

UltraEval: A Lightweight Platform for Flexible and Comprehensive Evaluation for LLMs

Chaoqun He¹, Renjie Luo^{2†}, Shengding Hu¹, Yuanqian Zhao^{3†}, Jie Zhou⁴
Hanghao Wu⁴, Jiajie Zhang^{5†}, Xu Han^{1*}, Zhiyuan Liu^{1*}, Maosong Sun¹

¹ Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua University, Beijing, China

² Institute of Artificial Intelligence, Beihang University, China

³ Renmin University of China ⁴ ModelBest Inc. ⁵ Northeastern University, China

{hcq21, hsd23}@mails.tsinghua.edu.cn, renjie.luo@outlook.com

Abstract

Evaluation is pivotal for refining Large Language Models (LLMs), pinpointing their capabilities, and guiding enhancements. The rapid development of LLMs calls for a lightweight and easy-to-use framework for swift evaluation deployment. However, considering various implementation details, developing a comprehensive evaluation platform is never easy. Existing platforms are often complex and poorly modularized, hindering seamless incorporation into research workflows. This paper introduces UltraEval, a user-friendly evaluation framework characterized by its lightweight nature, comprehensiveness, modularity, and efficiency. We identify and reimplement three core components of model evaluation (models, data, and metrics). The resulting composability allows for the free combination of different models, tasks, prompts, benchmarks, and metrics within a unified evaluation workflow. Additionally, UltraEval supports diverse models owing to a unified HTTP service and provides sufficient inference acceleration. UltraEval is now available for researchers publicly ¹.

1 Introduction

LLMs have been deployed in diverse domains, such as finance(Zhang and Yang, 2023), education(Kasneji et al., 2023), and law(Blair-Stanek et al., 2023), demonstrating their versatility and efficacy(Zhao et al., 2023). This advancement is significantly bridging the gap between the realization of the current state and Artificial General Intelligence (AGI)(Bubeck et al., 2023). Nevertheless, the expansion of model parameters and training datasets engenders increasing uncertainties and emergent capabilities, posing potential risks to humanity and

*Corresponding author: Xu Han (thu.hanxu13@gmail.com) and Zhiyuan Liu (liuzy@tsinghua.edu.cn)

†Work done during internship at ModelBest Inc.

¹The website of UltraEval is at <https://github.com/OpenBMB/UltraEval> and a demo video is at <https://youtu.be/C006BVzNAS8>.

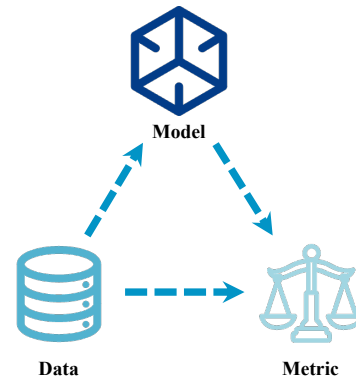


Figure 1: The three core modules of model evaluation.

challenges to stable training models (Chang et al., 2023; Bommasani et al., 2021; Wei et al., 2022a). Consequently, it is imperative to continuously and meticulously evaluate the evolving capabilities of LLMs throughout their development to ensure their responsible and beneficial applications.

Traditional benchmarks (Zellers et al., 2019; Suzgun et al., 2022; Austin et al., 2021; Clark et al., 2018) typically focus on evaluating model performance in a specific capability, making it challenging to assess the comprehensive abilities of a model. Additionally, these benchmarks generally do not include model deployment. Building pipelines from scratch for each combination of tasks and models for evaluation is highly inefficient and repetitive. Therefore, an integrated evaluation framework is crucial. Currently, some evaluation frameworks covering the entire pipeline from model deployment to model evaluation are proposed, and predominantly divided into two types: conversational websites, exemplified by platforms like Chatbot Arena ², and open-source evaluation tools, such as lm-evaluation-harness ³. The former effectively assesses the conversational abilities of a model

²<https://chat.lmsys.org/>

³<https://github.com/EleutherAI/lm-evaluation-harness>

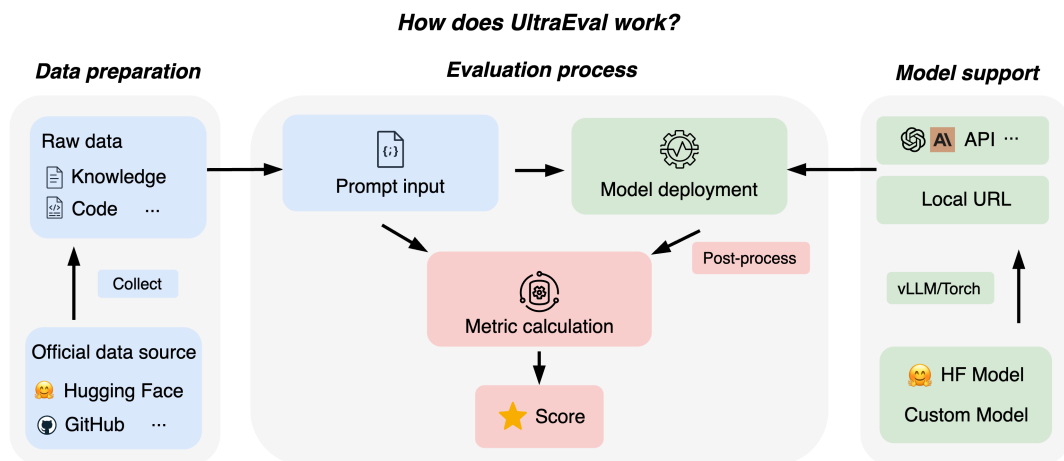


Figure 2: The overall pipeline of UltraEval, designed according to the three core modules of model evaluation.

but lacks comprehensive task coverage and transparency in the evaluation process. The latter open-source frameworks face challenges such as incomplete task coverage, complex code structure, heavy implementations, high difficulty to use, and tightly coupled functionalities. These issues hinder both convenient and comprehensive assessments.

In this paper, we identify three core components that form the evaluation process: models (or systems), task data, and metrics (i.e., evaluation methods), as illustrated in Figure 1. Rethinking the implementations in these three aspects would benefit the construction of a lightweight and easy-to-use evaluation framework, which covers mainstream tasks and complete evaluation pipeline, and can be easily expanded according to user customization. To this end, we introduce UltraEval, a lightweight and user-friendly open-source framework for LLMs evaluation. It stands out for its modular and scalable design, enabling thorough assessment of model capabilities. As illustrated in Figure 2, we segment the evaluation pipeline into three main modules: Data, Model, and Metrics, each operating independently and interacting through data exchange.

Specifically, UltraEval is characterized by the following features:

1. **Lightweight Usage Modes.** UltraEval is designed with minimal dependency requirements and features straightforward design and installation, complemented by detailed documentation. Users can initiate automated evaluations with just a few simple commands.
2. **Comprehensive Evaluation Tools.** UltraEval offers an extensive benchmarks suite, com-

prising over 50 commonly used benchmarks, and provides a customized prompt for each task. During the evaluation process, we replicated commonly used metrics and incorporated post-processing methods for more accurate metric calculation. We replicate some benchmarks from the LLaMA2 (Touvron et al., 2023), achieving consistent results, which demonstrates UltraEval’s reliability.

3. **Modular Architecture and Interfaces.** The three main modules are independent and have clear functions, enhancing the system stability of UltraEval. Moreover, its excellent scalability allows users to flexibly customize the evaluation workflow, such as by adding new models, tasks, metrics, and more.
4. **Efficient Inference Engines.** UltraEval deploys models as HTTP services, supporting the evaluation of LLMs from different sources, including the models deployed locally and the web-based API. When deployed locally, we also provide the interface to utilize vLLM⁴ (Kwon et al., 2023) and Gunicorn to enable multi-GPU acceleration.

Evaluation is currently in a phase of rapid and exploratory growth. UltraEval will be continuously updated and provide detailed tutorials to help researchers to efficiently deploy evaluation pipeline.

2 Related Work

The advancement of LLMs has led to the emergence of various evaluation frameworks, each with

⁴<https://github.com/vllm-project/vllm>

Framework	Bytes	Datasets	Acceleration	Model Types	Evaluation Method
Chatbot Arena	-	-	-	Chat	Human
AlpacaEval	3000k	-	-	Chat	GPT-4
FastChat	950k	1	-	Chat	GPT-4
HELM	3200k	10	-	All	Auto
FlagEval	760k	21	-	All	Auto
LLM harness	815k	50+	vLLM & HF Accelerate	All	Auto
OpenCompass	4000k	50+	Distributed Computing	All	Auto
InstructEval	480k	5	-	Chat	Auto
OpenAI-Evals	780k	-	Concurrent API Request	GPT	Auto
GPT-Fathom	445k	21	Concurrent API Request	GPT & LLaMA	Auto
UltraEval	315k	50+	vLLM & Gunicorn	All	Auto & GPT-4

Table 1: Comparison of Evaluation Frameworks. Bytes: the total bytes of Python and Jupyter Notebook code in each framework’s GitHub repository, Acceleration: tools or methods employed to expedite model inference, Model Types: Supported Models for Evaluation, GPT: GPT series models, LLaMA: LLaMA series models

its unique features. This section will provide a detailed overview of the current state of evaluation frameworks (also see Table 1).

Chatbot Arena (Zheng et al., 2024) offers a LLM evaluation platform where users vote on model responses, using a crowdsourced, anonymous Elo-rating system. Although innovative, its reliance on human judgment limits its suitability for fast, routine assessments. **AlpacaEval**⁵(Li et al., 2023) and **FastChat**⁶(Zheng et al., 2023a) conduct evaluation by employing GPT-4(Achiam et al., 2023) for automated judging. Yet, in evaluating complex reasoning tasks, they tend to favor verbose responses and face issues with robustness. Additionally, the scope of their evaluation capabilities is limited.

HELM⁷(Liang et al., 2022) streamlines language model evaluation but is constrained by its support solely for AutoModelForCausalLM⁸, excluding models without namespaces or stored locally. It lacks support for user models, demonstrates potential module coupling issues and absence of acceleration options. **FlagEval**’s⁹ “capability-task-indicator” framework is original but criticized for its closed-source approach and overly simplistic benchmark choices, raising data security and assessment depth concerns. Despite their innovations, both platforms fall short of the adaptability and comprehensiveness seen in more versatile frameworks like UltraEval.

LLM harness (Gao et al., 2023), used by HuggingFace’s Open LLM Leaderboard, and **OpenCompass**¹⁰(Contributors, 2023) have emerged as comprehensive solutions, offering extensive dataset support and rapid updates. These feature-rich environments, however, entail a trade-off: their complexity and dependency on specific software can complicate usage and customization. This underscores the importance of detailed documentation for those looking to adapt or extend these frameworks. Similarly, **InstructEval**¹¹ (Chia et al., 2023), leveraging the LLM harness infrastructure, caters specifically to models fine-tuned with instructions such as Alpaca and Flan-T5. Despite its targeted approach, InstructEval’s limitations in model and task coverage hint at its niche application rather than widespread utility. The adoption of such frameworks reflects the evolving landscape of model evaluation, where finding a balance between comprehensiveness and usability poses an ongoing challenge.

OpenAI Evals¹² and **GPT-Fathom**¹³(Zheng et al., 2023b). OpenAI Evals offers a straightforward, open-source framework for appraising OpenAI models, while GPT-Fathom expands upon this by analyzing the progression from GPT-3 to GPT-4 using a wider dataset array. Although it provides valuable insights into LLM development, GPT-Fathom shares OpenAI Evals’ limitations in supporting a diverse range of models.

⁵https://github.com/tatsu-lab/alpaca_eval

⁶<https://github.com/lm-sys/FastChat>

⁷<https://github.com/stanford-crfm/helm>

⁸<https://huggingface.co/docs/transformers/main>

⁹<https://flageval.baai.ac.cn>

¹⁰<https://github.com/open-compass/opencompass>

¹¹<https://github.com/declare-lab/instruct-eval>

¹²<https://github.com/openai/evals>

¹³<https://github.com/GPT-Fathom/GPT-Fathom>

First Level	Second Level	Dataset List
Knowledge	Disciplinary knowledge	MMLU, CMMLU, C-Eval, AGI-Eval JEC-QA, MEDMCQA, MEDQA-MCMLE MEDQA-USMLE, GAOKAO-Bench
	World knowledge	NQ-open, TriviaQA, TruthfulQA
Math	Math	GSM8K, MATH
Code	Code	HumanEval, MBPP
Reason	Logical reasoning	BBH
	Implicative relation	AX-B, AX-G, CB, CMNLI, OCNLI, RTE HellaSwag, OpenBookQA, ARC-c, ARC-e
	Commonsense reasoning	CommonsenseQA, COPA, PIQA, SIQA WinoGrande, Story Cloze, StrategyQA, TheoremQA
Language	Reading comprehension	BoolQ, C3, ChiD, DRCD, LAMBADA, MultiRC, QuAC RACE, RECORD, SQuAD, TyDiQA, SummEdits
	Translation	FLORES, WMT20-en-zh, WMT20-en-zh
	Semantic similarity	AFQMC, BUSTM
	Word sense disambiguation	CLUEWSC, WIC, Winogender, WSC
	Sentiment analysis	EPRSTMT
	News classification	TNEWS

Table 2: We compile a collection of 59 widely-used benchmarks and categorized them according to scenarios.

```

1 def transform_entry(row):
2     question, *choices, answer = row
3     target_scores = {
4         choice: int((ord(answer) -
5             ord("A")) == idx)
6         for idx, choice in enumerate
7             (choices)
8     }
9     return {
10         "passage": "",
11         "question": question,
12         "target_scores":
13         target_scores,
14         "answer": "",
15     }

```

Figure 3: The data formatting template for MMLU.

3 UltraEval

As illustrated in Figure 1, evaluation is a comprehensive process that integrates models, data, and metrics. With this in mind, the design philosophy considers both the independence and interconnectivity of each module. As shown in Figure 2, UltraEval encompasses the entire evaluation lifecycle (Chang et al., 2023) and segments the evaluation workflow into three main modules. In this section, we delve into the design and implementation of each component within UltraEval in detail.

3.1 Data Preparation

Data preparation involves transforming raw data into the final input format for the model, encompassing data preprocessing and prompt templates.

Data Preprocessing. We collect commonly

used benchmarks for evaluating LLMs, such as MMLU (Hendrycks et al., 2020), GSM8K (Cobbe et al., 2021), and Hellaswag (Zellers et al., 2019), covering multiple dimensions of capabilities. Currently, we have 59 benchmarks, listed in Table 2, and we plan to continually expand our collection of benchmarks.

To ensure the accuracy of the data, we source it from reputable platforms like Hugging Face¹⁴ and GitHub, rather than relying on data modified by third parties. Given the varying data formats across benchmarks, we design a set of templates to standardize these diverse formats into JSON, serving as the starting point for evaluations. As shown in Figure 3, different data items are categorized under unified attributes.

Prompt Templates. Prompts are used to guide models to generate specific outputs, and UltraEval supports prompt engineering (White et al., 2023), including few-shot and Chain of Thought (CoT) (Wei et al., 2022b), to enhance the model’s accuracy. The sensitivity of LLMs to prompts (Zhu et al., 2023) and the variability of prompts across different tasks often make it challenging for researchers to replicate results from papers, hindering research progress. UltraEval addresses this issue by providing customized, stable prompt templates for each task to facilitate result alignment. Figure 4 showcases an example of a prompt template for MMLU, demonstrating the rigorous process for forming the final prompt input.

¹⁴<https://huggingface.co/datasets>

```

1 question = f"Question:\n{question}]\n"
2 instruction = f"Requirement:\nChoose and respond with the letter of the
   correct answer, including the parentheses.\n"
3 options = "Options:\n"
4 for idx, item in enumerate(question_options):
5     options += f"({chr(65 + idx)}) {item}\n"
6 answer_prompt = f"Answer:\n"
7 final_input = question + instruction + options + answer_prompt

```

Figure 4: An example of a prompt template for an MMLU task.

3.2 Model Deployment

UltraEval employs a unique design approach, deploying models as HTTP online services and leveraging vLLM and gunicorn technologies to enable multi-GPU acceleration.

Http Service. In traditional evaluations, model deployment is closely integrated with task assessment, requiring models to be redeployed for each new task, which can lead to unnecessary consumption of time and hardware resources. To address this, we deploy models as HTTP services with Flask, facilitating modularization and efficient resource use. This approach has several advantages:

1. **Independence.** We provide a unified interface through which models receive task data and hyperparameters, returning results upon completing inference. This setup, which allows for adjustments via hyperparameters, ensures model independence.
2. **Comprehensiveness.** In UltraEval, we enable direct model loading from the Huggingface Transformers library. Given the independent deployment, UltraEval theoretically supports all models, including those from personal experimentation under different frameworks, greatly enhancing research and development flexibility. We utilize vLLM¹⁵ as the foundational acceleration framework, granting loading priority to models it supports.
3. **Scalability.** Thanks to its excellent scalability, users can easily extend models from language-based applications to multimodal models.

Multi-GPU Acceleration. We use the Gunicorn web server with Flask to deploy models via web endpoints, achieving a flexible and decoupled architecture for model deployment and evaluation. This setup allows for dynamic GPU acceleration, where

¹⁵<https://github.com/vllm-project/vllm>

the Gunicorn server, configured with environment-specific parameters, manages multiple worker processes. Each process, handling a slice of the available GPUs, executes inference tasks in parallel, significantly improving computational efficiency and throughput. A highlight of UltraEval’s performance is its ability to utilize 4 A800 GPUs to evaluate a test set of 41k data points in under 1.5 hours, showcasing remarkable efficiency¹⁶.

3.3 Evaluation Methods

Refined data and models are instantiated through the *Task* and *Model* classes, respectively, initiating the model inference process. The model performs inference based on the input data and its hyperparameters, generating prediction outcomes. Typically, between the model’s output and the final score calculation, there are intermediate steps including post-processing and metric calculation.

Post-Processing. Model outputs, influenced by task characteristics, prompt templates, and the model’s performance, often contain extraneous information beyond the answers needed (Park et al., 2024). As shown in Figure 7, when ChatGPT responds to HumanEval questions, the response may include both code and additional textual descriptions, complicating automatic evaluation. To more accurately assess the model, it is necessary to post-process the model’s outputs to extract the most crucial answers.

Post-processing is bifurcated into two dimensions: model and task. Variations in model training approaches result in different versions, such as chat and base, necessitating distinct processing methods. Additionally, certain tasks employ specific evaluation methodologies. Taking Figure 7 as an example, the initial step involves extracting the code segment from ChatGPT’s response. Subsequently, due to HumanEval (Chen et al., 2021)’s unique evaluation

¹⁶The model used in this experiment is Llama2-7B, and the benchmarks includes BBH, MMLU, C-Eval, CMMLU, HumanEval, MBPP, GSM8K and MATH. In total, they consist of 40,938 data points.

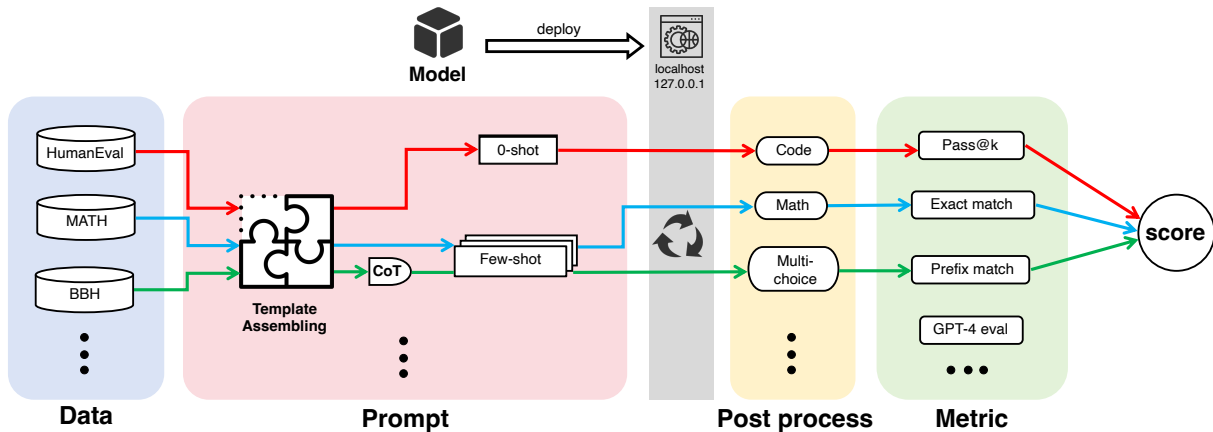


Figure 5: The combination of different modules within UltraEval.

criteria, it is necessary to extract the function body from the code while omitting the function name, yielding a cleaner and more precise answer. UltraEval develops several post-processing methods tailored to the tasks and models currently available.

Metric Calculation. Evaluation methods are categorized based on their ability to be computed automatically into automatic evaluation and human evaluation (Chang et al., 2023). For automatic evaluation, we implement common metrics such as exact match for text generation tasks, F1 score for binary classification tasks, ROUGE (Lin, 2004) for translation tasks, and pass@k (Chen et al., 2021) for coding tasks. Specifically, for exact match, we develop extensions like in match and prefix match to more effectively capture a wide range of scenarios.

Given that UltraEval is an automated evaluation framework, for human evaluation, we integrate GPT-4 (Achiam et al., 2023) as a discriminator to substitute for human evaluation. Moreover, all data results can be saved according to user preferences, allowing users to decide on human evaluation if desired, thus offering significant flexibility for a more objective assessment.

Benchmark	Llama2-7B		Llama2-13B		Mistral-7B	
	Official	UltraEval	Official	UltraEval	Official	UltraEval
ARC-C	45.9	43.2	45.9	47.4	55.5	50.8
HellaSwag	77.2	75.6	80.7	79.1	81.3	80.4
BBH	32.6	32.8	39.4	39.2	38.0*	40.4
MATH	2.5	2.8	3.9	4.8	13.1	10.2
GSM8K	14.6	14.8	28.7	22.6	52.1*	31.9
HumanEval	12.8	12.8	18.3	17.1	30.5	26.8
MBPP	20.8	20.8	30.6	29.0	47.5	47.3
MMLU	45.3	45.1	54.8	55.2	60.1	63.1

Table 3: Evaluation results on mainstream benchmarks (%). * The BBH score is not explicitly stated in the paper (Jiang et al., 2023), however, it is inferred to be 38.0 from the figures in the paper. The replicated result for GSM8K is 35.4 in Gemma paper (Team et al., 2024), which is close to our result.

4 Evaluation

UltraEval aims to provide a lightweight, comprehensive, and user-friendly evaluation framework to support research. As illustrated in Figure 5, UltraEval’s modular design effectively combines various models, tasks and metrics for evaluation. Using UltraEval, we evaluate models from the LLaMA2 series (Touvron et al., 2023) and Mistral (Jiang et al., 2023) on these widely-used benchmarks. As indicated in the Table 3, some reproduced results are higher, while others are lower, but within a certain margin of error, our reproduced results are consistent with the results reported in the papers, underscoring our framework’s reliability. The sources of error include hyperparameters (e.g., temperature, top-p) and hardware configurations. Since the evaluation details are not provided in the papers, verification is not possible. This highlights the importance of having an open and reproducible evaluation framework.

Furthermore, UltraEval supports innovative research efforts, research on predictable scaling (Hu et al., 2023), OlympiadBench (He et al., 2024) and model training, such as with MiniCPM (Hu et al., 2024).

5 Discussion and Future Work

In this section, we discuss future work following this study. Specifically, we focus on addressing data contamination and supporting a broader range of evaluation scenarios.

Data contamination refers to the phenomenon that examples from the evaluation set are also found in the training data (Li et al., 2024), causing inaccuracies in model evaluation. Common methods for contamination detection include n-gram over-

lap and embedding similarity search. However, none of these methods are perfect, making research on contamination detection still crucial. We will integrate these methods.

Supporting more evaluation scenarios, such as multimodal, long-text, and Retrieval-Augmented Generation (RAG), is crucial for meeting a broader range of evaluation needs. This will be an important direction for our future work.

6 Conclusion

We introduce UltraEval, a lightweight, user-friendly, and comprehensive framework for model evaluation. UltraEval establishes a unified structure with well-defined modules and flexible interactions, aiding researchers and developers in efficiently deploying evaluation workflows. Moving forward, we plan to continuously integrate new technologies and features into UltraEval, extending beyond large language models to support the evaluation of multimodal models, Retrieval-Augmented Generation (RAG), agents, and more, to advance the research on AGI. Additionally, we aim to expand our collection of representative benchmarks and also develop our own, exploring the capabilities and limits of large models.

Acknowledgements

This work is supported by the National Key R&D Program of China (No. 2022ZD0116312), Quan Cheng Laboratory (Grant No. QCLZD202301), the Postdoctoral Fellowship Program of CPSE (Grant No. GZB20230343), China Postdoctoral Science Foundation (Grant No. 2023M741945), and Institute Guo Qiang at Tsinghua University.

Limitations

Currently, our approach primarily utilizes text domain evaluation. However, we are looking to expand the scope of UltraEval by integrating multimodal and long-context evaluation datasets. This enhancement aims to facilitate more thorough and diverse assessments. Additionally, there is room for improvement in the visualization of our results. Future improvement will focus on enabling multi-dimensional visualization, thereby enriching the interpretability and depth of our evaluation results.

Ethical Considerations

In this paper, we present UltraEval, a lightweight, user-friendly, flexible, and comprehensive frame-

work for model evaluation. Adhering to the principles of modularity, UltraEval segments the evaluation process into three distinct modules: Data, Model, and Metrics. This approach enhances the framework’s extensibility and flexibility, allowing for the easy integration of new models and tasks. We offer an extensive benchmark suite and replicate commonly used models and benchmarks. Our results align with those reported in the corresponding papers, underscoring the stability and reliability of our framework. Committed to sustainable development, we publicly release all our code to minimize unnecessary carbon footprint. Throughout our experiments, we adhere to all licenses related to models and data.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Andrew Blair-Stanek, Nils Holzenberger, and Benjamin Van Durme. 2023. Can gpt-3 perform statutory reasoning? In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, pages 22–31.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large

- language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Yew Ken Chia, Pengfei Hong, Lidong Bing, and Soujanya Poria. 2023. Instructeval: Towards holistic evaluation of instruction-tuned large language models. *arXiv preprint arXiv:2306.04757*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- OpenCompass Contributors. 2023. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023. [A framework for few-shot language model evaluation](#).
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. 2024. Olympiad-bench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Shengding Hu, Xin Liu, Xu Han, Xinrong Zhang, Chaoqun He, Weilin Zhao, Yankai Lin, Ning Ding, Zebin Ou, Guoyang Zeng, et al. 2023. Unlock predictable scaling from emergent abilities. *arXiv preprint arXiv:2310.03262*.
- Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, Xinrong Zhang, Zheng Leng Thai, Kaihuo Zhang, Chongyi Wang, Yuan Yao, Chenyang Zhao, Jie Zhou, Jie Cai, Zhongwu Zhai, Ning Ding, Chao Jia, Guoyang Zeng, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. [Minicpm: Unveiling the potential of small language models with scalable training strategies](#).
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, et al. 2023. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.
- Yucheng Li, Frank Guerin, and Chenghua Lin. 2024. [Latesteval: Addressing data contamination in language model evaluation through dynamic and time-sensitive test construction](#).
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv: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.
- Peter S Park, Philipp Schoenegger, and Chongyang Zhu. 2024. Diminished diversity-of-thought in a standard large language model. *Behavior Research Methods*, pages 1–17.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutit Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022a. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.

Xuanyu Zhang and Qing Yang. 2023. Xuanyuan 2.0: A large chinese financial chat model with hundreds of billions parameters. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 4435–4439.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

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. 2023a. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).

Shen Zheng, Yuyu Zhang, Yijie Zhu, Chenguang Xi, Pengyang Gao, Xun Zhou, and Kevin Chen-Chuan Chang. 2023b. [GPT-Fathom: Benchmarking large language models to decipher the evolutionary path towards GPT-4 and beyond](#).

Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Neil Zhenqiang Gong, Yue Zhang, et al. 2023. Promptbench: Towards evaluating the robustness of large language models on adversarial prompts. *arXiv preprint arXiv:2306.04528*.

A More details

A.1 Post-process

We present the post-processing code in Figure 6 and explain the reasons necessitating post-processing in Figure 7.

A.2 UltraEval Usage

In the tutorial ¹⁷, we provide a detailed guide on UltraEval, including an introduction to its modules and instructions on how users can customize evaluations, such as adding their own new tasks and new models. This ensures that diverse evaluation needs are met.

A.3 Model as Judge

Currently, UltraEval supports ChatGPT as a substitute for human evaluation. However, since ChatGPT is a commercial model, users need to provide the relevant API key. Due to the modularity and flexibility of our framework, users can also use other closed-source models or highly performant open-source models as evaluators.

A.4 Multi-GPU Acceleration

Multi-GPU acceleration requires sufficient GPU resources. For larger models, such as Llama2-70B, deploying a single instance requires two A100 GPUs. Therefore, when using Multi-GPU acceleration, it is essential to consider both the model and the available GPU resources to ensure optimal configuration.

¹⁷<https://github.com/OpenBMB/UltraEval/blob/main/docs/tutorials/en/ultraeval.md>

```

1
2 def process_text(text):
3     triple_quotes_indices = [i for i, _ in enumerate(text) if text.
4                             startswith('\n\n\n', i)]
5
6     if len(triple_quotes_indices) % 2 == 0 and len(triple_quotes_indices) >
7         0:
8         for i in range(0, len(triple_quotes_indices), 2):
9             start_index = triple_quotes_indices[i]
10            end_index = triple_quotes_indices[i + 1]
11            if 'def' in text[end_index:]:
12                return text[:start_index].strip()
13            return text.strip()
14    elif len(triple_quotes_indices) > 0:
15        return text[:triple_quotes_indices[0]].strip()
16    else:
17        return text.strip()

```

Figure 6: An example of a post-processing method for the MBPP task.



You

```

from typing import List\n\n\ndef intersperse(numbers: List[int], delimiter: int) -> List[int]:\n
\"\"\" Insert a number 'delimiter' between every two consecutive elements of input list
`numbers`\n >>> intersperse([], 4)\n []\n >>> intersperse([1, 2, 3], 4)\n [1, 4, 2, 4, 3]\n \"\"\"

```



ChatGPT

Here's how you could implement the `intersperse` function. This function takes a list of integers, `numbers`, and an integer `delimiter`. It returns a new list where the `delimiter` is inserted between each pair of consecutive elements in `numbers`.

```

python Copy code

from typing import List

def intersperse(numbers: List[int], delimiter: int) -> List[int]:
    """
    Insert a number 'delimiter' between every two consecutive elements of input list
    >>> intersperse([], 4)
    []
    >>> intersperse([1, 2, 3], 4)
    [1, 4, 2, 4, 3]
    """
    if not numbers:
        return []

    # Start with the first element of numbers, as there's nothing before it to intersperse
    result = [numbers[0]]

    # Iterate over the numbers, starting from the second element.
    for number in numbers[1:]:
        result += [delimiter, number]

    return result

```

This function begins by checking if the `numbers` list is empty, in which case it immediately returns an empty list. Then, it initializes the `result` list with the first element of `numbers`, because there's no need to insert the `delimiter` before the first element. It iterates over the remaining elements in `numbers`, appending the `delimiter` followed by the current element to the `result` list for each iteration.



Figure 7: An example of ChatGPT to accomplish HumanEval. The figure illustrates that the responses from ChatGPT cannot be directly used for computation and require post-processing to extract the substantive content.