

Leveraging Large Language Models for Conversational Multi-Doc Question Answering: The First Place of WSDM Cup 2024

Yiming Li^{1,3,*}, Zhao Zhang^{2,3,*}

¹Beijing Key Laboratory of Mobile Computing and Pervasive Device,
Institute of Computing Technology, Chinese Academy of Sciences

²CAS Key Laboratory of Network Data Science and Technology,
Institute of Computing Technology, Chinese Academy of Sciences

³University of Chinese Academy of Sciences
{liyiming22s1,zhangzhao22s}@ict.ac.cn

Abstract

Conversational multi-doc question answering aims to answer specific questions based on the retrieved documents as well as the contextual conversations. In this paper, we introduce our winning approach for the “Conversational Multi-Doc QA” challenge in WSDM Cup 2024, which exploits the superior natural language understanding and generation capability of Large Language Models (LLMs). We first adapt LLMs to the task, then devise a hybrid training strategy to make the most of in-domain unlabeled data. Moreover, an advanced text embedding model is adopted to filter out potentially irrelevant documents, and several approaches are designed and compared for the model ensemble. Equipped with all these techniques, our solution finally ranked 1st place in WSDM Cup 2024, surpassing its rivals to a large extent. The source codes have been released at <https://github.com/zhangzhao219/WSDM-Cup-2024>.

1 Introduction

Conversational question answering (Zaib et al., 2022), which aims to generate correct and meaningful answers according to users’ intents identified from the dialog, plays a crucial role in modern search engines and conversational systems. However, it is still challenging, especially with current or trending topics as timely knowledge is inaccessible during the training stage of the language model. Although providing multiple relevant documents as contextual information seems feasible, the model is still at risk of being overwhelmed or misled by the massive input. Based on the real-world textual data, Xiaohongshu, the WSDM Cup 2024¹ presents a challenge of “Conversational Multi-Doc QA” to encourage further exploration of the problem.

Recently, LLMs, such as ChatGPT, have demonstrated impressive performance on several natural

language processing tasks. It is promising to solve the challenge by leveraging the understanding and reasoning abilities of LLMs. However, lots of factors, including the design of training configurations and the existence of irrelevant documents, still hinder the improvement of generation quality.

In this work, to unleash the power of LLMs, we first formulate the task as a multi-turn conditional generation problem with different LLMs. Then, a multi-stage hybrid training pipeline is conducted to incorporate an unlabeled eval set as an additional training corpus. To remove potentially irrelevant information, we implement certain strategies, including a state-of-the-art embedding model namely Nomic Embed (Nussbaum et al., 2024) to compute the similarity score between inputs and documents. Finally, several methods are considered to evaluate the quality of answers generated by various LLMs approximately before selecting the best response as the final answer for the model ensemble. Experimental results show that our solution achieves the highest score on each evaluation metric, far beyond the teams behind us, while ablation studies also suggest the effectiveness of the proposed techniques.

2 Preliminary

2.1 Dataset

For the competition, participants are required to train a model to produce answers corresponding to dialogue history (composed of sequential qa pairs), reference documents, and the final question. To this end, each training sample is organized as follows,

```
history: {q1}{a1}{q2}{a2}...{qn}{an}
documents: {d1}{d2}...{dn}
question: {q}
answer: {a}
```

But for the eval and test set, the answer field is invisible while an additional field, namely keywords is incorporated to comprehensively evaluate the

* Equal Contribution.

¹<https://sites.google.com/view/wsdm24-docqa>

generated answer, which will be detailed later.

2.2 Evaluation

Three metrics are involved to evaluate the lexical and semantic relatedness of generated answers. The definitions are detailed as follows,

ROUGE-L The intuition is that the longer the LCS (the longest common subsequence) of two sentences is, the more similar the two sentences are (Lin, 2004). As a result, LCS-based F-measure can be used to estimate the similarity between two sentences X of length m and Y of length n ,

$$R_{lcs} = \frac{LCS(X, Y)}{m}$$

$$P_{lcs} = \frac{LCS(X, Y)}{n}$$

$$F_{lcs} = \frac{(1 + \beta^2)R_{lcs}P_{lcs}}{R_{lcs} + \beta^2P_{lcs}}$$

Specifically, word-level ROUGE-L (abbreviated as W-ROUGE-L below) calculates F_{lcs} of the words in sentences, which is more focused on the accuracy of specific words, while character-level ROUGE-L (abbreviated as C-ROUGE-L below) pays more attention to word forms, grammar, and punctuation because of calculating F_{lcs} of the characters.

Keywords Recall Keywords Recall (abbreviated as KR below) focuses on whether the specific keywords of truth sentence X have appeared in the generated answer Y . If there are m reference keywords in X , and n of them appear in Y ,

$$KR = \frac{n}{m}$$

3 Methodology

3.1 Generation Baseline with LLMs

To adapt LLM to this task, we carefully design the input format and concatenate each textual component together in the following order,

$$u = \{\mathbf{q1}\}\{\mathbf{a1}\}\{\mathbf{q2}\}\{\mathbf{a2}\}\dots\{\mathbf{qn}\}\{\mathbf{an}\}\{\mathbf{q}\}\{\mathbf{d1}\}\{\mathbf{d2}\}\dots\{\mathbf{dn}\}\{\mathbf{a}\}$$

Note that we exclude special tokens (e.g., $\langle s \rangle$, [INST]) in above line for simplicity.

Then, the model θ can be trained by maximizing the log-likelihood over the whole sequence.

$$\mathcal{L}_{gen} = - \sum_{i=1}^{|\mathbf{u}|} m_i \log p(u_i | \mathbf{u}_{<i}; \theta)$$

where $p(u_i | \mathbf{u}_{<i}; \theta)$ is the probability to select a token u_i at step i given previous tokens $\mathbf{u}_{<i}$, and m_i is the loss mask for the i th token. Specifically, there are two modes to determine m_i : 1) the single-turn mode, which means that $m_i = 1$ if and only if u_i belongs to $\{\mathbf{a}\}$. 2) the multi-turn mode, $m_i = 1$ as long as u_i belongs to $\{\mathbf{a}\}$ or $\{\mathbf{ai}\}$. We conduct a toy experiment to examine them with the Llama2-13B-base model (Meta, 2023), the results are shown in Table 1. It can be seen that the multi-turn mode results in a better performance as it forces the LLM to pay more attention to contextual information.

Mode	C-ROUGE-L	KR
single-turn	0.5761	0.6291
multi-turn	0.5845	0.6359

Table 1: Performance of the specific model trained with single-turn and multi-turn modes on the eval dataset.

After deciding on the input format and mask mode, we compare lots of off-the-shelf LLMs, which are either pretrained only or instruction tuned. As shown in Table 2, the SOLAR-10.7B-Instruct model surpasses its counterparts a lot on the eval dataset, which uses depth up-scaling to scale LLM and fine-tuned for instruction-following capabilities (Kim et al., 2023). Therefore, it is chosen as our backbone in the subsequent experiments.

LLM	C-ROUGE-L	KR
ChatGPT (zero-shot) (OpenAI, 2022)	0.4815	0.5387
GPT-4 (zero-shot) (OpenAI et al., 2024)	0.5069	0.5537
Yi-6B (Yi, 2023)	0.5286	0.6464
ChatGLM3-6B (Zeng et al., 2023)	0.5740	0.6068
ChatGLM3-6B-Base (Zeng et al., 2023)	0.5782	0.6166
DeepSeek 7B Base (DeepSeek-AI, 2024)	0.5783	0.6184
Llama 2 Chat 13B (Meta, 2023)	0.5821	0.6266
Yi-6B-Chat (Yi, 2023)	0.5833	0.6365
Llama 2 13B (Meta, 2023)	0.5845	0.6359
Mistral 7B (Jiang et al., 2023)	0.6031	0.6489
Mistral 7B-Instruct (Jiang et al., 2023)	0.6048	0.6558
SOLAR 10.7B (Kim et al., 2023)	0.6099	0.6627
SOLAR 10.7B-Instruct (Kim et al., 2023)	0.6104	0.6691

Table 2: Performance of different LLMs on the eval dataset.

3.2 Hybrid Training

Appropriately labeled texts from a similar distribution may contribute a lot to the improvement of LLMs’ generation performance. During phase 2, we propose to utilize a well-trained model to produce (pseudo) answers for the eval dataset before

adding them to the original training set to finetune a new model from scratch. The intuition for the above hybrid training strategy is two-fold, on one hand, it can be viewed as the knowledge distillation process on in-domain unlabeled data, on the other hand, since we only generate the final target $\{\mathbf{a}\}$ in a pseudo labeling manner, $\{\mathbf{a}_i\}$ are still officially annotated, which may be beneficial for the multi-turn setting. Note that we do not further involve the test dataset for hybrid training as it may over-calibrate the model and thus weaken the model performance in the final evaluation, which is also examined by our empirical practices.

3.3 Noisy Document Filter

<p>Question: What are the advantages of Python? Noisy Doc: Listing the advantages of Python. <i>Almost rephrase the question</i> 😊</p> <hr/> <p>Question: So if I want to change the information of the order, such as the delivery address or delivery time, what do I need to do? Noisy Doc: No more worrying about your girlfriend getting angry! <i>Irrelevant information</i> 😞</p>
--

Figure 1: Two examples of noisy documents

There is no doubt that high-quality reference documents can not only help mitigate the hallucination phenomena but also enhance the inference quality of LLMs (Li et al., 2023). After manually scrolling through the whole dataset, we find there are mainly two types of noisy documents, the examples of which can be found in Figure 1.

- Documents almost rephrasing the question, which shares extremely high relevant scores with the document.
- Documents full of irrelevant information, so they share extremely low relevant scores with the question or history.

Therefore, it is vital to quantify the relevance without the existence of ground truth answers. From both semantic and lexical views, we come up with the following two indicators,

Embedding-level Cosine Similarity We adopt an advanced text embedding model, Nomic Embed, to compute the cosine similarity between documents and the corresponding question (or together with conversational history).

Word- or character-level ROUGE-L As illustrated before, the ROUGE-L scores can be viewed as lexical relatedness criteria.

Practically, we set a higher threshold τ_h and a lower one τ_l for each indicator separately, before screening out reference documents whose corresponding scores $\geq \tau_h$ or $\leq \tau_l$ for manual inspection. As a result, we filter out 193 noisy documents in phase 2.

Moreover, prior work (Liu et al., 2023) suggests that important passages placed at the beginning or end of the input can be better comprehended by LLMs. However, we find that there is a strong correlation between document indexes and their relative orders occurring at the officially annotated answers, which means that reranking the reference documents can result in severe performance degradation.

3.4 Model Ensemble

The model ensemble has proven to be effective in discriminative tasks, however, it is rarely explored under generative settings. In this work, we propose to approximately evaluate the quality of generated answers from different models and then select the best one as the final result. Suppose that given a test sample, we have M candidate responses to aggregate, for each candidate r_i , we calculate the relevance scores $s(r_i, r_j)$ between r_i and $r_j (j = 1, \dots, M, j \neq i)$ and add them together as the quality score q_i for $r_i (q_i = \sum_j s(r_i, r_j))$. Similarly, the relevance quantizers can be embedding-level cosine similarity (denoted as `emb_a_s`), word-level ROUGE-L (denoted as `word_a_f`), and character-level ROUGE-L (denoted as `char_a_f`). The motivation is that the final answer should be a representative who reaches a consensus with the most candidates.

4 Experiments

4.1 Experimental Settings

Thanks to the LLM fine-tuning framework provided by modelscope (Zhao et al., 2024), we can fine-tune and infer our model easily. All our experiments are carried out on NVIDIA A100 80G and V100 32G GPUs. As for the fine-tuning process, we employ configuration for the LoRA (Hu et al., 2021) fine-tuning, defined with the following parameters: `lora_rank` (rank of the low-rank matrices) set to 8, `lora_alpha` (scaling factor for learning rate) at 16, `lora_dropout` to manage overfitting set at 0.05, and `lora_target_modules` focused on all modules. The maximum length of input sequences was set to 3072. The AdamW (Loshchilov

Rank	Participant	W-ROUGE-L	C-ROUGE-L	KR
1	regtrh (ours)	0.46536	0.62084	0.69535
-	regtrh* (ours)	0.45548	0.61420	0.68359
2	wangkxu	0.44961	0.61183	0.67230
3	zhangmin186	0.45098	0.61048	0.66239
4	Ted	0.45013	0.61030	0.66513
5	tilbur	0.44554	0.60489	0.67718

Table 3: Top scores of the competition. We also list the results of our single model (marked with *) for comparison.

and Hutter, 2019) optimizer is employed for training with a learning rate set to $1e-4$ with a warm-up proportion of 0.03 and a batch size of 1 on each GPU. This model is trained for 4 epochs via deepspeed ZeRO-2, and we choose checkpoint 1700 as our best model. As for the inferring process, we set `do_sample=false` to use greedy decoding and ensure stable output. The repetition penalty is adjusted to be between 1.00 and 1.02, and `max_new_tokens` is set to 512. We use vLLM (Kwon et al., 2023) to accelerate our inference process, and it takes about 40 minutes to infer our final result on a V100 32G GPU.

4.2 Competition Results

Table 3 lists the final results of this competition. As shown, our solution, which aggregates outcomes from 8 different models, achieves 1.6%, 0.9%, and 2.3% absolute performance gains on W-ROUGE-L, C-ROUGE-L, and KR, respectively when compared to the 2nd place. Besides, it is noteworthy that our single model can also yield better performance than others, suggesting the effectiveness of our pipelines.

4.3 Ablation Studies

Ablation Study of the Noisy Document Filter.

Table 4 shows the experimental results of our single model inferred with and without the noisy document filter (abbreviated as NDF below). We find that it marginally improves the final scores since the provided documents are carefully chosen by the cup organizers and LLMs can somewhat tell apart the underlying distractors.

NDF	W-ROUGE-L	C-ROUGE-L	KR
-	0.45547	0.61368	0.68243
✓	0.45548	0.61420	0.68359

Table 4: Ablation study of the noisy document filter on the test dataset.

Ablation Study of the Hybrid Training Strategy.

We verify the effects of the proposed hybrid training strategy in Table 5. As seen, incorporating the

eval set with the corresponding pseudo targets can largely boost the generation quality, especially for the keywords recall score. But further incorporation of the test set has little effect, which validates our design choice.

Method	W-ROUGE-L	C-ROUGE-L	KR
-	0.45547	0.61368	0.68243
- hybrid training	0.45284	0.61308	0.67154
+ hybrid training	0.45476	0.61343	0.67626

Table 5: Ablation study of the hybrid training strategy on the test dataset.

Ablation Study of the Model Ensemble.

We first compare different ensemble methods as displayed in Figure 2 (a). Although the mentioned methods are all competitive on ROUGE scores, `emb_a_s` brings about many more improvements in keywords recall, thereby being selected as our final ensemble method. Then, a parameter analysis of the number of candidates for the ensemble is conducted. As seen from Figure 2 (b), more candidates generally lead to better performance. Due to the limited time and budget, we finally defined the number as 8.

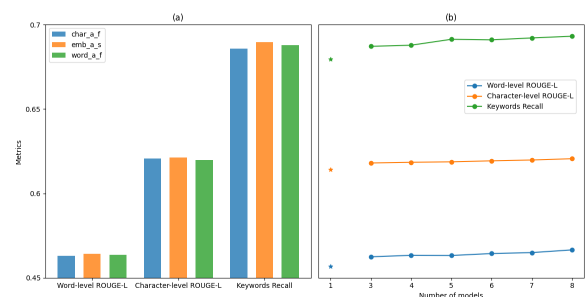


Figure 2: (a) generation performances on the test dataset with different ensemble approaches (5 candidates for ensemble); (b) generation performances on the test dataset with different numbers of candidates for ensemble (with the approach working best in (a));

5 Conclusion

In this paper, we detail our winning solution to the task of ‘‘Conversational Multi-Doc QA’’ in WSDM Cup 2024. Leveraging the abilities of LLMs, we

use the SOLAR-10.7B-Instruct model as the backbone, with the combination of hybrid training, noisy document filter, and selecting the best response by evaluating the quality from 8 results as our final submission. Our solution won 1st place on the public leaderboard.

Limitations

We would like to discuss our results on additional automated metrics, like METEOR (Banerjee and Lavie, 2005), BLEU (Papineni et al., 2002), BERTSCORE (Zhang* et al., 2020), and even LLM as a judge. However, the organizers have not made the standard answer public. As a result, we also can't compare with more RAG-based technique comparisons like REPLUG (Shi et al., 2024), RLHF-based tuning, or knowledge distillation.

Ethical Statement

In this research, it's crucial to acknowledge the potential limitations of LLMs, which may also produce biased information, especially on questions related to people. However, we do not adopt any additional processing for these answers, while other parts of the training data come from datasets by organizers. We keep fair and honest in our analysis of experimental results. Additionally, our method is extremely lightweight and allows the reproduction of the experiments on common GPUs. We have made our code accessible for future investigations.

Acknowledgments

We thank everyone who offers advice to us and everyone associated with organizing and sponsoring the WSDM Cup 2024. We also express our gratitude to the organizations that provided us with computing resources.

References

Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

DeepSeek-AI. 2024. [Deepseek llm: Scaling open-source language models with longtermism](#). *Preprint*, arXiv:2401.02954.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.

Dahyun Kim, Chanjun Park, Sanghoon Kim, Wonsung Lee, Wonho Song, Yunsu Kim, Hyeonwoo Kim, Yungi Kim, Hyeonju Lee, Jihoo Kim, Changbae Ahn, Seonghoon Yang, Sukyung Lee, Hyunbyung Park, Gyoungjin Gim, Mikyoung Cha, Hwalsuk Lee, and Sunghun Kim. 2023. [Solar 10.7b: Scaling large language models with simple yet effective depth up-scaling](#). *Preprint*, arXiv:2312.15166.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP '23*, page 611–626, New York, NY, USA. Association for Computing Machinery.

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.

Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. [Lost in the middle: How language models use long contexts](#). *Preprint*, arXiv:2307.03172.

Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *Preprint*, arXiv:1711.05101.

Meta. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.

Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. 2024. [Nomic embed: Training a reproducible long context text embedder](#). *Preprint*, arXiv:2402.01613.

OpenAI. 2022. [Introducing chatgpt](#).

- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeesh Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. [REPLUG: Retrieval-augmented black-box language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8371–8384, Mexico City, Mexico. Association for Computational Linguistics.
- Yi. 2023. A series of large language models trained from scratch by developers at 01-ai. <https://github.com/01-ai/Yi>.
- Munazza Zaib, Wei Emma Zhang, Quan Z. Sheng, Adnan Mahmood, and Yang Zhang. 2022. [Conversational question answering: a survey](#). *Knowl. Inf. Syst.*, 64(12):3151–3195.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, Weng Lam Tam, Zixuan Ma, Yufei Xue, Jidong Zhai, Wenguang Chen, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. [Glm-130b: An open bilingual pre-trained model](#). *Preprint*, arXiv:2210.02414.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. 2024. [Swift:a scalable lightweight infrastructure for fine-tuning](#). *Preprint*, arXiv:2408.05517.