

Meta-Cognitive Analysis: Evaluating Declarative and Procedural Knowledge in Datasets and Large Language Models

Zhuoqun Li^{1,3}, Hongyu Lin¹, Yaojie Lu¹, Hao Xiang^{1,3}, Xianpei Han^{1,2}, Le Sun^{1,2,✉}

¹Chinese Information Processing Laboratory ²State Key Laboratory of Computer Science

Institute of Software, Chinese Academy of Sciences, Beijing, China

³University of Chinese Academy of Sciences, Beijing, China

{lizhuoqun2021, hongyu, luyaojie, xianghao2022, xianpei, sunle}@iscas.ac.cn

Abstract

Declarative knowledge and procedural knowledge are two key parts in meta-cognitive theory, and these two hold significant importance in pre-training and inference of LLMs. However, a comprehensive analysis comparing these two types of knowledge is lacking, primarily due to challenges in definition, probing and quantitative assessment. In this paper, we explore from a new perspective by providing ground-truth knowledge for LLMs and evaluating the effective score. Through extensive experiments with widely-used datasets and models, we get conclusions: (1) In most tasks, benefits from declarative knowledge are greater than those from procedural knowledge. (2) Profits of procedural knowledge are larger than declarative knowledge only in reasoning tasks with simple logic. (3) As pre-training progresses and size increases, model ability to utilize both kinds of knowledge significantly improves, but in different speed. We do detailed analysis for the findings and this can provide primary guidance for evaluation and enhancement of large language models.

Keywords: large language model, declarative knowledge, procedural knowledge

1. Introduction

Recent advancements in large language models (LLMs) have been noteworthy, with models such as GPT4 (OpenAI, 2023), Llama (Touvron et al., 2023a) and Vicuna (Chiang et al., 2023) leading the way. These models are capable of solving various NLP tasks through an autoregressive approach and have show impressive performance (Ye et al., 2023; Bang et al., 2023). This evolution of LLMs has push NLP into a new era, moving away from traditional task-specific pre-train fine-tuning paradigm (Zhao et al., 2023).

According to metacognitive theories (Brown, 1987; Jacobs and Paris, 1987), there are two kinds of knowledge that critically contributes to the cognition of human beings: *declarative knowledge* and *procedural knowledge*. Declarative knowledge refers to “knowing that”. It encompasses facts, concepts, and information that can be explicitly verbalized or described (Ryle, 1945, 2009). For instance, knowing the capital of France is Paris or understanding the concept of gravity is declarative knowledge. Procedural knowledge, on the other hand, refers to “knowing how”. It pertains to skills and procedures that we might not be able to verbalize explicitly but can demonstrate through action (Ryle, 1945, 2009). For example, knowing how to ride a bike, play a musical instrument, or swim are all forms of procedural knowledge. This type of knowledge is typically acquired through practice and repetition and is believed to involve the basal ganglia and cerebellum in the brain. Declarative

and procedural knowledge are at the core of human cognition, underpinning our ability to function, learn, and adapt in various environments.

For large language models, many previous work has confirmed the existence of declarative and procedural knowledge in them. For declarative knowledge, it has been observed that LLMs are encoded with declarative knowledge across many domains (Zhong et al., 2023; Huang et al., 2023), which is learned from the pre-training on vast portions of Internet. For procedural knowledge, many researches have found that LLMs exhibit an understanding of processes or sequences of actions, to an extent. For instance, they can generate code snippets, describe step-by-step instructions, or help with problem-solving in certain domains, which significantly contribute to their ability to resolve complicated tasks (Gao et al., 2022; Wei et al., 2023; Zhou et al., 2023).

Despite acknowledgment of their existence and significance within LLMs (Wei et al., 2023; Shi et al., 2023; Fu et al., 2023), there is a marked absence of systematic exploration into how these two kinds of knowledge exert on the capabilities of LLMs. This oversight considerably hampers our comprehensive understanding of large model dynamics and effective improvement of LLM abilities in certain tasks, domain and scenarios. A potential explanation for this absence stems from inherent challenges associated with formulating and probing these two kinds of knowledge and quantitatively assessing their implications on LLMs, respectively because of the black box nature of LLMs. Therefore, it is challenging to determine

✉ Corresponding Author

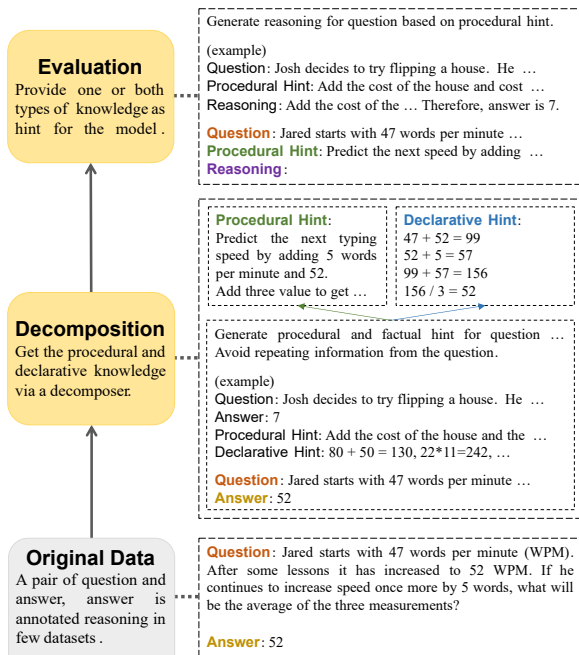


Figure 1: The illustration of overall processing. First decompose original reasoning of the question to procedural knowledge and declarative knowledge, then evaluate models by providing one or both type of knowledge after the question.

whether declarative or procedural knowledge is more critical for enhancement of current LLMs, as well as the impact of different types of knowledge on various models, tasks, and training phases. The lack of such studies significantly hampers efforts to target improvements in model capabilities.

In this paper, we investigate the impact of declarative and procedural knowledge on datasets and large language models from a fresh perspective. Instead of directly probing these types of knowledge from LLMs, we examine how introducing ground truth of these two knowledge types affects the model performance on specific tasks. For this purpose, we design an exploration method based on in-context knowledge injection. By providing the necessary ground truth declarative or procedural knowledge in the question input to the LLM, we compare the model performance with and without this knowledge. This helps us understand the potential benefits of adding specific types of knowledge to LLMs. In this setup, the performance results can be roughly seen as the maximum capability of LLMs for a given task when they are not missing certain knowledge.

Specifically, our method primarily uses an in-context approach, providing the LLMs with three types of information specific to the question: 1) declarative hints: This type of information includes all the declarative knowledge needed to solve the current problem; 2) procedural hints: This type pro-

vides step-by-step plans of how to solve the problem; 3) combined hints: This information includes both of the above hints. It can be seen as the model maximum ability to use and combine the provided information to complete the task when given all the necessary details. We conduct experiments on 32 openly available large language models and 13 evaluation datasets that cover kinds of different tasks including math, commonsense and reasoning¹. From our experiments, we find that:

- In most tasks, benefits from declarative knowledge are greater than those from procedural knowledge.
- Profits of procedural knowledge are larger than declarative knowledge only in reasoning tasks with simple logic.
- As pre-training progresses and size increases, model ability to utilize both kinds of knowledge significantly improves, but in different speed.

2. Related Work

Many works do illustrate declarative and procedural knowledge is important in LLM training and inference. LLMs can use chain of thought (Wei et al., 2023; Wang et al., 2023; Zhou et al., 2023) to solve complex tasks. Retrieval augmented LLMs (Ram et al., 2023; Shi et al., 2023) can use knowledge from knowledge bases or Internet to bolster model accuracy. LLMs can also be used to construct knowledge graph (Cohen et al., 2023). In addition, some works try to inject declarative knowledge (Kang et al., 2023) or procedural knowledge (Fu et al., 2023) during model training. However, there is no comprehensive analysis for these two types of knowledge in datasets and LLMs.

3. Meta-Cognitive Analysis Method

Procedural knowledge and declarative knowledge are two important aspects shared by most tasks. Exploring the difficulty of testing tasks and the capabilities of models from these two perspectives is of great significance for improving the training and testing of models. However, these two aspects, procedural knowledge and declarative knowledge, are coupled together in testing tasks. It requires a good method to decouple and quantify these two aspects, in order to analyze capabilities of models and the difficulty of test data.

In this paper, we address above challenges by a knowledge decomposition and injection method.

¹Our detailed source codes are openly available at <https://github.com/Li-Z-Q/meta-cognitive-analysis>

Dataset	Description	Procedural	Declarative	Combined	Delta
GSM8K	Elementary Arithmetic	5.14	2.15	6.68	-2.99
MultiArith	Elementary Arithmetic	5.15	3.70	7.69	-1.45
CommonsenseQA	Commonsense	1.17	-0.06	1.45	-1.23
ARC-Easy	Commonsense	0.41	2.05	2.87	1.64
ARC-Challenge	Commonsense	0.47	3.00	3.27	2.53
TruthfulQA	Commonsense	-0.01	2.61	2.50	2.62
MMLU-STEM	Understanding	0.72	0.80	1.26	0.08
MMLU-Humanities	Understanding	0.59	1.53	2.11	0.94
MMLU-Social	Understanding	0.64	3.46	3.56	2.82
MMLU-Other	Understanding	0.06	3.05	3.49	2.99
MATH-1	High School Mathematics	0.89	3.36	4.62	2.47
MATH-2	High School Mathematics	1.58	4.62	6.19	3.04
MATH-3	High School Mathematics	1.33	4.52	6.07	3.19

Table 1: Procedural, declarative and combined score of all datasets. The delta is how much larger the declarative score is than procedural score, it shows that procedural score is larger than declarative score only in GSM8K, MultiArith and CommonsenseQA. In addition, procedural scores of GSM8K and MultiArith are much larger than other datasets.

Firstly, we provide clear definitions of procedural and declarative knowledge, and utilize GPT4² to decouple the original reasoning into these two types of knowledge. Then, during the evaluation process, we provide questions and knowledge related to one aspect as hints to the model, observing the performance improvement of the model under these prompts. Subsequently, we transform these improvement value to scores and then do quantitative analysis on models and tasks.

3.1. Decomposition of Declarative and Procedural Knowledge

In order to quantify and analyze procedural and declarative knowledge, we first define these two types and then use GPT4 to decouple the original reasoning process into procedural knowledge and declarative knowledge.

Specially, we define declarative knowledge as fundamental facts crucial for resolving questions, each fact is independent with one another. Conversely, procedural knowledge embodies a generalized strategy essential for solving tasks, without any specific declarative details. Figure 1 shows one example, explaining the distinctions between declarative and procedural knowledge.

With clear define of procedural and declarative knowledge, leveraging remarkable capabilities of GPT4, we do decomposition and get declarative and procedural knowledge for different tasks. In detail, by providing two examples for GPT4 (each example including a question, answer, and corresponding declarative knowledge and procedural knowledge), we let GPT4 decouple the two

aspects of knowledge for unannotated question-answer data.

3.2. Evaluation via In-context Hint Knowledge Injection

After we obtain the decoupled data with procedural knowledge and declarative knowledge, we proceed to provide the model with a specific type of knowledge as a hint during the evaluation process. By observing the performance improvement resulting from these hints, we can further analyze the model capabilities in procedural and declarative aspects, the difficulty of tasks, and how effectively the hints are utilized.

Specially, during evaluation, models are given question with declarative knowledge, procedural knowledge, a combination of both, or neither, in contextual format to generate reasoning. One model input example of procedural type is shown in Figure 1, other types are similar with it. In a certain dataset-model pair, define e_o is the original error rate, e_p is the error rate after giving procedural knowledge, the procedural score is:

$$score_p = \frac{e_o - e_p}{e_o}$$

This score means effect of procedural knowledge. Via this metric, we can also get $score_d$ for declarative knowledge and $score_c$ for combination of both knowledge. With all dataset-model pairs score, to get a model score, we do average across all datasets and vice versa.

3.3. Datasets and Models

To derive more general conclusions regarding procedural and declarative aspects, our experimental

²<https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo>

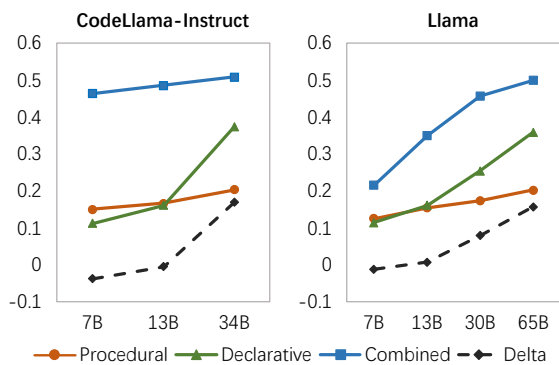


Figure 2: Procedural, declarative and combined score of different size models. The black line is the difference between procedural and declarative score. Above figure shows that the ability to utilize both knowledge becomes stronger as the model size increases, with different improvement rate.

dataset encompasses varying degrees of knowledge and reasoning difficulty, our model incorporates diverse pre-training processes and sizes.

In terms of datasets, for mathematics, we select easy mathematical datasets such as GSM8K (Cobbe et al., 2021) and MultiArith (Roy and Roth, 2015), and hard mathematical dataset MATH (Hendrycks et al., 2021b)—of which we opt for levels 1, 2, and 3, excluding levels 4 and 5 due to their extreme complexity. For commonsense reasoning, we choose datasets including CommonsenseQA (Talmor et al., 2019), ARC-Easy (Clark et al., 2018), ARC-Challenge (Clark et al., 2018), and TruthfulQA (Lin et al., 2022). Additionally, we choose MMLU (Hendrycks et al., 2021a) benchmark, which assesses many LLMs across four subcategories: humanities, social, STEM, and other. Note that while we focus on these selected datasets, our method is general and can be applied across all kinds of datasets.

In terms of large language models, we select widely-used models including Llama (Touvron et al., 2023a), Vicuna-v1.3 (Chiang et al., 2023), Llama-2 (Touvron et al., 2023b), Llama-2-Chat (Touvron et al., 2023b), Vicuna-v1.5 (Zheng et al., 2023), Vicuna-v1.5-16K (Zheng et al., 2023), CodeLlama-Instruct (Rozière et al., 2023), and GPT3.5 API³. Furthermore, we assess 11 checkpoints from Baichuan-2-7B (Yang et al., 2023), each separated by an interval of 220 steps. Note that we only show a portion of models result in Figure 2 and Figure 3.

It is noteworthy that in some datasets, declarative hints may inherently encompass procedural el-

³<https://platform.openai.com/docs/guides/text-generation/chat-completions-api>

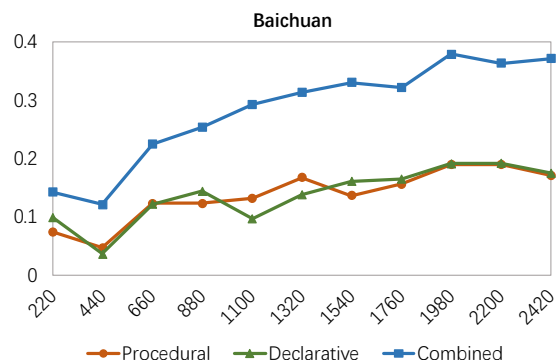


Figure 3: Scores of different checkpoints. The 220 means Baichuan-2-7B-00220, the model after 220 steps pre-training. It shows that the ability to utilize both type of knowledge becomes stronger as the pre-training step increases.

ements due to sequencing. To address this, we collect noise declarative hints, mixing them with declarative hints, followed by randomization.

4. Findings

Finding 1. *In most tasks, benefits from declarative knowledge are greater than those from procedural knowledge.*

As shown in Table 1, we calculate the procedural score, declarative score, combined score and difference between procedural and declarative score, we observe that declarative score is larger than procedural score in 9 datasets. Thus we get conclusion that benefits from declarative knowledge is greater in most tasks.

Finding 2. *Procedural knowledge profits are larger than declarative knowledge only in reasoning tasks with simple logic.*

As shown in Table 1, it shows that in simple mathematical datasets such as GSM8K and MultiArith, and basic commonsense reasoning dataset like CommonsenseQA, benefits from procedural knowledge are more than other datasets. This meets our expectations, as tasks like mathematics and commonsense reasoning often require more logical knowledge, commonly covered within procedural knowledge, hence golden procedural knowledge can improve a lot in these tasks.

On the other hand, for tasks with more complex logic, such as high school mathematics (MATH), we do not observe a significantly higher benefit from procedural knowledge than declarative knowledge, and gains from procedural knowledge are lower compared with simpler logic tasks. We speculate that this may be due to the limited capabilities of our test models. When the logic of question be-

comes overly complex, the model might struggle to understand and utilize logical information from procedural knowledge, resulting in the introduction of procedural knowledge not providing larger additional benefits.

Finding 3. *As pre-training step and size increases, model ability to utilize both kinds of knowledge significantly improves, but in different speed.*

To find out effects of different model size and pre-training step, we do experiments in Llama, Code-Llama-Instruct and Baichuan model. Specifically, to observe impacts of model size in knowledge utilization, we compare performance of Llama 7B, 13B, 30B, 65B, and Code-Llama-Instruct 7B, 13B, 34B. To observe effects of pre-training step, we examine the performance of Baichuan-2-7B at pre-training step from 220 to 2420.

In terms of model size, models with more parameters show clear improvement in capturing both declarative and procedural knowledge, as shown in Figure 2. Simultaneously, improvement in capturing declarative knowledge is significantly higher than procedural knowledge. This indicates that larger models are easier to utilize external declarative information, while relevant procedural abilities might more rely on model inherent capabilities.

In terms of pre-training steps, as shown in Figure 3, it shows a steady improvement in both type of knowledge as pre-training progresses. Note that basic performance of the model is continuously improving with more pre-training steps, we find that even above this baseline, the benefits of introducing additional knowledge continue to increase. This suggests that benefits of deeper training likely go beyond just enhancing the model's internal knowledge capabilities, mainly improving the model ability to utilize knowledge.

5. Conclusion

In this paper, we conduct a comprehensive analysis of declarative and procedural knowledge, via a novel perspective. For each pair of dataset and model, we provide ground truth knowledge and then evaluate this knowledge effect score. Our experiments get insightful conclusions regarding the significance of these two types of knowledge across diverse datasets, and efficacy with which different models utilize them. These findings provide primary guidance in enhancing evaluation and improvement processes of LLMs.

6. Acknowledgements

We sincerely thank all anonymous reviewers for their insightful comments and valuable sugges-

tions. This work is supported by the Strategic Priority Research Program of Chinese Academy of Sciences under Grant XDA27020200 and the Natural Science Foundation of China (No. 62122077 and 62106251).

7. Bibliographical References

Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Love-nia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity](#). *ArXiv preprint*, abs/2302.04023.

Ann L Brown. 1987. Metacognition, executive control, self-regulation, and other more mysterious mechanisms. *Metacognition, motivation, and understanding*, pages 65–116.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#).

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#).

Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. 2023. [Crawling the internal knowledge-base of language models](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1856–1869, Dubrovnik, Croatia. Association for Computational Linguistics.

Ruth Dannenfelser, Jeffrey Zhong, Ran Zhang, and Vicky Yao. 2023. [Into the single cell multiverse: an end-to-end dataset for procedural knowledge extraction in biomedical texts](#).

Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. [Specializing smaller language models towards multi-step reasoning](#).

- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2022. [Pal: Program-aided language models](#). *ArXiv preprint*, abs/2211.10435.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring mathematical problem solving with the math dataset](#).
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. [C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models](#).
- Janis E Jacobs and Scott G Paris. 1987. Children’s metacognition about reading: Issues in definition, measurement, and instruction. *Educational psychologist*, 22(3-4):255–278.
- Minki Kang, Seanie Lee, Jinheon Baek, Kenji Kawaguchi, and Sung Ju Hwang. 2023. [Knowledge-augmented reasoning distillation for small language models in knowledge-intensive tasks](#). *ArXiv preprint*, abs/2305.18395.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, Edouard Grave, Yann LeCun, and Thomas Scialom. 2023. [Augmented language models: a survey](#).
- OpenAI. 2023. [Gpt-4 technical report](#).
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhl-gay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [In-context retrieval-augmented language models](#).
- Subhro Roy and Dan Roth. 2015. [Solving general arithmetic word problems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752, Lisbon, Portugal. Association for Computational Linguistics.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. [Code llama: Open foundation models for code](#).
- Gilbert Ryle. 1945. Knowing how and knowing that: The presidential address. In *Proceedings of the Aristotelian society*, volume 46, pages 1–16. JS-TOR.
- Gilbert Ryle. 2009. *The concept of mind*. Routledge.
- Gregory Schraw and David Moshman. 1995. Metacognitive theories. *Educational psychology review*, 7:351–371.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Min-joon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. [Replug: Retrieval-augmented black-box language models](#).
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucu-rull, 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, Rannan 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, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#).
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#).
- Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, Fan Yang, Fei Deng, Feng Wang, Feng Liu, Guangwei Ai, Guosheng Dong, Haizhou Zhao, Hang Xu, Haoze Sun, Hongda Zhang, Hui Liu, Jiaming Ji, Jian Xie, JunTao Dai, Kun Fang, Lei Su, Liang Song, Lifeng Liu, Liyun Ru, Luyao Ma, Mang Wang, Mickel Liu, MingAn Lin, Nuolan Nie, Peidong Guo, Ruiyang Sun, Tao Zhang, Tianpeng Li, Tianyu Li, Wei Cheng, Weipeng Chen, Xiangrong Zeng, Xiaochuan Wang, Xiaoxi Chen, Xin Men, Xin Yu, Xuehai Pan, Yanjun Shen, Yiding Wang, Yiyu Li, Youxin Jiang, Yuchen Gao, Yupeng Zhang, Zenan Zhou, and Zhiying Wu. 2023. [Baichuan 2: Open large-scale language models](#).
- Junjie Ye, Xuanning Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, Zeyang Zhou, Chao Gong, Yang Shen, Jie Zhou, Siming Chen, Tao Gui, Qi Zhang, and Xuanjing Huang. 2023. [A comprehensive capability analysis of gpt-3 and gpt-3.5 series models](#). *ArXiv preprint*, abs/2303.10420.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Z. Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jianyun Nie, and Ji rong Wen. 2023. [A survey of large language models](#). *ArXiv preprint*, abs/2303.18223.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).
- Wanjuan Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. [AgiEval: A human-centric benchmark for evaluating foundation models](#).
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#).