different LLMs. By engaging two distinct models, we likely enhanced the robustness and accuracy of the labels through a form of cross-validation, thereby mitigating potential biases or errors that might arise from relying on a single LLM.

Table 6: Comparing silver labels with LLM labels on PopQA.

Approach	EM	F1	K	DS
Sparse RAG	67.17	71.16	7.84	12.28
w/ silver label	66.97	71.05	8.26	11.99

A.2 Ablation on Foundation Model Size

We applied Sparse RAG to different sizes of LLMs, specifically testing it on Gemini XS and Gemini XXS. The results of these experiments are presented in Table 7. The findings demonstrate that Sparse RAG is compatible with various foundation models, effectively adapting to different model sizes. Notably, with a reduced amount of decoding caches, Sparse RAG is capable of achieving the highest quality results. This indicates that Sparse RAG maintains its efficiency and effectiveness across different foundation models, making it a versatile approach for various LLM configurations.

Table 7: Ablation on Different LLM model sizes.

Approach	Model Size	EM	F1	K
RAG	XS	66.52	70.87	20
PCW RAG	XS	65.75	70.37	20
Sparse RAG	XS	68.26	72.26	6.27
RAG	XXS	65.43	69.99	20
PCW RAG	XXS	65.04	69.95	20
Sparse RAG	XXS	67.17	71.16	7.84

A.3 Using Golden Context Label during inference

Since QMSum provides golden per-context labels, we leverage these labels during inference to evaluate the upper bound performance under the condition of perfect per-context assessment. By utilizing the highest quality labels available, we aim to determine the best possible outcomes that our model can achieve. The results of this experiment, which highlight the performance ceiling under ideal labeling conditions, are presented in Table 8. This approach allows us to understand the maximum potential of our model when supplied with optimal input data, thereby offering insights into its ultimate capabilities.

Table 8: Trying golden labels on QMSum.

Approach	F1	RougeLSum	K	DS
Sparse RAG	23.96	20.1	4.45	16.05
+ golden label	26.76	21.93	1.13	21.16

B Prompts and Instructions

B.1 Prompts Used for LLMs

We share the prompt used for calling LLMs to get per context assessment in Table 9.

Table 9: The zero-shot prompts for LLM labeling and critique.

Round 1 prompt
<p>You are now doing a reading comprehension task. It is important that you be as thorough, detail-oriented, and accurate as possible in your response.</p> <p>You are given a question, a set of accepted answers, a document and its title. The document does not necessarily contain the right answer to the question.</p> <p>You should read the title and the document and then check if they provide one of the correct answers to the question.</p> <p>If the title and document together contain the correct answer to the question, output a score of 1.0, otherwise output a score of 0.0.</p> <p>question: <question> accepted answers: <answers> title: <title> document: <document> output:</p>
Round 2 prompt
<p>Your job is to correct another model’s performance on a reading comprehension task.</p> <p>The model was given a question, a set of accepted answers, a document and its title. The document and title do not necessarily contain the right answer. The model was instructed to output a score of 1.0 if the document contains the answer, and a score of 0.0 otherwise.</p> <p>You will be given the same information as the other model along with its output. You should read the title and document and then check if they provide one of the correct answers to the question.</p> <p>Then check if you agree with the previous model’s output.</p> <p>If you agree, output the same score unchanged.</p> <p>If you disagree, output the corrected score.</p> <p>Your output should be as accurate as possible.</p> <p>question: <question> accepted answers: <answers> title: <title> document: <document> previous model’s score: <score> output:</p>

B.2 Rater Guidelines

We share the instructions provided to the human labelers in Table 10.

Table 10: Instructions for raters creating ground-truth relevance dataset.

Human Rater Instructions
<p>Please read the question, the answer and the context. Please answer if the context can help answer the question. If it can, select 1. Otherwise select 0.</p> <p>1: good 0: bad</p> <p>Please use the answers as a hint. However, do not use "is the answer in the context?" as a heuristic for making the decision.</p>

C Dataset Analysis

During the human-labeling process, several raters flagged documents and questions that were difficult to label. In total, 23 out of 500 documents were flagged and 15 out of 50 questions were flagged.

We observed trends in questions and contexts shared in Table 11 that raised concerns about whether and how a human would be able to assess the relevance of the context. These concerns extend to expectations of how well LLMs would be able to do that as well.

Table 11: Overview of trends, datasets, and examples with associated comments.

Trend	Dataset	Example question	Comments
Questions with time-dependent answers	NQ	who is the president of usa right now	Depends on when the question is asked.
	NQ	who is the current director of the us mint	Depends on when the question is asked.
	NQ	when is the next dead-pool movie being released	Depends on when the question is asked.
	NQ	total number of death row inmates in the us	Fluctuates over time.
	PopQA	What is Prague the capital of?	The borders in this region and the name of the country have changed several times in the 20th century.
	PopQA	What is Dennis Rodman's occupation?	The accepted answers are "actor, actress, actors, actresses". He was an actor later in his career, but he rose to prominence as a professional basketball player.
Missing synonyms, abbreviations, and aliases, combined with unclear granularity	NQ	in which regions are most of Africa petroleum and natural gas found	"Region" can refer to different levels of granularity (e.g. Sub-Saharan Africa vs. Ethiopia), but the only accepted answer is "Nigeria".
	NQ	what type of car is a jeep	The accepted answers are only "off-road vehicles", "light utility vehicles", "sport utility vehicles", but "SUV" is clearly a correct answer as well.
Non-exhaustive list of answers	NQ	cast of law & order special victim unit	The accepted answers include 16 cast members, but the show went on for 25 seasons with many cast changes and guest stars not included in the list.
Oddly phrased question	NQ	right to property according to the constitution of India is a	The only correct answer is "constitutional right", but that is included in the question. It's not clear what type of answer would be appropriate here.
Overly specific answers expected	NQ	where does the story the great gatsby take place	The only accepted answer here is "Long Island of 1922", but the place is Long Island and the question does not ask about when the story is set.
Question refers to an entity with a common name without disambiguation	PopQA	What genre is Frances?	There is a musician and a film called "Frances", and both of those could arguably have a genre associated with them.
	PopQA	Who was the producer of Hurt?	The question is referring to a song performed by Christina Aguilera but there are many other songs, movies, and other entities that share the name and also have a producer.
	PopQA	What is the capital of Cherokee County?	There are many different Cherokee Counties in different states in the USA.

We explored several ways of filtering our human-labeled subsample of Natural Questions to determine how they impacted context assessment F-scores overall and for each auto-rater. We provide two additional filtered versions of the human-labeled RAG relevance dataset as alternatives. See Table 12 for statistics on the filtered datasets and Table 13 for the auto-rater F-scores for each. Both statistical filtering approaches (e.g. removing contexts with non-unanimous labels) and targeted filtering approaches (e.g. removing questions or contexts flagged by human raters) lead to some improvement in F-scores for relevance labels.

Table 12: Summary of Dataset Evaluations

Dataset	# Questions	# Docs	Total % Relevant	Average # Docs per Question	Std Dev # Docs per Question	Average % relevant docs per question	Std Dev % relevant docs per question
Original 3-rater labels	50	500	0.30	10	0	0.30	0.23
Revisited (released)	50	500	0.31	10	0	0.31	0.26
Filter non-unanimous (released)	50	351	0.23	7.02	2.59	0.27	0.29
Filter flagged docs	48	467	0.31	9.729	0.57	0.31	0.27
Filter flagged questions	35	350	0.29	10	0	0.29	0.26
Filter flagged docs and question (released)	34	330	0.29	9.7	0.58	0.29	0.27
Filter 4-rater split	50	444	0.28	8.88	1.49	0.29	0.27

Table 13: Evaluation of Labeling Methods w/ Filtering

Dataset Filters	# Docs	Rater Model	Critic Model	Average F1
Specialized rater corrections	500	PALM2 XL	n/a	0.729
		PALM2 XL	PALM2 XL	0.781
		Gemini Ultra	n/a	0.761
		Gemini Ultra	Gemini Ultra	0.704
		PALM2 XL	Gemini Ultra	0.728
		Gemini Ultra	PALM2 XL 340B	0.821
Filter non-unanimous docs	351	PALM2 XL	n/a/	0.741
		PALM2 XL	PALM2 XL	0.811
		Gemini Ultra	n/a	0.792
		Gemini Ultra	Gemini Ultra	0.739
		PALM2 XL	Gemini Ultra	0.763
		Gemini Ultra	PALM2 XL	0.856
Filter flagged docs and questions	330	PALM2 XL	n/a	0.750
		PALM2 XL	PALM2 XL	0.797
		Gemini Ultra	n/a	0.782
		Gemini Ultra	Gemini Ultra	0.741
		PALM2 XL	Gemini Ultra	0.753
		Gemini Ultra	PALM2 XL	0.833

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We lay out the main scope and contribution in both abstract and introduction: introducing "sparse" into RAG that makes it achieve best quality and latency tradeoff.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We put limitations in Section 7.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not have theoretical result.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We explained our experimental settings well in Section 5.3 and 5.4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: Our dataset are public available. However the Gemini LLMs is not open sourced LLMs.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We share our full settings in Section 5.3 and 5.4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Calculating Statistical Significance introduce too much TPU/GPU resource usages, given our models needs to be trained.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Section 5.3 shares the details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We follow the guidelines.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: We don't see very obvious social impact of this work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We don't see any potential of misuse of our approach.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the related dataset papers or technique reports.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We introduced a testset with details in section 4 and appendix.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: We shared the guideline details in appendix.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The main contribution of the paper does not involves human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.