Do Retrieval-Augmented Language Models Adapt to Varying User Needs?

Anonymous ACL submission

Abstract

Existing evaluations for Retrieval-Augmented Language Models (RALMs) often assume a single, optimal use of retrieved information, neglecting real-world scenarios where user needs dictate how to prioritize external context versus internal memory. We address this gap by introducing user-Need Instruction Compliance Evaluation (NICEval), a user-need-aware evaluation framework that separates factual correctness from faithfulness-to-instruction. The framework systematically assesses RALMs in three user need cases: **Context-Exclusive**, **Context-First**, and **Memory-First**, and three context settings: **Context Matching**, **Knowledge Conflict**, and **Information Irrelevant**, where the correct answer depends on both the instruction and retrieval quality. Through extensive experiments on multiple QA datasets, including our synthetic User-focused Retrieval-Augmented QA (URAQ) dataset, we find that restricting memory usage improves robustness in adversarial retrieval conditions but decreases peak performance with ideal retrieval results and model family dominates behavioral differences. Our findings highlight the necessity of user-centric evaluations in the development of retrieval-augmented systems and provide insights into optimizing model performance across varied retrieval contexts.

1 Introduction

Retrieval Augmented Language Models (RALMs) are deployed across diverse, knowledge-intensive tasks (Lewis et al., 2020), including Real-time Question Answering (Wang et al., 2024b), Educational Tutoring (Han et al., 2024), and Personal Assistants (Wang et al., 2024c). However, the optimal use of the retrieved information is not monolithic but varies with user needs expressed via instructions on whether to prioritize an external context (Es et al., 2023) or internal memory (Yu et al., 2023), or adhere strictly to specified contexts (Pipitone and Alami, 2024) that can be even counterfactual (Zhou et al., 2023). The ground truth changes accordingly. For instance, Real-time QA may rely heavily on updated external facts, whereas tutoring may draw more on the model's conceptual understanding. In addition, current RALMs still struggle to identify relevant references (Laban et al., 2024), resolve knowledge conflicts (Wang et al., 2024a), and reason effectively (Islam et al., 2024), when having low quality retrieval results (Park and Lee, 2024). These challenges show the need for robust evaluation about how RALMs adapt to various user need instructions under retrievals of different quality.

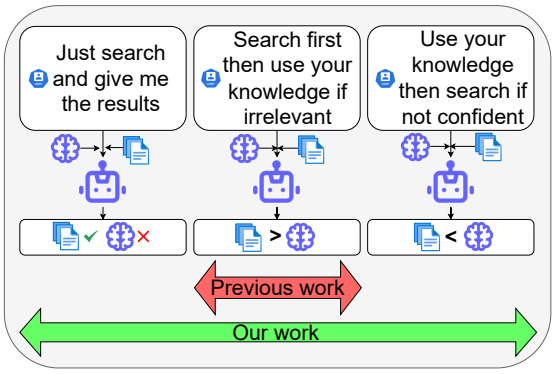

Figure 1: Users may have different directions on how to use retrieved context and internal memory as knowledge sources and most of the previous work only focused on a small portion of them.

The central gap in the current RALM evaluation is the assumption of singular optimality. Typically, the retrieved context should always be trusted and prioritized (Es et al., 2023; Chen et al., 2024; Ming et al., 2025). While benchmarks for multi-scenario evaluation exist (Friel et al., 2024; Zhu et al., 2024), they rarely condition the ground truth on user instructions that explicitly alter source prioritization. This obscures the critical distinction between *factual correctness* and *faithfulness to instruction*. Similarly, work on knowledge conflicts (Longpre et al., 2021; Xu et al., 2024b) has

proposed methods to enforce context-faithfulness (Shi et al., 2024; Zhou et al., 2023), but evaluations often overlook the hybrid strategies (e.g., memory as a fallback) or challenging conditions such as irrelevant context that are common in real applications. This narrow perspective overlooks how diverse user instructions can dramatically alter the desired model's behavior within the same scenario. In medical fact-checking, for instance, one user might demand answers derived only from peer-reviewed studies, while another relies on the model's internal knowledge—even if these sources conflict (Miao et al., 2024). Further related work discussion is in Appendix A.

This paper addresses these issues with NICEval, an evaluation framework that measures RALMs' ability to adapt to varying user needs. We consider three generic cases—**(1) Context-Exclusive**, **(2) Context-First**, and **(3) Memory-First**—to capture different degrees of reliance on context vs. memory usage. We also vary the **context settings—(a) Context Matching**, **(b) Knowledge Conflict**, and **(c) Information Irrelevant**—to represent scenarios where context may align with, contradict, or be unrelated to the query. By intersecting user cases with context settings, we closely mirror the complexities of real-world applications, where both the user's priorities and the reliability of retrieved information can shift dramatically.

We conduct extensive experiments on two public datasets, DisentQA (Neeman et al., 2023) and HotpotQA (Yang et al., 2018), along with our curated dataset, URAQ, evaluating two model families, Llama3.1 (Grattafiori et al., 2024) and Qwen2.5 (Qwen et al., 2025), across various model sizes and numbers of retrieved contexts. Our findings reveal that: 1) **Current LMs struggle to satisfy diverse user needs**, achieving below 50% accuracy across all datasets, with Llama-3.1-8B-Instruct occasionally nearing 0%. 2) **Contextual restriction alters performance**: Restricting models to rely solely on retrieved context improves LMs performance when external context content is different from internal memory by up to 23% accuracy difference. Increasing distracting context also abnormally increases the performance, possibly due to model laziness or trade-off between reasoning and instruction following capability. 3) **Model family dominate behavioral differences**: Model family contributes the majority of behavioral differences, which emphasize the importance of choosing the correct model for different user needs through proper evaluations.

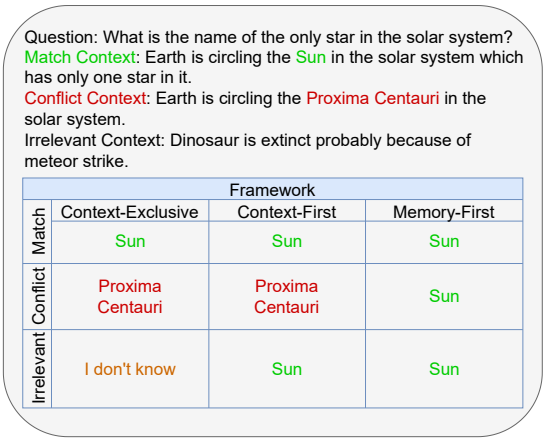

Figure 2: An illustration of the framework with an example question with its possible retrieved context and the ground truth answer under each situation.

2 NICEval Framework

In this section, we present NICEval framework. We first describe the design of three abstract **user need cases** in Section 2.1 representing different user needs expressed by context usages. Then, we describe the three **context settings** in Section 2.2 motivated by practical usage conditions in which the relevancy of the context varies and may conflict with the LMs' memory.

2.1 User Need Cases

To evaluate RALMs with varying *user needs*, we define a spectrum based on reliance of context versus memory. This spectrum, illustrated in Figure 2, consists of three distinct cases, determined by how LMs are instructed. Example prompts are in Appendix B. 1) **Context-Exclusive**: LMs must strictly base answers on retrieved context regardless of actual correctness, responding "I don't know" if context is unrelated. 2) **Context-First**: LMs prioritize the retrieved context but fall back on memory when no relevant context exists. 3) **Memory-First**: LMs rely on internal memory unless uncertain, in which case they defer to the retrieved context.

2.2 Context Settings

To better analyze RALMs under real-world situations with suboptimal retrieval results, we also consider the spectrum of context quality on top of each user case. For any context retrieved in

an RAG system, we can assess its quality based on two primary dimensions: 1) **Relevance to the Task or Question**, 2) **Alignment with LM's Internal Knowledge**. These two dimensions are reduced to the following three distinct context settings (relevant × match/conflict + irrelevant). a) **Context Matching**: There is at least one retrieved context *relevant* to the question and *matches* with the LM's memory. This is an ideal situation for RALMs as correct knowledge is presented in both the external context and the internal memory. b) **Knowledge Conflict**: There is at least one retrieved context *relevant* to the question but *conflicts* with the LM's memory. This setting simulates context-memory knowledge conflicts (Xu et al., 2024b) and tests the model's ability on generation with strictly following instructions regarding context usages. c) **Information Irrelevant**: All retrieved contexts are unrelated to the question. This setting evaluates the model's ability to recognize and disregard entirely irrelevant information, testing its knowledge selection when faced with a complete retrieval failure.

2.3 Evaluation Metrics

To rigorously assess user need awareness with different retrieval content, we test each user need with identical questions but varying the guidance on context usage, spanning three levels: 1) **Overall Accuracy**: The model must satisfy *all situations* simultaneously. Specifically, each test sample can be counted as correct only if the model can answer the same question under *all user need cases and context settings*. In this way, we can evaluate the LMs in a generic setting. 2) **Case-Level Accuracy**: For each user need case, a test sample is considered correct only if the model consistently provides the correct answer *across all context settings under that specific case*. This method ensures that the model demonstrates reliability in addressing a given requirement, independent of the context variations presented. 3) **Setting-Specific Accuracy**: In each context setting, a test sample is considered correct if the model's answer is the same as the ground truth. To ensure the robustness of the evaluation, for suboptimal **Knowledge Conflict** and **Information Irrelevant** settings, we also require that the test sample is correct for the same question under the optimal **Context Matching** setting. The formulas for the accuracy calculation are in the Appendix C.

3 Experiment & Analysis

3.1 Experiment Setup

This experiment employs three QA datasets: HotpotQA (Yang et al., 2018), DisentQA (Neeman et al., 2023), and our synthetic User-focused Retrieval-Augmented QA (URAQ) designed with controlled difficulty and high number of retrieval context. A detailed explanation of the URAQ curation pipeline, which can be reused in other fields, is in Appendix E. We conduct experiments on two distinct open-source LLM families: Llama-3.1 (Grattafiori et al., 2024), and Qwen2.5 (Qwen et al., 2025)—which vary in model size. We set the maximum context length to 128k, the temperature to 0, and Top-p to 1, while leaving all other configurations at their default values which defers to the Appendix D.

To examine how performance changes with varying amounts of retrieved context, rather than using a fixed retrieval count as in previous work (Zhu et al., 2024), we evaluate LM performance by exponentially increasing the retrieval count across different datasets, shown in Table 1. To assess the models' tolerance to distracting or irrelevant contexts, we ensure that only one relevant context is present for both the context-matching and conflicting settings, randomly positioned within the prompt. A detailed description of the prompt formatting and example is in the Appendix B. All other contexts are selected from a pool of *original* and *manipulated* knowledge that excludes any information directly related to the current question.

3.2 Overall Performance

We start our analysis on the overall performance across all three user cases by using the overall accuracy. The results are shown in Figure 3.

LMs struggle across all datasets No model surpasses 50% accuracy across different user needs, with Llama-3.1-8B-Instruct reaching nearing 0%. Among all datasets, URAQ is more challenging than DisentQA and HotpotQA. The best-performing model, Qwen2.5-72B-Instruct, scores up to 44.4% lower on URAQ. This result emphasizes the need for stronger reasoning capabilities to handle complex real-world user needs.

LMs behave differently at the model-family level Overall, we observe distinct patterns in LMs across different model families. On DisentQA and HotpotQA, the Qwen2.5 models ex-

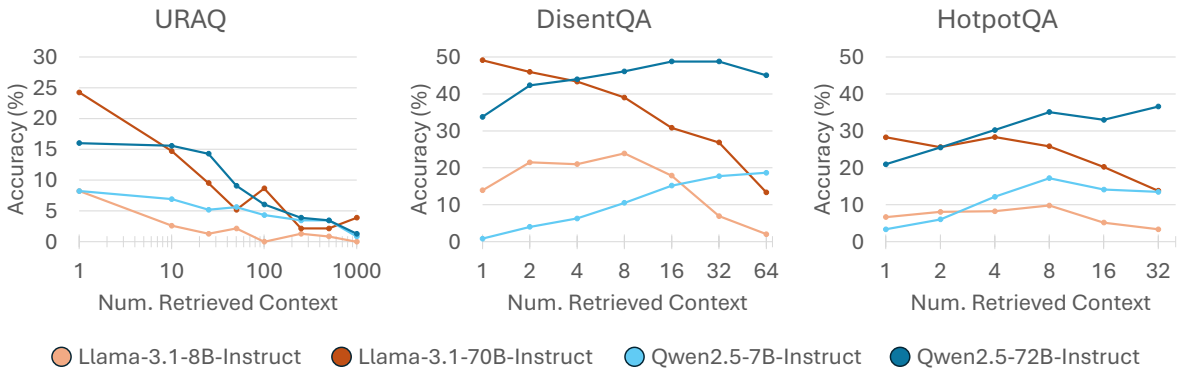

Figure 3: Overall user need performance curve of all models on each dataset.

hibit an increasing trend in accuracy as the number of retrieved contexts grows, whereas the Llama-3.1 models follow a decreasing trend. This difference likely stems from model-specific behavioral tendencies and a potential trade-off between instruction-following capability and multi-hop reasoning ability, which we further discuss in Section 3.3. On URAQ, the Llama-3.1 models experience a steeper drop in performance compared to the Qwen models. For example, the performance gap from 1 to 10 retrieved contexts in the Qwen family is around relative accuracy 1.5%, whereas for the Llama-3.1 family, it is 9.1%, indicating a more pronounced decline.

3.3 Performance for Each User Need Case

To further analyze the behavior of LMs on each user need, we measure the curve of *Case-Level Accuracy* versus number of retrieved context, as shown in Figure 5 in the Appendix F. We also conduct additional analysis on Setting-level in Appendix G, which further support our conclusions.

Restricting memory usage improves real-world performance. We find that the model's accuracy increased from *Context or Memory-First* to *Context-Exclusive* case, meaning that limiting the usage of internal memory improves the lower limit of general performance, possibly because *Context-Exclusive* strategy forces strict reliance on retrieved evidence and prevents hallucinations. This trend is particularly evident in Qwen2.5 models on HotpotQA dataset that maintain at least 7.7% increase in accuracy. However, as the number of context increases, the performance gap gradually shrinks and may even be inverted on Llama-3.1 models where *Context-Exclusive* accuracy drops by up to 12.5% when the number of retrieved context increases to 32.

Accuracy Increase as Num. of Retrieval Increase. To investigate the counterintuitive pattern in which the accuracy of *Context or Memory-First* cases increases as the number of retrieved contexts grows across all models, we analyze the impact of different context settings in both cases, as shown in Figure 8. Interestingly, the *Information Irrelevant* setting appears to contribute to this upward trend. By randomly sampling 100 cases across different retrieval context lengths, we observe that models are easily influenced by irrelevant information, often generating responses such as "no," "none," or "0." However, as more context is retrieved, models exhibit emergent reasoning capabilities. This phenomenon may stem from a form of "lazy" behavior, where models, instead of identifying the correct context, increasingly rely on their own memory as the context length grows. This can also be seen as the trade-off between reasoning and instruction following (Zhao et al., 2025; Fu et al., 2025). Because of the introduced noise from long context, models failed to follow the instructions about giving short and concise answers, which delimit their ability of reasoning as it allows more space for explicit Chain-of-Thought output and thus improves performance. A case study example is in Appendix H.

4 Conclusion

We introduced NICEval, a user need-ware framework to address a critical gap in RAG evaluation: the conflict of factual correctness with instruction faithfulness on user needs. By testing models across a 3×3 grid of user need cases and context settings, our framework and metrics provide a more nuanced understanding of RALM capabilities. This approach is a necessary step towards developing more reliable, adaptable, and truly user-centric retrieval-augmented systems.

Limitations

While our study provides a structured evaluation framework for Retrieval-Augmented Language Models (RALMs) under diverse user needs and retrieval conditions, several limitations remain. Our experiments rely on three datasets: HotpotQA, DisentQA, and the synthetic URAQ dataset. While these datasets cover various knowledge retrieval challenges, they may not fully capture the diversity of real-world retrieval scenarios, particularly in highly specialized domains such as medical or legal applications. Additionally, the synthetic URAQ dataset, although designed to control retrieval complexity, may not generalize perfectly to naturally occurring retrieval conflicts found in real-world settings. In addition, our results are based on evaluations of two model families, Llama-3.1 and Qwen2.5, across different sizes. While these models are representative of current state-of-the-art retrieval-augmented systems, our conclusions may not generalize to other architectures, such as retrieval-heavy fine-tuned transformers or proprietary models with distinct retrieval and reasoning mechanisms. Future work should extend this analysis to a broader range of models.

Ethics Statement

Our framework is designed to assess how well RALMs adhere to different user instructions, reflecting real-world applications where users may have distinct expectations regarding knowledge usage. However, models may still exhibit disparities in their ability to satisfy certain user needs, especially in adversarial retrieval settings. We recommend further research on mitigating disparities and enhancing fairness in retrieval-augmented systems. The datasets used in our experiments include HotpotQA, DisentQA, and the newly introduced synthetic URAQ dataset. While these datasets contain diverse question-answer pairs, we acknowledge that biases may be present in both retrieved and internally generated content. We have taken measures to minimize biases by curating synthetic data with balanced question difficulty and by evaluating model performance under varying retrieval conditions. However, residual biases in training corpora or retrieval mechanisms may influence the observed model behavior. One of our primary motivations is to analyze how models handle conflicting or irrelevant retrieved information. While our evaluation reveals scenarios where models fail to distinguish misinformation or exhibit hallucination tendencies, our work does not actively promote the generation or dissemination of false information. Instead, we highlight the need for more robust mechanisms to ensure factual consistency, particularly in knowledge-conflict scenarios. By conducting this study, we aim to advance the ethical design of retrieval-augmented models while encouraging further research on mitigating biases, improving factual robustness, and ensuring alignment with diverse user needs.

We acknowledge the use of the GPT-4o language model provided by OpenAI in the final stages of manuscript preparation. This tool was employed exclusively for identifying and correcting typographical and grammatical errors, ensuring clarity and precision in the written presentation. Its use was strictly limited to linguistic refinement and did not impact the study's conceptual framework, research methodology, data analysis, or conclusions. All intellectual contributions and substantive content remain those of the authors.

We intend to release our evaluation framework, code, and the URAQ dataset under a permissive license upon publication. Our use of existing public datasets, HotpotQA and DisentQA, is fully consistent with their original licenses and intended research purposes. These datasets are derived from public knowledge sources and are not expected to contain personally identifiable information (PII). For our synthetic URAQ dataset, the generation process was designed to create general and neutral knowledge statements. We have also conducted reviews to ensure the dataset is free from any PII or offensive content.

References

Dhananjay Ashok and Barnabas Poczos. 2024. Controllable text generation in the instruction-tuning era.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, T. W. Hennigan, Saffron Huang, Lorenzo Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and L. Sifre. 2021. Improving language models by

retrieving from trillions of tokens. In *International Conference on Machine Learning*.

Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'24/IAAI'24/EAAI'24. AAAI Press.

Shahul Es, Jithin James, Luis Espinosa-Anke, and Steven Schockaert. 2023. Ragas: Automated evaluation of retrieval augmented generation.

Robert Friel, Masha Belyi, and Atindriyo Sanyal. 2024. Ragbench: Explainable benchmark for retrieval-augmented generation systems.

Tingchen Fu, Jiawei Gu, Yafu Li, Xiaoye Qu, and Yu Cheng. 2025. Scaling reasoning, losing control: Evaluating instruction following in large reasoning models.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn,

Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. The llama 3 herd of models.

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: retrieval-augmented language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org.

Zifei FeiFei Han, Jionghao Lin, Ashish Gurung, Danielle R. Thomas, Eason Chen, Conrad Borchers, Shivang Gupta, and Kenneth R. Koedinger. 2024. Improving assessment of tutoring practices using retrieval-augmented generation.

Shayekh Bin Islam, Md Asib Rahman, K S M Tozammel Hossain, Enamul Hoque, Shafiq Joty, and Md Rizwan Parvez. 2024. Open-RAG: Enhanced retrieval augmented reasoning with open-source large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14231–14244, Miami, Florida, USA. Association for Computational Linguistics.

Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: few-shot learning with retrieval augmented language models. *J. Mach. Learn. Res.*, 24(1).

Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, Xiaojian Jiang, Jiexin Xu, Qiuxia Li, and Jun Zhao. 2024. Tug-of-war between knowledge: Exploring and resolving knowledge conflicts in retrieval-augmented language models.

Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.

Philippe Laban, Alexander R. Fabbri, Caiming Xiong, and Chien-Sheng Wu. 2024. Summary of a haystack: A challenge to long-context llms and rag systems.

Patrick 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 *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.

Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu, and Sanjiv Kumar. 2023. Large language models with controllable working memory. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1774–1793, Toronto, Canada. Association for Computational Linguistics.

Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. Entity-based knowledge conflicts in question answering. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Jing Miao, Charat Thongprayoon, Supawadee Suppadungsuk, Oscar A. Garcia Valencia, and Wisit Cheungpasitporn. 2024. Integrating retrieval-augmented generation with large language models in nephrology: Advancing practical applications. *Medicina*, 60.

Yifei Ming, Senthil Purushwalkam, Shrey Pandit, Zixuan Ke, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. 2025. Faitheval: Can your language model stay faithful to context, even if "the moon is made of marshmallows".

Ella Neeman, Roee Aharoni, Or Honovich, Leshem Choshen, Idan Szpektor, and Omri Abend. 2023. DisentQA: Disentangling parametric and contextual knowledge with counterfactual question answering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10056–10070, Toronto, Canada. Association for Computational Linguistics.

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao

Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi, Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Cunninghman, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. 2024a. Gpt-4o system card.

OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich, Andrey Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghorbani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botao Hao, Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, Camillo Lugaresi, Cary Bassin, Cary Hudson, Chak Ming Li, Charles de Bourcy, Chelsea Voss, Chen Shen, Chong Zhang, Chris Koch, Chris Orsinger, Christopher Hesse, Claudia Fischer, Clive Chan, Dan Roberts, Daniel Kappler, Daniel Levy, Daniel Selsam, David Dohan, David Farhi, David Mely, David Robinson, Dimitris Tsipras, Doug Li, Dragos Oprica, Eben Freeman, Eddie Zhang, Edmund Wong, Elizabeth Proehl, Enoch Cheung, Eric Mitchell, Eric Wallace, Erik Ritter, Evan Mays, Fan Wang, Felipe Petroski Such, Filippo Raso, Florencia Leoni, Foivos Tsimpourlas, Francis Song, Fred von Lohmann, Freddie Sulit, Geoff Salmon, Giambattista Parascandolo, Gildas Chabot, Grace Zhao, Greg Brockman, Guillaume Leclerc, Hadi Salman, Haiming Bao, Hao Sheng, Hart Andrin, Hessam Bagherinezhad, Hongyu Ren, Hunter Lightman, Hyung Won Chung, Ian Kivlichan, Ian O'Connell, Ian Osband, Ignasi Clavera Gilaberte, Ilge Akkaya, Ilya Kostrikov, Ilya Sutskever, Irina Kofman, Jakub Pachocki, James Lennon, Jason Wei, Jean Harb, Jerry Twore, Jiacheng Feng, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joaquin Quiñonero Candela, Joe Palermo, Joel Parish, Johannes Heidecke, John Hallman, John Rizzo, Jonathan Gordon, Jonathan Uesato, Jonathan Ward, Joost Huizinga, Julie Wang, Kai Chen, Kai Xiao, Karan Singhal, Karina Nguyen, Karl Cobbe, Katy Shi, Kayla Wood, Kendra Rimbach, Keren Gu-Lemberg, Kevin Liu, Kevin Lu, Kevin Stone, Kevin Yu, Lama Ahmad, Lauren Yang, Leo Liu, Leon Maksin, Leyton Ho, Liam Fedus, Lilian Weng, Linden Li, Lindsay McCallum, Lindsey Held, Lorenz Kuhn, Lukas Kondraciuk, Lukasz Kaiser, Luke Metz, Madelaine Boyd, Maja Trebacz, Manas Joglekar, Mark Chen, Marko Tintor, Mason Meyer, Matt Jones, Matt Kaufer, Max Schwarzer, Meghan Shah, Mehmet Yatbaz, Melody Y. Guan, Mengyuan Xu, Mengyuan Yan, Mia Glaese, Mianna Chen, Michael Lampe, Michael Malek, Michele Wang, Michelle Fradin, Mike McClay, Mikhail Pavlov, Miles Wang, Mingxuan Wang, Mira Murati, Mo Bavarian, Mostafa Rohaninejad, Nat McAleese, Neil Chowdhury, Neil Chowdhury, Nick Ryder, Nikolas Tezak, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Patrick Chao, Paul Ashbourne, Pavel Izmailov, Peter Zhokhov, Rachel Dias, Rahul Arora, Randall Lin, Rapha Gontijo Lopes, Raz Gaon, Reah Miyara, Reimar Leike, Renny Hwang, Rhythm Garg, Robin Brown, Roshan James, Rui Shu, Ryan Cheu, Ryan Greene, Saachi Jain, Sam Altman, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Santiago Hernandez, Sasha Baker, Scott McKinney, Scottie Yan, Shengjia Zhao, Shengli Hu, Shibani Santurkar, Shraman Ray Chaudhuri, Shuyuan Zhang, Siyuan Fu, Spencer Papay, Steph Lin, Suchir Balaji, Suvansh Sanjeev, Szymon Sidor, Tal Broda, Aidan Clark, Tao Wang, Taylor Gordon, Ted Sanders, Tejal Patwardhan, Thibault Sottiaux, Thomas Degry, Thomas Dimson, Tianhao Zheng, Timur Garipov, Tom Stasi, Trapit Bansal, Trevor Creech, Troy Peterson, Tyna Eloundou, Valerie Qi, Vineet Kosaraju, Vinnie Monaco, Vitchyr Pong, Vlad Fomenko, Weiyi Zheng, Wenda Zhou, Wes McCabe, Wojciech Zaremba, Yann Dubois, Yinghai Lu, Yining Chen, Young Cha, Yu Bai, Yuchen He, Yuchen Zhang,

Yunyun Wang, Zheng Shao, and Zhuohan Li. 2024b. Openai o1 system card.

Seong-Il Park and Jay-Yoon Lee. 2024. Toward robust RALMs: Revealing the impact of imperfect retrieval on retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 12:1686–1702.

Nicholas Pipitone and Ghita Houir Alami. 2024. Legalbench-rag: A benchmark for retrieval-augmented generation in the legal domain.

Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. Qwen2.5 technical report.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

Sagi Shaier, Ari Kobren, and Philip V. Ogren. 2024. Adaptive question answering: Enhancing language model proficiency for addressing knowledge conflicts with source citations. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17226–17239, Miami, Florida, USA. Association for Computational Linguistics.

Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Wen-tau Yih. 2024. Trusting your evidence: Hallucinate less with context-aware decoding. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 783–791, Mexico City, Mexico. Association for Computational Linguistics.

Fei Wang, Xingchen Wan, Ruoxi Sun, Jiefeng Chen, and Sercan Ö. Arık. 2024a. Astute rag: Overcoming imperfect retrieval augmentation and knowledge conflicts for large language models.

Yuhao Wang, Ruiyang Ren, Junyi Li, Xin Zhao, Jing Liu, and Ji-Rong Wen. 2024b. REAR: A relevance-aware retrieval-augmented framework for open-domain question answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5613–5626, Miami, Florida, USA. Association for Computational Linguistics.

Zheng Wang, Zhongyang Li, Zeren Jiang, Dandan Tu, and Wei Shi. 2024c. Crafting personalized agents through retrieval-augmented generation on editable memory graphs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4891–4906, Miami, Florida, USA. Association for Computational Linguistics.

Zhepei Wei, Wei-Lin Chen, and Yu Meng. 2024. Instructrag: Instructing retrieval-augmented generation via self-synthesized rationales.

Rongwu Xu, Brian Lin, Shujian Yang, Tianqi Zhang, Weiyan Shi, Tianwei Zhang, Zhixuan Fang, Wei Xu, and Han Qiu. 2024a. The earth is flat because...: Investigating LLMs' belief towards misinformation via persuasive conversation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16259–16303, Bangkok, Thailand. Association for Computational Linguistics.

Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024b. Knowledge conflicts for LLMs: A survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8541–8565, Miami, Florida, USA. Association for Computational Linguistics.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.

Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2023. Generate rather than retrieve: Large language models are strong context generators. In *International Conference for Learning Representation (ICLR)*.

Weixiang Zhao, Xingyu Sui, Jiahe Guo, Yulin Hu, Yang Deng, Yanyan Zhao, Bing Qin, Wanxiang Che, Tat-Seng Chua, and Ting Liu. 2025. Trade-offs in large reasoning models: An empirical analysis of deliberative and adaptive reasoning over foundational capabilities.

Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2023. Context-faithful prompting for large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14544–14556, Singapore. Association for Computational Linguistics.

Kunlun Zhu, Yifan Luo, Dingling Xu, Ruobing Wang, Shi Yu, Shuo Wang, Yukun Yan, Zhenghao Liu, Xu Han, Zhiyuan Liu, and Maosong Sun. 2024. Rageval: Scenario specific rag evaluation dataset generation framework.

Appendix

A Related Work

Our work intersects with four key research areas: (1) Retrieval-Augmented Generation Systems (Section A.1), (2) Knowledge Conflict Resolution (§A.2), and (3) RAG Evaluation Benchmarks (§A.3). We situate our framework within this landscape and highlight critical gaps in current approaches.

A.1 RAG Systems

Modern RAG systems built on foundational architectures like REALM (Guu et al., 2020) and DPR (Karpukhin et al., 2020), which first demonstrated the value of integrating neural retrieval with language modeling. Subsequent work improved context utilization through better attention mechanisms (RETRO (Borgeaud et al., 2021)) and multi-stage reasoning (Atlas (Izacard et al., 2023)). While these systems demonstrate impressive performance on knowledge-intensive tasks, they primarily optimize for single objective functions under the implicit assumption that retrieved context should always be prioritized. Recent work on controllable generation (Li et al. 2023; Ashok and Poczos 2024; Wei et al. 2024) begins to address this limitation but focuses on content style rather than source prioritization. We aim to raise the attention to diversified objectives of RAG system by this work about evaluating performance under different *user needs*.

A.2 Knowledge Conflict

The challenge of resolving conflicts between internal knowledge and external context has gained attention as LMs and RAG systems mature (Xu et al., 2024b). Early work by Longpre et al. (2021) identified context-memory conflicts as a key failure mode of LMs through evaluation on QA dataset. Subsequent works proposed multiple solutions, including but not limit to various fine-tuning, prompting, or decoding methods, to context-memory conflicts that require LM to be faithful to context in order to ignore outdated knowledge (Shi et al., 2024; Zhou et al., 2023) or faithful to memory in order to discriminate misinformation are rarely explored (Xu et al., 2024a). However, the hybrid strategies that utilize both context and memory with prioritization, although commonly appeared in real-world applications, are rarely explored. In addition, there also exists applications that require LMs and RAG systems to work along or accept fictitious information or knowledge, which are commonly ignored by the previous works. Our framework includes the hybrid strategies that stem from the fundamental *user needs*, providing a wider coverage of evaluating RALMs performance under context-memory conflict situations.

A.3 Recent RAG Benchmark

Previous RAG benchmarks like RAGAS (Es et al., 2023) and RGB (Chen et al., 2024) have facilitated progress by quantifying performance across various scenarios. However, many of these benchmarks focused on a single type of optimal setting in terms of context usages (for instance, always prioritizing the context), overlooking how different user instructions may drastically affect model behaviors and performances. Moreover, previous multi-scenario evaluations (Friel et al. 2024; Zhu et al. 2024), while covering a wide range of specific tasks and purpose abundant metrics for evaluating different aspects of RAG systems, also tend to follow the paradigm of focusing on singular optimality, neglecting that different user needs can actually happen in the same scenario, ultimately hindering the comprehensiveness of benchmark. FaithEval (Ming et al., 2025) proposes a benchmark to evaluate the faithfulness of the RAG system. Our work diverges by decoupling evaluation criteria from predefined singular optimality and measuring model capability to *adapt* to dynamic *user needs* by using different instructions. This mirrors real-world deployments where systems must honor diverse users' requirements rather than optimize for monolithic accuracy.

B Example Prompts and Formatting

This section provides full prompt templates for the three user need cases and details on input formatting.

B.1 Context-Exclusive Instruction

```
You are a helpful AI assistant tasked
    with answering the given question
    ONLY based on the provided
    information. Here are the
    requirements to answer the question:

1. The answer should be a numeric value,
     a boolean ("yes" or "no"), or an
    entity.
```

```
2. You MUST directly provide the final
    answer within an <output> XML tag,
    without including any units if the
    answer is numeric.

3. You MUST utilize the RELEVANT
    knowledge contained in the provided
    information to answer the question,
    even if the knowledge is INCORRECT.
    If NONE of the provided information
    is RELEVANT to the question, you
    MUST output "I don't know".
```

B.2 Context-First Instruction

```
You are a helpful AI assistant tasked
    with answering the given question by
     referring to the provided
    information. Here are the
    requirements to answer the question:

1. The answer should be a numeric value,
     a boolean ("yes" or "no"), or an
    entity.

2. You MUST directly provide the final
    answer within an <output> XML tag,
    without including any units if the
    answer is numeric.

3. If the provided information contains
    RELEVANT knowledge that can be used
    to answer the question, you MUST
    utilize the provided information,
    even if the knowledge is INCORRECT.

4. If NONE of the provided information
    is RELEVANT to the question, you
    MUST utilize your own knowledge to
    answer the question.
```

B.3 Memory-First Instruction

```
You are a helpful AI assistant tasked
    with answering the given question by
     referring to the provided
    information. Here are the
    requirements to answer the question:

1. The answer should be a numeric value,
     a boolean ("yes" or "no"), or an
    entity.

2. You MUST directly provide the final
    answer within an <output> XML tag,
    without including any units if the
    answer is numeric.

3. You MUST utilize your own knowledge
    to answer the question if you are
    certain of the accuracy (e.g.,
    factual information you are sure
    about). If you are UNSURE about your
     knowledge, you MUST use the
    relevant knowledge from the given
    information instead.
```

B.4 Input Prompt Formatting

The input prompt is organized as (I, C, Q) or (I_f, I_u, C, Q), where I is the instruction and can be separated into formatting instruction I_f and user needs instructions I_u, $C = \{c_1, c_2, ..., c_n\}$ is a series of retrieved context with retrieval number of n, and Q is the question. Given an input (I_f, I_u, C, Q), we have the following prompting template:

$$\text{<sys>}I_f \oplus I_u\text{</sys><user>}C \oplus Q\text{</user>} \quad (1)$$

where <sys></sys> and <user></user> denote the system prompt and the user prompt. Among all data samples, the I_u and C may change according to the **user case** and **context setting**, while the I_f remaining the same by instructing models to directly output a simple answer that is either a numeric value, a boolean ("yes" or "no"), or an entity.

We introduce an example input prompt that is designed for **Context-Exclusive and Context Matching** with 2 total retrieved context following the abstract input (I_f, I_u, C, Q). The prompt is formatted with XML for both input and output. Specifically, the formatting instructions I_f are separated into two parts: 1) The first and second instructions in the system prompt describing that the answer should be as simple as possible with XML format. 2) The instruction in the user prompt about format of context with an reinforcement of output format. The *user need* instruction I_u is at the third instruction in the system prompt. The retrieved context C is all the sentences in the user prompt within <info> xml tag. Finally, the question Q is at the last of user prompt within <question> xml tag.

```
System Prompt:
You are a helpful AI assistant tasked
    with answering the given question
    ONLY based on the provided
    information. Here are the
    requirements to answer the question:
1. The answer should be a numeric value,
     a boolean ("yes" or "no"), or an
    entity.
2. You MUST directly provide the final
    answer within an <output> XML tag,
    without including any units if the
    answer is numeric.
3. You MUST utilize the RELEVANT
    knowledge contained in the provided
    information to answer the question,
    even if the knowledge is INCORRECT.
    If NONE of the provided information
    is RELEVANT to the question, you
    MUST output 'I don't know'.
```

```
User Prompt:
Given the following information enclosed
     in xml tag <info></info>, answer
    the question enclosed in xml tag <
    question></question>.
**
<info>2 + 2 equals 4.</info>
<info>The sun is the only star in our
    solar system.</info>
**
<question>If you have two apples and you
     get two more, how many apples do
    you have?</question>
```

C Accuracy Calculation Formulas

C.1 Overall Accuracy

Let U be the set of all user need cases and S be the set of all context settings. For a given question i, the model's output for user case $u \in U$ and context setting $s \in S$ is $LM_{i,u}(C_s)$, and the ground truth is $A_{i,u,s}$. The Overall Accuracy (Acc_{oa}) is:

$$Acc_{oa} = \frac{1}{N} \sum_{i=1}^{N} \bigwedge_{u \in U} \bigwedge_{s \in S} \mathbb{I}(LM_{i,u}(C_s) = A_{i,u,s})$$

Where:

- N is the total number of test samples.
- $\mathbb{I}(\cdot)$ is the indicator function, which is 1 if the condition is true and 0 otherwise.
- $\bigwedge$ represents the logical AND operator, ensuring the condition is met for all combinations of user cases and context settings.

C.2 Case-Level Accuracy

Formally, given the language model as $LM()$, the case-level accuracy for the combination Acc_{case} can be written as:

$$Acc_{case} = \frac{1}{N} \sum_{i=1}^{N} \Big(\mathbb{I}(\text{LM}_i(C_a) = A_i^a) \wedge \mathbb{I}(\text{LM}_i(C_b) = A_i^b) \wedge \mathbb{I}(\text{LM}_i(C_c) = A_i^c) \Big) \quad (2)$$

where C_a, C_b, C_c are the series of external context for question i under **Setting (a) Context Matching, (b) Knowledge Conflict, (c) Information Irrelevant** respectively; A_i^a, A_i^b, A_i^c is the ground-truth answer for question i under **Setting (a) Context Matching, (b) Knowledge Conflict, (c) Information Irrelevant** respectively.

Dataset	Num. of Context Sequence	Size	Max. Token
Synthetic	1, 10, 25, 50, 100, 250, 500, 1000	231	25k
DisentQA	1, 2, 4, 8, 16, 32, 64	1415	59k
HotpotQA	1, 2, 4, 8, 16, 32	1274	35k

Table 1: Basic information of the three datasets used in the experiment. The number of retrieved context is increased in a exponential way until the average number of tokens at the highest number of each sequence reaches around 20k in order to balance the effectiveness of the experiment on long context and the consumption of computational resources. The number of maximum tokens among all samples for a dataset may vary based on context retrieved.

C.3 Setting-Specific Accuracy

Formally, given the language model as $LM()$, the accuracy for **Setting (a) Context Matching** Acc_a can be written as:

$$Acc_a = \frac{\sum_{i=1}^{N} \mathbb{I}(\text{LM}_i(C_a) = A_i^a)}{N} \quad (3)$$

where N is the total number of test samples; $\mathbb{I}(\cdot)$ is the indicator function (1 if the model generates the correct answer, 0 otherwise); C_a is the series of external context for question i under **Setting (a) Context Matching**; A_i^a is the ground-truth answer for question i under **Setting (a) Context Matching**.

Because for suboptimal **Knowledge Conflict** and **Information Irrelevant** settings, we also require that the test sample is correct for the same question under the optimal **Context Matching** setting, the accuracy for **Setting (b) Knowledge Conflict** Acc_b and **Setting (c) Information Irrelevant** can be written as:

$$Acc_b = \frac{1}{N} \sum_{i=1}^{N} \Big(\mathbb{I}(\text{LM}_i(C_a) = A_i^a) \wedge \mathbb{I}(\text{LM}_i(C_b) = A_i^b) \Big) \quad (4)$$

$$Acc_c = \frac{1}{N} \sum_{i=1}^{N} \Big(\mathbb{I}(\text{LM}_i(C_a) = A_i^a) \wedge \mathbb{I}(\text{LM}_i(C_c) = A_i^c) \Big) \quad (5)$$

D Detailed Experiment Setups

To assess RALMs' real-world performance, we use HotpotQA and DisentQA versions augmented with conflicted knowledge by Shaier et al. (2024)

for the retrieval-content knowledge conflict setting. While valuable, these benchmarks lack controlled knowledge boundaries and have varying question difficulty, limiting evaluation. They also rely on long-document contexts only, thereby restricting retrieval diversity in terms of document length. In addition, the nature of factual-based for these datasets makes them may not be fully aligned with the evaluation of under needs. URAQ complements these by providing uniformly difficult questions and numerous concise modified contexts, specifically to isolate instruction-following and conflict-resolution capabilities when adapting to varied user needs, distinct from general comprehension over long or highly complex factual texts. While specialized domain datasets (e.g., medical, real-time QA) would be ideal for demonstrating our three user needs, we opted for these known and synthetic benchmarks to ensure reproducibility, broader comparability, and generalizability within budget constraints. The framework itself remains applicable to more domain-specific evaluations.

E URAQ Dataset Curation Procedure

Below, we provide a step-by-step description of how we constructed the URAQ dataset:

E.1 Knowledge Generation

We used *gpt-4o-mini* (OpenAI et al., 2024a) to produce an initial list of short, simple knowledge statements. These statements are general facts (e.g., "A hummingbird can hover in mid-air" or "Blue whales are the largest animals on Earth") rather than domain-specific or specialized knowledge. The generated statements were deliberately kept concise and straightforward to facilitate subsequent manipulation and question generation.

E.2 Redundancy Filtering

Since GPT-based generators can produce highly similar or paraphrased statements, we employed *SentenceBERT* (Reimers and Gurevych, 2019) to measure the semantic similarity between all knowledge statements. Any pair of statements with a cosine similarity above 0.5 was considered near-duplicate and therefore removed to ensure diversity in the final knowledge set.

E.3 Manipulated Knowledge Creation

For every remaining "original" knowledge statement, we prompted *gpt-4o* to generate a *manipulated* variant. The manipulation involved either substituting key elements (e.g., entities, numerical values, or critical details) or adding a negation that changes the statement's truth value (e.g., "A hummingbird cannot hover in mid-air"). Each pair of statements (original vs. manipulated) thus serves as a pairwise contrast for subsequent question-answer (QA) creation.

E.4 Question-Answer (QA) Generation

From each pair of original and manipulated knowledge statements, we prompted *gpt-o1* (OpenAI et al., 2024b) to generate a question that requires between 1 to 5 *reasoning steps* to arrive at an answer. The reasoning steps typically involve either numerical computation, logical inference, or entity comparison. Each question was tied to both the original and the manipulated knowledge. The resulting QA format consists of one question and two different answers: one correct answer derived from the original statement, and a second answer derived from the manipulated statement. For multi-hop questions, where other knowledge is needed, these knowledge are included inside the question itself through explicit prompting the o1 model so that they will not interfere with the knowledge retrieval process.

E.5 Answer Format and Difficulty Selection

We constrained valid answers to be either (i) a numeric value, (ii) a boolean ("yes" or "no"), or (iii) a single entity. Among the generated questions, those requiring 4-hop reasoning were chosen for the final dataset, as a full manual inspection on the generated QA pairs suggested these exhibited higher quality and clearer multi-step logic compared to simpler or more complex variants.

E.6 Final Ground Truth Assignment

For each question, we designated the correct ground truth answer to be the one aligned with the *original* knowledge statement. An example illustrating how this ground truth is integrated into the evaluation framework is provided in Figure 2 of the main paper.

By following these steps, we ensure that the URAQ dataset offers well-defined pairs of knowledge (original vs. manipulated) and corresponding multi-step questions designed to differentiate between factual and altered information. This framework supports a diverse range of potential use

cases, from fact-checking systems to more elaborate multi-step reasoning models.

F Accuracy Curves of All Datasets

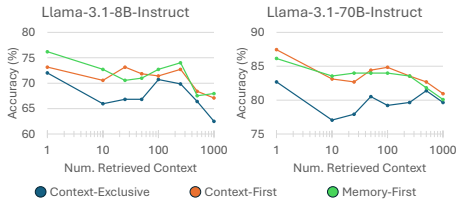

(1) Accuracy curve of Llama-3.1 on URAQ dataset under *Context Matching* setting.

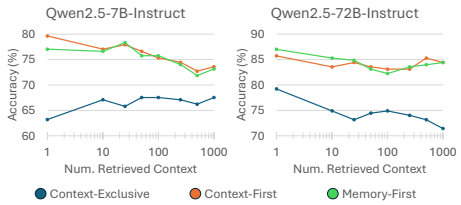

(2) Accuracy curve of Qwen2.5 on URAQ dataset under *Context Matching* setting.

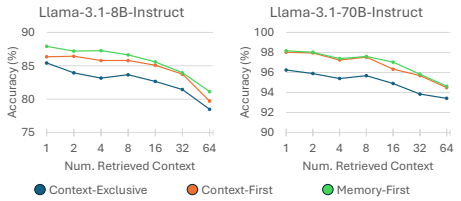

(3) Accuracy curve of Llama-3.1 on DisentQA dataset under *Context Matching* setting.

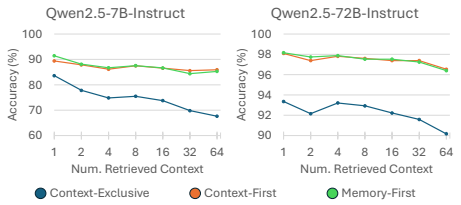

(4) Accuracy curve of Qwen2.5 on DisentQA dataset under *Context Matching* setting.

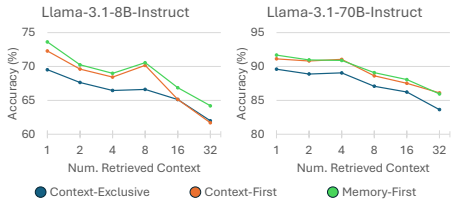

(5) Accuracy curve of Llama-3.1 on HotpotQA dataset under *Context Matching* setting.

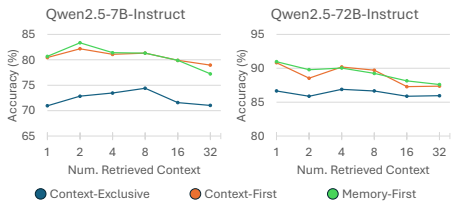

(6) Accuracy curve of Qwen2.5 on HotpotQA dataset under *Context Matching* setting.

Figure 4: Accuracy curve of all models under *Context Matching* setting.

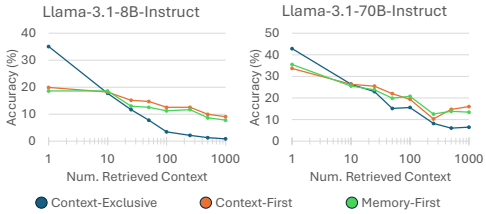

(1) Case-Level Accuracy curve of Llama-3.1 on URAQ dataset.

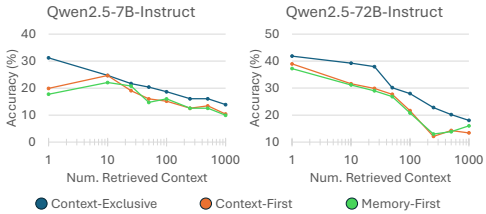

(2) Case-Level Accuracy curve of Qwen2.5 on URAQ dataset.

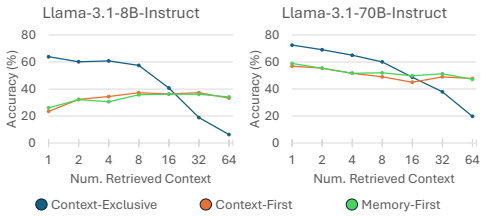

(3) Case-Level Accuracy curve of Llama-3.1 on DisentQA dataset.

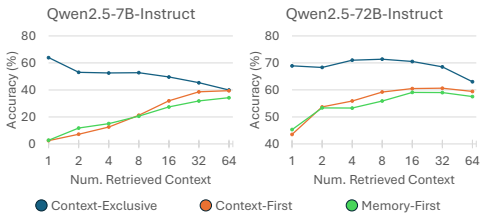

(4) Case-Level Accuracy curve of Qwen2.5 on DisentQA dataset.

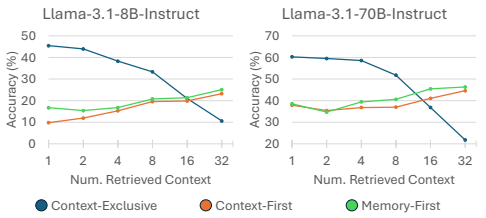

(5) Case-Level Accuracy curve of Llama-3.1 on HotpotQA dataset.

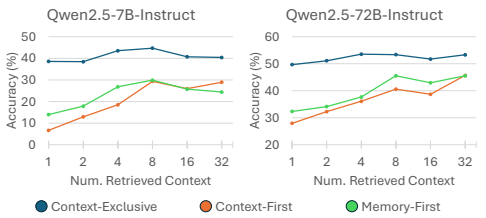

(6) Case-Level Accuracy curve of Qwen2.5 on HotpotQA dataset.

Figure 5: Case-Level Accuracy of all models.

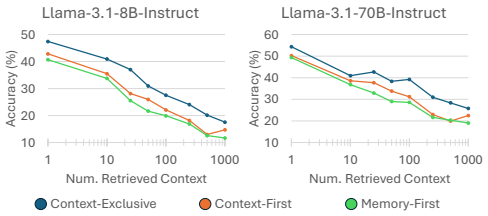

(1) Accuracy curve of Llama-3.1 on URAQ dataset under *Context Matching & Knowledge Conflict* setting.

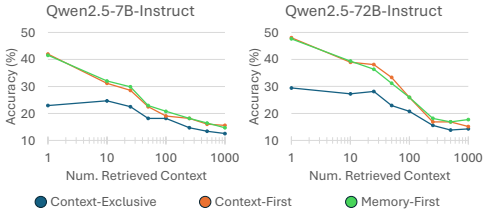

(2) Accuracy curve of Qwen2.5 on URAQ dataset under *Context Matching & Knowledge Conflict* setting.

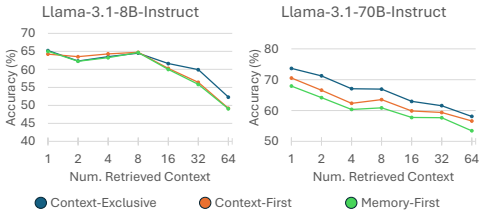

(3) Accuracy curve of Llama-3.1 on DisentQA dataset under *Context Matching & Knowledge Conflict* setting.

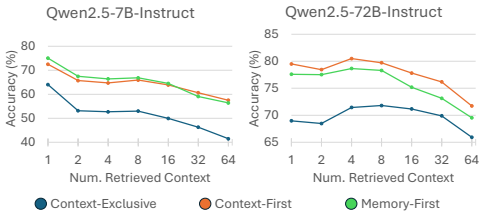

(4) Accuracy curve of Qwen2.5 on DisentQA dataset under *Context Matching & Knowledge Conflict* setting.

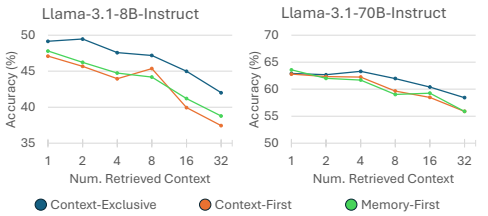

(5) Accuracy curve of Llama-3.1 on HotpotQA dataset under *Context Matching & Knowledge Conflict* setting.

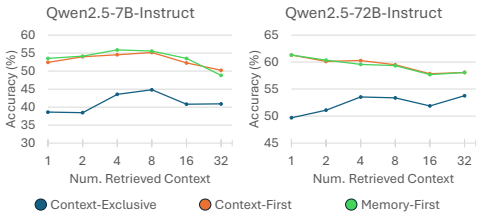

(6) Accuracy curve of Qwen2.5 on HotpotQA dataset under *Context Matching & Knowledge Conflict* setting.

Figure 6: Accuracy curve of all models under *Context Matching & Knowledge Conflict* setting.

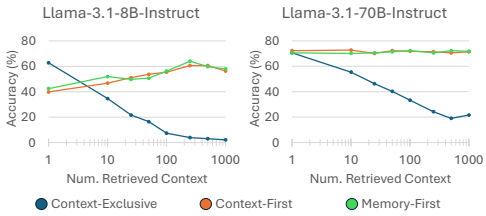

(1) Accuracy curve of Llama-3.1 on URAQ dataset under *Context Matching & Information Irrelevant* setting.

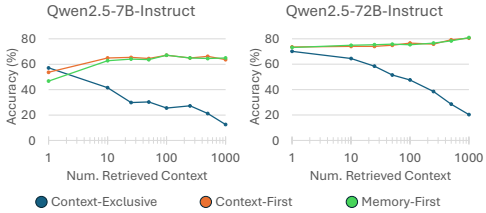

(2) Accuracy curve of Qwen2.5 on URAQ dataset under *Context Matching & Information Irrelevant* setting.

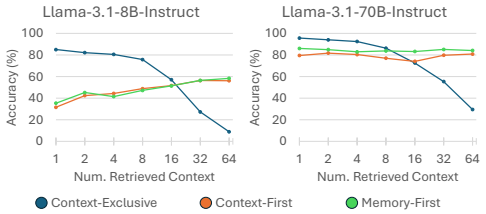

(3) Accuracy curve of Llama-3.1 on DisentQA dataset under *Context Matching & Information Irrelevant* setting.

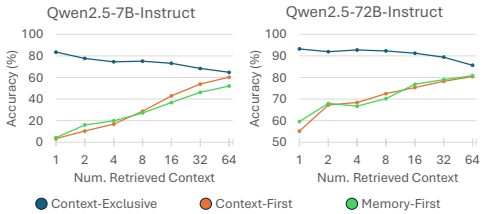

(4) Accuracy curve of Qwen2.5 on DisentQA dataset under *Context Matching & Information Irrelevant* setting.

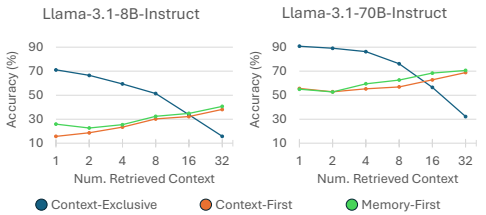

(5) Accuracy curve of Llama-3.1 on HotpotQA dataset under *Context Matching & Information Irrelevant* setting.

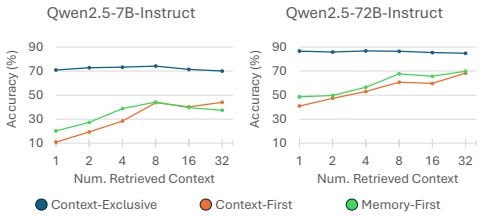

(6) Accuracy curve of Qwen2.5 on HotpotQA dataset under *Context Matching & Information Irrelevant* setting.

Figure 7: Accuracy curve of all models under *Context Matching & Information Irrelevant* setting.

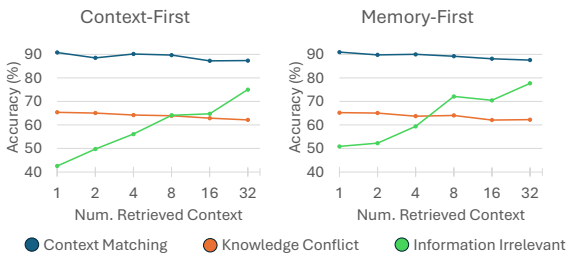

Figure 8: Accuracy curve of Qwen2.5-72B-Instruct on HotpotQA dataset under all context settings with *Context-First* and *Memory-First*.

Dataset	Llama-3.1-Instruct		Qwen2.5-Instruct	
	8B (%)	70B (%)	7B (%)	72B (%)
URAQ	52	74	85	97
DisentQA	70	84	92	98
HotpotQA	63	76	84	95

Table 2: Percentage of errors that is "I don't know" among the shortest 100 randomly selected samples that under *Context Matching* setting that is **incorrect** for *Context-Exclusive* user need and **correct** for *Context or Memory-First*.

G Individual Setting Performance

To provide more detailed analysis on models' behavior on the context setting-level, we measure the *Setting-Specific Accuracy* curve for each user need case, categorizing them into two groups: **Optimal Context**, where the provided context aligns with the model's memory, and **Challenging Context**, where the context is conflicting or irrelevant.

G.1 Performance on Optimal Context

Under the *Context Matching* setting, where the model receives fully relevant and correct context, we assess its maximum potential performance. This defines an **optimal performance**, isolating the model's ability to utilize ideal context without retrieval constraints.

Restricting memory usage limits optimal performance. Based on the results in Figure 4, we observe that models' accuracy declines when internal memory is restricted under the *Context-Exclusive* strategy. This effect is more pronounced in the Qwen2.5 family, where Qwen2.5-7B-Instruct experiences up to a 12.1% accuracy drop from *Context or Memory-First* to *Context-Exclusive*, whereas the Llama-3.1 family shows only a slight decrease, with Llama-3.1-8B-Instruct losing up to 4.1%.

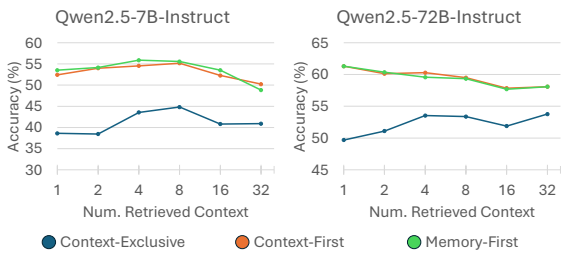

(1) Setting-Specific Accuracy curve of Qwen2.5 model family on HotpotQA dataset with knowledge conflict.

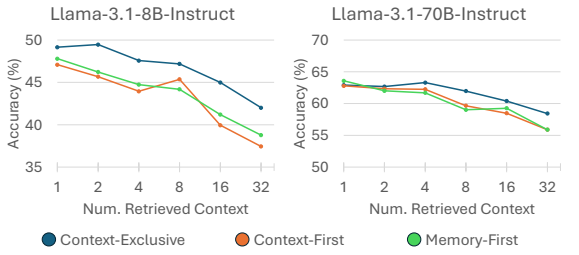

(2) Setting-Specific Accuracy curve of Llama-3.1 model family on HotpotQA dataset with knowledge conflict.

Figure 9: Setting-Specific Accuracy curve of Qwen2.5 and Llama-3.1 model family on HotpotQA dataset with knowledge conflict. While two models have similar accuracy on *Context or Memory-First* case, Llama models has higher accuracy on *Context-Exclusive* compared with *Context or Memory-First* and Qwen models has lower accuracy.

LLMs exhibit self-protective conservatism. To examine the accuracy drop under the *Context-Exclusive* setting, we analyze 100 randomly selected cases with up to four retrieved context segments, where the model provides an incorrect answer under *Context-Exclusive* but a correct one under *Context or Memory-First*. Errors are categorized into two types: (1) the model refuses to answer by stating, "I don't know," and (2) the model generates an incorrect hallucinated response. Table 2 reports the percentage of refusals.

We observe that models overwhelmingly prefer rejection over hallucination when they struggle to locate relevant context, with refusal rates exceeding 50% across all models and datasets. This tendency is particularly strong in the Qwen2.5 family, where the 7B and 72B models reject answers in over 85% of cases, with Qwen2.5-72B-Instruct reaching a 98% rejection rate on DisentQA. Similarly, the Llama-3.1 models exhibit high rejection rates, ranging from 70% to 84% on DisentQA. This conservative behavior may stem from its training objectives or alignment strategies prioritizing answer correctness over speculative responses.

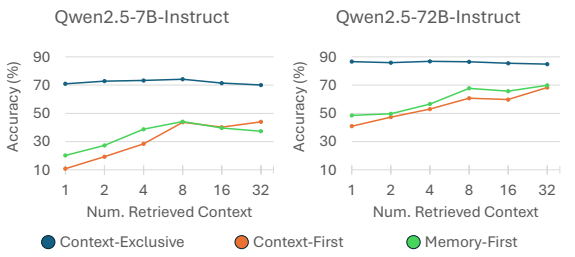

(1) Setting-Specific Accuracy curve of Qwen2.5 model family on HotpotQA dataset with irrelevant context.

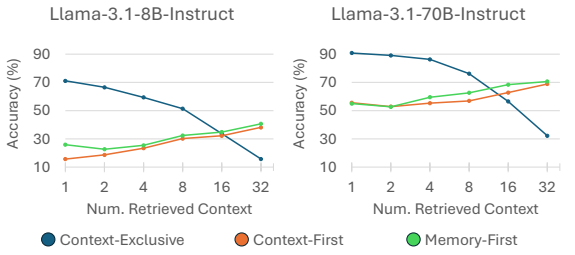

(2) Setting-Specific Accuracy curve of Llama-3.1 model family on HotpotQA dataset with irrelevant context.

Figure 10: Setting-Specific Accuracy curve of Qwen2.5 and Llama-3.1 model family on HotpotQA dataset with irrelevant context.

G.2 Performance with Challenging Context

For performance under *Knowledge Conflict* or *Irrelevant Context*, we realize that evaluating only the performance of single context setting in isolation can introduce bias and skewed interpretations due to LMs preference on using memory than context or vise versa (Longpre et al., 2021; Jin et al., 2024), resulting performing perfectly in one setting but failed in other. For example, succeeding in *Irrelevant Context* but failing in *Matching Context* may suggest that the model is prone always relying on memory without actually complying with the instructions to use retrieved context. Therefore, we measure the *Setting-Specific Accuracy* for Challenging Context in a way that the same question need to be also answered correctly in *Context Matching* settings, ensuring the robustness of evaluation. Such measuring method is applied to all experiments in this section shown in Figure 9 and 10.

Model family dominates behavioral difference. Model families still exhibit distinct behavioral patterns: When knowledge conflict exists as Figure 6, Llama3.1 models show an increase of performance from *Context-First* and *Memory-First* to *Context-Exclusive* case for up to 10.2% accuracy, while Qwen2.5 models demonstrate the opposite trend with a decrease close to 20%.

	Context-First (%)	Memory-First (%)
Qwen2.5-72B-Instruct	84	77
Llama-3.1-70B-Instruct	56	65

Table 3: Percentage of testing samples that answered with single negative output for short input but correct output with explicit reasoning, among 100 randomly selected samples that the question answered incorrectly with 1 retrieved context and correctly with 1000 retrieved context.

This behavior suggests fundamental differences in knowledge reliance—Llama3.1 appears more context-dependent, struggling to effectively integrate memory, whereas Qwen2.5 leverages its parametric knowledge more effectively when permitted. Such difference also appears in the as Figure 7 with *Information Irrelevant* setting, Llama models exhibit significant decreasing accuracy on *Context-Exclusive* strategy with increasing context length for up to 60.1%, whereas Qwen exhibit almost no loss in performance, for the same reason as discussed in Section 3.3.

H Case Study

```
Question: How do plants create their
    food, photosynthesis or moonthesis?

Answer with 1 context: <output>No</
    output>

Answer with 1000 context: I think there
    might be a small mistake in your
    question. The correct term is "
    photosynthesis" not "moonthesis".
    Photosynthesis is the process by
    which plants, algae, and some
    bacteria make their own food using
    sunlight, water, and carbon dioxide.
     So, to answer your question: Plants
     create their food through
    photosynthesis.
```

Above question-answer pair from our synthetic URAQ dataset is an example of model relying on their own memory on long context and acquire better result than the result from shorter input. For answer with only 1 context, the model outputs a single word "No" with correct output format. For answer with 1000 context, the model provides clear thinking path with correct, affirmative answer without the desinated XML format for output. We also calculate the percentage of 100 randomly selected testing samples that has similar behaviors using Qwen2.5-72B-Instruct and Llama-3.1-70B-Instruct as shown in Table 3.