

WisdomBot: Tuning Large Language Models with Artificial Intelligence Knowledge

Jingyuan Chen^a, Tao Wu^b, Wei Ji^c, Fei Wu^b

^a College of Education, Zhejiang University, Hangzhou 310058, China

^b College of Computer Science and Technology, Zhejiang University, Hangzhou 310058, China

^c School of Computing, National University of Singapore, Singapore 119077, Singapore

© Higher Education Press 2024

Abstract Large language models (LLMs) have emerged as powerful tools in natural language processing (NLP), showing a promising future of artificial generated intelligence (AGI). Despite their notable performance in the general domain, LLMs have remained suboptimal in the field of education, owing to the unique challenges presented by this domain, such as the need for more specialized knowledge, the requirement for personalized learning experiences, and the necessity for concise explanations of complex concepts. To address these issues, this paper presents a novel LLM for education named WisdomBot, which combines the power of LLMs with educational theories, enabling their seamless integration into educational contexts. To be specific, we harness self-instructed knowledge concepts and instructions under the guidance of Bloom's Taxonomy as training data. To further enhance the accuracy and professionalism of model's response on factual questions, we introduce two key enhancements during inference, i.e., local knowledge base retrieval augmentation and search engine retrieval augmentation during inference. We substantiate the effectiveness of our approach by applying it to several Chinese LLMs, thereby showcasing that the fine-tuned models can generate more reliable and professional responses.

Keywords artificial intelligence (AI), large language models (LLMs), intelligent education

1 Introduction

In recent years, the field of artificial intelligence (AI) has witnessed significant advancements, giving rise to a powerful class of large language models (LLMs). At the forefront of these language models is GPT-4 (OpenAI, 2023). With an impressive number of parameters in the trillions, GPT-4 is one of the largest language models ever created, capable of understanding complex language patterns and generating responses that are often indistinguishable from those of a human. At the same time, open-source communities have developed several alternatives, including LLaMA (Touvron et al., 2023). Training on trillions of tokens based on publicly available data, LLaMA achieves competitive performance with a relatively small number of parameters and affordable training costs. Thus, it emerges as a valuable open-source tool for a wide range of applications.

While large language models have demonstrated high capacity, they often encounter distinct challenges when applied to educational tasks. We outline the key limitations they face as follows and provide an illustration in Figure 1.

First, inadequacy of basic cognitive capacities. Although LLMs possess exceptional capabilities in general domain, their fundamental educational capabilities, such as retention and comprehension, remain limited due to constraints imposed by the training data. Primarily, owing to the vastness of knowledge, LLMs often exhibit constrained comprehension of specialized expertise that extends beyond their training data, yielding inaccurate responses. Additionally, academic knowledge is incessantly evolving, especially within practical subjects. The information contained in their training data may become outdated and obsolete, thus constraining the

Received May 15, 2024; accepted June 20, 2024

Jingyuan Chen (✉)

E-mail: jingyuanchen@zju.edu.cn

Fei Wu (✉)

E-mail: wufei@zju.edu.cn

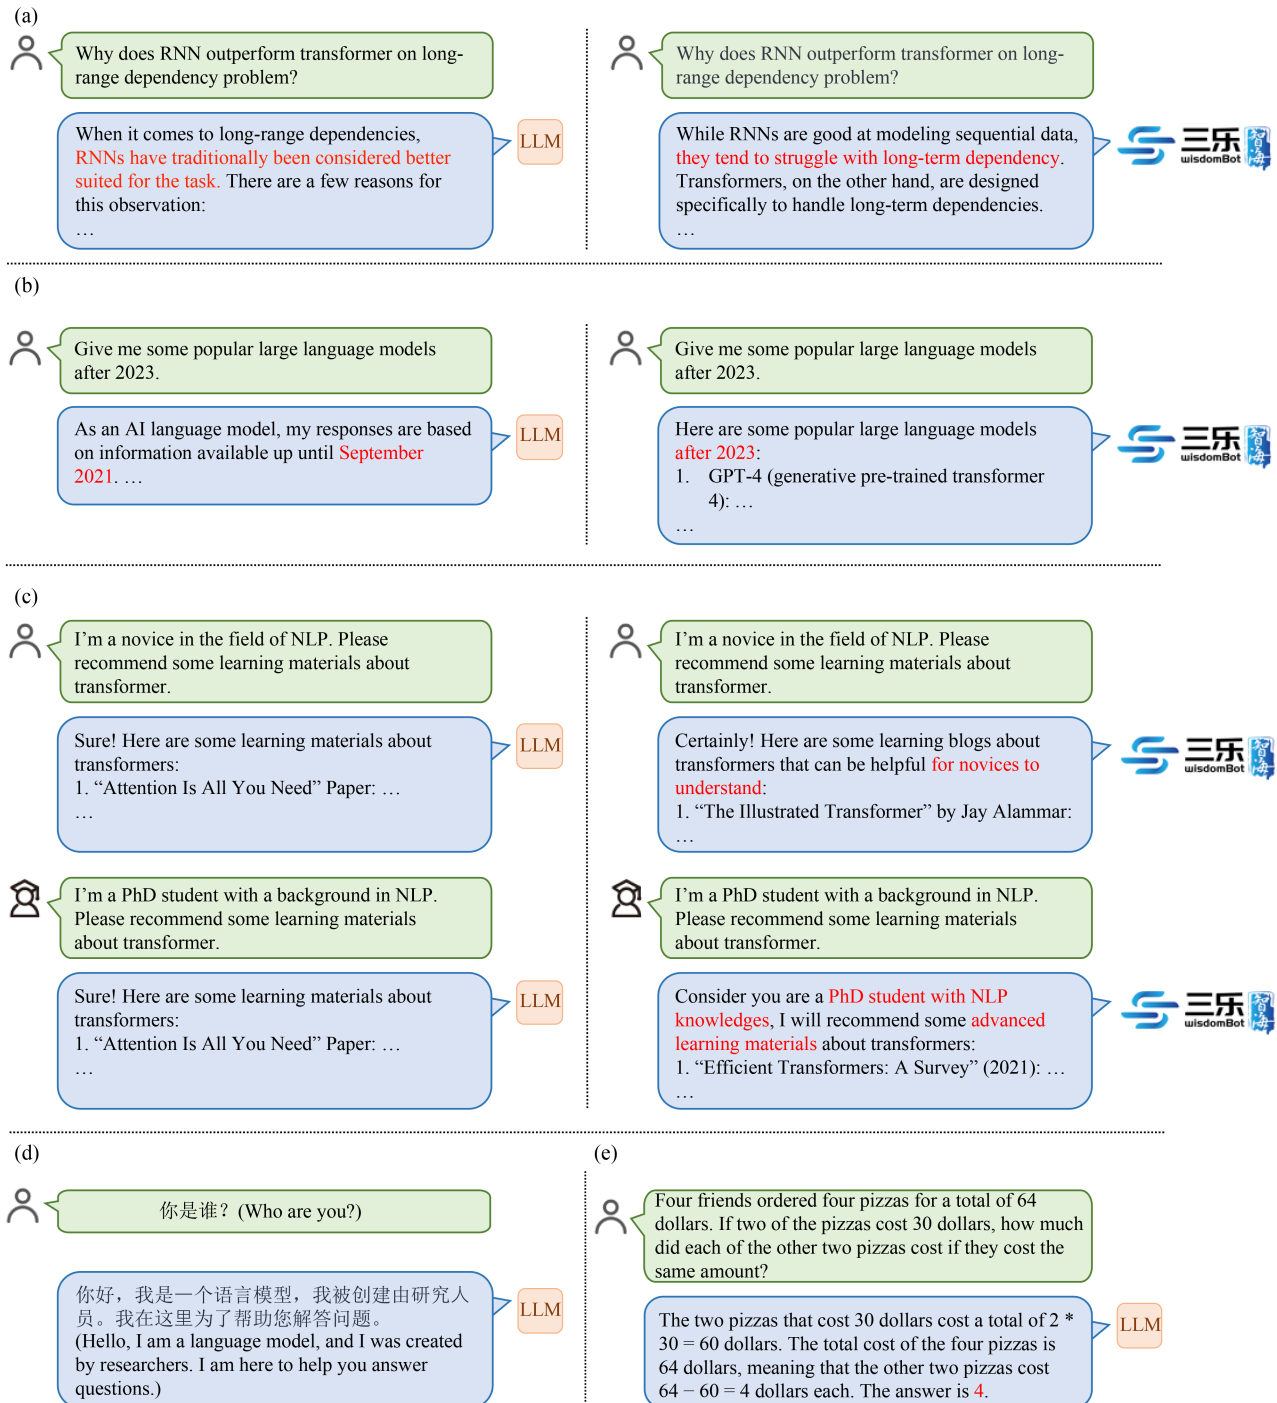


Figure 1 Limitations of general LLMs in education: (a) comprehension ability, (b) out-of-date knowledge, (c) personalized ability, (d) Chinese proficiency, (e) logical reasoning ability.

basic ability of LLMs to generate factually accurate responses when faced with inquiries that necessitate up-to-date awareness of post-training occurrences and knowledge (Cao et al., 2021; Liu et al., 2023; Wang et al., 2021; Yang et al., 2023).

Second, lack of advanced cognitive capacities. Proficiency in advanced cognitive capacities, such as analysis, evaluation and innovation, is vital for tackling challenging tasks. Existing studies have demonstrated

that LLMs lack these capabilities, which can lead to failures. For instance, Liu et al. (2023) have highlighted the persisting difficulty in logical reasoning for ChatGPT and GPT-4, especially when confronted with unfamiliar data and natural language inference datasets. Baidoo-Anu and Owusu Ansah (2023) discovered that generative models were limited to generating responses solely based on patterns present in their training data, thereby restricting the creativity and originality of their

outputs. Moreover, Baidoo-Anu and Owusu Ansah (2023) provided evidence that ChatGPT and other generative AI models may offer general information and assistance but lack the ability to personalize instruction to cater to the specific needs of individual students. The insufficiency of advanced skills in LLMs hinders their broader utilization in educational contexts.

Third, limited Chinese proficiency. While several large language models, such as LLaMA, have been made available to the public, their primary focus has been on English corpora, with limitation in applicability to other languages. Cui et al. (2023) showed that vocabularies of LLaMA or Alpaca (Taori et al., 2023) contained only a few hundred Chinese tokens, substantially hindering their efficiency in encoding and decoding Chinese text.

The key to rectify the shortcomings and adapt LLMs to the realm of education lies in the amalgamation of LLMs with educational theories, thereby equipping LLMs with varying levels of abilities. Among a large amount of educational theories, Bloom's Taxonomy, as delineated in Anderson et al. (2001), proffers a framework for categorizing the diverse objectives and proficiencies that educators aspire to instill into their students. The new taxonomy embraces a two-dimensional framework encompassing "knowledge" and "cognitive processes." "Knowledge" pertains to the relevant content involved in learning, while "cognitive processes" refer to the academic behaviors and manifestations of learning that need to be mastered. Bloom's Taxonomy has provided ample scope for guiding teaching practices and helping

learners progress to higher-order thinking (Ramirez, 2017), henceforth, it ought to be utilized to boost the abilities of LLMs in the realm of education.

In accordance with Bloom's Taxonomy, the authors propose a method that transfers general domain LLMs into educational field by simultaneously learning the "knowledge" and "cognitive processes" dimension in this paper, as illustrated in Figure 2. Firstly, to expand the breadth of the model's knowledge, we manually summarize coarse-grained knowledge concepts drawn from authentic textbooks, and meticulously utilize strong LLMs to generate fine-grained knowledge concepts that are aligned with the "knowledge" category. This approach facilitates a comprehensive coverage of detailed knowledge concepts spanning multiple levels and complexities, which help to enhance the model's basic cognitive capacities. Subsequently, we employ self-instruction, as demonstrated in (Wang et al., 2023), to construct over forty thousands Chinese instructions based on the "cognitive processes" dimension with educational tasks including professional knowledge question answering, test problem generation, and intelligent tutoring. These instructions not only embody a multitude of educational capacities at various levels, particularly those advanced ones, but also encompass all aforementioned fine-grained knowledge concepts. With them as instruction-tuning data, the proficiency of models can be significantly enhanced. Finally, in order to enhance model's awareness to the knowledge lies beyond training data, we utilize two retrieval augmentation strategies during inference, namely local knowledge base retrieval and search engine retrieval, to serve as extra knowledge sources.

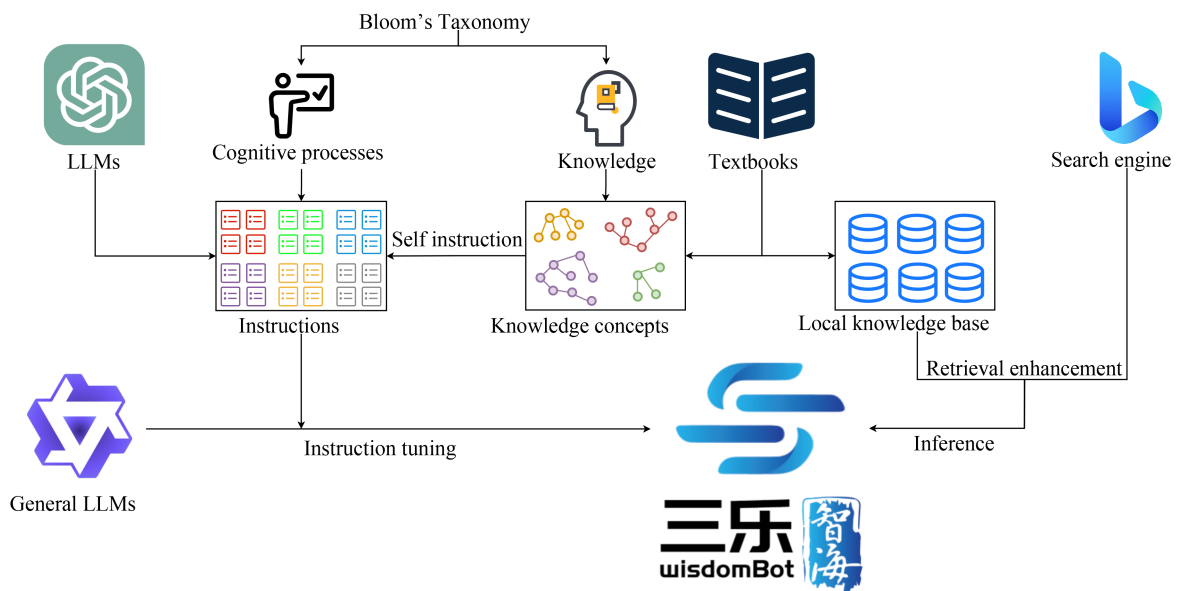


Figure 2 Training pipeline. We collect knowledge concepts and instructions under the guidance of textbooks, Bloom's Taxonomy and strong LLMs, serving as instruction-tuning data to transform general LLMs to educational LLMs. During inference, we construct a local knowledge base based on the textbook, incorporating search engine capabilities for retrieval enhancement.

We conduct our method on two open-source Chinese language models, Chinese-LLaMA and Alpaca (Cui et al., 2023) and Qwen-7B-Chat (Bai et al., 2023). Experiments have demonstrated the superiority of our fine-tuned models compared to the original models across various educational tasks, assessed from a diverse range of evaluation perspectives. It is worthy of note that our experiments are conducted specifically within the domain of Chinese AI instruction. In summary, our contributions are three-fold:

(1) We devise instructions based on textbooks and the guidance of Bloom’s Taxonomy to transfer general LLMs to educational domain;

(2) We utilize retrieval augmentation strategies during inference to expand the width of the model’s knowledge and enhance the quality of responses to factual inquiries;

(3) We conduct evaluation on various education tasks, demonstrating the superiority of our finetuned models compared to origin models.

2 Related Work

2.1 | Large Language Models

LLMs have revolutionized the field of AI and natural language processing, opening up new possibilities for human-computer interaction and advancing our understanding of language and its applications. With the ability to process and analyse vast amounts of textual data, large language models have showcased remarkable capabilities in tasks such as text generation, question answering, summarization, translation, sentiment analysis, and more.

Among all the models, ChatGPT (Ouyang et al., 2022) and GPT-4 (OpenAI, 2023) are two prominent iterations of large language models developed by OpenAI. Despite their excellent performance on general tasks, they are not open-source models and have a huge amount of parameters, which hinders personal deployment and research. LLaMA (Touvron et al., 2023) is an open-source substitute for GPT, with number of parameters ranging from 7 billion to 65 billion. Alpaca (Taori et al., 2023) conducts instruction tuning on LLaMA with 52K instruction data, achieving comparable performance with ChatGPT in English at an affordable cost.

Despite their great performance in English, LLaMA has shown weakness in Chinese due to the lack of Chinese corpus in the training data. To tackle this problem, Chinese-LLaMA-Alpaca (Cui et al., 2023) augment LLaMA with capabilities for understanding and generating Chinese text and its ability to follow instructions by extending LLaMA’s existing vocabulary with additional Chinese tokens and further fine-tuning

the model with Chinese instruction datasets. Unlike Chinese-LLaMA-Alpaca directly fine-tuning on LLaMA, ChatGLM (Zeng et al., 2022) is a new model based on the existing GLM structure. With training on large Chinese-English bilingual corpus and aligning with human intentions through supervised fine-tuning, ChatGLM has shown powerful Chinese language capabilities.

2.2 | Bloom’s Taxonomy

Bloom’s Taxonomy (Anderson et al., 2001) is a widely recognized framework that categorizes educational objectives and cognitive processes. It was first developed in the 1950s by Benjamin Bloom and his colleagues, and it has since become a fundamental tool in the field of education. Bloom’s Taxonomy provides a structured way to understand and organize different levels of thinking and learning (Bloom et al., 1956).

In its 2001 revision, the new taxonomy embraces a two-dimensional framework encompassing “knowledge” and “cognitive processes.” Knowledge pertains to the relevant content involved in learning, encompassing four categories ranging from concrete to abstract: factual knowledge, conceptual knowledge, procedural knowledge, and metacognitive knowledge. Cognitive processes refer to the academic behaviors and manifestations of learning that need to be mastered, including six categories: remember, understand, apply, analyse, evaluate, and create, which are arranged in ascending order of cognitive complexity. This theory has been used to explore the weaknesses of ChatGPT in the field of education (Elsayed, 2023).

3 Methods

The details of our finetuning method will be introduced in this section. Firstly, we collect fine-grained concepts corresponding to categories in “knowledge” dimension of Bloom’s Taxonomy. Then the instruction templates are designed to align with “cognitive processes” dimension of Bloom’s Taxonomy. By combining knowledge concepts and instruction templates, we obtain instruction-output pairs as our training data. Besides, two retrieval augmentation methods are developed during inference, enhancing the accuracy and professionalism of model’s response.

3.1 | Knowledge Concepts

Knowledge concepts, as the fundamental unit for transmitting instructional information in teaching activities, plays an important role in both teaching and learning. It can be regarded as the basic component of subject knowledge and serves as the cornerstone for

constructing a systematic knowledge system. We collect knowledge concepts with 2 steps, corresponding to different levels of granularity. Coarse-grained knowledge concepts have fewer quantities and are easy to obtain, so we manually extract them from the textbooks in the first step. In this way, we collect 117 coarse-grained knowledge concepts about AI, which are used for fine-grained knowledge concepts generation.

In the second step, we aim to collect intricate fine-grained knowledge concepts that are derived from the coarse ones. However, acquiring these detailed concepts manually is a time-consuming task and requires significant human effort. Therefore, we employ a self-instruction approach (Wang et al., 2023) to acquire them. To be specific, for each coarse concept and knowledge category within Bloom's Taxonomy, we employ ChatGPT to act as an AI learner. We prompt it to provide a series of questions that it may encounter during the learning process and to summarize the corresponding fine-grained knowledge concepts related to each category. Following careful manual extraction, cleaning, and filtering of the responses, we obtain a total of 981 fine-grained concepts and 1,196 questions. These concepts encompass various levels and diverse subjects.

3.2 | Knowledge-Based Instruction Tuning

Instruction tuning (Wei et al., 2021) is a simple method to improve the ability of language models to respond to NLP instructions, demonstrating promising abilities of language models to perform tasks described purely via instructions. Inspired by the automatic generation of instruction, we design templates and construct instructions using concepts above.

Specifically, we develop 39 distinct templates, in which fine-grained concepts or questions will be filled. These templates use three common educational tasks, i.e., subject knowledge QA, test problem generation, and intelligent tutoring as carriers, comprehensively encompassing all the learning abilities described in the cognitive processes of Bloom's Taxonomy. Subsequently, these templates are merged with concepts or questions to formulate original instructions. However, these instructions often suffer from substandard quality and a dearth of diversity, which might affect the performance of the model (Wang et al., 2023). Therefore, we employ ChatGPT to assess the coherence of each instruction. Only those instructions deemed of high quality are selected and revised to reduce similarity with others. Through this process, we acquire a set of instructions that exhibit exceptional quality. These instructions are then submitted to human experts utilizing GPT as an assistant to generate corresponding answers. Each instruction, along with its respective answer, is organized following the Stanford Alpaca

(Taori et al., 2023) style. It consists of indispensable Instruction and Output fields, while the Input field remains optional. Ultimately, we obtain 38,784 pairs of instruction and output, which serve as the foundation for supervised fine-tuning. We illustrate some data samples from our dataset in Figure 3.

3.3 | Retrieval Enhancement

As mentioned in (Cao et al., 2021; Liu et al., 2023; Wang et al., 2021; Yang et al., 2023), LLMs have limited performance in producing factually accurate answers. To tackle this problem, we utilize two retrieval enhancement methods during inference, namely local knowledge base retrieval and search engine retrieval.

Local knowledge base retrieval primarily addresses factual information contained in textbooks, which provides advantages when the model responds to queries involving obscure knowledge or highly professional language. To establish such a local library, we follow the standard long chain procedure. Initially, we import our textbooks to constitute unstructured textual contents, dividing them into multiple text chunks using a text splitter. Subsequently, a text embedding model is employed to transform these text chunks into a vector space while preserving textual coherence and similarity. Through this approach, each query can be converted into the same vector space, enabling the retrieval of the most similar k text segments from the textbooks to serve as reference materials.

To address inquiries that exceed the scope of textbooks, we employ a search engine to augment our proficiency. This is accomplished by dispatching each inquiry to the Azure Bing Search API, which provides us with a collection of search outcomes serving as pertinent resources. Whether retrieving information from a local library or through the search engine, these resources are amalgamated with the user's query, thereby forming an input for the model to generate a more accurate, comprehensive, and professional response. Please note that we set the retrieval enhancement as an optional feature, which means users need to decide whether to use the retrieval enhancement feature according to their needs.

4 Experiments

4.1 | Baselines

We conduct experiments on two Chinese baselines:

(1) Chinese-LLaMA-Alpaca (Cui et al., 2023) continues training on LLaMA with Chinese data. We conduct our training on Chinese-Alpaca-7B and Chinese-Alpaca-13B.

Instruction:

如何教授没有贝叶斯网络原理及应用基础知识的学生？应该从哪些方面入手？
(How do you teach students who lack basic knowledge of Bayesian network principles and applications? Where should you start?)

Answer:

1. 引入基本概念：首先，教授诸如条件概率、独立性、联合概率等基本概念是很重要的。正确理解这些概念是奠定贝叶斯网络学习基础的前提。
2. 介绍贝叶斯定理：贝叶斯网络的核心原理是贝叶斯定理，因此教授学生贝叶斯定理的概念和实践，打好基础。
3. 理解概率图模型：强调贝叶斯网络是一种概率图模型，需要深入了解概率图模型的定义、组成、优势等方面的知识。
4. 探讨模型学习：了解贝叶斯网络的学习方法也是重要的。通过介绍极大似然估计和基于贝叶斯方法的参数学习，帮助学生熟悉贝叶斯网络的模型学习。
5. 应用案例：通过实际的案例应用，让学生了解贝叶斯网络在实际问题中的应用和效果，进一步培养学生的学习和实践能力。
6. 提供示例代码：最后，提供相关的贝叶斯网络的学习实践代码给学生，引导学生自主进行实践，进一步巩固和提高理论学习的实效性。

(1. Introducing basic concepts: Firstly, it is important to teach basic concepts such as conditional probability, independence, and joint probability. Correctly understanding these concepts is a prerequisite for learning Bayesian networks.

2. Introducing Bayes' theorem: The core principle of Bayesian networks is Bayes' theorem. Therefore, teaching students the concept and practice of Bayes' theorem lays a solid foundation.

3. Understanding probabilistic graphical models: Emphasize that Bayesian networks are a type of probabilistic graphical model. It is necessary to deeply understand the definition, components, and advantages of probabilistic graphical models.

4. Exploring model learning: Understanding the learning methods of Bayesian networks is also important. By introducing maximum likelihood estimation and Bayesian-based parameter learning, help students become familiar with model learning in Bayesian networks.

5. Applying case studies: Through practical case applications, let students understand the application and effectiveness of Bayesian networks in real-world problems, further cultivating their interest and practical skills.

6. Providing example code: Finally, provide relevant example code for learning Bayesian networks. Guide students to practice independently, further consolidating and enhancing the effectiveness of theoretical learning.)

Instruction:

请用简单易懂的语言解释逻辑斯蒂回归模型，同时讲解其参数的意义。
(Please explain the logistic regression model in simple and easy-to-understand language, and also explain the meaning of its parameters.)

Answer:

逻辑斯蒂回归是一种用于分类问题的机器学习模型。其原理基于线性模型，但是通过 sigmoid 函数将结果映射到 0 到 1 之间，即可用于二分类问题。模型的核心是从训练数据中学习到一个最佳权重系数，使得模型的预测结果与实际结果的误差最小化。逻辑斯蒂回归的参数包括：

1. 正则化系数：控制模型的复杂度以避免过度拟合。正则化系数越大，惩罚系数就越大，模型的复杂度就会降低。
2. 截距项：用于调整模型的基准预测值，即将所有特征都设置为 0 时的预测结果。截距项也可以被视为特征向量中的一个额外的变量。
3. 特征系数：对每个特征变量分别设置的权重系数，用于衡量该变量对分类结果的贡献。特征系数越大，意味着该特征对分类结果的影响越大。

通过调整这些参数，可以最优化模型的性能，以获得更准确、可靠的分类结果。

(Logistic regression is a machine learning model used for classification problems. Its principle is based on a linear model, but it uses the sigmoid function to map the results between 0 and 1, making it suitable for binary classification problems. The core of the model is to learn an optimal weight coefficient from the training data to minimize the error between the model's predicted results and the actual results.

The parameters of logistic regression include:

1. Regularization coefficient: This controls the complexity of the model to avoid overfitting. The larger the regularization coefficient, the greater the penalty, and the lower the complexity of the model.
2. Intercept term: This adjusts the baseline prediction value of the model, which is the prediction result when all features are set to 0. The intercept term can also be seen as an additional variable in the feature vector.
3. Feature coefficients: These are weight coefficients set for each feature variable, used to measure the contribution of that variable to the classification result. The larger the feature coefficient, the greater the impact of that feature on the classification result. By adjusting these parameters, the performance of the model can be optimized to achieve more accurate and reliable classification results.)

Figure 3 Demonstrations of the dataset.

(2) Qwen (Bai et al., 2023) is an optimized dialogue model specifically designed for the Chinese chatting scenario. We conduct our training on Qwen-7B-Chat.

4.2 | Experiment Detail

We adopt AdamW optimizer with an initial learning rate of $2e-5$. The models are trained on 8 A100 GPUs and the batch size of each GPU is set to 16. We use Low-Rank Adaptation (Hu et al., 2021) training strategy to reduce training parameters when training on Chinese-LLaMA-Alpaca. The lora rank of Chinese-Alpaca-7B is set to 8, while the lora rank of Chinese-Alpaca-13B is set to 32. For Qwen-7B-Chat, we utilize full-parameter finetuning. For inference, we set the temperature to 1, top-p to 0.7, beam size to 1, and maximum generation length to 1,024.

4.3 | Testset Detail

4.3.1 Self-Constructed Dataset

Our testset consist of 2 parts: educational functions and cognitive capacities. There are 70 instances in the educational functions part, encompassing professional question answering, test problem generation and intelligent tutoring. Each function is evaluated from diverse perspectives. The remaining 60 instances are dedicated to testing cognitive capacities, corresponding to the 6 cognitive processes described in Bloom's

Taxonomy. Figure 4 illustrates the distribution of our test data.

4.3.2 Public Dataset: C-Eval

C-Eval (Huang et al., 2024) is a comprehensive Chinese evaluation suite designed to assess the advanced knowledge and reasoning abilities of foundation models in a Chinese context. It includes 13,948 multiple-choice questions across four difficulty levels (middle school, high school, college, and vocational) and spans 52 disciplines. C-Eval also features a subset called C-Eval Hard, focusing on particularly challenging subjects. Evaluation results show that even the most advanced models like GPT-4 have significant room for improvement, highlighting the suite's ability to benchmark the capabilities and limitations of current language models. We conduct evaluation on the validation set of C-Eval.

4.4 | Results on Self-Constructed Dataset

We conduct model inference on our self-constructed dataset, aiming to compare the results generated by the original LLMs (i.e., Chinese-Alpaca-7B and Chinese-Alpaca-13B) and our WisdomBot. We utilize both human and GPT-4 evaluation when comparing performance, ensuring the accuracy and diversity of evaluation. For human evaluation, we recruit ten experts in the field of AI to compare the responses of

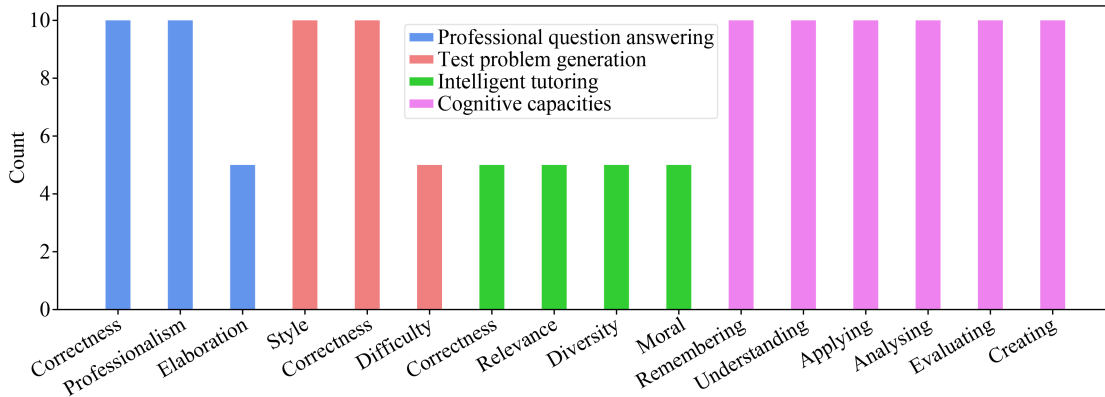


Figure 4 Distribution of test data.

two models. For GPT-4 evaluation, we conduct two evaluation processes for each question. The order of responses from the two models differs in the prompt of each evaluation processes because GPT-4 favors response that come first in the sequence.

The results are shown in Figures 5–8. From these pictures we can observe that for each evaluation part, WisdomBot has a winning rate of at least 63%. WisdomBot even reaches a 100% winning rate on professional question answering part comparing with

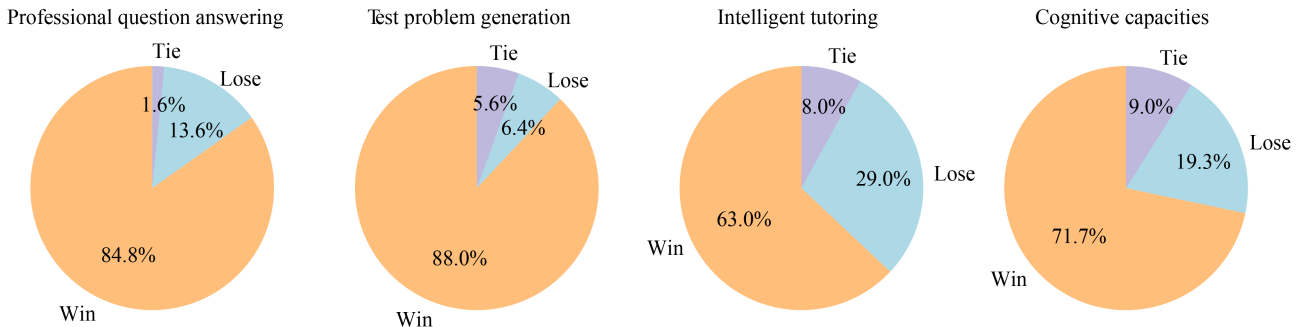


Figure 5 Human evaluation of whether WisdomBot outperforms Chinese-Alpaca-7B.

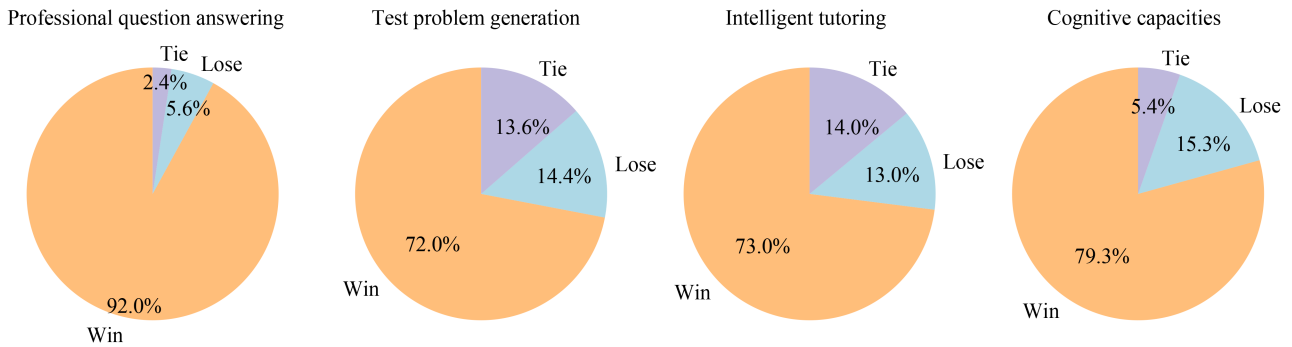


Figure 6 Human evaluation of whether WisdomBot outperforms Chinese-Alpaca-13B.

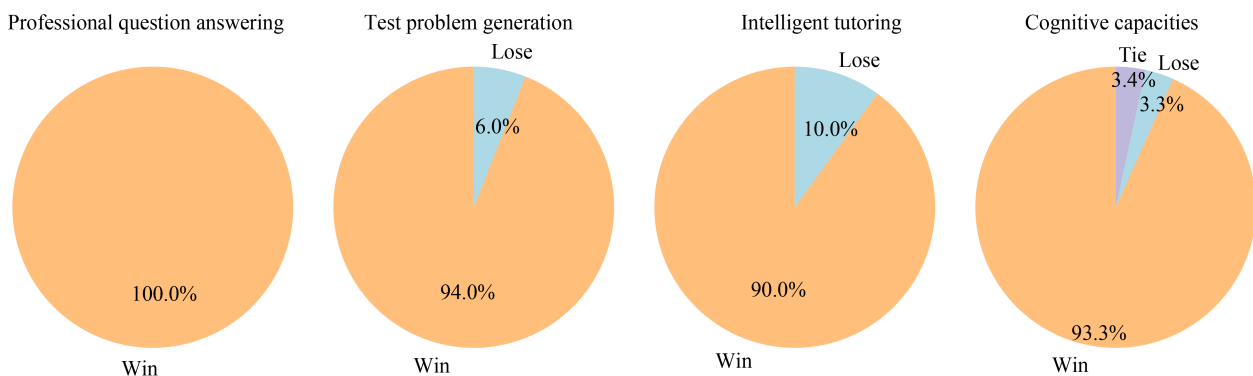


Figure 7 GPT-4 evaluation of whether WisdomBot outperforms Chinese-Alpaca-7B.

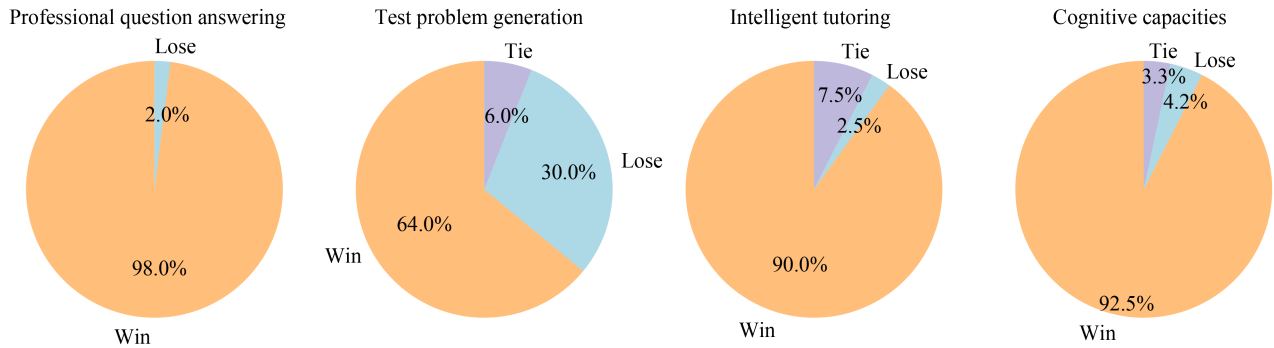


Figure 8 GPT-4 evaluation of whether WisdomBot outperforms Chinese-Alpaca-13B.

Chinese-Alpaca-7B. These results demonstrate that WisdomBot can provide more accurate responses.

4.5 | Results on C-Eval

We conduct zero-shot evaluation on the validation set

of C-Eval benchmark of all the baselines and WisdomBot. We list the overall performance of the three models in Table 1, and the performance on each subset of C-Eval benchmark in Tables 2–5. From these tables we can observe that WisdomBot outperforms the baselines in most subjects, especially in subjects related

Table 1 Results on the validation set of C-Eval benchmark

Model	STEM	Social science	Humanities	Other	Hard	Average
Chinese-Alpaca-7B	35.45	51.53	47.67	41.87	28.28	42.49
Qwen-7B-Chat	51.61	72.64	66.94	53.83	35.14	59.37
WisdomBot	59.17	72.01	65.38	54.96	49.26	62.06

Table 2 Results on the STEM subset within the validation set of C-Eval benchmark

Model	Computer network	Operating system	Computer architecture	College programming
Chinese-Alpaca-7B	36.84	52.63	38.1	43.24
Qwen-7B-Chat	42.11	42.11	52.38	64.86
WisdomBot	52.63	57.89	57.14	62.16
Model	College physics	College chemistry	Advanced mathematics	Probability and statistics
Chinese-Alpaca-7B	31.58	16.67	21.05	33.33
Qwen-7B-Chat	31.58	54.17	10.53	22.22
WisdomBot	57.89	58.33	26.32	33.33
Model	Discrete mathematics	Electrical engineer	Metrology engineer	High school mathematics
Chinese-Alpaca-7B	43.75	37.84	50	16.67
Qwen-7B-Chat	18.75	24.32	75	33.33
WisdomBot	37.5	35.14	70.83	33.33
Model	High school physics	High school chemistry	High school biology	Middle school mathematics
Chinese-Alpaca-7B	31.58	31.58	42.11	21.05
Qwen-7B-Chat	57.89	52.63	73.68	63.16
WisdomBot	78.95	68.42	68.42	63.16
Model	Middle school biology	Middle school physics	Middle school chemistry	Veterinary medicine
Chinese-Alpaca-7B	47.62	47.37	40	26.09
Qwen-7B-Chat	85.71	84.21	100	43.48
WisdomBot	90.48	84.21	95	52.17

Table 3 Results on the social science subset within the validation set of C-Eval benchmark

Model	College economics	Business administration	Marxism	Mao Zedong Thought
Chinese-Alpaca-7B	32.73	45.45	52.63	54.17
Qwen-7B-Chat	45.45	54.55	73.68	75.00
WisdomBot	38.18	54.55	84.21	62.50
Model	Education science	Teacher qualification	High school politics	High school geography
Chinese-Alpaca-7B	37.93	59.09	57.89	42.11
Qwen-7B-Chat	65.52	84.09	94.74	63.16
WisdomBot	72.41	81.82	94.74	57.89
Model	Middle school politics	Middle school geography		
Chinese-Alpaca-7B	66.67	66.67		
Qwen-7B-Chat	95.24	75.00		
WisdomBot	90.48	83.33		

Table 4 Results on the humanities subset within the validation set of C-Eval benchmark

Model	Modern Chinese history	Ideological and moral cultivation	Logic	Law
Chinese-Alpaca-7B	52.17	52.63	54.55	20.83
Qwen-7B-Chat	78.26	84.21	36.36	41.67
WisdomBot	69.57	94.74	59.09	37.50
Model	Chinese language and literature	Art studies	Professional tour guide	Legal professional
Chinese-Alpaca-7B	34.78	48.48	51.72	39.13
Qwen-7B-Chat	56.52	66.67	79.31	43.48
WisdomBot	47.83	69.70	68.97	43.48
Model	High school Chinese	High school history	Middle school history	
Chinese-Alpaca-7B	47.37	50.00	72.73	
Qwen-7B-Chat	78.95	80.00	90.91	
WisdomBot	57.89	75.00	95.45	

Table 5 Results on the other subset within the validation set of C-Eval benchmark

Model	Civil servant	Sports science	Plant protection	Basic medicine
Chinese-Alpaca-7B	40.43	57.89	36.36	47.37
Qwen-7B-Chat	48.94	47.37	68.18	63.16
WisdomBot	53.19	52.63	59.09	68.42
Model	Clinical medicine	Urban and rural planner	Accountant	Fire engineer
Chinese-Alpaca-7B	36.36	52.17	36.73	38.71
Qwen-7B-Chat	45.45	63.04	51.02	48.39
WisdomBot	50.00	60.87	53.06	45.16
Model	Environmental impact assessment engineer	Tax accountant	Physician	
Chinese-Alpaca-7B	45.16	34.69	34.69	
Qwen-7B-Chat	48.39	53.06	55.10	
WisdomBot	58.06	44.90	59.18	

to information and computer science. We attribute the performance enhancement to our training data, which is highly relevant to these subjects. For other subjects

such as social science and humanities, WisdomBot does not exhibit a significant performance decrease. The performance on “other” subset is even increased

compared with the baselines. The results demonstrate the superiority of our WisdomBot model.

4.6 | Advanced Cognitive Ability Comparisons

We compare WisdomBot with baseline models to evaluate their advanced cognitive abilities, encompassing creativity, personalized ability, and logical reasoning ability. We curate 50 test samples for each ability test. For the creativity and personalization tests, we ask GPT-4 to score each model's response on a scale from 1 to 5, with higher scores indicating stronger abilities. For the logical reasoning test, we directly assess the outputs and calculate each model's accuracy. The results are reported in Table 6, which demonstrates the superiority of WisdomBot in terms of advanced cognitive abilities.

4.7 | Experiments on Retrieval Enhancement

We evaluate the effectiveness of two retrieval enhancements: local knowledge base retrieval and search engine retrieval. For the knowledge base retrieval, we curate 50 professional AI questions and assess the professional level of the answers using GPT-4. For the search engine retrieval, we curate 30 factual questions and evaluate the correctness of the answers. The results, presented in Table 7, show that the local knowledge base retrieval enables the model to generate more professional answers, while the search engine retrieval improves the model's accuracy on factual questions.

Table 6 Comparisons on three advanced cognitive abilities

Model	Creativity	Personalized ability	Logical reasoning (%)
Chinese-Alpaca-7B	2.78	3.56	8
Qwen-7B-Chat	2.86	3.34	46
WisdomBot	3.28	3.80	52

Table 7 Comparisons on retrieval enhancements

Model	Local knowledge base (%)	Search engine (%)
w/o retrieval	30	35
w retrieval	70	93


4.8 | Case Study

As illustrated in Figure 9, WisdomBot demonstrates enhanced creativity (a), personalized ability (b), and logical reasoning ability (c) compared to baseline models. Additionally, local knowledge base retrieval (d) and search engine retrieval (e) respectively improve the professional level and accuracy of WisdomBot's response.


5 Conclusion

The general large language models lack basic cognitive abilities and advanced cognitive abilities. We propose a novel tuning approach, using high-quality textbook-level corpora as the basis, focusing on knowledge concepts to construct training data, migrating open-source large language models to the education field, and forming the educational large language model WisdomBot. Experiments show that WisdomBot has achieved excellent performance in different educational scenarios and various subjects.


(a)



How can AI be utilized to create personalized reading lists that adapt to students' evolving interests and comprehension levels?




AI can be used to create personalized reading lists by analyzing the student's reading habits, interests, and comprehension level. This can be done through machine learning algorithms which analyze the student's previous reading behavior and generate a list of books tailored to their interests and abilities. Additionally, AI can also be used to track the student's progress as they read each book on the list and adjust the list accordingly.



AI can be utilized to create personalized reading lists that adapt to students' evolving interests and comprehension levels by using natural language processing (NLP) and machine learning algorithms. Here are the steps that can be taken:

1. Collect data: Collect data on students' reading habits, interests, and comprehension levels.
2. Analyze data: Use NLP algorithms to analyze the data and identify patterns in students' reading preferences and comprehension levels.
3. Develop algorithms: Develop machine learning algorithms that can use the data to create personalized reading lists for each student...
4. Implement algorithms: Implement the algorithms in a reading management system that can be used by teachers and students...
5. Monitor progress: Monitor students' progress as they read the books on their personalized reading lists...
6. Provide feedback: Provide feedback to students on their reading progress and suggest...



AI can be utilized to create personalized reading lists that adapt to students' evolving interests and comprehension levels by using natural language processing and machine learning algorithms. Here are the steps that can be taken:

1. Collect data: Collect data on students' reading levels, interests, and previous reading lists...
2. Analyze data: Use natural language processing and machine learning algorithms to analyze the data and identify patterns and trends in students' reading preferences and comprehension levels.
3. Create personalized reading lists: Based on the analysis, create personalized reading lists that are tailored to each student's interests and comprehension levels. The lists can include books, articles, and other reading materials that are appropriate for each student's current level.
4. Adapt to evolving interests: As students' interests change, the AI system can adapt the reading lists accordingly...
5. Monitor progress: The AI system can monitor students' progress and adjust the reading lists accordingly. For example, if a student is struggling with a particular book, the AI system can suggest alternative reading materials that are more appropriate for the student's current level.



Figure 9 Case examples generated by WisdomBot and baselines: (a) creativity, (b) personalized ability, (c) logical reasoning. WisdomBot with two retrieval enhancement methods: (d) local knowledge library retrieval, (e) search engine retrieval.

Acknowledgments This work was supported by the National Science and Technology Major Project, China (Grant No. 2022ZD0117104), the National Natural Science Foundation of China (Grant Nos. 62037001 and 62307032), and the Starry Night Science Fund at Shanghai Institute for Advanced Study (SN-ZJU-SIAS-0010).

Conflict of Interest Fei Wu is a member of the Editorial Board of *Frontiers of Digital Education*, who was excluded from the peer-review process and all editorial decisions related to the acceptance and publication of this article. Peer-review was handled independently by the other editors to minimise bias.

References

Krathwohl, D. R. (2002). A Revision of Bloom’s Taxonomy: An Overview. *Theory Into Practice*, 41(4), 212–218.

Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., Hui, B., Ji, L., Li, M., Lin, J., Lin, R., Liu, D., Liu, G., Lu, C., Lu, K., Ma, J., Men, R., Ren, X., Ren, X., Tan, C., Tan, S., Tu, J., Wang, P., Wang, S., Wang, W., Wu, S., Xu, B., Xu, J., Yang, A., Yang, H., Yang, J., Yang, S., Yao, Y., Yu, B., Yuan, H., Yuan, Z., Zhang, J., Zhang, X., Zhang, Y., Zhang, Z., Zhou, C., Zhou, J., Zhou, X., & Zhu, T. (2023). Qwen technical report. *arXiv Preprint*, arXiv:2309.16609.

Baidoo-Anu, D., & Owusu Ansah, L. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7(1), 52–62.

Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). *Handbook I: Cognitive domain*. New York: David McKay.

Cao, B., Lin, H., Han, X., Sun, L., Yan, L., Liao, M., Xue, T., & Xu, J. (2021). Knowledgeable or educated guess? Revisiting language models as knowledge bases. In: *Proceedings of the*

- 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). Online: Association for Computational Linguistics, 1860–1874.
- Cui, Y., Yang, Z., & Yao, X. (2023). Efficient and effective text encoding for Chinese LLaMA and Alpaca. *arXiv Preprint*, arXiv:2304.08177.
- Elsayed, S. (2023). Towards mitigating ChatGPT’s negative impact on education: Optimizing question design through Bloom’s Taxonomy. In: *Proceedings of 2023 IEEE Region 10 Symposium (TENSYMP)*. Canberra, 1–6.
- Hu, E., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2021). LoRA: Low-rank adaptation of large language models. *arXiv Preprint*, arXiv:2106.09685v2.
- Huang, Y., Bai, Y., Zhu, Z., Zhang, J., Zhang, J., Su, T., Liu, J., Lv, C., Zhang, Y., Lei, J., Fu, Y., Sun, M., & He, J. (2024). C-Eval: A multi-level multi-discipline Chinese evaluation suite for foundation models. *Advances in Neural Information Processing Systems*, 36.
- Liu, H., Ning, R., Teng, Z., Liu, J., Zhou, Q., & Zhang, Y. (2023). Evaluating the logical reasoning ability of ChatGPT and GPT-4. *arXiv Preprint*, arXiv:2304.03439.
- OpenAI. (2023). GPT-4 technical report. *arXiv Preprint*, arXiv:2303.08774.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P. F., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730–27744.
- Ramirez, T. V. (2017). On pedagogy of personality assessment: Application of Bloom’s Taxonomy of educational objectives. *Journal of Personality Assessment*, 99(2), 146–152.
- Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., & Hashimoto, T. B. (2023). *Stanford Alpaca: An instruction-following LLaMA model*. Retrieved from GitHub.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M. A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., & Lample, G. (2023). LLaMA: Open and efficient foundation language models. *arXiv Preprint*, arXiv:2302.13971.
- Wang, C., Liu, P., & Zhang, Y. (2021). Can generative pre-trained language models serve as knowledge bases for closed-book QA? In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, 3241–3251.
- Wang, Y., Kordi, Y., Mishra, S., Liu, A., Smith, N. A., Khashabi, D., & Hajishirzi, H. (2023). Self-instruct: Aligning language models with self-generated instructions. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto: Association for Computational Linguistics, 13484–13508.
- Wei, J., Bosma, M., Zhao, V., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., & Le, Q. V. (2021). Finetuned language models are zero-shot learners. *arXiv Preprint*, arXiv:2109.01652v2.
- Yang, L., Chen, H., Li, Z., Ding, X., & Wu, X. (2023). ChatGPT is not enough: Enhancing large language models with knowledge graphs for fact-aware language modeling. *arXiv Preprint*, arXiv:2306.11489.
- Zeng, A., Liu, X., Du, Z., Wang, Z., Lai, H., Ding, M., Yang, Z., Xu, Y., Zheng, W., Xia, X., Tam, W. L., Ma, Z., Xue, Y., Zhai, J., Chen, W., Liu, Z., Zhang, P., Dong, Y., & Tang, J. (2022). GLM-130B: An open bilingual pre-trained model. *arXiv Preprint*, arXiv:2210.02414v2.