

Embedding-Informed Adaptive Retrieval-Augmented Generation of Large Language Models

Chengkai Huang¹, Yu Xia², Rui Wang³, Kaige Xie⁴, Tong Yu⁵, Julian McAuley², Lina Yao^{1,6}

¹University of New South Wales, ²University of California San Diego, ³Duke University,
⁴Georgia Institute of Technology, ⁵Adobe Research, ⁶CSIRO's Data61

Correspondence: chengkai.huang1@unsw.edu.au

Abstract

Retrieval-augmented large language models (LLMs) have been remarkably competent in various NLP tasks. However, it was observed by previous works that retrieval is not always helpful, especially when the LLM is already knowledgeable on the query to answer. Motivated by this, Adaptive Retrieval-Augmented Generation (ARAG) studies retrieving only when the knowledge asked by the query is absent in the LLM. Previous works of ARAG either require accessing the pre-training corpus or prompting with additional model inferences. Aiming to avoid such drawbacks, we propose to determine whether the model is knowledgeable on a query via inspecting the (contextualized) pre-trained token embeddings of LLMs. We hypothesize that such embeddings capture rich information on the model's intrinsic knowledge base, which enables an efficient way of judging the necessity to retrieve from an external corpus. Extensive experiments demonstrate our ARAG approach's superior performance across various benchmarks.

1 Introduction

Retrieval-augmented LLMs have excelled in various NLP tasks (Li et al., 2022; Yasunaga et al., 2023; Lin et al., 2022). However, it was observed (Mallen et al., 2023) that the retrieved knowledge might not necessarily improve the quality of a generated response, especially when the model is already knowledgeable on the input query from pre-training. Moreover, the retrieval process can incur additional computational costs and latency, e.g., by significantly increasing the context length.

To solve this, Adaptive Retrieval-Augmented Generation (ARAG) dynamically determines whether the LLM has already acquired the knowledge to answer the query during pre-training, then only to retrieve from external corpus when the knowledge required is absent. The pilot work of ARAG (Mallen et al., 2023) only retrieves if an

entity from the query has low frequency in the pre-trained corpus. The LLMs are deemed not knowledgeable on the query if the extracted entities have low frequencies. An obvious drawback of such a heuristic approach is that it only works on *entity-centric* question answering, e.g., "What is the capital of Louisiana?", where the question is mostly about an entity (a person, country, etc.) that can be identified by existing entity extraction tools like Spacy (Honnibal and Montani, 2017). Additionally, the pre-trained corpus could be proprietary and not readily accessible to compute the frequency. More recent approaches of ARAG, e.g. (Zhang et al., 2024), prompt the LLMs for a retrieval decision with extra LLM calls, assuming that the LLMs is aware of its knowledge boundary. This assumption is challenged by studies on LLM factuality, i.e., LLMs are overly confident in their ability to answer a question (Ren et al., 2023a).

In this work, we propose to decide whether to retrieve by analyzing the LLM pre-trained parameters, which does not require accessing the pre-trained data or extra LLM calls. Specifically, we predict whether the external knowledge from retrieval could help, by inspecting the pre-trained embeddings of tokens from the input query. We hypothesize that the pre-trained token embeddings capture rich information on the model's intrinsic knowledge base. This stems from previous analysis (Cai et al., 2020), claiming that tokens of different frequencies during pre-training tend to exhibit different layout in the embeddings space. In Section C, we also show that the pre-trained embeddings are discriminative by the frequency information, consistent with (Cai et al., 2020). Considering that such frequencies are indicative of whether the LLM is knowledgeable (Mallen et al., 2023), the token embeddings should contain sufficient information to judge the necessity to augment with external knowledge via retrieval. We term our approach **EI-ARAG** (Embedding-Informed ARAG).

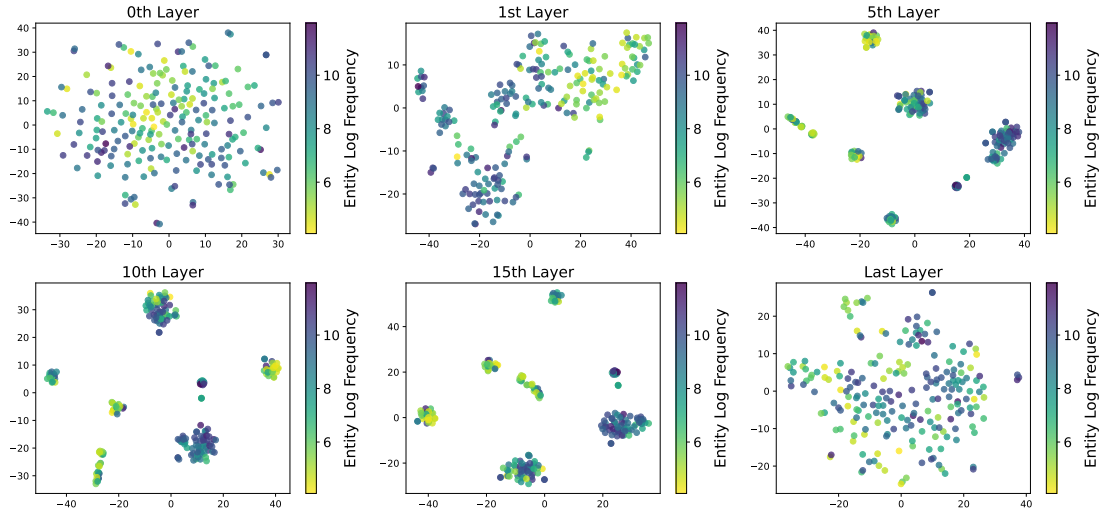


Figure 1: Visualization of embeddings at different layers of LLaMA 2 7B for director-related questions in PopQA. Darker color indicates the question contains entities with higher frequency in pre-training data.

The main contributions of this work are as follows:

- We hypothesize that the pre-trained LLMs token embeddings capture rich information on the model’s intrinsic knowledge base, and propose EI-ARAG to predict the necessity to retrieve based on these embeddings.
- Instead of requiring extra LLM calls to judge whether the LLM is knowledgeable on the input, our proposed EI-ARAG shows that we can do it efficiently with the pre-trained token embedding.
- Extensive experiments and in-depth analyses demonstrate the superiority of our approach on various benchmarks.

To determine when to retrieve, existing methods either need to access pre-training data to calculate entity frequency (e.g., how many times ‘Louisiana’ appears in pre-training corpus (Mallen et al., 2023)), or require extra inference prompting LLMs to make decision themselves (e.g., “Do you need to retrieve external knowledge to answer the question correctly?” (Zhang et al., 2024)). However, (i) relying on the entity frequency restricts the approach, making it suitable only for entity-centric question answering, (ii) the access to the pre-training data is not often available and (iii) the extra prompting will double the computational cost due to and additional LLM call. To circumvent the above drawbacks, we hypothesize that the pre-trained token embeddings capture rich information

on the model’s intrinsic knowledge (Section C), proposing to utilize such embeddings to efficiently judge the necessity to retrieve from external knowledge base.

2 Embedding-Informed Retrieval

While the input embedding layer (i.e. 1st layer) contains rich semantic information of each token, later layers of LLMs provide embeddings that are more contextualized on current question. To capture more precisely LLMs’ intrinsic knowledge given a specific question, we extract embeddings for each token from the contextualized layers. We then obtain a sentence embedding for the given question via average pooling following (Li et al., 2020).

Formally, given a question $q \sim Q$, we tokenize it with tokenizer T and subsequently convert into a token embedding sequence, $\text{embed}_{1\text{st}}(T(q))$. We hypothesize that $\text{embed}_{1\text{st}}(T(q))$ encapsulates intrinsic knowledge of LLMs on the question q . To predict whether the LLM is already knowledge on q , we train a classifier $C : Q \rightarrow Y \in \{0, 1\}$ whose prediction is $y \in Y$, where $y = 1$ indicates the need for retrieval augmentation, and $y'_i = 0$ indicates otherwise. Specifically, we have,

$$y = C(\text{embed}_{1\text{st}}(T(q))).$$

where C is a neural network. We denoted its training set as $D = \{q_i, y_i\}_{i=1}^N$. The question q_i is randomly sampled from the original training split of datasets. The label y_i is determined by inferencing twice with the LLM on q_i , one with retrieval and one without retrieval. We determine $y_i = 1$ if

the result with retrieval is better than not. During the test, we only retrieve when C predicts 1.

To validate our hypothesis that the pre-trained embeddings contains rich information on the intrinsic knowledge, we visualize in Figure 1 the embeddings of entities in different layers of LLaMA 2 7B and color them based on the logarithmic entity frequency associated in each question following (Cai et al., 2021). We can observe that at the 0th layer, i.e., the input embedding layer, the embeddings is less discriminative in terms of the frequency. While on higher layers, we observe clear patterns for questions with entities of different frequency, which shows that the (contextualized) pre-trained embeddings captures the entity frequency information during pre-training, thus show be indicative on whether the LLM is knowledge about the given question (Mallen et al., 2023). Such embeddings-frequency correlation are also observed in prior analysis (Cai et al., 2020). While the embedding patterns appear from the first layer, embeddings from later layers do not bring obvious improvement to our classifier as shown in Table 4. Thus, we simply extract sentence embeddings from the first contextualized layer to ensure computational efficiency compared to prompt-based methods that require full inference. More details are provided in Appendix C. Utilizing only the first contextualized layer, our embedding-informed method offers an efficient way to determine the need for retrieval augmentation compared to methods that need full inference (Zhang et al., 2024). In contrast to the method requiring the access of the pre-training data (Mallen et al., 2023), our method focuses on the analysis of embeddings, enhancing the applicability and scalability in real-world systems.

3 Experimental Setup

In this section, we describe the datasets, baselines, metrics, and implementation details of our experiments. The experiments are designed to answer the following research questions:

RQ1: Can our method effectively decide when to retrieve for ARAG? How does it perform compared to state-of-the-art baselines?

RQ2: Is using embeddings more computationally efficient than prompting-based methods?

RQ3: What kinds of knowledge are captured from the embeddings to decide when to retrieve? The extensive discussion in Section C has already answered RQ3 where we find contextualized em-

beddings to reflect LLM intrinsic knowledge similarly as entity frequency.

Datasets. Following the evaluation in (Mallen et al., 2023; Zhang et al., 2024), we verify the effectiveness of our proposed framework on two open-domain QA datasets (Entity QA dataset and Non-entity QA dataset) to answer our research questions. Please refer to Appendix E for details. **Baselines.** We compare our EI-ARAG against relevant baselines, including: (1) **Simple Full Retrieval** (Mallen et al., 2023); (2) **DARAG** (Mallen et al., 2023); (3) **PARAG-Vanilla** (Zhang et al., 2024); (4) **PARAG-TAARE** (Zhang et al., 2024); (5) **Adaptive-RAG w/ Oracle:** using the Oracle classifier with our Adaptive-RAG, which includes the best performance (i.e., the upper-bound performance) in an ideal scenario. See Appendix D for more details.

Evaluation Metrics. Following similar settings as (Mallen et al., 2023; Zhang et al., 2024), we report the results with (i) retrieval *Accuracy* (ACC), which evaluates how well LLMs can perform adaptive retrieval and (ii) *Percentage of Retrieval* (POR) which assesses the efficiency of implementing adaptive retrieval. *Accuracy* measures the proportion of the predicted answer containing the ground-truth answer. *Percentage of Retrieval* quantifies the proportion of instances in which ARAG methods opt to activate retrieval across all test samples.

4 Experimental Results

Entity-centric QA (RQ1). As shown in Table 1, we can observe that: (i) The simple No Retrieval method yields the lowest accuracy, highlighting the importance of incorporating external knowledge for enhanced response quality. Conversely, Full Retrieval shows a significant improvement in accuracy, affirming the benefit of external data. Among adaptive retrieval methods, DARAG, which is a prompting-based method, shows varied performance. Unfortunately, since GPT-Neo is not a chat model, DARAG tends to default to full retrieval on this platform, not achieving the desired adaptive retrieval effect. This limitation is evident as it performs well on the LLaMA 2 model but fails on the GPT-Neo models. (ii) Our proposed EI-ARAG approach stands out by not only achieving high accuracy but also maintaining lower retrieval rates. This indicates efficient use of external information, retrieving only when the LLM lacks the necessary knowledge for a query. To help answer RQ3, in

Table 1: Experiment results on the entity-centric PopQA dataset queries with different LLMs. We emphasize our results in bold for easy comparisons. *ACC(%)* denotes retrieval accuracy and *POR(%)* denotes the percentage of retrieval. Please refer to Section 3.

Types	Methods	GPT-Neo (1.3B)		GPT-Neo (2.7B)		LLaMA 2 (7B)	
		ACC(%)	POR(%)	ACC(%)	POR(%)	ACC(%)	POR(%)
Simple	No Retrieval	11.19	0	12.56	0	24.64	0
	Full Retrieval	20.44	100	22.99	100	29.55	100
Adaptive	DARAG (Mallen et al., 2023)	20.32	95.93	22.76	94.89	31.99	69.80
	PARAG-Vanilla (Zhang et al., 2024)	11.85	0	12.56	0	27.78	88.98
	PARAG-TAARE (Zhang et al., 2024)	11.85	0	12.56	0	29.21	95.15
	EI-ARAG (Ours)	21.47	84.33	24.00	94.37	33.08	57.89
Oracle	Adaptive-RAG w/ Oracle	22.34	88.81	24.87	87.44	37.62	75.36

Table 2: Experiment results on the non-entity centric TriviaQA dataset queries with LLaMA 2 (7B). We emphasize our results in bold for easy comparisons. *ACC(%)* denotes retrieval accuracy and *POR(%)* denotes the percentage of retrieval. Please refer to Section 3.

Methods	ACC (%)	POR(%)
No Retrieval	47.33	0
Full Retrieval	62.33	100
DARAG (Mallen et al., 2023)	N/A	N/A
PARAG-Vanilla (Zhang et al., 2024)	61.78	97.67
PARAG-TAARE (Zhang et al., 2024)	62.33	98.56
EI-ARAG (Ours)	62.67	92.11
Adaptive-RAG w/ Oracle	68.56	52.67

Table 4 and Figure 2, we also run additional experiments and report the results on PopQA, with varying layers when extracting the latent embeddings from LLaMA 2 7B.

Non Entity-centric QA (RQ1). In this subsection, we show the overall experimental results and offer in-depth analyses of our method in the non-entity-centric dataset TriviaQA.

Table 3: Comparison of ARAG methods across different dimensions in LLaMA2 7B model (detailed in Section 4 and 4).

	DARAG (Mallen et al., 2023)	PARAG (Zhang et al., 2024)	EI-ARAG (Ours)
Entity QA	✓	✓	✓
Non-entity QA	✗	✓	✓
Latency (ms)	1740	2328	1744
ACC(%)	31.99	29.21	33.08
POR(%)	69.80	95.15	57.89

The main results as shown in Table 2 and the error analysis of our EI-ARAG are shown in Figure 3. It is worth noting that DARAG is not suitable for sentence-level experiments because it is an entity-centric method, and the TriviaQA dataset contains questions involving multi-hop entities, resulting in an N/A outcome for DARAG. We ob-

serve that EI-ARAG is better than all baselines both for ACC and POR. Our approach reveals higher confidence in the LLM’s intrinsic knowledge than PARAG, resulting in fewer retrievals and more precise answers, demonstrating our method’s superior decision-making on when to engage retrieval compared to the baseline model.

Comparison of Computational Cost to Determine When to Retrieve (RQ2). To better demonstrate the efficiency of our methods, we record the inference time needed by our methods to extract the first layer embedding of LLaMA 2 (7B). In comparison, we also record the time needed for TAARE (Zhang et al., 2024) to prompt LLaMA 2 to generate “Yes” or “No” for retrieval. We use the same instruction prompt as Zhang et al. (2024) and limit the maximum generated new tokens to 5 to avoid extra inference cost. On the test sets of PopQA and TrivialQA, our method extracts the first layer embeddings with only 0.0443 seconds on average for each question. For the prompting-based method, it takes an average of 0.3885 seconds per question for the LLM to answer whether it needs retrieval. The results further demonstrate the advantage of our methods for more efficient adaptive RAG. To summarize, our method achieves the highest accuracy and the lowest proposition retrieval, while maintaining a relatively shorter average inference time.

5 Conclusion

In this work, we introduce a novel approach for Adaptive Retrieval-Augmented Generation (ARAG), which strategically employs retrieval only when the LLM lacks the necessary knowledge for a query. Unlike previous methods that require access to the pre-training corpus or additional model inferences, our method leverages the rich information

captured by the pre-trained token embeddings of LLMs. This approach allows us to efficiently determine the necessity of external retrieval. Extensive experiments demonstrate that our proposed ARAG method achieves superior performance across various benchmarks.

6 Limitations

The effectiveness of retrieval-augmented strategies depends on the quality and relevance of the external data sources. If the retrieved content is of low quality or irrelevant, it not only fails to aid in answering queries but might also introduce noise that can mislead the model and impair its decision-making process. Ensuring that retrieval mechanisms access high-quality and pertinent information is crucial for enhancing model performance and reliability. While this is out of the scope of this paper, future works are required to make RAG systems more robust and effective.

7 Ethics Statement

It is important to note that LLMs can still generate incorrect (hallucination) or biased outputs, even when they are retrieval-augmented. Therefore, it is always important to verify the outputs of language models with other sources of information.

References

- 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.
- Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, USVSN Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. [Gpt-neox-20b: An open-source autoregressive language model](#). *CoRR*, abs/2204.06745.
- Xingyu Cai, Jiayi Huang, Yuchen Bian, and Kenneth Church. 2020. Isotropy in the contextual embedding space: Clusters and manifolds. In *International conference on learning representations*.
- Xingyu Cai, Jiayi Huang, Yuchen Bian, and Kenneth Church. 2021. [Isotropy in the contextual embedding space: Clusters and manifolds](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Zhangyin Feng, Weitao Ma, Weijiang Yu, Lei Huang, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. Trends in integration of knowledge and large language models: A survey and taxonomy of methods, benchmarks, and applications. *CoRR*, abs/2311.05876.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. [Representation degeneration problem in training natural language generation models](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- 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.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- 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.
- 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.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1601–1611. Association for Computational Linguistics.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. 2022. Language models (mostly) know what they know. *CoRR*, abs/2207.05221.

- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: a benchmark for question answering research. *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. [On the sentence embeddings from pre-trained language models](#). *CoRR*, abs/2011.05864.
- Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lemao Liu. 2022. [A survey on retrieval-augmented text generation](#). *CoRR*, abs/2202.01110.
- Bill Yuchen Lin, Kangmin Tan, Chris Miller, Beiwen Tian, and Xiang Ren. 2022. [Unsupervised cross-task generalization via retrieval augmentation](#). In *NeurIPS*.
- 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.
- Ruiyang Ren, Yuhao Wang, Yingqi Qu, Wayne Xin Zhao, Jing Liu, Hao Tian, Hua Wu, Ji-Rong Wen, and Haifeng Wang. 2023a. Investigating the factual knowledge boundary of large language models with retrieval augmentation. *arXiv preprint arXiv:2307.11019*.
- Ruiyang Ren, Yuhao Wang, Yingqi Qu, Wayne Xin Zhao, Jing Liu, Hao Tian, Hua Wu, Ji-Rong Wen, and Haifeng Wang. 2023b. Investigating the factual knowledge boundary of large language models with retrieval augmentation. *CoRR*, abs/2307.11019.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrubti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura,
- Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Richard James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen-Tau Yih. 2023. [Retrieval-augmented multimodal language modeling](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 39755–39769. PMLR.
- Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. 2023. Do large language models know what they don’t know? In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 8653–8665. Association for Computational Linguistics.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, Weng Lam Tam, Zixuan Ma, Yufei Xue, Jidong Zhai, Wenguang Chen, Zhiyuan Liu, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. [GLM-130B: an open bilingual pre-trained model](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Zihan Zhang, Meng Fang, and Ling Chen. 2024. [Retrievalqa: Assessing adaptive retrieval-augmented generation for short-form open-domain question answering](#). *CoRR*, abs/2402.16457.

A Related Work

Retrieval-Augmented Generation. Retrieval-augmented Large Language Models (RALM) have been remarkably competent in various NLP tasks (Li et al., 2022), which augments the input space of LMs with retrieved text passages (Guu et al., 2020), leading to large improvements in knowledge-intensive tasks (Izacard et al., 2022). However, standard RAG methods indiscriminately retrieve information regardless of the query, potentially degrading performance and increasing costs, as LLMs can often handle straightforward queries on their own, but noisy retrieved context might impede their performance. To alleviate the limitations of RAG mentioned above, recent studies advocate for Adaptive RAG (ARAG), which dynamically determines retrieval necessity and relies only on LLMs’ parametric knowledge when considered unnecessary (Mallen et al., 2023). Mallen et al. (Mallen et al., 2023) propose to heuristically retrieve when the popularity of an entity on Wikipedia is below a certain threshold. Jiang et al. (Jiang et al., 2023) trigger retrieval if any token in the temporarily generated sentence has low confidence. (Feng et al., 2023; Asai et al., 2023; Ren et al., 2023b) directly prompt LLMs for retrieval decisions, given the observation that LLMs can acknowledge their knowledge boundaries to some extent (Yin et al., 2023; Kadavath et al., 2022). Recently, Zhang et al. (Zhang et al., 2024) propose the TA-ARE, a simple yet effective method to help LLMs assess the necessity of retrieval, obviating the need for calibration or additional training.

Parametric and Non-parametric Knowledge. Previous research demonstrates that large pre-trained language models (LMs) such as BERT, and GPT (Gao et al., 2019; Li et al., 2020; Cai et al., 2021) encapsulate a significant amount of world knowledge within their parameters. However, depending exclusively on their parameters to store extensive world knowledge necessitates an impractically large number of parameters, and the knowledge can quickly become outdated. Recent studies (Mallen et al., 2023) show that enhancing LMs with non-parametric memories (i.e., retrieved text chunks) allows for smaller models to achieve performance levels comparable to those of larger models.

Table 4: EI-ARAG accuracy with different embedding layers of LLaMA 2 7B averaged on all subsets of PopQA.

Methods	ACC (%)
0th Layer	38.54
1st Layer	40.98
5th Layer	40.97
10th Layer	41.13
15th Layer	41.09
Last Layer	40.86

B Embedding-Informed ARAG

In the context of open-domain QA, the primary objective of the ARAG method is to ascertain whether a given question (e.g., “*What is the capital of Louisiana?*”) requires retrieval augmentation. The core of this task is to determine whether language models already possess knowledge related to the question, thereby deciding if there is a need to retrieve external knowledge bases to enhance the model prediction. This adaptive retrieval approach can effectively save context length and thus reduce latency during LLM inference. Besides, it can also mitigate performance degradation caused by redundant retrievals in LLMs (Mallen et al., 2023).

C Why the Pre-Trained Embeddings Help?

To validate our hypothesis that the pre-trained embeddings contain rich information on the intrinsic knowledge and are sufficient to help judge the necessity to retrieve, we visualize in Figure 1 the embeddings of entities in different layers of LLaMA 2 7B and color them based on the logarithmic entity frequency associated with each question following (Cai et al., 2021). Questions are the entire director-related subset of PopQA dataset (Mallen et al., 2023) selected as a representative challenging subset.

From Figure 1, we can observe that at the 0th layer, i.e., the input embedding layer, the embeddings is less discriminative in terms of the frequency. While on higher layers, we observe clear patterns for questions with entities of different frequencies. In Table 4, we show the testing results of using the embeddings from different layers. We find that the 0th layer is also producing detect results. Therefore, we speculate that embeddings from the 0th layer should also be correlated with frequency, but require a more complex visualization method to exhibit in a 2D plot. The results in Figure 1 show that the (contextualized) pre-trained

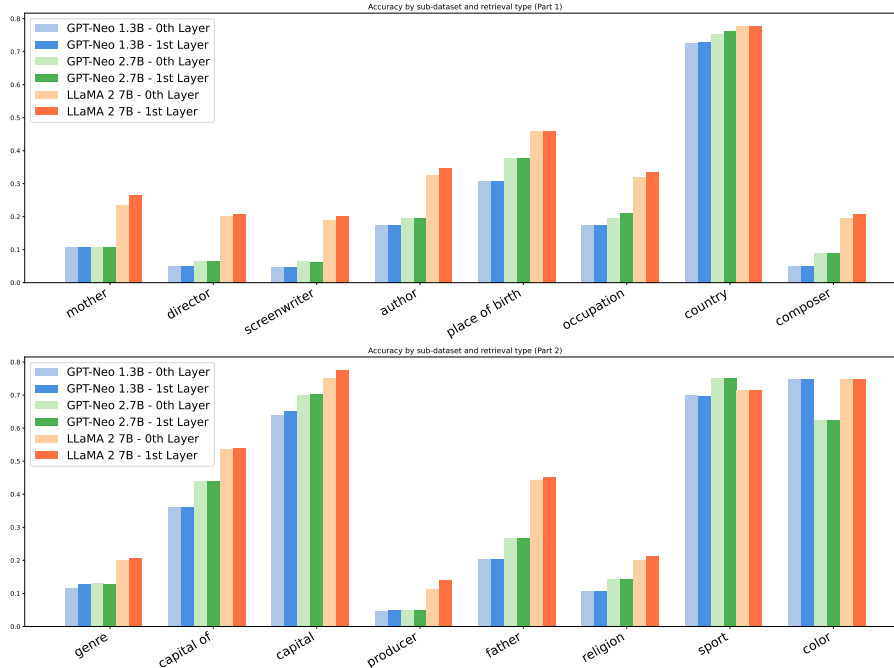


Figure 2: Per-relationship type results on PopQA by different models, showing overall accuracy of EI-ARAG using 0th and 1st layer embeddings based on BM25 RALM.

embeddings capture the entity frequency information during pre-training, thus indicative of whether the LLM is knowledgeable about the given question (Mallen et al., 2023). Such embeddings-frequency correlations are also observed in prior analysis (Cai et al., 2020).

While the embedding patterns appear from the first layer, embeddings from later layers do not bring obvious improvement to our classifier as shown in Table 4. Thus, in our method, we simply extract sentence embeddings from the first contextualized layer to ensure computational efficiency compared to prompt-based methods that require full inference. We also show in Figure 2 additional results on the full PopQA dataset comparing the performance of utilizing the 0th layer and 1st layer in Table 4.

D Baselines

We compare our EI-ARAG against relevant baselines, including two retrieval-augmented LLM strategies (No retrieval and full retrieval) and the adaptive retrieval approaches (Mallen et al., 2023; Zhang et al., 2024), which can be grouped into one of two categories: Simple and Adaptive. Specifically, baseline approaches include: **Simple No Retrieval**: Directly input the question prompt into LLMs to generate answers. **Simple Full Retrieval** (Mallen et al., 2023): For all questions, we use

the retriever to retrieve related external knowledge, which is then added to the prompts before inputting them into large language models to generate answers. **DARAG** (Mallen et al., 2023): Data-aware Adaptive RAG determines the complexity level questions based on the frequency of its entities. It suggests using the retrieval modules only when the entity frequency falls below a certain threshold. In the PopQA dataset, this threshold for each relation type is determined using brute force search. **PARAG-Vanilla** (Zhang et al., 2024): Prompting Adaptive RAG (Vanilla) uses vanilla prompt to ask LLMs whether retrieval is necessary to determine “Yes” or “No”. If retrieval is needed, proceed with it; if not, do not perform retrieval. This adaptive retrieval method requires two complete inferences. **PARAG-TAARE** (Zhang et al., 2024): Prompting Adaptive RAG (Time-aware) incorporates time-related prompts to ask LLMs whether retrieval is needed to determine “Yes” or “No”. If retrieval is required, it will be conducted; otherwise, it will not. This adaptive retrieval method also requires two complete inferences. **Adaptive-RAG w/ Oracle**: using the Oracle classifier with our Adaptive-RAG, which includes the best performance (i.e., the upper-bound performance) in an ideal scenario. In the ideal setting, Adaptive approaches are expected to be more effective than those using Full Retrieval, while also being more efficient than the

No Retrieval approach.

For the fair comparison, we evaluate our framework across models of varying capacities: GPT-Neo (1.3 billion), GPT-Neo (2.7 billion) (Black et al., 2022), LLaMA 2 (7 billion) (Touvron et al., 2023). In this evaluation, we do not perform any fine-tuning on models.

E Dataset Construction

PopQA (Mallen et al., 2023): An entity-centric open-domain QA dataset about entities with a wide variety of popularity, consisting of 14k questions to cover factual information in the long tail that might have been missed in other popular QA datasets (Kwiatkowski et al., 2019). Following the setting of Mallen et al. (2023), we choose 75% of data for training, then evaluate our model and baselines on the remaining 25% of PopQA.

TriviaQA (Joshi et al., 2017): This Non-entity QA dataset is designed for reading comprehension tasks and consists of question-answer-evidence triples, where each question incorporates both entity and non-entity information, challenging models to navigate and integrate diverse types of data within a single query, thus providing a more complex test of the models’ comprehension and reasoning abilities. Since the TriviaQA-unfiltered (open) test set is not publicly available, we derived from a random sampling of 3,600 entries from the original TriviaQA-unfiltered dataset, with a training and testing split of 2,700 and 900, respectively. For the RAG task, we utilize the top-5 documents retrieved from Google Search API to gather relevant passages.

F Details of Data Constructions

In this paper, for entity-centric experiments, we adopt the 16 relationship types identified in the PopQA dataset as outlined in (Mallen et al., 2023). The authors of this reference manually annotated templates to transform knowledge triples into natural language questions. The complete list of templates utilized to generate EI-ARAG is presented in Table 5.

G Implementation Details

For conducting experiments, we use a single NVIDIA RTX A5000 GPU equipped with 24GB of GPU memory. Besides, we use int8bit (Zeng et al., 2023) quantization with LLaMA2 7 billion to make them fit our GPUs. Similar to Mallen

Relationship	Template
occupation	What is [subj] ’s occupation?
place of birth	In what city was [subj] born?
genre	What genre is [subj]?
father	Who is the father of [subj] ?
country	In what country is [subj] ?
producer	Who was the producer of [subj] ?
director	Who was the director of [subj] ?
capital of	What is [subj] the capital of?
screenwriter	Who was the screenwriter for [subj] ?
composer	Who was the composer of [subj] ?
color	What color is [subj] ?
religion	What is the religion of [subj] ?
sport	What sport does [subj] play?
author	Who is the author of [subj] ?
mother	Who is the mother of [subj] ?
capital	What is the capital of [subj] ?

Table 5: Full list of the manually annotated templates used for our creations. [subj] denotes a placeholder for subject entities.

et al. (2023), in our experiments with GPT-Neo 1.3 and 2.7 billion parameters, we did not detect a notable performance decline when employing quantization. Besides, for a fair comparison, by configuring the number of beams as 1 and setting sampling to False, we ensured that the LLMs consistently produced fixed inference results following (Mallen et al., 2023).

For the prompts design, we follow the previous work’s experiment setting (Mallen et al., 2023). For the PopQA dataset, we sample few-shot examples stratified by relationship type to diversify the samples: for each of the 15 relationship types other than the one in the test question, we sample one random question-answer pair to include in the context (detailed relation type templates can be found in Appendix A of (Mallen et al., 2023)). Besides, since the quality of the retrieved documents is not the focus of this paper, we use the off-the-shelf BM25 retrieval model and author-provided top-5 documents (Mallen et al., 2023) extracted from Wikipedia where possible. For the TriviaQA dataset, we adopt the standard template following as (Zhang et al., 2024). Consistency across experiments is maintained by employing the same random seed as used in (Mallen et al., 2023), guaranteeing that all models are tested against a standardized set of exemplars. Consequently, the 15 exemplars provided to each model are uniformly generated using this identical seed. Besides, all questions are formatted using the prompt template ‘Q:

Table 6: Case study with LLaMA 2 (7B). We represent the factual error in red and the accurate information in blue.

Dataset	Question	PARAG-TAARE (Zhang et al., 2024)	EI-ARAG (Ours)
PopQA	Who is the mother of Melissa Benn ?	Query Type: (Non Retrieval) Generation: Hilary Mantel .	Query Type: (Retrieval) Retrieved Passages: 'Melissa Ann Benn (born 1957) is a British journalist and writer. She is the only daughter of Tony and Caroline Benn .' Generation: Caroline Benn .
PopQA	What is Peter W. Barca 's occupation?	Query Type: (Retrieval) Retrieved Passages: 'Fever/Dream is a play by Sheila Callaghan which premiered in 2009 at Woolly Mammoth Theatre Company in Washington, DC. It is a reinvention of Pedro Calderón de la Barca's play Life is a Dream.' Generation: actor .	Query Type: (Non Retrieval) Generation: politician .
TriviaQA	Against which South American country was Bolivia fighting during 'The War of the Pacific'?	Query Type: (Retrieval) Retrieved Passages: Bolivia then declared war on Chile and called upon Peru for help. Chile declared war on both Peru and Bolivia (April 5, 1879). Generation: Peru .	Query Type: (Non Retrieval) Generation: Bolivia fought against Chile during the War of the Pacific.
TriviaQA	Who was the first man to orbit the Earth?	Query Type: (Retrieval) Retrieved Passages: John Glenn , the First American to Orbit the Earth aboard ... Generation: John Glenn .	Query Type: (Non Retrieval) Generation: The first man to orbit the Earth was Yuri Gagarin , a Soviet cosmonaut.

<question> A:' to facilitate generative prediction. A 15-shot prompting method is adopted to maintain uniformity across both the GPT-Neo and LLaMA2 models. In our EI-ARAG, for the sentence classifier, we use and train a three-layer MLP, by considering a good tradeoff between the accuracy and latency. Specifically, the classifier is trained using the epoch that shows the best performance until 50 training iterations from the validation set, with the learning rate of 1e-3 and the Adam (Kingma and Ba, 2015) as an optimizer.

H Error Analysis on TriviaQA

We plot the error analysis on LLaMA2 7B using our EI-ARAG in Figure 3, which illustrates the ARAG cases in TriviaQA. The green area indicates when ARAG models activate retrieval, whereas the red area shows when they assess no external information is needed. The yellow area is where retrieval may or may not occur but predictions are incorrect predictions.

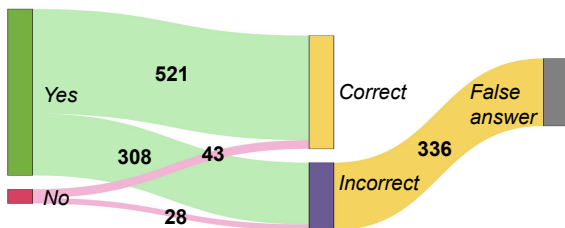


Figure 3: Sankey Diagram for our EI-ARAG method on the TriviaQA dataset.

I Case Study

We conduct a case study to qualitatively compare our Adaptive-RAG against previous Adaptive Retrieval (Zhang et al., 2024). Table 6 shows the classification complexity and the query handling patterns for both entity-centric (PopQA) and sentence-level (TriviaQA) questions.

Example 1: For this entity-centric question, TAARE Adaptive-RAG (Zhang et al., 2024) identifies that it is answerable by only using the LLM's parametric knowledge about 'Melissa Been' and generating the hallucination information. In contrast, our model fetches additional documents, leading to producing correct responses about 'Melissa Been'.

Example 2: In the sentence-level scenario involving historical conflict, PARAG-TAARE (Zhang et al., 2024) seeks out relevant information, including details like 'Peru', but incorrectly associates the event with 'Peru' due to an error in retrieving contextually relevant information. While our Adaptive Retrieval identifies this sentence knowledge have been stored in LLMs not need to request such information from external sources, resulting in accurate answers.

In the additional two examples in Table 6 we have similar observations, which validate the advantages of our approach.