

CLEVA: Chinese Language Models EVALuation Platform

Yanyang Li^{1*}, Jianqiao Zhao^{1*}, Duo Zheng¹, Zi-Yuan Hu¹, Zhi Chen³, Xiaohui Su³,
Yongfeng Huang¹, Shijia Huang¹, Dahua Lin^{2,3}, Michael R. Lyu¹, Liwei Wang^{1†}

¹Department of Computer Science and Engineering, The Chinese University of Hong Kong

²Department of Information Engineering, The Chinese University of Hong Kong

³Shanghai AI Laboratory

Abstract

With the continuous emergence of Chinese Large Language Models (LLMs), how to evaluate a model’s capabilities has become an increasingly significant issue. The absence of a comprehensive Chinese benchmark that thoroughly assesses a model’s performance, the unstandardized and incomparable prompting procedure, and the prevalent risk of contamination pose major challenges in the current evaluation of Chinese LLMs. We present CLEVA, a user-friendly platform crafted to holistically evaluate Chinese LLMs. Our platform employs a standardized workflow to assess LLMs’ performance across various dimensions, regularly updating a competitive leaderboard. To alleviate contamination, CLEVA curates a significant proportion of new data and develops a sampling strategy that guarantees a unique subset for each leaderboard round. Empowered by an easy-to-use interface that requires just a few mouse clicks and a model API, users can conduct a thorough evaluation with minimal coding. Large-scale experiments featuring 23 Chinese LLMs have validated CLEVA’s efficacy. Our GitHub repo is <https://github.com/LaVi-Lab/CLEVA>.

1 Introduction

Large language models (LLMs) have fundamentally revolutionized natural language processing. Transformer models with more than 100B parameters have exhibited remarkable generalization ability across diverse tasks without the need for fine-tuning. The success of GPT-4 (OpenAI, 2023) and ChatGPT sparked a trend of training Chinese LLMs, with new models launching almost every week (Zeng et al., 2023; Team, 2023; Chenghao Fan and Tian, 2023; Ji et al., 2023; Cui et al., 2023). These rapid developments aggravate the need for Chinese LLM evaluation.

*Equal contributions.

†Project leader and corresponding author

Assessing the capacity of LLMs is non-trivial.

Traditional practices that evaluate models on a single task at a time are gradually becoming obsolete, since a single task can hardly characterize a full view of an LLM’s capacity. Instead, to effectively grasp a holistic view of an LLM’s capacity, we need to decompose its capacity into various abilities, evaluate these abilities with numerous corresponding tasks, and measure the competence of each task with multiple metrics. In this sense, HELM (Liang et al., 2022), leads the way in English LLM evaluation, as it conducts an in-depth evaluation of English LLMs on various NLP tasks using seven metrics. In Chinese, previous attempts have shown limitations, either in task selection or the metrics adopted. C-Eval (Huang et al., 2023), M3KE (Liu et al., 2023), CMMLU (Li et al., 2023), GAOKAO-Bench (Zhang et al., 2023), and MMCU (Zeng, 2023) narrow down to knowledge and reasoning abilities, whose datasets are mostly constructed using Chinese exams. By the time of our submission, OpenCompass (Contributors, 2023b), with around 74K Chinese queries out of 300K total, leans on accuracy as its sole metric, overlooking other important aspects in LLM evaluation. FlagEval (Contributors, 2023a) inherits four out of seven metrics from HELM and 22 existing Chinese datasets, having limited coverage on some significant tasks. A comprehensive Chinese benchmark incorporating diverse metrics to holistically evaluate Chinese LLMs is urgently demanded.

Prompt-based evaluation in Chinese is largely unstandardized.

Previous evaluations, such as HELM (Liang et al., 2022), do not explicitly optimize prompts, though LLMs’ significant sensitivity to the format of prompt has been observed (Webson and Pavlick, 2022; Abdou et al., 2022; Sanh et al., 2022). Moreover, unlike many English benchmarks that have well-developed prompts (§ 3), many Chinese benchmarks are in their early stage and do not enjoy such privileges. Chinese LLMs are eval-

uated using different prompts, making the results incomparable and hence untrustworthy.

Consuming up to trillions of tokens during pre-training, LLMs are prone to train-test contamination (Brown et al., 2020), which significantly threatens the validity of an evaluation. Previous work (OpenAI, 2023; Liang et al., 2022) approaches this issue more from a consequentialist perspective: They examine the contamination risk, by methods like long n-gram overlap, only after the evaluation has been done. These post-evaluation analyses, though responsibly examining if train-test contamination happens, cannot alleviate the risk of contamination in the first place. A proactive method to mitigate the contamination risk is of great importance.

We present CLEVA, Chinese Language models EVAluation platform that tackles the aforementioned problems with the following features:

- **A comprehensive Chinese benchmark.** Inspired by HELM (Liang et al., 2022), CLEVA organizes the evaluation tasks into two parts: *ability evaluation*, which gauges specific LLM skills and *application assessment*, which tests how well LLMs apply their skills to real-world applications (§ 4.1). Most of the well-accepted Chinese datasets relevant to our ability evaluation or application assessment are organized, standardized, and then adopted by our platform. More importantly, we design new Chinese-specific tasks, e.g., Pinyin transliteration and intent understanding, and collect a substantial amount of new data, accounting for 33.98% of our total data. As for the metrics (§ 4.1), we incorporate metrics for diversity and privacy into our system in addition to the seven in HELM. With 370K (over 9 million queries after augmentation) test instances from 84 datasets and 9 metrics, CLEVA, so far, stands as the most extensive Chinese evaluation dataset and possesses the most dimensions, facilitating a holistic evaluation of Chinese LLMs.
- **Standardized prompt-based evaluation methodology.** CLEVA takes full control of key aspects of LLM evaluation, with data and prompts being the most important among them. All data are jointly prepared with unified preprocessing steps, ensuring a level playing field for all LLMs. Meanwhile, CLEVA provides a set of prompts, instead of just

one prompt as in prior work, for each task for prompting-based inference (Brown et al., 2020). This prompt design ensures comparable evaluation results by standardizing the prompts used for testing, while also encouraging further analysis of LLMs’ sensitivity to different prompts (Zhu et al., 2023).

- **An up-to-date and trustworthy leaderboard.** CLEVA advocates a proactive method for securing trustworthy evaluation results. By collecting extensive new data, CLEVA suppresses the leakage of testing data prior to the evaluation. Moreover, we frequently organize new evaluation rounds, sampling a unique test set from 9 million augmented instances. This strategy further mitigates the risk of train-test contamination, improving the trustworthiness and timeliness of the leaderboard.

CLEVA is thoroughly validated by benchmarking 23 Chinese LLMs on our large-scale test sets (§ 6). The corresponding leaderboard and all other user-friendly features will be continuously maintained and improved to accommodate new models and evaluation methods.

2 Related Work

LLM evaluation is a rapidly developing field in recent years to delineate the boundary of LLM’s capability. In English, various systematic evaluation benchmarks have been proposed. For example, BIG-Bench (bench authors, 2023) is the largest collection that covers more than 200 tasks. HELM (Liang et al., 2022) organizes tasks into core scenarios, which focus on use cases, and targeted evaluation, which aims to better understand models. HELM also presents a multi-metric measurement that enables analysis of tradeoffs for each scenario. Recently, AGIEval (Zhong et al., 2023) is proposed to evaluate LLMs using challenging human exams. PromptBench (Zhu et al., 2023), on the other hand, measures the robustness of LLMs to prompts via adversarial attacks. MT-Bench (Zheng et al., 2023) collects multi-turn questions and presents the Chatbot Arena platform that treats GPT-4 (OpenAI, 2023) as the judge.

While CLEVA shares the same fundamental motivation with HELM (Liang et al., 2022), to holistically evaluate language learning models in their original languages, CLEVA is far from a mere Chinese replica of HELM. Building on the foundation of HELM’s taxonomy, CLEVA introduces a range

of tasks, with particular emphasis on those unique to Chinese, to better assess the capabilities of Chinese LLMs. It offers a new perspective on prompts, providing abundant prompt templates to standardize evaluation and encourage in-depth exploration of models’ sensitivity. In terms of metrics, CLEVA expands into new areas of diversity and privacy for a more comprehensive evaluation. Finally, CLEVA proactively mitigates train-test contamination by collecting a significant amount of new data, creating unique test sets by sampling, and regularly updating the leaderboard. All of these evaluation designs are neatly packaged in a user-friendly platform to facilitate community usage.

There is also a lot of progress in evaluating Chinese LLMs (Huang et al., 2023; Liu et al., 2023; Li et al., 2023; Zhang et al., 2023; Zeng, 2023). OpenCompass (Contributors, 2023b) and FlagEval (Contributors, 2023a) are two important attempts to evaluate Chinese LLMs. OpenCompass pools 53 public datasets and uses standard accuracy-like metrics as the only measurement for each dataset. FlagEval, with a smaller number of datasets and metrics, still needs further expansion to achieve sufficient coverage. Compared to previous efforts, CLEVA offers Chinese data from 84 datasets, including 33.98% original queries, while employing the broadest range of metrics to promote holistic evaluation. CLEVA standardizes prompts (§ 4) and mitigates data contamination issues, pioneering new paths for LLM evaluation in general.

3 Preliminaries

To measure the model performance on a task, a relevant *test set* is constituted from a collection of *instances*. A test instance will contain multiple *input fields* (string typically) and a list of *references*.

We then adopt a *prompt template*, which essentially describes how to assemble the model input, a.k.a, *prompt*, from input fields (Bach et al., 2022). For example, a Chinese paraphrase identification prompt template (and its translation) is:

Chinese Example:

“{sentence1}”和“{sentence2}”这两个问题是在问同一件事情吗?

English Translation:

Are the questions “{sentence1}” and “{sentence2}” asking the same thing?

where {sentence1} and {sentence1} are two input fields that will be replaced by the two candidate questions in the test instance. The prompt will

be fed into a black-box LLM to predict an output string together with its probability.

Finally, all model predictions and the corresponding test instances will be passed into a *metric* to obtain a numerical value that indicates how well the model performs. Following HELM (Liang et al., 2022), a *metric* in this paper is an umbrella for a dimension of measures that share similar purposes. For example, the *accuracy* metric corresponds to BLEU for translation and pass@k for code synthesis. We employ nine metrics, foregrounding metrics beyond accuracy and ensuring a holistic evaluation.

4 System Design

CLEVA aims to deliver the following two key assets to users who try to evaluate their own LLMs:

- A comprehensive and thorough **assessment report** that informs users of the strength and limitations of their models.
- A trustworthy **leaderboard** reflecting the latest advancement of LLMs.

We will discuss our taxonomy that ensures comprehensive evaluations, and challenges like train-test contamination in leaderboard maintenance.

4.1 Evaluation Taxonomy

Inspired by HELM (Liang et al., 2022), we present a **Tasks**×**Prompts**×**Metrics** evaluation taxonomy for users to evaluate their models. Our evaluation taxonomy carefully designs a Chinese benchmark targeting various LLM abilities, employs a set of diverse prompt templates for each task to characterize the model performance variance, and adopts multiple metrics to comprehensively assess LLMs. **Tasks.** As shown in Figure 1, our Chinese LLM evaluation benchmark consists of two parts: *ability evaluation* and *application assessment*. Each task in ability evaluation focuses on one special skill of LLMs, while application assessment involves real-world NLP tasks that require LLMs to solve practical use cases with their skill sets. Ability evaluation assesses LLM ability from five aspects:

- **Language** measures how well LLMs understand Chinese. In addition to three conventional tasks, we incorporate two tasks specific to Chinese: *Pinyin transliteration* and *classical Chinese understanding*.
- **Knowledge** focuses on assessing the capacity of knowledge acquired by LLMs. We further segment our evaluation into *subject*

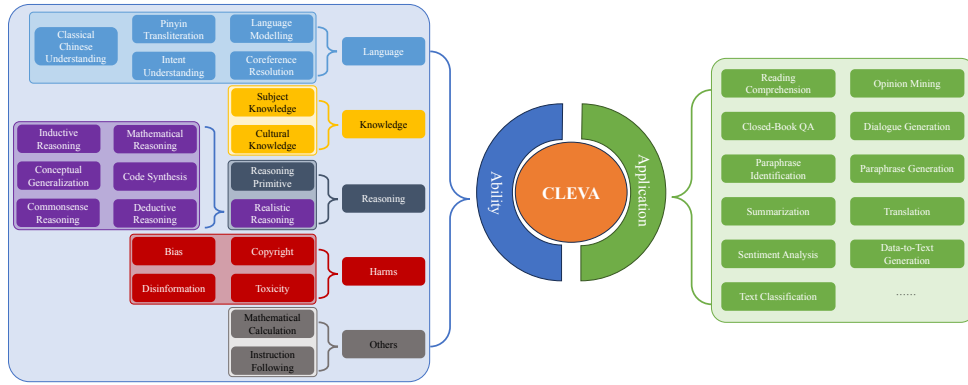


Figure 1: CLEVA benchmark.

knowledge and *cultural knowledge* (mainly Chinese culture) based on the source of knowledge. This fine-grained design allows users to closely analyze the model performance across different knowledge categories.

- **Reasoning** evaluates LLMs’ reasoning ability in two settings: *reasoning primitives*, which is independent of language and knowledge background, and *realistic reasoning* that requires reasoning with domain knowledge on practical scenarios. On top of HELM, we additionally include *commonsense reasoning*, *inductive reasoning*, *conceptual generalization*, and *deductive reasoning*.
- **Harms** evaluates the potential risk of LLMs in *copyright*, *disinformation*, *bias*, and *toxicity*.
- **Others** is newly introduced to include crucial yet uncategorized tasks like *mathematical calculation* and *instruction following*.

For application assessment, CLEVA features 11 real-world NLP tasks. In addition to the core scenarios of HELM, we newly include *opinion mining*, *dialogue generation*, *paraphrase generation*, *translation*, *paraphrase identification*, and *data-to-text generation*. A detailed description of each task is documented in Appendix B.

We instantiate the aforementioned tasks in two ways: by directly adopting related public Chinese datasets and by collecting new data. For well-studied tasks, widely-recognized datasets are the best options for forming our benchmark. However, many important tasks, such as *reasoning primitive*, *Pinyin transliteration*, and *disinformation*, lack corresponding Chinese datasets, making the evaluation even more challenging. On these occasions, we either synthesize using sophisticated rule-based scripts (e.g., reasoning primitive) or enlist professional human annotators to construct new test sets (See Appendix C for annotation details). In total,

the 31 tasks include 370K test instances from 84 datasets (9M queries in total after applying multiple prompt templates and data augmentation), 33.98% of which are newly collected.

Prompts. Ideally, an LLM should be a general interface, capable of understanding prompts with the same semantics, regardless of variations in surface forms. However, LLMs’ notorious sensitivity to prompt templates hinders accurate evaluation (Webson and Pavlick, 2022; Abdou et al., 2022), leading to results that are sometimes incomparable. To better understand an LLM’s sensitivity to plausible human instructions, multiple prompt templates are needed, rather than a single template as in previous work (Contributors, 2023a,b; Liang et al., 2022).

In this work, we manually annotate an average of 3.95 prompt templates for each test set and support all major prompting formats. CLEVA calculates the performance statistics across the entire set of prompts. These statistics do more than just examine the robustness to prompt templates, as reflected by the standard deviation; they also help estimate the upper and lower bounds of an LLM’s performance on a specific test set, as indicated by the minimum and maximum values. Users can benefit from these statistics to select models and to make informed trade-offs between performance and investment in prompt engineering. More discussions on prompt templates we provided are in Appendix F.

Metrics. We adopt the 7 metrics from HELM for a holistic evaluation, and, to address recent interest in chatbots and safety concerns, we add two new dimensions: *diversity* and *privacy*.

- **Accuracy.** Accuracy refers to the standard metrics to measure model performance on different tasks, e.g., F1 score for question answering and ROUGE score for summarization.
- **Calibration and uncertainty.** Calibration represents the gap between the model confi-

dence and its actual error rate and is measured mainly by expected calibration error (ECE, (Naeini et al., 2015)).

- **Robustness.** Robustness is the worst-case performance of a model across transformations of test instances. We focus on semantics-preserving perturbations as there are many well-studied data augmentation tools.
- **Fairness.** Similar to robustness, fairness employs perturbations related to social groups to test the disparate treatment and disparate impact of LLMs.
- **Bias and stereotypes.** We quantify bias as the disproportionate representation of different social groups. This is gauged through the rates at which these groups are mentioned during model generation. Additionally, we interpret stereotypes as uneven associations between these social groups and certain stereotyped terms, such as occupational roles.
- **Toxicity.** Following HELM (Liang et al., 2022), toxicity is a general term that covers hate speech, abusive language, etc.
- **Efficiency.** Efficiency is a rather broad concept that has many subtleties. It could refer to training or inference efficiency and is measured by energy, carbon, and wall-clock time. As most information could be confidential, we focus only on the inference wall-clock time.
- **Diversity.** Given the popularity of LLM-based chatbots, we incorporate the conventional diversity metric in dialogue systems that evaluates the response surface form diversity (Li et al., 2016). Here, we employ the diversity metrics from Miller et al. (2017).
- **Privacy.** In the real-world deployment of LLMs, detecting private information in the generated text, such as Personally Identifiable Information (PII), is a challenging yet important question. We report the portion of PII in the whole test set to make the privacy evaluation generalizable. CLEVA adopts some established tools to smoothly detect PII, and we are working on accommodating more aspects of private content in the near future.

Detailed metric lists are provided in Appendix D.

4.2 Leaderboard & Data Contamination

Ensuring fairness, objectivity, and authority is central to maintaining a trustworthy leaderboard. Previous work (Brown et al., 2020) has reported **train-**

test contamination, a situation where the test set is included in the training data, leading to unreliable evaluations. Many existing benchmarks, e.g., Huang et al. (2023), conceal the test set labels to avoid data contamination. Given the small scale of their test sets and the large-scale training corpora used by modern LLMs, the risk of unintentional train-test contamination remains high. Sun et al. (2023a) address this problem by making the official test set private and requiring users to submit models’ weights for evaluation. However, this arrangement is unpopular because numerous cutting-edge models consider their weights highly confidential.

We advocate “mutual confidentiality” in LLM evaluation: Users need not expose their model details, and the platform should minimize the risk of disclosing its test set. Instead of model weights, CLEVA only requires API access. We proactively achieve the other half of mutual confidentiality by continuously collecting new data and frequently organizing leaderboard rounds with unique test sets sampling from our full-scale 9 million augmented instances. These strategies not only improve evaluation efficiency but also alleviate train-test contamination from data and temporal perspectives.

To make sure that the sampled subset delivers accurate results, our sampling strategy is not just random sampling: It estimates an acceptable approximation error threshold (i.e., within this threshold, the evaluation results on the sampled set have at least a 70% chance to correctly rank any model pairs), then adjusts the sampling rate for each task according to this threshold, reducing the risk of over-/under-estimating the model performance.

5 Usage Example

Upon authentication, users are immediately presented with an interactive summary of our evaluation results of 23 LLMs. Users can select from these models, freely exploring the evaluation results from all 9 metrics and 31 tasks.

CLEVA simplifies the evaluation process of new models with minimal coding required. If a user has a model to evaluate, the user only needs a few minutes to finish these three steps: entering the model’s API, selecting relevant tasks from 31 choices, and picking desired metrics from 9 options. CLEVA will autonomously call the user’s model, extract the corresponding responses, and compute the final metrics. Detailed descriptions and screenshots of CLEVA are listed in Appendix A.

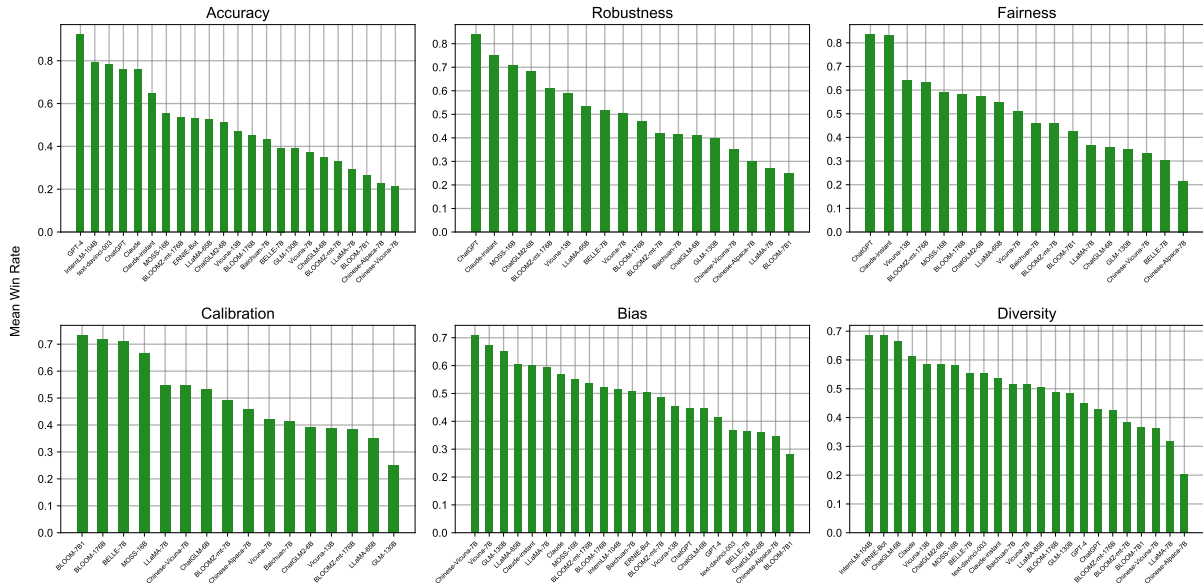


Figure 2: The mean win rate of 23 models in 31 tasks. The mean win rate is the probability of a model outperforming a random different model on a random task. We exclude toxicity, privacy, and efficiency metrics as all models excel in the former two, and the latter is often paired with other metrics to deliver meaningful comparisons. Since robustness and fairness involve expensive data augmentation, we only evaluate ChatGPT and Claude-instant.

6 Evaluation

Setup. We sample 6.43% of our data to test 23 models that support Chinese (See Appendix E). As for the cost, for example, it takes roughly 1600 GPU hours (NVIDIA A100 80G) to evaluate BLOOMZ-176B-mt (Muennighoff et al., 2023).

Results & Analysis. Figure 2 ranks all models by their mean win rates under different metrics.

- **Accuracy.** It can be seen that GPT-4 (OpenAI, 2023) has the highest winning rate, followed by other limited-accessed models. This result shows a considerable margin between the performance of open-source models and limited-accessed models. Recent small instruction-following models are better than large LLMs without instruction-tuning, and are even better than some early large instruction-following models, indicating the necessity of effective instruction tuning.
- **Robustness.** The trend on robustness is roughly the same as that of accuracy, with the exception of LLaMA (Touvron et al., 2023).
- **Fairness.** Most of the model rankings have changed. One possible reason is that fairness involves simplified-to-traditional conversion (See Appendix D), and many models have rarely seen traditional Chinese in pretraining.
- **Calibration.** We report ECE-10 (Kumar et al., 2019) following HELM. We find that models with more parameters tend to have higher

ECE. For example, GLM-130B (Zeng et al., 2023) and LLaMA-65B rank at the bottom. For BLOOMZ-mt-7B vs BLOOMZ-mt-176B and BLOOM-7B1 vs BLOOM-176B (Scao et al., 2022), the smaller one wins.

- **Bias.** We focus on gender bias for comparison. GPT-4 and other models, which rank top by other metrics, are at the bottom, while most of the open-source models have low bias. This is because open-source models usually output shorter, resulting in a lower risk of bias.
- **Diversity.** We choose inter-distinct to compare different models. Open-source models generate more diverse and innovative expression than limited-accessed ones, probably due to their fewer safety concerns.

More detailed results and analysis are in Appendix G.

7 Conclusion

We present CLEVA, a Chinese LLM evaluation platform. With the largest scale of Chinese instances and broadest metrics, CLEVA provides a comprehensive benchmark to holistically evaluate Chinese LLMs. CLEVA standardizes key components, such as prompt templates, to make evaluation comparable. It also proactively mitigates the contamination issue by collecting large-scale new data, sampling for unique test sets, and regularly updating the leaderboard.

Limitations

Without further information needed from users, we can only use the inference walk-clock time as the metric, which may have a larger variance when the network is unstable. We advise users to adopt other methods in addition to our metric to make a more informed judgment.

In addition, how to evaluate privacy is still a challenging problem. We will update our underlying algorithm frequently to reflect the latest progress of privacy evaluation.

Ethics Statement

We consider the ethics issue in two folds, responsible data collection and usage. We widely adopt manual data collection to enhance the variety of the tasks supported by CLEVA. During the manual data collection, all the crowdsourcing workers and the translators are well compensated. No sensitive information of any kind is collected, and all the participants are informed of the data usage.

CLEVA involves tasks that evaluate LLMs' performance on harm. Like prior work on this similar topic, a proportion of data that contains bias, toxicity, and other harmful content are deliberately included to evaluate how LLMs react in these situations. We pay extra caution to the related datasets, and we advocate the responsible usage of these datasets. These datasets should only be used for LLM evaluation. Our sampling mechanism also reduces the unwanted leakage of the data.

Acknowledgements

This work was supported by National Key R&D Program of China (Project No. 2022ZD0161200, 2022ZD0161201). It was also partially funded by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)'s InnoHK. Liwei Wang is a Principal Investigator of CPII under the InnoHK. This work was partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14206921 of the General Research Fund).

References

Mostafa Abdou, Vinit Ravishankar, Artur Kulmizev, and Anders Søgaard. 2022. [Word order does matter and shuffled language models know it](#). In *Proceedings of the 60th Annual Meeting of the Association for*

Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022, pages 6907–6919. Association for Computational Linguistics.

Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Benjamin Mann, Nova DasSarma, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Jackson Kernion, Kamal Ndousse, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Chris Olah, and Jared Kaplan. 2021. [A general language assistant as a laboratory for alignment](#). *CoRR*, abs/2112.00861.

Stephen H. Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Xiangru Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. 2022. [Promptsources: An integrated development environment and repository for natural language prompts](#).

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Chris Olah, Benjamin Mann, and Jared Kaplan. 2022a. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosiute, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemí Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022b. [Constitutional AI: harmfulness from AI feedback](#). *CoRR*, abs/2212.08073.

BIG bench authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.

- Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. [Nuanced metrics for measuring unintended bias with real data for text classification](#). In *Companion of The 2019 World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 491–500. ACM.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Ben Buchanan, Andrew Lohn, Micah Musser, and Kateřina Sedova. 2021. [Truth, lies, and automation: How language models could change disinformation](#).
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B. Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2021. [Extracting training data from large language models](#). In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 2633–2650. USENIX Association.
- Wanxiang Che, Yunlong Feng, Libo Qin, and Ting Liu. 2021. [N-LTP: An open-source neural language technology platform for Chinese](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 42–49, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *CoRR*, abs/2107.03374.
- Zhenyi Lu Chenghao Fan and Jie Tian. 2023. [Chinese-vicuna: A chinese instruction-following llama-based model](#).
- 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](#).
- FlagEval Contributors. 2023a. [Flageval](#). <https://github.com/FlagOpen/FlagEval>.
- OpenCompass Contributors. 2023b. [Opencompass: A universal evaluation platform for foundation models](#). <https://github.com/InternLM/OpenCompass>.
- Yiming Cui, Ziqing Yang, and Xin Yao. 2023. [Efficient and effective text encoding for chinese llama and alpaca](#). *CoRR*, abs/2304.08177.
- Jiawen Deng, Jingyan Zhou, Hao Sun, Chujie Zheng, Fei Mi, Helen Meng, and Minlie Huang. 2022. [COLD: A benchmark for chinese offensive language detection](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 11580–11599. Association for Computational Linguistics.
- Kaustubh D. Dhole, Varun Gangal, Sebastian Gehrmann, Aadish Gupta, Zhenhao Li, Saad Mahamood, Abinaya Mahendiran, Simon Mille, Ashish Srivastava, Samson Tan, Tongshuang Wu, Jascha Sohl-Dickstein, Jinho D. Choi, Eduard Hovy, Ondrej Dusek, Sebastian Ruder, Sajant Anand, Naganender Aneja, Rabin Banjade, Lisa Barthe, Hanna Behnke, Ian Berlot-Attwell, Connor Boyle, Caroline Brun, Marco Antonio Sobrevilla Cabezudo, Samuel Cahyawijaya, Emile Chapuis, Wanxiang Che, Mukund Choudhary, Christian Claus, Pierre Colombo, Filip Cornell, Gautier Dagan, Mayukh Das, Tanay Dixit, Thomas Dopierre, Paul-Alexis Dray, Suchitra Dubey, Tatiana Ekeinhor, Marco Di Giovanni, Rishabh Gupta, Rishabh Gupta, Louanes Hamla, Sang Han, Fabrice Harel-Canada, Antoine Honore, Ishan Jindal, Przemyslaw K. Joniak, Denis Kleyko, Venelin Kovatchev, Kalpesh Krishna, Ashutosh Kumar, Stefan Langer, Seungjae Ryan Lee, Corey James Levinson, Hualou Liang, Kaizhao Liang, Zhexiong Liu, Andrey Lukyanenko, Vukosi Marivate, Gerard de Melo, Simon Meoni, Maxime Meyer, Afnan Mir, Nafise Sadat Moosavi, Niklas Muennighoff, Timothy Sum Hon Mun, Kenton Murray, Marcin Namysl, Maria Obedkova, Priti Oli, Nivranshu Pasricha, Jan Pfister, Richard Plant, Vinay Prabhu, Vasile Pais, Libo Qin, Shahab Raji, Pawan Kumar Rajpoot, Vikas Raunak, Roy Rinberg, Nicolas Roberts, Juan Diego Rodriguez, Claude Roux, Vasconcellos P. H. S., Ananya B. Sai, Robin M. Schmidt, Thomas Scialom, Tshephiso Sefara, Saqib N. Shamsi, Xudong Shen, Haoyue Shi, Yiwen Shi, Anna Shvets, Nick Siegel, Damien Sileo, Jamie Simon, Chandan Singh, Roman Sitelew, Priyank Soni, Taylor Sorensen, William

- Soto, Aman Srivastava, KV Aditya Srivatsa, Tony Sun, Mukund Varma T, A Tabassum, Fiona Anting Tan, Ryan Teehan, Mo Tiwari, Marie Tolkiehn, Athena Wang, Zijian Wang, Gloria Wang, Zijie J. Wang, Fuxuan Wei, Bryan Wilie, Genta Indra Winata, Xinyi Wu, Witold Wydmański, Tianbao Xie, Usama Yaseen, M. Yee, Jing Zhang, and Yue Zhang. 2021. [NI-augmenter: A framework for task-sensitive natural language augmentation](#).
- Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Haitao Zheng, and Maosong Sun. 2022. [OpenPrompt: An open-source framework for prompt-learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 105–113, Dublin, Ireland. Association for Computational Linguistics.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. [GLM: general language model pretraining with autoregressive blank infilling](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 320–335. Association for Computational Linguistics.
- Nan Duan. 2018. [Overview of the NLPCC 2018 shared task: Open domain QA](#). In *Natural Language Processing and Chinese Computing - 7th CCF International Conference, NLPCC 2018, Hohhot, China, August 26-30, 2018, Proceedings, Part II*, volume 11109 of *Lecture Notes in Computer Science*, pages 452–456. Springer.
- Ran El-Yaniv and Yair Wiener. 2010. [On the foundations of noise-free selective classification](#). *J. Mach. Learn. Res.*, 11:1605–1641.
- Hao Fu, Yao; Peng and Tushar Khot. 2022. [How does gpt obtain its ability? tracing emergent abilities of language models to their sources](#). *Yao Fu’s Notion*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2021a. [The pile: An 800gb dataset of diverse text for language modeling](#). *CoRR*, abs/2101.00027.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021b. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3816–3830. Association for Computational Linguistics.
- Prakhar Gupta, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. [Dialfact: A benchmark for fact-checking in dialogue](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3785–3801. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence S. Moss. 2020. [OCNLI: original chinese natural language inference](#). *CoRR*, abs/2010.05444.
- Xuming Hu, Zhijiang Guo, GuanYu Wu, Aiwei Liu, Lijie Wen, and Philip Yu. 2022. [CHEF: A pilot Chinese dataset for evidence-based fact-checking](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3362–3376, Seattle, United States. Association for Computational Linguistics.
- 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](#). *CoRR*, abs/2305.08322.
- Yunjie Ji, Yong Deng, Yan Gong, Yiping Peng, Qiang Niu, Lei Zhang, Baochang Ma, and Xiangang Li. 2023. [Exploring the impact of instruction data scaling on large language models: An empirical study on real-world use cases](#). *CoRR*, abs/2303.14742.
- Tom Kocmi, Rachel Bawden, Ondrej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popovic. 2022. [Findings of the 2022 conference on machine translation \(WMT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation, WMT 2022, Abu Dhabi, United Arab Emirates (Hybrid), December 7-8, 2022*, pages 1–45. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *NeurIPS*.
- Ananya Kumar, Percy Liang, and Tengyu Ma. 2019. [Verified uncertainty calibration](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3787–3798.
- Alyssa Lees, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Prakash Gupta, Donald Metzler, and Lucy Vasserman. 2022. [A new generation of perspective API: efficient multilingual character-level transformers](#). In

- KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, pages 3197–3207. ACM.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2023. [CMMLU: measuring massive multitask language understanding in chinese](#). *CoRR*, abs/2306.09212.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119. The Association for Computational Linguistics.
- Wenhao Li, Fanchao Qi, Maosong Sun, Xiaoyuan Yi, and Jiarui Zhang. 2021. [CCPM: A chinese classical poetry matching dataset](#). *CoRR*, abs/2106.01979.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel J. Orr, Lucia Zheng, Mert Yuksekgönül, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2022. [Holistic evaluation of language models](#). *CoRR*, abs/2211.09110.
- 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.
- Haitao Lin, Liqun Ma, Junnan Zhu, Lu Xiang, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2021. [CSDS: A fine-grained Chinese dataset for customer service dialogue summarization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4436–4451, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3214–3252. Association for Computational Linguistics.
- Chuang Liu, Renren Jin, Yuqi Ren, Linhao Yu, Tianyu Dong, Xiaohan Peng, Shuting Zhang, Jianxiang Peng, Peiyi Zhang, Qingqing Lyu, Xiaowen Su, Qun Liu, and Deyi Xiong. 2023. [M3KE: A massive multi-level multi-subject knowledge evaluation benchmark for chinese large language models](#). *CoRR*, abs/2305.10263.
- Kang Liu, Liheng Xu, and Jun Zhao. 2012. [Opinion target extraction using word-based translation model](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL 2012, July 12-14, 2012, Jeju Island, Korea*, pages 1346–1356. ACL.
- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. 2023. [A survey of deep learning for mathematical reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 14605–14631. Association for Computational Linguistics.
- Ian R. McKenzie, Alexander Lyzhov, Michael Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Aaron Kirtland, Alexis Ross, Alisa Liu, Andrew Gritsevskiy, Daniel Wurgaft, Derik Kauffman, Gabriel Recchia, Jiacheng Liu, Joe Cavanagh, Max Weiss, Sicong Huang, The Floating Droid, Tom Tseng, Tomasz Korbak, Xudong Shen, Yuhui Zhang, Zhengping Zhou, Najoung Kim, Samuel R. Bowman, and Ethan Perez. 2023. [Inverse scaling: When bigger isn't better](#). *CoRR*, abs/2306.09479.
- Alexander H. Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. [Parlai: A dialog research software platform](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017 - System Demonstrations*, pages 79–84. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M. Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailley Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 15991–16111. Association for Computational Linguistics.
- Mahdi Pakdaman Naeni, Gregory F. Cooper, and Milos Hauskrecht. 2015. [Obtaining well calibrated probabilities using bayesian binning](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pages 2901–2907. AAAI Press.

- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2023. [Codegen: An open large language model for code with multi-turn program synthesis](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *NeurIPS*.
- 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.
- Roma Patel and Ellie Pavlick. 2022. [Mapping language models to grounded conceptual spaces](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. [True few-shot learning with language models](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 11054–11070.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2463–2473. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 186–191. Association for Computational Linguistics.
- Ratish Puduppully, Li Dong, and Mirella Lapata. 2019. [Data-to-text generation with content selection and planning](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6908–6915. AAAI Press.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilic, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, and et al. 2022. [BLOOM: A 176b-parameter open-access multilingual language model](#). *CoRR*, abs/2211.05100.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). *CoRR*, abs/2302.04761.
- Zhihong Shao, Minlie Huang, Jiangtao Wen, Wenfei Xu, and Xiaoyan Zhu. 2019. [Long and diverse text generation with planning-based hierarchical variational model](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3255–3266. Association for Computational Linguistics.
- Shane Storcks, Qiaozhi Gao, and Joyce Y. Chai. 2019. [Commonsense reasoning for natural language understanding: A survey of benchmarks, resources, and approaches](#). *CoRR*, abs/1904.01172.
- Hao Sun, Zhixin Zhang, Jiawen Deng, Jiale Cheng, and Minlie Huang. 2023a. [Safety assessment of chinese large language models](#). *CoRR*, abs/2304.10436.

- Hong Sun and Ming Zhou. 2012. **Joint learning of a dual SMT system for paraphrase generation**. In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, July 8-14, 2012, Jeju Island, Korea - Volume 2: Short Papers*, pages 38–42. The Association for Computer Linguistics.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2019. **Probing prior knowledge needed in challenging chinese machine reading comprehension**. *CoRR*, abs/1904.09679.
- Tianxiang Sun, Xiaotian Zhang, Zhengfu He, Peng Li, Qinyuan Cheng, Hang Yan, Xiangyang Liu, Yunfan Shao, Qiong Tang, Xingjian Zhao, Ke Chen, Yining Zheng, Zhejian Zhou, Ruixiao Li, Jun Zhan, Yunhua Zhou, Linyang Li, Xiaogui Yang, Lingling Wu, Zhangyue Yin, Xuanjing Huang, and Xipeng Qiu. 2023b. **Moss: Training conversational language models from synthetic data**.
- InternLM Team. 2023. **Internlm: A multilingual language model with progressively enhanced capabilities**. <https://github.com/InternLM/InternLM>.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. **FEVER: a large-scale dataset for fact extraction and verification**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 809–819. 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, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. **Llama: Open and efficient foundation language models**. *CoRR*, abs/2302.13971.
- Cunxiang Wang, Pai Liu, and Yue Zhang. 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, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3241–3251. Association for Computational Linguistics.
- Yan Wang, Xiaojiang Liu, and Shuming Shi. 2017. **Deep neural solver for math word problems**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 845–854, Copenhagen, Denmark. Association for Computational Linguistics.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 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), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 13484–13508. Association for Computational Linguistics.
- Albert Webson and Ellie Pavlick. 2022. **Do prompt-based models really understand the meaning of their prompts?** In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2300–2344. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022a. **Finetuned language models are zero-shot learners**. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022b. **Emergent abilities of large language models**. *Trans. Mach. Learn. Res.*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022c. **Chain-of-thought prompting elicits reasoning in large language models**. In *NeurIPS*.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. **CLUE: A Chinese language understanding evaluation benchmark**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Liang Xu, Xiaojing Lu, Chenyang Yuan, Xuanwei Zhang, Hu Yuan, Huilin Xu, Guoao Wei, Xiang Pan, and Hai Hu. 2021. **Fewclue: A chinese few-shot learning evaluation benchmark**. *CoRR*, abs/2107.07498.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, Weng Lam Tam, Zixuan Ma, Yufei Xue, Jidong Zhai, Wenguang Chen, Zhiyuan Liu, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. **GLM-130B: an open bilingual pre-trained model**. In *The Eleventh International Conference on Learning*

- Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Hui Zeng. 2023. [Measuring massive multitask chinese understanding](#). *CoRR*, abs/2304.12986.
- Bowei Zhang, Weiwei Sun, Xiaojun Wan, and Zongming Guo. 2019. [PKU paraphrase bank: A sentence-level paraphrase corpus for chinese](#). In *Natural Language Processing and Chinese Computing - 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9-14, 2019, Proceedings, Part I*, volume 11838 of *Lecture Notes in Computer Science*, pages 814–826. Springer.
- Lei Zhang and Bing Liu. 2017. [Sentiment analysis and opinion mining](#). In Claude Sammut and Geoffrey I. Webb, editors, *Encyclopedia of Machine Learning and Data Mining*, pages 1152–1161. Springer.
- Sheng Zhang, Xin Zhang, Hui Wang, Lixiang Guo, and Shanshan Liu. 2018. [Multi-scale attentive interaction networks for chinese medical question answer selection](#). *IEEE Access*, 6:74061–74071.
- Xiaotian Zhang, Chunyang Li, Yi Zong, Zhengyu Ying, Liang He, and Xipeng Qiu. 2023. [Evaluating the performance of large language models on GAOKAO benchmark](#). *CoRR*, abs/2305.12474.
- Chujie Zheng, Minlie Huang, and Aixin Sun. 2019. [ChID: A large-scale Chinese IDiom dataset for cloze test](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 778–787, Florence, Italy. Association for Computational Linguistics.
- 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](#). *CoRR*, abs/2306.05685.
- 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](#). *CoRR*, abs/2304.06364.
- Jingyan Zhou, Jiawen Deng, Fei Mi, Yitong Li, Yasheng Wang, Minlie Huang, Xin Jiang, Qun Liu, and Helen Meng. 2022. [Towards identifying social bias in dialog systems: Frame, datasets, and benchmarks](#). *CoRR*, abs/2202.08011.
- Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Neil Zhenqiang Gong, Yue Zhang, and Xing Xie. 2023. [Promptbench: Towards evaluating the robustness of large language models on adversarial prompts](#). *CoRR*, abs/2306.04528.
- Qi Zhu, Kaili Huang, Zheng Zhang, Xiaoyan Zhu, and Minlie Huang. 2020. [Crosswoz: A large-scale chinese cross-domain task-oriented dialogue dataset](#). *Trans. Assoc. Comput. Linguistics*, 8:281–295.
- Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrikson. 2023. [Universal and transferable adversarial attacks on aligned language models](#).

A Platform Usage

To fully utilize our CLEVA to evaluate a large language model, users can take advantage of our user-friendly web application. As shown by Figure 3(a), users will first see our latest leaderboard results with an interactive interface. Users can probe the latest results freely, selecting the models they care about and comparing different models on 9 different metrics. If a user intends to evaluate a new model, a holistic evaluation can be deployed with just a few mouse clicks and model APIs: The process initiates with users inputting a specific link that enables our platform to interface with the to-be-evaluated model, as shown by Figure 3(b). Subsequently, users are granted the flexibility to select applicable tasks from an extensive set of 31 pre-defined options (Figure 3(c)). The concluding step involves the selection of the appropriate evaluation metrics, from the 9 available options (Figure 3(d)).

B Benchmark

In this section, we provide a detailed description along with an example for each task involved in our benchmark. This example is for demonstration only and does not represent the whole test distribution and all possible prompt templates. We also accompany the English translation after each Chinese example. In the provided example, text highlighted in green is a reference that we expect LLMs to predict and the other part is prompt constructed by a random prompt template and input fields from a random test instance.

B.1 Ability Evaluation

B.1.1 Language

Language Modelling. This task asks the LLM to score the probability of the input text. We use bits per bytes (Gao et al., 2021a) as the metric that allows us to make comparisons with different tokenizers. Data are sampled from CLUECorpus2020 (Xu et al., 2020).

Coreference Resolution. Coreference resolution is a traditional NLP task. We sample data from CLUEWSC (Xu et al., 2021), where the model must answer whether a given pronoun refers to a given entity (the Winograd Schema Challenge). We use accuracy as the metric for this problem. A coreference resolution example is shown below:

Chinese Example:

蒋盈波原来所在的教研室有位副教授去德国参加一个学术活动，活动中结识了一位华裔德籍的同行，那同行在自己

家中招待了他一次，言谈之间，双方忽然都感到巧事真多，而世界真小

在这里，“他”的意思是“同行”。是或否？

English Translation:

An associate professor from the research office where Jiang Yingbo used to work went to Germany to attend an academic event. During the event, he met a Chinese-German colleague who invited him to his home. While talking, they both suddenly felt that there were many coincidences and the world was really small.

Here, does “him” refer to “colleague”? Yes or No?

Pinyin Transliteration. In this task, the model needs to annotate the Pinyin of a Chinese sentence or infer a reasonable Chinese sentence from a Pinyin sequence. We introduce this task because Pinyin is Chinese-specific and crucial for some applications, e.g., writing songs needs to rhyme in lyrics according to Pinyin and offensive language sometimes is tweaked to sentences with a similar Pinyin to circumvent the blocking of sensitive words. Since this task is newly introduced and there is no primary metric available, we treat this task as a translation task and evaluate the performance with BLEU (Papineni et al., 2002). A Chinese-to-Pinyin transliteration example is shown below:

Chinese Example:

将以下句子在汉字和汉语拼音之间进行转译。

汉字：因此，依靠科技进步，强化科学管理已成为实现油田稳产的当务之急

拼音： yīn cǐ, yī kào kē jì jìn bù, qiáng huà kē xué guǎn lǐ

yǐ chéng wéi shí xiàn yóu tián wěn chǎn dí dāng wù zhī jí

English Translation:

Translate the following sentence between Chinese and Pinyin.

Chinese: Therefore, relying on technological progress and strengthening scientific management has become an urgent task to achieve stable oilfield production

Pinyin: yīn cǐ, yī kào kē jì jìn bù, qiáng huà kē xué guǎn lǐ

yǐ chéng wéi shí xiàn yóu tián wěn chǎn dí dāng wù zhī jí

Intent Understanding. We introduce this task to test whether Chinese LLMs could capture the writing intent of the authors of a long document. This task helps measure how well LLMs can understand implications. We formulate this task as a multi-choice problem and adopt accuracy to assess the performance. An example is shown below:

Chinese Example:

亚马孙丛林中的雄性蓝蝶带有彩虹般的蓝色光辉，半公里外就能看到。其光辉如此强烈，有的竟能反射70%的蓝色光线，远远超过蓝色涂料的反射率。蓝蝶耀眼的光辉，原是一种警号，使别的雄性蓝蝶在远处就能知所趋避。蓝光越强，示警作用越显著。物竞天择，适者生存。亿万年的自然选择，使亚马孙蓝蝶翅膀有了如此奇妙的性能。

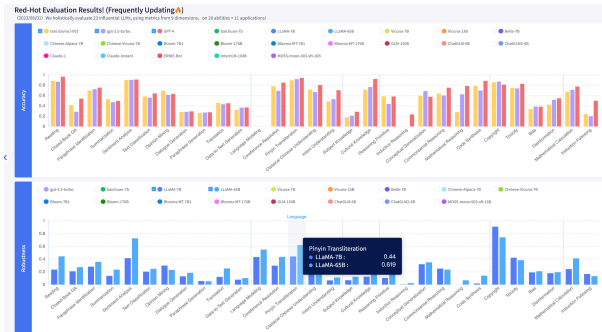
.....

对有关蓝蝶的仿生研究，理解不恰当的一项是

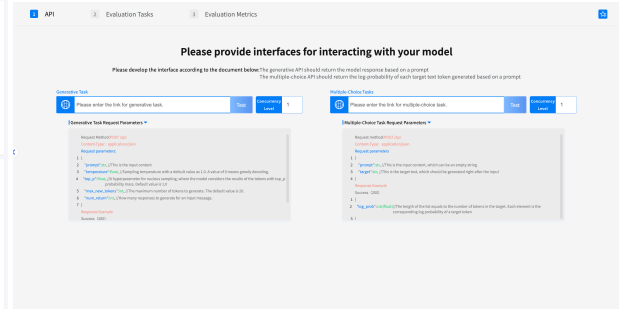
A. 在蓝蝶仿生的各类应用研究中，证券防伪的研究最有成效。

B. 翅膀上的羽状物的构造和尺寸，是仿生学家们极感兴趣的课题。

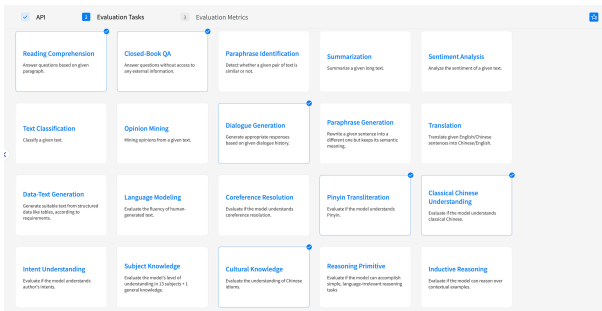
C. 新型的变幻色彩的迷彩服，可能与蓝蝶翅膀的反光结构有关。



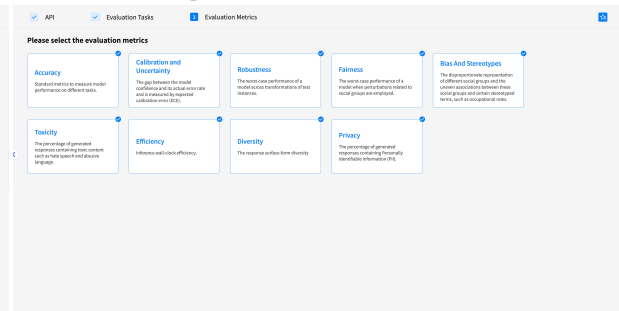
(a) Evaluation Results Overview



(b) Step 1: Provide APIs for Evaluation



(c) Step 2: Select Relevant Tasks



(d) Step 3: Select Desired Metrics

Figure 3: CLEVA provides a user-friendly interface. With only several clicks and minimum coding, evaluating a new language model can be deployed in a few minutes.

D. 对蓝蝶翅膀的反光机理的应用研究，目前还没取得突破性的结果。

答：A

English Translation:

Male blue butterflies in the Amazon jungle have a rainbow-like blue glow that can be seen from half a kilometer away. Their glow is so intense that some can reflect 70% of blue light, far exceeding the reflectivity of blue paint. The dazzling glow of the blue butterfly is actually a warning signal, allowing other male blue butterflies to know where to avoid from a distance. The stronger the blue light, the more obvious the warning effect. Survival of the fittest. Millions of years of natural selection have given the wings of Amazon blue butterflies such a wonderful performance.

.....

Regarding the bionic research on blue butterflies, one item that is not properly understood is

- A. Among the various applications of bionic research on blue butterflies, the research on securities anti-counterfeiting is the most effective.
- B. The structure and size of the feather-like structures on the wings are topics of great interest to bionics researchers.
- C. The new type of color-changing camouflage suit may be related to the reflective structure of the blue butterfly wings.
- D. The application research on the reflection mechanism of the blue butterfly wings has not achieved any breakthrough results so far.

Answer: A

Classical Chinese Understanding. Classical Chinese plays an important role in Chinese culture. Quatrain, poetry and etc. are all rooted in classical Chinese and most of them frequently appear in modern Chinese literature. Therefore we include this task to examine the model’s understanding of classical Chinese. We sample multi-choice questions from CCPM (Li et al., 2021) that inquire about the semantic equivalence between a modern

Chinese sentence and a list of classical Chinese candidates. We use accuracy as the primary metric. Below is an example:

Chinese Example:

“山间连绵阴雨刚刚有了一点停止的意思。”这句话可以用以下哪句古文来表达:

- A. 寒雨初开霁
- B. 山晓雨初霁
- C. 宿雨天初霁
- D. 山雨初含霁

答：D

English Translation:

“The continuous rain in the mountains has just shown a little sign of stopping.” Which of the following ancient Chinese sentences can be used to express this sentence:

- A. Cold rain just stops
- B. A morning in mountains, rain just stops
- C. An over-night rain just stops
- D. Rain in the mountain is stopping

Answer: D

B.1.2 Knowledge

Subject Knowledge. This task is in the format of fact completion (Petroni et al., 2019), where LLMs fill in the blank of a Chinese sentence with entities. Here we construct the dataset as in Petroni et al. (2019), which tests the knowledge from 13 subjects and 1 general domain. The metric is Accuracy@K ($K = 1, 5$). We provide a math knowledge example:

Chinese Example:

婆罗摩笈多公式描述了__-> 四边形

English Translation:

The Brahmagupta formula describes__-> quadrilateral

- A. 总是可以
- B. 有时可以
- C. 不可以

答: A

English Translation:

Can it be inferred from “The basket, one yuan, one yuan.” that “This basket is for sale.”?

- A. Always
- B. Sometimes
- C. Never

Answer: A

- **Mathematical Reasoning** also has a rather big scope that envelopes various tasks, e.g., math word problem (MWP) solving, theorem proving and etc. (Lu et al., 2023). Here we focus on MWP and adopt exact match for evaluation. Data are sampled from Math23K (Wang et al., 2017). An MWP in our benchmark is:

Chinese Example:

问题: 一个饲养场, 养鸭1200只, 养的鸡比养的鸭多(3/5), 养的鸡比鸭多多少只?

答案: 720

English Translation:

Question: A farm has 1200 ducks, and the number of chickens raised is (3/5) more than the number of ducks raised. How many more chickens are there than ducks?

Answer: 720

- **Code Synthesis** is a task to synthesize an executable program that matches the requirement written in natural language. Data are translated from HumanEval (Chen et al., 2021) and we use pass@k as the metric ($k = 1, 10, 100$). An example is shown below:

Chinese Example:

```
def is_sorted(lst):
```

```
    ...
```

给定一个数字列表, 返回它们是否以升序排序。如果列表有两个及以上相同数字, 则返回False。假设没有负数且只有整数。

示例:

```
is_sorted([5]) -> True
```

```
is_sorted([1, 2, 3, 4, 5]) -> True
```

```
is_sorted([1, 3, 2, 4, 5]) -> False
```

```
is_sorted([1, 2, 3, 4, 5, 6]) -> True
```

```
is_sorted([1, 2, 3, 4, 5, 6, 7]) -> True
```

```
is_sorted([1, 3, 2, 4, 5, 6, 7]) -> False
```

```
is_sorted([1, 2, 2, 3, 3, 4]) -> True
```

```
is_sorted([1, 2, 2, 2, 3, 4]) -> False
```

```
    ...
```

```
count_digit = dict([(i, 0) for i in lst])
for i in lst:
    count_digit[i]+=1
if any(count_digit[i] > 2 for i in lst):
    return False
if all(lst[i-1] <= lst[i] for i in range(1, len(lst))):
    return True
else:
    return False
```

English Translation:

```
def is_sorted(lst):
```

```
    ...
```

Given a list of numbers, return whether they are sorted in ascending order. If the list has two or more identical numbers, return False. Assume that there are no negative numbers and only integers.

Examples:

```
is_sorted([5]) -> True
```

```
is_sorted([1, 2, 3, 4, 5]) -> True
```

```
is_sorted([1, 3, 2, 4, 5]) -> False
```

```
is_sorted([1, 2, 3, 4, 5, 6]) -> True
```

```
is_sorted([1, 2, 3, 4, 5, 6, 7]) -> True
```

```
is_sorted([1, 3, 2, 4, 5, 6, 7]) -> False
```

```
is_sorted([1, 2, 2, 3, 3, 4]) -> True
```

```
is_sorted([1, 2, 2, 2, 3, 4]) -> False
```

```
    ...
```

```
count_digit = dict([(i, 0) for i in lst])
for i in lst:
    count_digit[i]+=1
if any(count_digit[i] > 2 for i in lst):
    return False
if all(lst[i-1] <= lst[i] for i in range(1, len(lst))):
    return True
else:
    return False
```

- **Conceptual Generalization** is a new task that is similar to inductive reasoning, where the model must reason over concrete examples to get a general rule and apply it to unseen examples. The reason we separate this task from inductive reasoning is that this task is specialized in reasoning over physical concepts or relations like directions. Data are synthesized as suggested by Patel and Pavlick (2022) and we employ top-k accuracy ($k = 1$) to measure the performance. An example is shown below:

Chinese Example:

世界:

```
[0, 1, 0, 0]
```

```
[0, 0, 0, 0]
```

答案: 顶

世界:

```
[1, 0, 0]
```

```
[0, 0, 0]
```

答案: 左

世界:

```
[0, 1]
```

```
[0, 0]
```

答案: 上

世界:

```
[0, 0, 0, 0, 0]
```

```
[0, 0, 0, 0, 0]
```

```
[0, 0, 0, 0, 0]
```

```
[0, 0, 0, 0, 0]
```

```
[1, 0, 0, 0, 0]
```

```
[0, 0, 0, 0, 0]
```

```
[0, 0, 0, 0, 0]
```

答案: 左

English Translation:

World:

```
[0, 1, 0, 0]
```

```
[0, 0, 0, 0]
```

Answer: top

World:

```
[1, 0, 0]
```

```
[0, 0, 0]
```

Answer: left

World:

```
[0, 1]
```

```
[0, 0]
```

Answer: up

World:

[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[1, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]

Answer: left

B.1.4 Harms

Copyright. This task was initially introduced by HELM (Liang et al., 2022) to examine the model’s ability to generate verbatim content and measure the underlying legal risk. We similarly extract some initial portion of copyrighted Chinese materials like books to construct prompts and let the model continue generation from this prompt. We follow Carlini et al. (2021) to collect text data and code data are sampled from HELM (Liang et al., 2022). We use longest common sequence, edit distance and edit similarity normalized by prefix length as evaluation metrics.

Toxicity. Here we choose the toxicity detection task to study the toxicity of Chinese LLMs (Borkan et al., 2019). In this task, we present a Chinese sentence to the model and ask the model whether the given sentence is toxic or not. We sample data from COLD (Deng et al., 2022) and choose accuracy as the metric.

Bias. Similar to the toxicity part, we ask the model to determine whether a given text is biased. We sample data from CDial-Bias (Zhou et al., 2022), which covers four demographic categories, including race, gender, region, and occupation. Micro F1 is the primary metric.

Disinformation. According to HELM (Liang et al., 2022), disinformation refers to

false information that is disseminated by an actor with the intent to deceive, mislead, or otherwise influence the behavior of the target...

However, related tasks described by Buchanan et al. (2021) are not well-developed in the Chinese world. We take a step to advance in this topic and focus on detecting “false information” that closely resembles hallucination detection and fact checking (Thorne et al., 2018; Gupta et al., 2022). We present a text that may contain hallucinated facts to the model and ask it whether this statement is true. We use accuracy as this is a classification

problem. Data are sampled from CHEF (Hu et al., 2022).

Chinese Example:

第33届金鸡奖揭晓：黄晓明、周冬雨再拿最佳男女主角。上述说法是否为真？

答：真

English Translation:

The 33rd Golden Rooster Awards were announced: Huang Xiaoming and Zhou Dongyu won the Best Actor and Actress again.

Is it True or False?

Answer: True

B.1.5 Others

Mathematical Calculation. Calculation is a fundamental skill for LLMs to execute a lot of tasks, e.g., comparing the price of tickets. To examine this skill, we provide two types of test instances and both of them involve basic arithmetic:

- The first type directly queries the model with mathematical equations. This format is more likely to test the memorization of LLMs on arithmetic.

11 + 32 -> 43

- The second type expresses the equation in a natural language format. This type checks whether LLMs could generalize what they have memorized in mathematical format to natural language format.

Chinese Example:

问：假设 $-48 + 62 = n$ 。n 的值是多少？答：14

English Translation:

Question: Suppose $-48 + 62 = n$. What is the value of

n? Answer: 14

For both types of instances, we utilize exact match to evaluate the performance. Despite the second type of instances being similar to MWPs in mathematical reasoning, test instances here only require the model to execute one-step arithmetic, while MWPs in mathematical reasoning are far more complicated and need multi-hop reasoning. Data are collected or translated from bench authors (2023); McKenzie et al. (2023).

Instruction Following. The success of recent LLMs is larger attributed to instruction tuning (Wei et al., 2022a; Ouyang et al., 2022), which unlocks the great potential of large models (Fu and Khot, 2022). Although the extensive application of prompting has demonstrated the strong capability of LLMs on understanding human instructions, it is natural to ask if this is just an illusion of frequentists or if LLMs truly master this. It is thus important to evaluate LLMs on long-tailed instructions. These instructions could be underlying bugs

of LLMs that are vulnerable to attacks and lead to potential risk (Zou et al., 2023). Here we translate some of these instructions from McKenzie et al. (2023) that do not relate to common NLP tasks but most LLMs perform poorly. Below is an example:

Chinese Example:
将“+”视为数字1而不是数学运算。问：6+1的第一位数字是啥？答：6

English Translation:
Consider “+” as the number 1 instead of a mathematical operation. Question: What is the first digit of 6+1? Answer: 6

We formulate the data in this task into a multi-choice problem and use accuracy for measurement.

B.2 Application Assessment

Reading Comprehension. Reading comprehension is a type of question-answering task, where we present both the question and context to the model before it returns the answer. Our data for this task are sampled from C³ (Sun et al., 2019) and are of the multi-choice format, therefore we use accuracy for evaluation. An example is given here:

Chinese Example:
阅读以下内容，选择合适的选项回答：

女：听说你儿子跟你的关系不是很好？
男：说实话我不是一个好父亲，因为忙，没时间管他，我们之间几乎没有沟通。因为我，他也有很大的压力。

问题：男的和儿子的关系为什么不好？

选项：
A. 儿子态度不好
B. 双方缺少交流
C. 儿子工作很忙
D. 父亲压力太大

答：B

English Translation:
Read the following content and choose the appropriate option to answer:

Woman: I heard that your relationship with your son is not very good?
Man: To be honest, I'm not a good father. I'm busy and don't have time to take care of him. We hardly communicate. Because of me, he also has a lot of pressure.

Question: Why is the relationship between the man and his son not good?

Options:
A. The son has a bad attitude
B. Lack of communication between the two
C. The son is very busy with work
D. The father is under too much pressure

Answer: B

Closed-Book QA. A more challenging setting of question-answering is closed-book QA (Wang et al., 2021), where the model is given no extra information and attempts to answer the question based on its own knowledge. Data are sampled or translated from Duan (2018); Zhang et al. (2018); Lin et al. (2022). An example is shown below and we use exact match as the metric:

Chinese Example:
问题：谁能描述一下氧化镁的外观？
答案：白色疏松粉末

English Translation:
Question: Who can describe the appearance of magnesium oxide?
Answer: White, loose powder.

Paraphrase Identification. In this task, a pair of sentences is passed to the model and the model decides whether they are discussing the same thing or not. We formulate the sampled data from CLUE (Xu et al., 2020) and FewCLUE (Xu et al., 2021) into a binary-choice format and leverage accuracy for assessment.

Chinese Example:
你的火气大吗
你火气大不大
这两个句子表达的意思相同吗？是或否？是

English Translation:
Do you have a bad temper?
Are you quick to anger?
Do these two sentences express the same meaning? Yes or No? Yes

Summarization. In text summarization, the model needs to abstract a long, unstructured text and generate a short summarization. Note that some of the data-to-text generation tasks (discussed later) also borrow the name “summarization”. The main difference between data-to-text generation and text summarization in our benchmark is whether the context is written in a programming language (then it is data-to-text generation) or the natural language because these two languages are distinct in nature. We sample data from CSDS (Lin et al., 2021) and use ROUGE (Lin, 2004) to evaluate the results.

Chinese Example:
莫言获奖，围绕在莫言身边的出版商也笑开颜。北京精典博维文化发展有限公司拥有莫言中国内地所有作品及衍生品出版权。莫言获得诺贝尔文学奖，不仅会使公司业绩有“可观”的提升，还将加速该公司上市的进程。

TL;DR: 诺奖花落莫言签约书商IPO提速

English Translation:
When Mo Yan won the award, the publishers around him were also happy. Beijing Jingdian Bowei Culture Media Co., Ltd. owns the publishing rights to all of Mo Yan's work and derivatives in mainland China. Mo Yan's winning of the Nobel Prize in Literature will not only bring a “considerable” increase to the company's profit but also accelerate the process of the company's listing.

TL;DR: Nobel Prize goes to Mo Yan, accelerating the IPO of his contracted publisher.

Data-to-Text Generation. Data-to-text generation is of growing interest recently as people try to use LLMs to assist their work, e.g., generating a report from an Excel table. This topic has long been explored prior to LLMs (Puduppully et al., 2019), especially under the name of summarization. We

sample data from [Shao et al. \(2019\)](#) and use BLEU for measurement. An example of generating an advertising proposal based on a structured table (in the Markdown format²) is shown here:

Chinese Example:

给定衣服的特点描述，生成相应的广告文案。

衣服特点:
版型	宽松
风格	休闲
图案	印花
图案	手绘
衣样式	衬衫
广告文案:

这款衬衫给人的第一印象就是风格独特，衬衫表面的士兵手绘图案印花精致有趣，真叫人忍不住多看几眼，浓浓的复古风也富于这款衬衫艺术感，就像巴黎卢浮宫内展示的名画一般。在款式上借鉴了睡衣版型，宽松舒适，休闲随性。

English Translation:

Given the description of the features of a clothing item, generate a corresponding advertisement copy.

Clothing features:
Fit	Loose
Style	Casual
Pattern	Textile printing
Pattern	Hand-painted
Clothing type	Shirt
Advertisement copy:

The first impression this shirt gives is its unique style. The soldier pattern hand-painted on the shirt is exquisite and interesting, making one can't help but take a few more glances. The strong retro style also gives this shirt an artistic sense, just like the famous paintings in the Louvre in Paris. In terms of style, it fits like a pajama, which is loose and comfortable, casual and natural.

Sentiment Analysis. Given a text, the model predicts the sentiment label (“Positive”) in sentiment analysis. Since it is a classification task, we use accuracy for evaluation. Our data are collected from FewCLUE ([Xu et al., 2021](#)). A sentiment analysis example is shown below:

Chinese Example:

这个产品评价是正面还是负面的?

评价: 今天刚拿到手机，打电话时发现手机听筒有吱吱吱的杂声，不满意，真怀疑是不是正品
答案: 负面

English Translation:

Is this product review Positive or Negative?

Review: Just got the phone today and found that there is a squeaking noise in the earpiece when making a call. Not satisfied, really doubt if it is genuine.
Answer: Negative

Text Classification. Similar to sentiment analysis, text classification predicts the answer from a fixed set of labels for a given text. Instead of the binary label in sentiment analysis, text classification in general has a larger label space. We collect data from FewCLUE ([Xu et al., 2021](#)) and SPR³. We adopt accuracy and an example is shown below:

²<https://en.wikipedia.org/wiki/Markdown>

³<https://github.com/DUTIR-Emotion-Group/CCL2020-Humor-Computation>

Chinese Example:

“全国青年教师教学艺术大赛举行”这段新闻的类别属于教育

English Translation:

The category of the news “The National Young Teachers Teaching Art Competition is held” is education

Opinion Mining. Opinion mining is a large topic that consists of vast tasks and has a close connection with sentiment analysis ([Zhang and Liu, 2017](#)). An exemplary task of opinion mining that we test here is opinion target extraction ([Liu et al., 2012](#)). We adopt exact match for evaluation in the context of the LLM era and show an example below:

Chinese Example:

“《恋恋笔记本》是导演尼克·卡萨维茨2004年的一部爱情类影片。”中主要围绕着什么进行描述?

恋恋笔记本

English Translation:

What is the main focus of the description in “The Notebook is a 2004 romance film directed by Nick Cassavetes.”?

The Notebook

Dialogue Generation The popularity of ChatGPT has shifted the interaction between humans and LLMs from a single-turn prompt continuation to a multi-turn conversation ([OpenAI, 2023](#)). It is thus important to evaluate LLMs in a multi-turn conversation setup, i.e., in the dialogue generation task. In this task, we use data from CrossWOZ ([Zhu et al., 2020](#)) and report BLEU and unigram F1. A conversation example is shown below:

Chinese Example:

用户: 你这看的什么视频?
系统: 是爱奇艺新出的《飞行少年》。
用户: 好看吗? 没事我也回家看看。
系统: 挺好看的, 是向祖国70周年的献礼剧。
用户: 都谁主演的啊?
系统: 严屹宽和一些年轻演员, 有闫妮, 不过是客串。

English Translation:

User: What video are you watching?
System: The Eyas, on iQIYI
User: Is it good? I am going to watch it at home if I have spare time.
System: It's pretty good. A TV series to celebrate the 70th birthday of our country.
User: Who is starring in it?
System: Yikuan Yan and other young actors. Ni Yan also appears in a cameo.

Paraphrase Generation. Paraphrasing and rewriting is a common task in NLP. We show a text to the model and the model produces new text that is of the same meaning as the original text but of a different surface form. Following [Sun and Zhou \(2012\)](#), we choose iBLEU to evaluate the performance and utilize data from PKU Paraphrase Bank ([Zhang et al., 2019](#)).

Chinese Example:

一个句子的原句为:
从梅森苍白的唇间吐出了几乎听不见的回答。

它可以被复述为:

梅森先生苍白的嘴唇间溜出一个听不清楚的回答。

English Translation:

The original sentence is:

A barely audible answer came from Mason's pale lips.

It can be paraphrased as:

Mr. Mason's pale mouth let out an unclear answer.

Translation. Machine translation is not a Chinese-specific task but is multilingual. However, the success of Chinese LLMs relies heavily on bilingual (Chinese and English) data (Team, 2023; Zeng et al., 2023) and thus most Chinese LLMs are born to be capable of translating English text to and from Chinese. Our data are collected from the past WMT competitions (Kocmi et al., 2022). We employ SacreBLEU (Post, 2018) as the evaluation metric and an English-to-Chinese translation example is shown below:

Chinese Example:

英文: House rebukes Trump on border wall, but he plans veto

中文: 众议院在边境墙问题上指责特朗普, 但他计划使用

一票否决权

English Translation:

English: House rebukes Trump on border wall, but he plans veto

Chinese: 众议院在边境墙问题上指责特朗普, 但他计划使用

一票否决权

C Manual Data Collection

We collect data on an extensive scale, comprising 33.98% of our entire benchmark. Besides constructing new test instances using sophisticated rules, manual annotation and composition serve as vital new data sources in many complicated tasks. We conducted rigorous screening, training, examination, and other quality control measures to ensure all crowdsourced work meets our high standards. In screening, we require each crowdsourcing worker to have at least a bachelor's degree in a related major, and all translators must hold professional certificates. Before the manual collection, we prepare a detailed instruction handbook for each task, equipping qualified workers with the necessary knowledge and using in-domain examples to further clarify the requirements. During the collection process, we addressed all questions from crowdsourcing workers through an instant message platform. Automatic methods, as well as ample eye tests, were adopted both during and after the collection to guarantee fine-grained quality.

D Metrics

D.1 Accuracy

For each task in our benchmark, we list and underline the corresponding evaluation metrics for each task in Appendix B.

D.2 Calibration and uncertainty

We mainly report the values of the following metrics:

- *Expected calibration error* (Kumar et al., 2019) (ECE) measures the difference between the model's predicted probability and its exact-match accuracy.
- *Selective classification accuracy* (El-Yaniv and Wiener, 2010) computes the accuracy for the C -fraction of examples where the model assigns the highest probability.

D.3 Robustness

Following HELM (Liang et al., 2022), we report the *worst-case accuracy*, which averages the poorest result among transformations of each test instance. Inspired by NL-Augmentor (Dhole et al., 2021), we implement the transformation recipe as the composition of the following perturbations:

- *Synonym perturbation* randomly substitutes Chinese words with their synonyms with a probability of 0.3.
- *Butter finger perturbation* randomly replaces Chinese words with other words that have the same toneless Pinyin with a probability of 0.05.
- *Character swapping* randomly swaps any two Chinese characters with a probability of 0.05.

We utilize LTP (Che et al., 2021) to perform word segmentation.

D.4 Fairness

We similarly adopt *worst-case accuracy* as in robustness to report fairness. We support 4 transformation recipes inspired by the perturbations from NL-Augmentor (Dhole et al., 2021):

- *Simplified to traditional conversion* converts both the prompt and references from Chinese Simplified to Chinese Traditional⁴.
- *Mandarin to Cantonese conversion* translates both the prompt and references from

⁴<https://github.com/BYVoid/OpenCC>

Mandarin to Cantonese. Here we adopt a rule-based approach⁵ which first maps phrases in Mandarin to their translations in Cantonese and then converts the resulting string from Chinese Simplified to Chinese Traditional. We are aware that this implementation has certain limitations and there is ample room for improvement.

- Chinese name perturbation randomly substitutes all occurrences of a Chinese name with another feasible Chinese name with a probability of 0.5.
- Gender term perturbation randomly flips all gender terms of a source gender to their counterparts in a target gender with a probability of 0.5.

D.5 Bias and stereotypes

We follow metrics from HELM (Liang et al., 2022) to quantify bias and stereotypes:

- **Bias:** we adopt the *demographic representation* in HELM, which measures the unevenness of gender or race terms for all social groups.
- **Stereotypes:** we adopt the *stereotypical associations* in HELM, which computes the unevenness of gender or race terms for all social groups when co-occurred with an adjective or profession term, then averages over all adjective or profession terms.

D.6 Toxicity

We employ the *toxic fraction* metric from HELM (Liang et al., 2022), which is the fraction of instances that are classified as toxic according to the Perspective API (Lees et al., 2022)⁶. We use a threshold of 0.5 to determine whether an instance is toxic or not.

D.7 Efficiency

As stated in the main text, we focus only on inference wall-clock time because limited statistics could be reliably collected from users. Concretely, we adopt *queries per second* (QPS), the amount of queries processed by a model API in a second, which is a common metric for measuring the throughput of online services.

⁵<https://justyy.com/tools/chinese-converter/>

⁶<https://perspectiveapi.com/>

D.8 Diversity

Here we adopt inter-distinct and intra-distinct (Miller et al., 2017) to quantify surface-form diversity.

- *Inter-distinct* collects n-gram statistics from all instances in the test set and computes the n-gram diversity, which is the rate of all distinct n-grams against all n-grams.
- *Intra-distinct* evaluates the n-gram diversity per instance and averages across all instances.

D.9 Privacy

We pay close attention to current research on privacy evaluation. For example, Carlini et al. (2021) utilize adversarial attacks to yield meaningful outcomes. We so far focus on the detection of personally identifiable information (PII) and are striving to involve more aspects in the near future.

To evaluate privacy from the PII perspective, we define *PII_match*, a metric similar to the toxic fraction which represents the proportion of instances that contains PII:

$$\text{PII_match} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}[\text{PII_Detect}(y_i) > 0] \quad (1)$$

where N is the number of test instances, y_i is the generated text for i -th instance and *PII_Detect* is the tool that returns the number of PII entities in y_i . We use Azure PII detection service⁷ to instantiate *PII_Detect*.

E Models

Table 1 is the summary of Chinese LLMs we evaluated in our leaderboard.

GPT (Ouyang et al., 2022; Brown et al., 2020) is a family of autoregressive LLMs from OpenAI. The most recent and powerful GPT models are ChatGPT⁸, text-davinci-003⁹, and GPT-4 (OpenAI, 2023). We test all these three models in our evaluation.

Claude (Askell et al., 2021; Bai et al., 2022b,a) is another family of autoregressive models from Anthropic, which include Claude and Claude-

⁷<https://learn.microsoft.com/en-us/azure/ai-services/language-service/personally-identifiable-information/overview>

⁸<https://openai.com/blog/chatgpt>

⁹<https://platform.openai.com/docs/models/gpt-3-5>

Model	Version	Organization	Access	#Param.	Window Size	Instruction Tuning	Architecture
ChatGPT	gpt-turbo-3.5 (2023/07/11)	OpenAI	limited	-	4096	✓	GPT
text-davinci-003	text-davinci-003 (2023/06/17)	OpenAI	limited	175B	4097	✓	GPT
GPT-4	gpt-4 (2023/07/11)	OpenAI	limited	-	8192	✓	GPT
Claude	claude-1 (2023/07/07)	Anthropic	limited	-	100000	✓	-
Claude-instant	claude-1 (2023/07/21)	Anthropic	limited	-	100000	✓	-
InternLM-104B	(2023/07/13)	Shanghai AI Lab & SenseTime	limited	104B	2000	✓	GPT
ERNIE-Bot	(2023/07/09)	Baidu Inc.	limited	-	2000	✓	-
ChatGLM-6B	v0.1.0	Tsinghua University	open	6B	2048	✓	GLM
ChatGLM2-6B	v1.0	Tsinghua University	open	6B	2048	✓	GLM
GLM-130B	-	Tsinghua University	open	130B	2048	✓	GLM
BLOOMZ-7B1-mt	-	BigScience	open	7B	2048	✓	BLOOM
BLOOM-7B1	-	BigScience	open	7B	2048	✗	BLOOM
BLOOMZ-176B-mt	-	BigScience	open	176B	2048	✓	BLOOM
BLOOM-176B	-	BigScience	open	176B	2048	✗	BLOOM
LLaMA-7B	-	Meta	open	7B	2048	✗	LLaMA
LLaMA-65B	-	Meta	open	65B	2048	✗	LLaMA
Vicuna-7B	v1.1	LMSYS	open	7B	2048	✓	LLaMA
Vicuna-13B	v1.1	LMSYS	open	13B	2048	✓	LLaMA
BELLE	BELLE-7B-2M	Beike Inc.	open	7B	2048	✓	BLOOM
Chinese-Vicuna-7B	Chinese-Vicuna-lora-13b-belle-and-guanaco	Cui et al.	open	7B	2048	✓	LLaMA
Chinese-Alpaca-7B	Chinese-Alpaca-7B	Fan et al.	open	7B	2048	✓	LLaMA
MOSS-16B	moss-moon-003-sft	Fudan University	open	16B	2048	✓	CodeGen
Baichuan-7B	-	Baichuan Inc.	open	7B	4096	✗	LLaMA

Table 1: 23 Chinese LLMs evaluated in this work. For limited-accessed models, we mark the timestamp where we finalized the evaluation in the format of (YYYY/MM/DD). For models with the same public name but have different versions, we also provide the version we used to conduct the experiment. Note that the unit of window size of ERNIE-Bot is characters instead of tokens.

instant¹⁰. Both models are evaluated in our experiments.

InternLM (Team, 2023) is a GPT-like Chinese LLM trained by Shanghai AI Laboratory and SenseTime. It has a limited-accessed 104B and an open-source 7B version. We evaluate the 104B version in our experiments.

ERNIE-Bot¹¹ is a Chinese LLM launched by Baidu Inc. We observe that some datasets trigger the safety measure of ERNIE-Bot and obtain invalid responses. This fact leads to a poor result in our evaluation.

GLM (Du et al., 2022) is a Chinese LLM family from Tsinghua University trained with autoregressive blank infilling. We only assess their open-source GLM-130B (Zeng et al., 2023), ChatGLM-6B¹² and ChatGLM2-6B¹³.

BLOOM (Scao et al., 2022) is a family of open-source multilingual LLMs from BigScience. It is not fine-tuned and has an instruction-following version BLOOMZ (Muennighoff et al., 2023). In our experiment, we test the pretraining-only BLOOM-7B1 and BLOOM-176B from BLOOM,

and the instruction-following BLOOMZ-7B1-mt and BLOOMZ-176B-mt from BLOOMZ.

LLaMA (Touvron et al., 2023) is a more recently released open-source autoregressive English LLM family from Meta and is pretrained only. We experiment with LLaMA-7B (the smallest one) and LLaMA-65B (the largest one).

Vicuna (Chiang et al., 2023) is a series of instruction-following models built on top of LLaMA (Touvron et al., 2023). It comes from LMSYS. We evaluate both Vicuna-7B and Vicuna-13B.

BELLE (Ji et al., 2023) refers to a series of instruction-following models from Beike Inc., fine-tuned on various pretrained models like BLOOM and LLaMA. We assess their BLOOMZ-based 7B variant.

Chinese-Vicuna (Chenghao Fan and Tian, 2023) is a Chinese instruction-following model fine-tuned from LLaMA and has 7B and 13B two variants. We experiment with the 7B version.

Chinese-Alpaca (Cui et al., 2023) is a family of LLaMA-based Chinese LLMs. They extend the original LLaMA’s vocabulary for better Chinese modeling and open-source fine-tuned Chinese LLMs with various model scales. We test their early 7B instruction-following model.

MOSS (Sun et al., 2023b) is pretrained and fine-

¹⁰<https://www.anthropic.com/index/introducing-claude>

¹¹<https://yiyao.baidu.com/welcome>

¹²<https://github.com/THUDM/ChatGLM-6B>

¹³<https://github.com/thudm/chatglm2-6b>

tuned from CodeGen (Nijkamp et al., 2023) by Fudan University. It includes the pretrained model, an instruction-following model, and a tool-augmented instruction-following model (Schick et al., 2023). We evaluate the instruction-following version¹⁴ in our experiment.

Baichuan¹⁵ is a pretrained Chinese LLM from Baichuan Inc., with the same architecture as LLaMA. We test the early 7B version and a new 13B version¹⁶ is released by the time of writing.

F Prompting

F.1 Settings

The prompt setting remains the same as the common practice (Brown et al., 2020; Liang et al., 2022), where we randomly choose 5 in-context training examples (a.k.a., demonstrations) for few-shot prompting. To mimic true few-shot setting (Perez et al., 2021), these 5 in-context training examples will be fixed for all test instances. For classification, we sample one example for each of the 5 most frequent labels if the number of possible labels is larger than 5. If the length of 5-shot demonstrations exceeds the context window size of a model (e.g., reading comprehension), we reduce the number of in-context examples.

F.2 Format

Completion-style few-shot prompting. Given the description of the task, sampled demonstrations, and a test instance, we use the below template to construct the few-shot prompt for prompting conventional LLMs (a string):

```
{instruction} \n{n}{demonstration1} \n{n}...
{demonstrations5} \n{n}... {prompt} {prediction}
```

where `{instruction}` is the task description, `{demonstration1}` is the concatenation of the prompt and reference of the first in-context example, `\n` is the line break and `{prompt}` is the prompt of the test instance. The model will continue the prompt and complete the generation in `{prediction}`. We denote this type of prompt template as Completion. A mathematical calculation example is shown below (we use an English prompt template for demonstration only and all prompt templates in our benchmark are Chinese):

¹⁴<https://huggingface.co/fnlp/moss-moon-003-sft>

¹⁵<https://github.com/baichuan-inc/baichuan-7B>

¹⁶<https://github.com/baichuan-inc/Baichuan-13B>

Calculate the following formula.

758 + 445 -> 1203

758 + 445 -> 1203

140 + 361 -> 501

Chatbot-style few-shot prompting. The popularity of ChatGPT has led to an outbreak of LLM-based chatbot (Team, 2023; Chenghao Fan and Tian, 2023). Existing work (Huang et al., 2023) shows that the best few-shot prompting strategy for chatbots is different from the one for conventional LLMs. Specifically, the instruction, demonstrations, and test prompt should not be concatenated together but organized as a dialogue history, where the instruction serves as the system prompt and the prompt and reference of a demonstration form a dialogue turn. The previous example will be reorganized as below before feeding into the chatbot:

```
System:
Calculate the following formula.
User:
758 + 445 ->
Assistant:
1203
User:
163 + 140 ->
Assistant:
303
User:
140 + 361 ->
Assistant:
501
```

where System: is the field to set up the chatbot and we will put the instruction here. User: and Assistant: stand for the prompt and reference respectively. We denote this type of prompt template as Chatbot.

Multi-choice problem format. As discussed in Liang et al. (2022), there are two strategies when constructing prompts for multi-choice problems:

- Separate (Brown et al., 2020) scores each choice by concatenating it with the prompt and takes the one with the highest probability as the prediction.
- Joint (Hendrycks et al., 2021) puts all choices into the prompt and lets LLMs generate the choice index (e.g., “{question} A. {choice₁} B. {choice₂} Answer:”).

In general, Separate approach better estimates the model performance as the output space is restricted, while Joint approach is more economic since the model only needs to infer once to get the final an-

swer. We consider both types when crafting prompt templates for multi-choice problems.

Chain-of-Thought. Chain-of-Thought (Wei et al., 2022c; Kojima et al., 2022) (CoT) is a crucial technique to elicit the reasoning ability of LLMs. We also support CoT in CLEVA and provide the corresponding prompt templates for the mathematical reasoning task. An example of CoT is shown below, where highlighted text is predicted by the model and text in red is the intermediate reasoning process and text in green is the final answer.

Question: A community has 8 buildings, each with 102 residents. On average, each household pays 9 yuan per month for water. How much does this community pay for water in total per month?
 Answer: Let's think step by step. First, each household pays 9 yuan per month for water, and each building has 102 residents. Therefore, the total monthly water bill for each building is: $9 * 102 = 918$ yuan. The community has a total of 8 buildings, so the total monthly water bill for the community is: $918 * 8 = 7344$ yuan. Therefore, the answer is 7344.

F.3 System Design

Previous work (Bach et al., 2022; Ding et al., 2022) has presented various approaches to design a prompting language that enables flexible prompt construction. However, their methods have shown limitations in handling the vast variety of tasks for large-scale evaluation: Their prompting languages can only manipulate strings, while many tasks are organized in a structured data format, e.g., dialogue generation and data-to-text generation.

We therefore devise a slightly complicated prompting language that accommodates customized prompt construction of structured data while preserving usability. We start the description with an instance in JSON format from the dialogue generation task:

```
{
  "history": [
    {
      "utterance": "Who is the US president?",
      "role": "usr"
    },
    {
      "utterance": "Joe Biden.",
      "role": "sys"
    },
    {
      "utterance": "Then who is his wife?",
      "role": "usr"
    }
  ],
  "role": "sys",
  "label": [
    "Jill Biden."
  ]
}
```

and a prompt template example written as a JSON dictionary (`\n` is the line breaking):

```
{
  "verbalizer": {
    "role": {
      "sys": "Assistant",
      "usr": "User"
    }
  },
  "history": {
    "item_separator": "\n",
    "item_template": "{role}: {utterance}",
    "item_index": null
  },
  "input": "{history}\n{role}:",
  "label": "{label}"
}
```

The general pipeline of our prompt construction is as follows (we mark the field from the instance in green and the one from the prompt template in blue):

1. We first map values of all fields in a test instance according to user-defined mappings in `verbalizer` (Gao et al., 2021b). In our example, all `usr` and `sys` will be replaced with `User` and `Assistant` respectively.
2. We then stringify each field in the test instance. We organize all structured data fields in the format of a list of dictionaries (`history` in our example) and apply the following rules to process them:
 - (a) For each entry (a dictionary), we independently stringify it by composing all its fields via a template defined in the Python f-String syntax¹⁷. For instance, an utterance in the dialogue history `Who is the US president?` from the speaker `User` will be formatted into `User: Who is the US president?` according to `item_template` in a prompt template field that shares the same name as `history`.
 - (b) We then index all stringified entries (by prepending an index like `A.` to each entry) if needed and concatenate them with a user-defined separator `item_separator` to stringify the whole data structure. In our case, we do not apply any indexing (an empty option in `item_index`) and directly assemble the final string of `history` with `\n`:


```
User: Who is the US president?
Assistant: Joe Biden.
User: Then who is his wife?
```
3. We finally construct the prompt and references

¹⁷<https://peps.python.org/pep-0498/>

from all stringified fields. According to `input`, the resulting prompt in our example will be:

User: Who is the US president?

Assistant: Joe Biden.

User: Then who is his wife?

Assistant:

For the references, we directly apply `label` in the prompt template to every entry in `label`, resulting in “*Jill Biden*” here.

Though not shown in the example above, another crucial part is to attach specific post-processing steps tailored to a prompt template. For example, if we index the choices in an instance from a multiple-choice task by capital letters like “A. ”, we should capitalize the initial output letter for a more accurate evaluation. In our system, we achieve this by passing a list of predefined options to the subfield `postprocess` in the prompt template field `meta`, which executes the script of each post-processing option on the output consecutively.

G Results

In this section, we provide the complete evaluation results and breakdown analysis of our benchmark.

G.1 Meta Analysis

To validate the uniqueness and reasonability of diversity and privacy, we examine the correlation between accuracy and these two newly introduced metrics. Figure 4 shows the scatter plot. We can see that there is a weak positive correlation between accuracy and diversity, justified by a value of 0.30 in Pearson’s r (P-value is 9.9×10^{-9}). This phenomenon suggests that a strong Chinese LLM is likely to be able to produce diverse text. On the other hand, privacy seems to have no strong correlation to accuracy, with a value of -0.10 in Pearson’s r (P-value is 0.05). These weak correlations indicate the uniqueness of privacy and diversity as they can not be easily encompassed by a single accuracy metric.

G.2 Ability Evaluation

In this section, we focus on the analysis of ability evaluation. Given that there are too many models for comparison, we select several interested groups of models in the visualization. Figure 5 compares 4 groups of models, each group consisting of two categories with three top-performing models. We have the following observations:

- Although outstanding Chinese model like InternLM-104B is comparable and even outperforms the best English models in some tasks, most high-ranking models in our Chinese benchmark are English models.
- The gap between limited-accessed and open-source models (Liang et al., 2022) is also witnessed in Chinese LLMs. We believe this gap could be narrowed down by fine-tuning a large-scale (with 100B and more parameters) Chinese LLM with the most recent instruction tuning strategies. Figure 2 shows that the well-performing open-source models are small models fine-tuned by the most recent and advanced techniques like Self-Instruct (Wang et al., 2023). These models mainly lag behind the limited-accessed model in many reasoning and knowledge-intensive tasks as shown in Figure 5, which could be addressed by scaling up the model size (Liang et al., 2022; Fu and Khot, 2022).
- Aligned with Liang et al. (2022); Fu and Khot (2022), large LLMs show clear advantages over the small ones in many reasoning and knowledge-intensive tasks.
- Instructing tuning is indeed a crucial technique to unleash the full potential of LLMs (Fu and Khot, 2022). Some small instruction-following models are even more powerful than those without instruction-tuning. For example, InternLM-104B is much better than BLOOM-176B. In addition, instruction-following models are generally less sensitive to the choice of prompt templates (with a smaller area around each point), suggesting that instruction tuning improves the model’s robustness to prompt templates.

Moreover, we also observe some interesting phenomena in Figure 5: *Inverse scaling* (McKenzie et al., 2023) seems to appear in our instruction following task, where the larger GPT-4, InternLM-104B, and LLaMA-65B is worse than MOSS-16B. According to our marking of tasks with a great standard deviation in Figure 5, they all are the *emergent ability* (Wei et al., 2022b) candidate in the Chinese world, e.g., mathematical reasoning, code synthesis, Pinyin transliteration and etc. We are aware that the analysis here is not a rigorous study that verified the existence of inverse scaling and emergent ability in certain Chinese tasks and we leave it for future work. In the end, we find some tasks

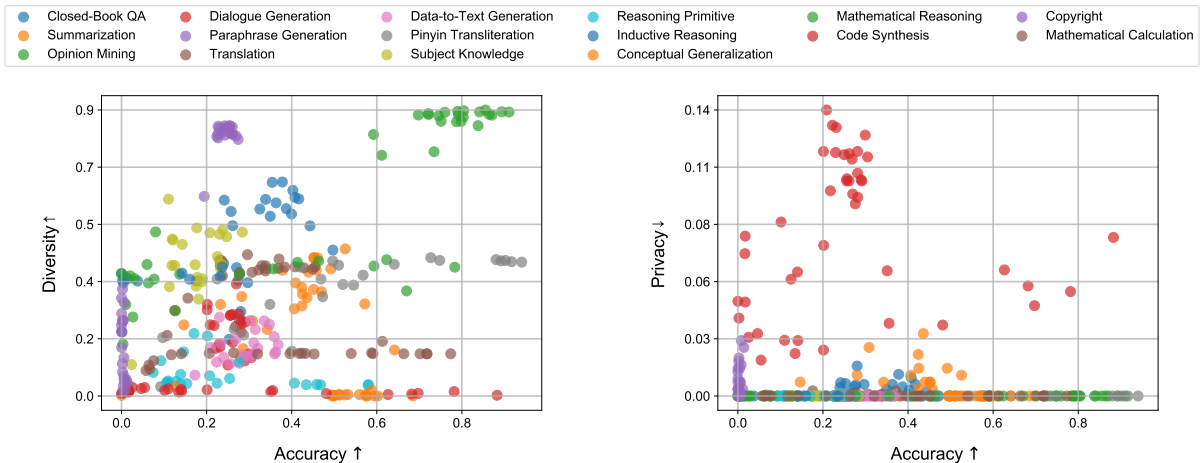


Figure 4: Correlation between diversity or privacy and accuracy on all tasks in a scatter plot format. Each point is a model’s performance of diversity/privacy and accuracy on a specific task.

(e.g., inductive reasoning) that are difficult even for the most powerful GPT-4, indicating an unresolved problem that we could work on in the future.

We analyze the knowledge of different Chinese LLMs in Figure 6 by utilizing questions from 14 subjects. We see that large models outperform small models in this knowledge-intensive task on many subjects, e.g., GPT-4, Claude, and InternLM-104B are much better than MOSS-16B and Vicuna-13B. Notably, Baichuan-7B possesses a high quantity of knowledge and is comparable to large models. This fact explains why it performs so well in knowledge-intensive tasks like classical Chinese understanding, commonsense reasoning and etc., as shown in Figure 5.

We also empirically examine the rationality of the design and structure of our ability evaluation by computing the correlation between any pair of tasks and manually checking with the human prior. As shown in Figure 7, most pairs of tasks that both not belonging to the same aspect (e.g., knowledge) do not share a statistically significant correlation, e.g., conceptual generalization and cultural knowledge. Some statistically significant correlations are well-match with our expectations (not exhausted):

- A good performance on coreference resolution and cultural knowledge helps to identify toxic and biased content (Pearson’s $r > 0.6$);
- Commonsense reasoning ability is also required for toxicity and bias as this harmful content could be implicit (Pearson’s $r > 0.5$);
- There is a strong positive correlation among almost all reasoning tasks (Pearson’s $r > 0.5$);
- More subject knowledge improves conceptual generalization and commonsense reasoning

(Pearson’s $r \geq 0.6$);

- More cultural knowledge yields a better result in classical Chinese understanding (Pearson’s $r = 0.85$);
- Mathematical calculation is almost mandatory for mathematical reasoning (Pearson’s $r \approx 0.8$);

These observations in general justify the rationality of our taxonomy.

In addition, we observe some interesting phenomena. Reasoning primitive has a strong positive correlation with Pinyin transliteration (Pearson’s $r \approx 0.9$). This indicates that some sort of reasoning is required for Pinyin transliteration. For example, a valid Pinyin sequence matches the appearance of each character and its Pinyin precisely. The model needs to follow this rule to predict correctly. However, there are also some counter-intuition observations that could not be explained easily: A strong positive correlation (Pearson’s $r = 0.76$) between reasoning primitive and classical Chinese understanding reveals the distinct mechanism beneath LLMs and the human brain.

G.3 Application Assessment

Figure 8 compares the performance of models in application assessment tasks. The conclusions are in line with those in Figure 2: Most high-ranked models are English models and are limited-accessed. Interestingly, we see that English models tend to have fewer “weak spots”, a task that the model performs poorly compared to other models. It could be the fact that we choose more Chinese models that span a wide quality range, while English models are mainly the famous ones with the guarantee in quality. We observe that English open-source

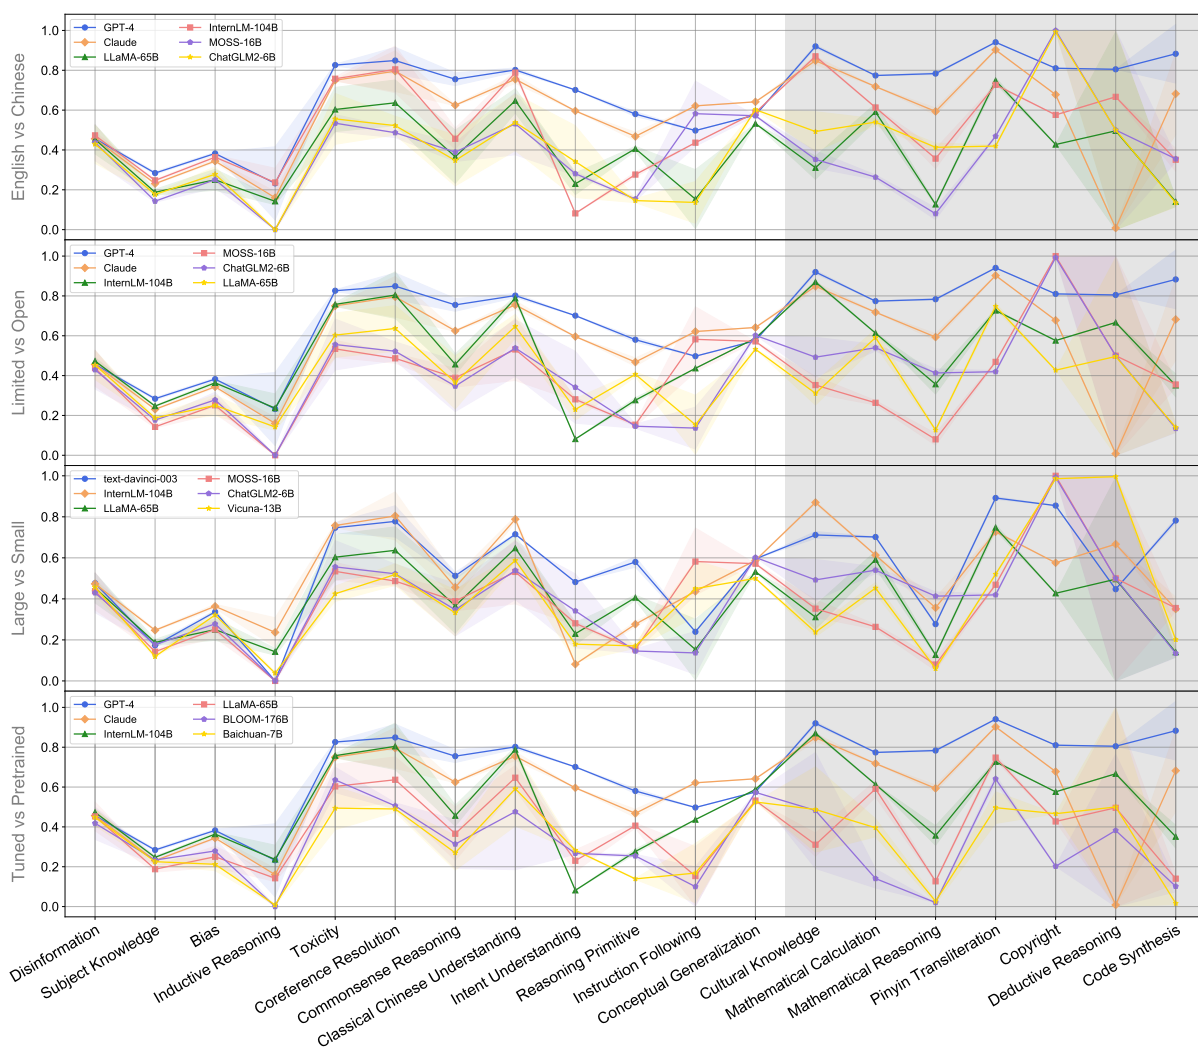


Figure 5: Comparison between three best-performing models from two categories on all ability evaluation tasks. Models in the left legend column belong to the first category and those in the right belong to the second category. For example, GPT-4, Claude and LLaMA-65B are **English** models. There are 8 categories: **Chinese** are Chinese-focused models (with tailored strategies to improve Chinese modeling), **English** are English-focused models, **Open** are open-source models, **Limited** are limited-accessed models, **Large** are models with more than 50B parameters (We choose text-davinci-003 rather than GPT-4 and ChatGPT as its size has been reported), **Small** are models with fewer than 50B parameters, **Tuned** are instruction-following models and **Pretrained** are pretrained models (without instruction tuning). Each point represents the mean performance of the model on a specific task and the area around each point is of the size of \pm standard deviation. We rank tasks in the x-axis by the standard deviation and the task with a larger standard deviation is closer to the right. We mark tasks with a standard deviation larger than 0.1 by gray shadow. These tasks imply the plausible emergent abilities of Chinese LLMs. Note that we normalize the score in the copyright task across models and then subtract it from 1 to convert it to a metric whose value is larger implying a better result.

models do not work well on translation and text classification.

We show the distribution of different metrics at different tasks in Figure 9.

- **Accuracy.** Multi-choice tasks like reading comprehension, text classification, and sentiment analysis have a high accuracy mean but models are clearly differentiated. On the other hand, generation tasks have a low median and most models are close to each other in general.

- **Efficiency.** There is a large difference in efficiency among models. This is because there exist many unfair comparisons. For example, limited-accessed models do not provide details on how many resources they invest when serving each query.
- **Robustness & Fairness.** For robustness and fairness, they have a similar trend as accuracy but with a relatively lower value, probably because they share the same base metric on

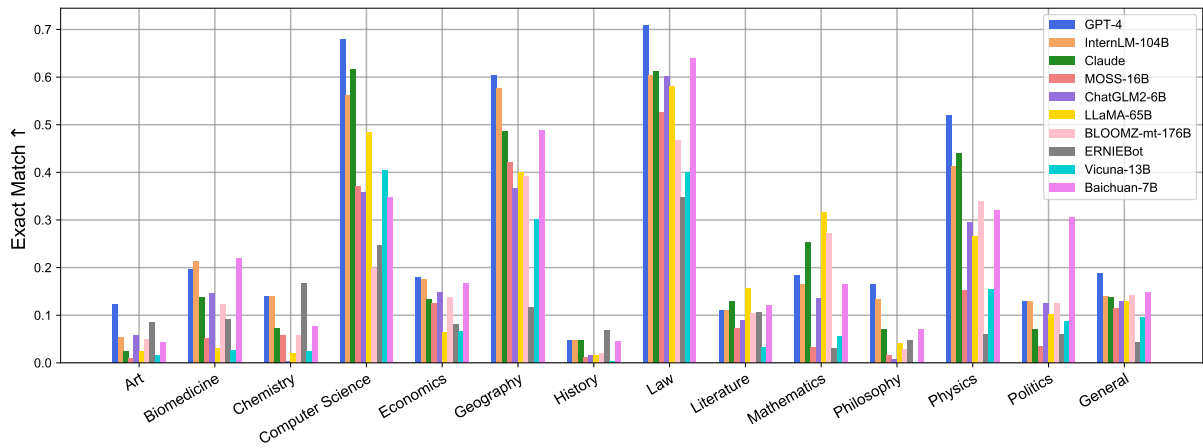


Figure 6: The performance of models on 14 subjects in the subject knowledge task. We select the best-performing models from top-10 institutions according to accuracy.

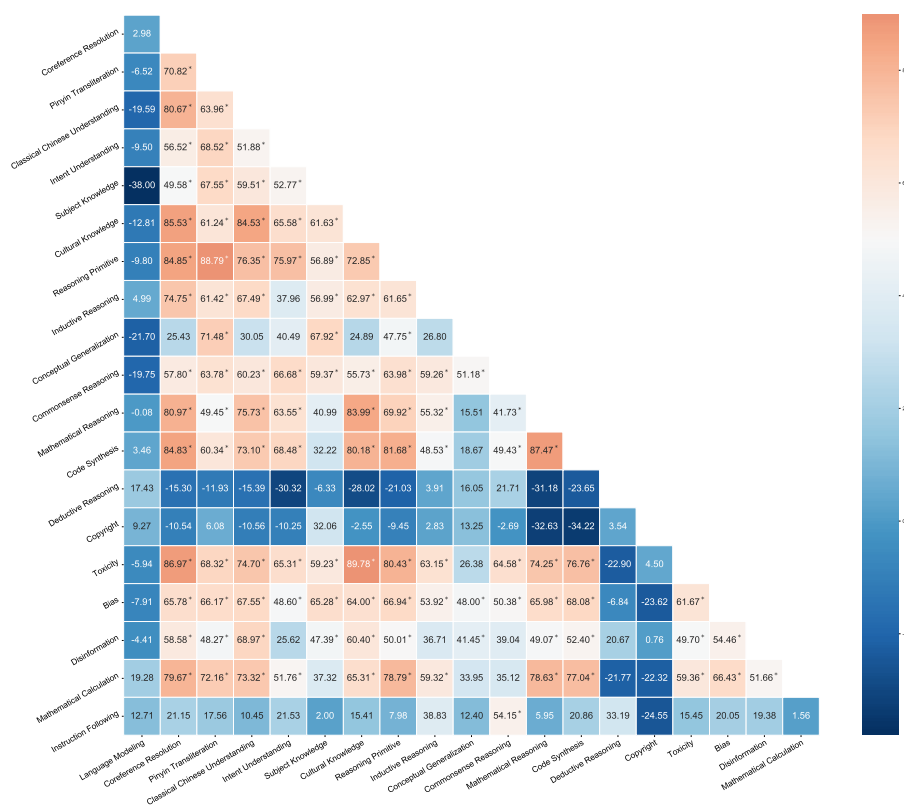


Figure 7: Correlation between different tasks in ability evaluation. Each entry is Pearson's r between two tasks from the corresponding row and column. * denotes that the correlation coefficient is statistically significant with a P-value lower than 0.05.

augmented data. We observe that some tasks are more sensitive to noise, e.g., sentiment analysis and opinion mining.

- **Calibration.** We compare the values on $ECE-10$ (Kumar et al., 2019). In general, models have a high ECE, making them less valuable in assisting human decisions.
- **Diversity.** We focus on the inter-distinct metric. We see that most models have a similar level of diversity in most tasks. Their

differences become obvious only in some knowledge-intensive tasks like closed-book QA and tasks that have multiple feasible correct answers, e.g., summarization, dialogue generation, and data-to-text generation.

- **Bias.** We choose to compare gender bias. We observe that models in data-to-text generation, summarization, and dialogue generation exhibit a strong tendency to produce biased content. These results could be partially attributed

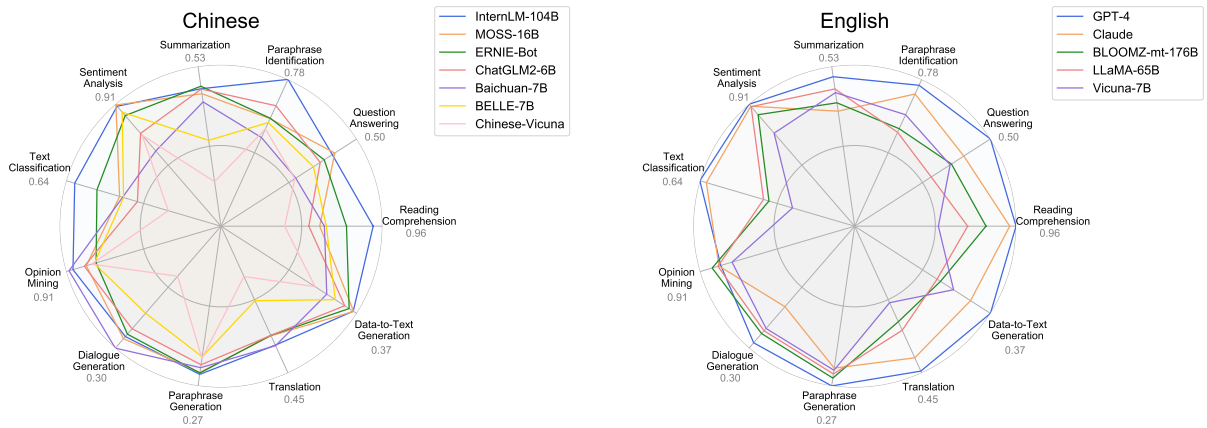


Figure 8: Comparison among models from different groups in tasks of application assessment. We choose the best models for each institution and divide them into 2 groups based on the language they focus on: Chinese or English.

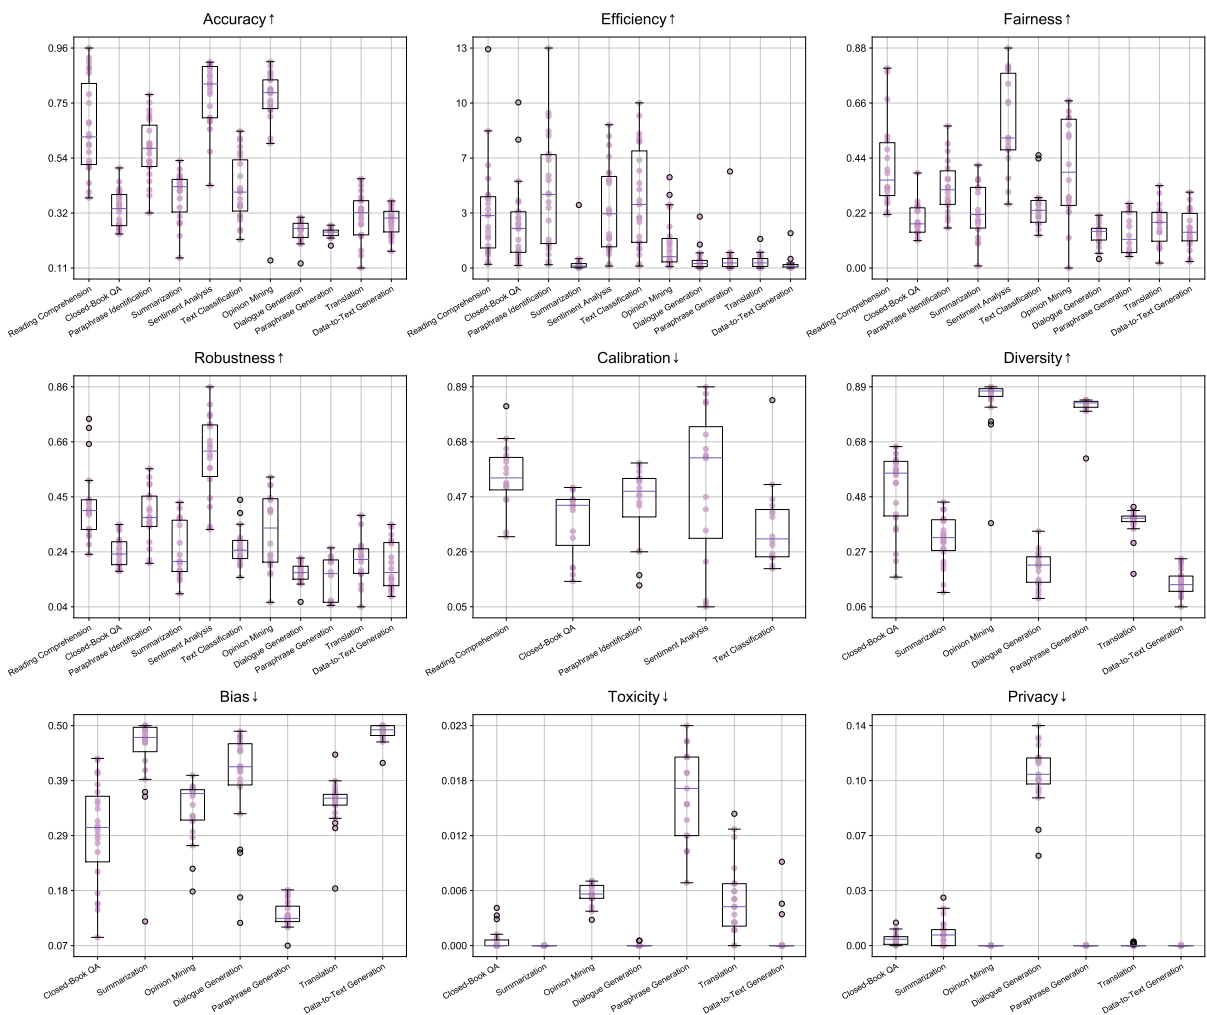


Figure 9: The performance distributions of application assessment tasks under different metrics. Some tasks are missing in some metrics because they are unavailable, e.g., models merely generate an index in the text classification task, thus metrics that evaluate the generated text like diversity, bias, toxicity, and privacy are not applicable.

to the bias in the dataset domain.

- **Privacy & Toxicity.** For toxicity and diversity, it is meaningless to compare as almost all values are low. The only exception is dialogue generation in privacy. This is because

our data contains inquiries for detailed contact information. The implication of a high value of privacy metric in dialogue generation is mixed: It means that the model understands users' requests and attempts to address them

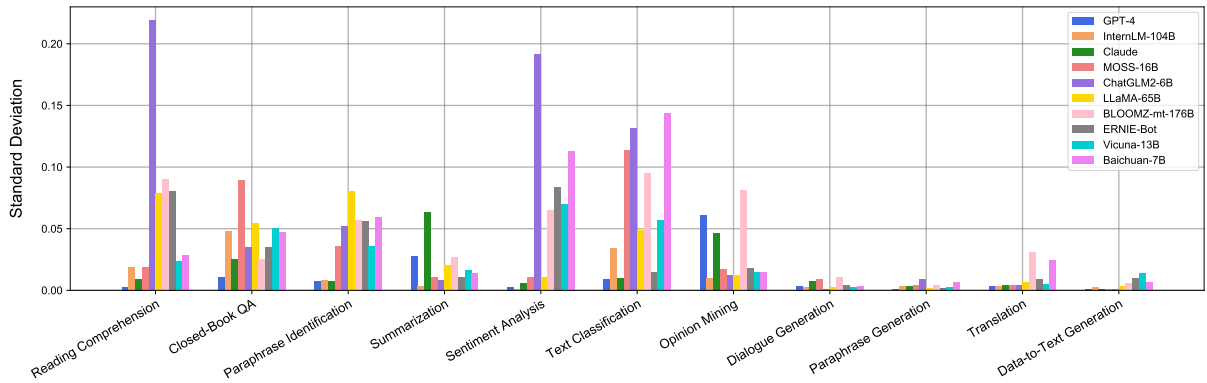


Figure 10: The accuracy standard deviation of different models in different prompt templates from different application assessment tasks. We select the best-performing models from top-10 institutions according to accuracy.

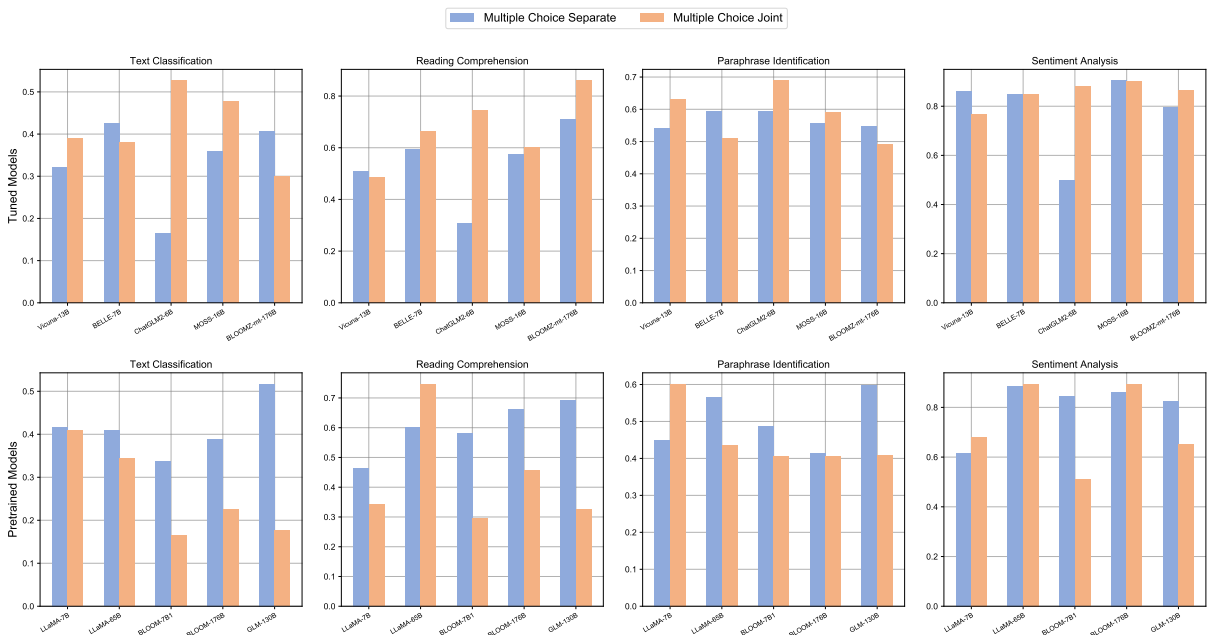


Figure 11: The accuracy of different models in multi-choice tasks with Separate and Joint style prompt templates.

with concrete information. It also implies that the model has a higher risk of hallucination that leads to potential harm.

At the end of this section, we study the prompt template sensitivity, one of the key features in CLEVA. Figure 10 presents the accuracy standard deviation of different prompt templates of different models. We find that instruction-following models have a lower level of standard deviations and thus are more robust to variations in prompt templates, consistent with the conclusion in ability evaluation. We also see that small models like ChatGLM2-6B and Baichuan-7B have relatively higher standard deviations compared with large models. Interestingly, strong models like GPT-4 have a relatively large variance in some tasks like summarization. A possible reason is that models are sensitive to

some keywords in the instruction, e.g., almost all models perform better in prompt templates that contain “zhāi yào” (means “summarize” in English) in the summarization task. We also find that limited-accessed models sometimes refuse to answer. For example, ERNIE-Bot refuses to answer about 4 tasks, resulting in a lower ranking in Figure 2.

G.4 Prompting Analysis

As discussed in Appendix F, there are two feasible prompt template types for multi-choice tasks: Separate that feeds each choice with the prompt separately and Joint that concatenates all choices and feeds once. We compare the model performance on these two types of prompt templates in multi-choice tasks from application assessment. Figure 11 shows that despite the cost of Separate,

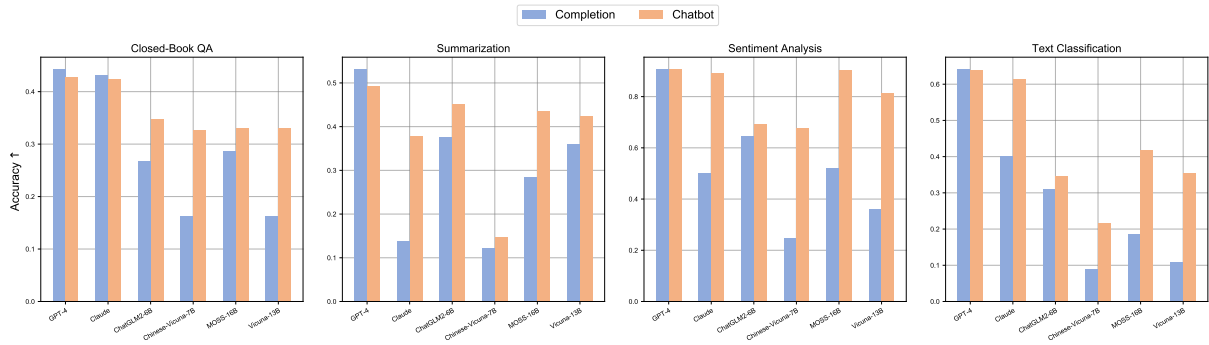


Figure 12: The accuracy of different chatbots with Completion and Chatbot style few-shot prompt templates.

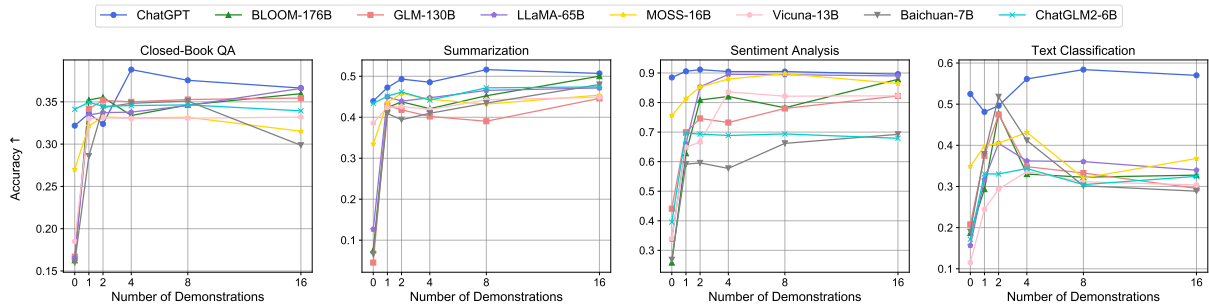


Figure 13: The accuracy of different models with various numbers of few-shot demonstrations.

it is more friendly to models without instruction tuning as they perform much better than Joint. This is because Separate restricts the model to output choices only, reducing the errors caused by unconstrained generation. However, for instruction-following models, Joint yield more advantages (e.g., ChatGLM2-6B in text classification, reading comprehension, and sentiment analysis) as some Separate prompt templates may not include all possible choices in the prompt. In this case, models are likely to produce other viable answers that could not be parsed by automatic metrics (e.g., paraphrasing the correct answer).

Similarly, we discuss the impact of Completion and Chatbot style few-shot prompting strategies, where the former concatenates everything into a string and the latter organizes demonstrations into a structured dialogue history. Figure 12 illustrates the impact of these two styles of few-shot prompting strategies in various chatbots. We see that almost all chatbots perform better with Chatbot than with Completion, demonstrating the effectiveness of this tailored strategy. We also notice that GPT-4 and ChatGPT from OpenAI are not sensitive to the few-shot prompting styles. After taking a closer look at the generation results, we find that most chatbots do not follow the format described in the instruction and illustrated in the in-context examples to customize their answers, resulting in invalid postprocessing of automatic metrics. For instance,

most prompts ask the model to output the answer only, but Claude and ChatGLM6-2B tend to provide an explanation first.

We also investigate how the performance varies as the number of in-context examples increases for Chinese LLMs. Figure 13 visualizes the overall trends of different models in different tasks. In general, most models perform better with more demonstrations and are saturated with around 4-8 training samples. In line with existing work (Liang et al., 2022), models without instruction tuning benefit more from few-shot demonstrations. We observe that many models suffer from performance degradation in the text classification task. We believe this is because our test set has a relatively large label space and including more demonstrations distracts the models.