

From REAL to SYNTHETIC: Synthesizing Millions of Diversified and Complicated User Instructions with Attributed Grounding

Chiwei Zhu^{1,2§}, Benfeng Xu^{1,2†}, Xiaorui Wang², Zhendong Mao¹

¹University of Science and Technology of China

²Metastone Technology

{tanz, benfeng}@mail.ustc.edu.cn

Abstract

The pursuit of diverse, complex, and large-scale instruction data is crucial for automatically aligning large language models (LLMs). While there are methods capable of generating synthetic instructions at scale, they either suffer from limited grounding sources, leading to a narrow distribution, or rely on trivial extensions that fail to produce meaningful trajectories in terms of complexity. In contrast, instructions that benefit efficient alignment are typically crafted with cognitive insights and grounded in real-world use cases. In this paper, we synthesize such instructions using attributed grounding, which involves 1) a top-down attribution process that grounds a selective set of real instructions to situated users, and 2) a bottom-up synthesis process that leverages web documents to first generate a situation, then a meaningful instruction. This framework allows us to harvest diverse and complex instructions at scale, utilizing the vast range of web documents. Specifically, we construct a dataset of 1 million instructions, called SYNTHQUESTIONS, and demonstrate that models trained on it achieve leading performance on several common benchmarks, with improvements that continually scale with more web corpora. Data, models and codes will be available at <https://github.com/Ignoramus0817/SynthQuestions>.

1 Introduction

Alignment training (Wei et al., 2022) has become an essential technique for instruction-following large language models (Ouyang et al., 2022; OpenAI et al., 2024; Yang et al., 2024; Dubey et al., 2024; GLM et al., 2024), which aims to align language models behaviors with human when given certain instructions through training on instruction-response pairs. Researchers have been studying how to achieve such alignment effectively, on which educational psychology provides us with wisdom. Vygotsky proposed in his Zone of Proximal Development theory (Vygotsky, 1978) that

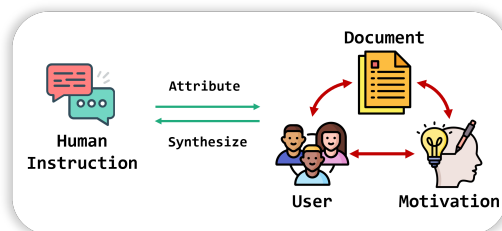


Figure 1: Human instructions can be attributed to documents, users and motivations. Conversely, instructions can also be synthesized from them.

tasks that are just beyond the learner’s capabilities promote maximum cognitive growth, which is suitable for alignment as well. Showing testimonies to such arguments, plenty of studies have shown that to obtain strong instruction-following and reasoning capability, instruction data that are sufficiently **diversified, complicated** and **scaled** are required (Kaplan et al., 2020; Hoffmann et al., 2022; Lu et al., 2023; Köpf et al., 2023; Chiang et al., 2023; Zhao et al., 2024).

However, collecting such instructions is an intractable mission, relying on massive use cases and brain labor from human users. As a result, increasing number of works seek to synthesize instructions with language models. Typical existing methods involve augmenting seed tasks (Taori et al., 2023; Xu et al., 2023b), generating instructions according to real-world concepts (Ding et al., 2023), or training model mimicking human to ask questions (Kong et al., 2024). While these approaches succeed in the automatic generation of scaled instructional data, they are constrained by the design of their synthesizing methodologies and inevitably fall into certain local distributions (e.g., knowledge and concepts from Wikipedia, limited seed instructional patterns and so on). Consequently, they fall short

[§]Work done during the internship in Metastone Technology.

[†]Corresponding author.

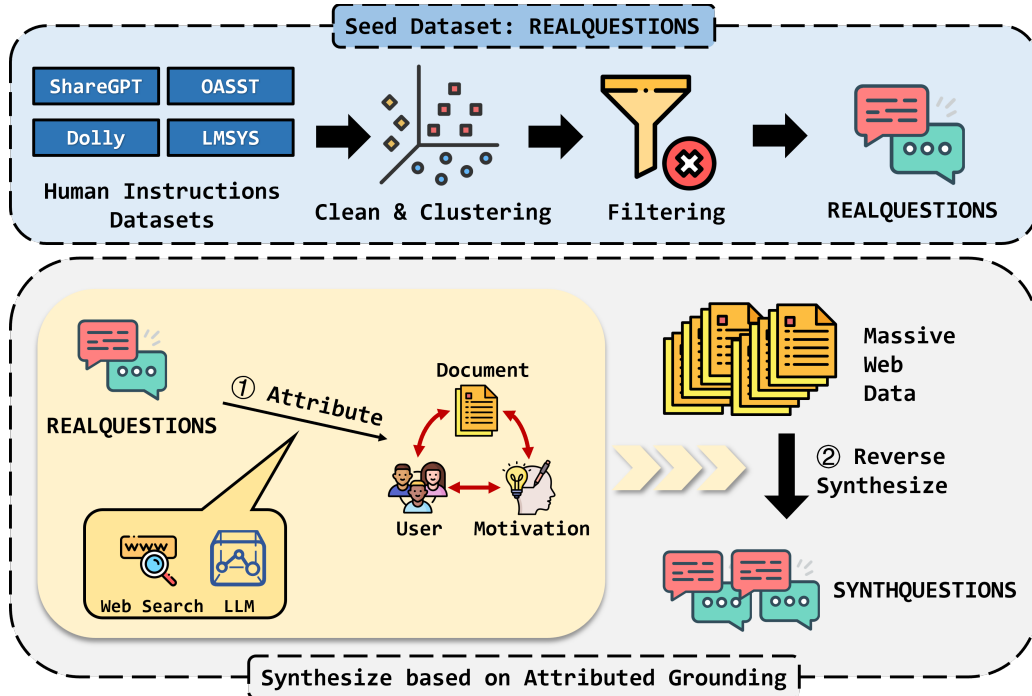


Figure 2: Overview of our synthesizing framework.

in generating lifelike and complex instructions that accurately reflect the diversity and intricacy of real-world tasks and queries.

Various linguistic and social studies have pointed out that language understanding is based on world knowledge, which is situated, being in part a product of the activity, context, and culture in which it is developed and used (Brown et al., 1989; Epley et al., 2004; Bisk et al., 2020). Given this idea, we believe it is critical to assure the generated instructions are grounded to real world. In this work, we propose a synthesizing framework based on the core idea of **attributed grounding**, which consists of two main parts: top-down **attributing** and bottom-up **synthesizing**. From our perspective, a human instruction can be attributed to three key factors (as shown in Figure 1): **(1) Document:** the background knowledge involved in the instruction. **(2) User:** who proposes this instruction. **(3) Motivation:** why the users need an LLM to do the task for them. Through these factors, an instruction is grounded to the real world. Conversely, we can build situations including users and motivations from documents and synthesize instructions accordingly. As massive and diverse web documents are accessible without much effort, it is possible to generate pre-training-level instruction data with high complexity and diversity.

We conduct an in-context-learning driven im-

plementation for our framework, as can be seen in Figure 2. We first build a seed dataset for the attributing process. We collect commonly-used human-labeled SFT datasets, clean and deduplicate the instructions, and keep the ones with the highest quality. We call the resulting seed dataset **REALQUESTIONS**. In the attributing step, we recall web documents for each instruction in **REALQUESTIONS**, based on which we build a lifelike situation with users and motivations leveraging an advanced LLM. In the synthesizing step, we start with existing web documents and prompt an LLM to generate grounding situations along with new instructions. Above process is done in an in-context-learning style, driven by demonstrations from **REALQUESTIONS**. Figure 3 showcases an example of web document, the corresponding user/motivation and the synthesized instruction, which is well grounded and complicated (see Appendix K for more cases).

With above process we harvest a 1M-size dataset, **SYNTHQUESTIONS**, which exhibits higher diversity than other synthesized datasets. Model fine-tuned on **SYNTHQUESTIONS** achieves leading results on various benchmarks and is comparable with models trained with 10 times more data and preference training, which demonstrate the effectiveness of our method.

Our contributions can be concluded as follows:

- We construct **REALQUESTIONS**, an instruc-

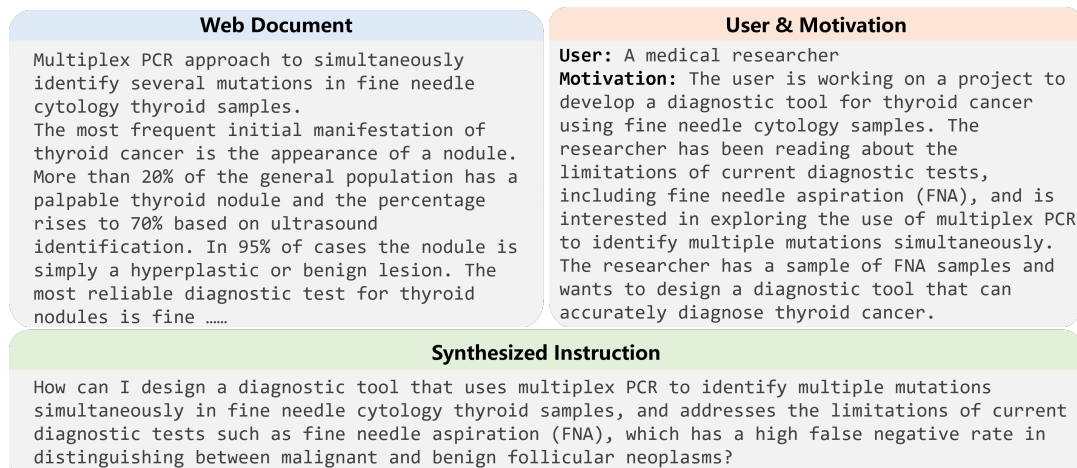


Figure 3: A random sample of web document, the corresponding grounding situation and generated instruction.

tion dataset of cleaned and filtered human instructions.

- We propose and implement a new data synthesizing framework, which can generate pre-training-scale data with high quality. Based on the framework we construct a 1M scale dataset, SYNTHQUESTIONS.
- We demonstrate the effectiveness of our methods through intense experiments on series of open-ended and closed-form benchmarks.

2 Related Works

2.1 Alignment Training

Since first appearance in (Weller et al., 2020), researchers have found that models can show surprising generalization ability on understanding human intents when fine-tuned with instruction-response pairs (Wei et al., 2022; Chung et al., 2022; Mishra et al., 2022; Ouyang et al., 2022; Xu et al., 2023a). Nowadays supervised fine-tuning (SFT), also referred to as instruction tuning, has become essential for aligning large language models with human behaviors, which is invariably applied to almost every instruction following LLMs (Chiang et al., 2023; OpenAI et al., 2024; Dubey et al., 2024; Yang et al., 2024; GLM et al., 2024). The broad application of SFT raises a high demand on instruction-response data. A number of works have put effort on collecting SFT data from human. ShareGPT (Chiang et al., 2023) and OpenAssistant (Köpf et al., 2023) collect user conversations with proprietary LLMs like GPT-4. Chatbot Arena (Chiang et al., 2024) is a benchmarking platform for users to chat with

different LLMs and rate their responses which collects human instructions in the same time, resulting in several conversation datasets (Zheng et al., 2023, 2024). However, collecting conversations from users or crowd-sourcing to annotate instruction data is rather expensive, which leads to the limited scale of most human-labeled SFT datasets.

2.2 Instruction Data Synthesizing

As the generating ability of LLMs become stronger, more recent works seek to synthesize SFT data automatically to break through the scale limit of human annotating. Self-Instruct (Wang et al., 2023a) firstly introduces the idea of generating instructions with LLMs themselves. Following this idea, Alpaca (Taori et al., 2023) generates 52K instructions from 175 seed tasks with OpenAI’s text-davinci-003. WizardLM (Xu et al., 2023b) prompts gpt-3.5-turbo to evolve a seed dataset to generate more complicated instructions. Hump-Back (Li et al., 2023) constructs a translation model that back-translates web documents into instructions. PlatoLM (Kong et al., 2024) directly train a model on existing SFT datasets to simulate users and raise questions. To improve the diversity and complexity of the instructions, later works begin to inject more real-world information to the generating process. UltraChat (Ding et al., 2023) leverage Wikipedia entities to improve field coverage. MAMMO2 (Yue et al., 2024) directly extract Q-A pairs from web documents and refine them to construct a 10M SFT dataset.




Figure 4: Scores distribution of human instructions, based on Arena Hard Pipeline. The instructions scored 7 go into the REALQUESTIONS dataset.

3 REALQUESTIONS

REALQUESTIONS is a high-quality human instruction dataset that we construct as the seed dataset, which will later be attributed to grounding situations and drive the synthesizing process as demonstrations. REALQUESTIONS is built with the following steps:

Data Collection. We collect conversation data from 7 commonly used human-labeled instruction datasets, namely Chatbot Arena Conversations (Chiang et al., 2024), Databricks-dolly-15k (Conover et al., 2023), LMSYS-Chat-1M (Zheng et al., 2024), OpenAssistant (Köpf et al., 2023), ShareGPT (Chiang et al., 2023), UltraChat (Ding et al., 2023) and WildChat (Zhao et al., 2024). We collect a total of 1.92M raw conversation data, details of which are shown in Figure 9.

Data Cleaning and Deduplication. There is massive noise in the raw conversation data and a variety of measures are applied to reduce it:

- Conversations that are not complete or do not appear in English are discarded.
- Remove conversations where the user instructions are similar to that of our evaluation benchmarks, Alpaca Eval 2.0 and Arena Hard.
- As the tasks of user instructions in the conversations exhibit high degree of duplication, we deduplicate them according to their semantics.

For the detail of deduplication, we apply a community detection algorithm* on the instructions,

*https://sbert.net/docs/package_reference/util.html#sentence_transformers.util.community_detection

where data points whose embedding similarity exceeds a specified threshold (0.85 in our implementation) are considered to constitute a cluster. For each cluster standing for a certain task or topic, we keep only one instruction to maximize the diversity of the resulted dataset. For outlier instructions that do not belong to any clusters, we keep them all as they are not similar to any other ones and can be regarded as unique samples. After the cleaning and deduplication, we get a total of 690K conversations. We discard all the responses from the conversations and only keep the user instructions of the first rounds for later use.

Data Filtering. As mentioned previously, challenging tasks are particularly beneficial for model performance. To filter the most challenging instructions, we modify criteria from Arena Hard (Li* et al., 2024) to score them, where instructions are assessed on 7 dimensions (details in Appendix I).

We use LLaMA-3-70B-Instruct as the judging model. For each dimension that an instruction satisfies, 1 score will be added. Detailed statistics of scored instructions are shown in Figure 4, where we refer the 29K instructions with full score as our seed dataset REALQUESTIONS. We fine-tune a LLaMA-3-8B on REALQUESTIONS and experimental results show its superior quality compared to existing instruction datasets (see Table 6).

4 Attributed Grounding

4.1 Attributing

In the top-down attributing step, instructions from REALQUESTIONS will be attributed to documents, users and motivations.

Documents. Attributing starts by collecting documents, i.e. relevant real-world information, for instructions in REALQUESTIONS, which is done with web search in our implementation. We utilize LLaMA-3-70B-Instruct to extract key concepts of each instruction, and recall web pages from Google using the key concepts as queries. We keep the top-1 result as the document for each instruction.

Users and Motivations. With documents that provide background knowledge about real-world, we can further simulate the situation where the instruction appears. We provide LLaMA-3-70B-Instruct with the document along with the instruction and prompt it to conceive a situation where a user interacts with the document

and brings up the instruction out of certain motivation. To improve the grounding process, we conduct the prompting with manually crafted demonstrations, which are shown in Appendix J.1. Following above process, we get attributed SYNTHQUESTIONS which we refer to as:

$$RQ^\alpha = \{(i, d, u, m)\} \quad (1)$$

where i represents instructions from SYNTHQUESTIONS, d , u and m refers to the attributed factors documents, users and motivations respectively.

4.2 Synthesizing

In the bottom-up synthesizing process, we reverse the attributing process, constructing situations with users and motivations from existing web documents, based on which new instructions are synthesized. Attributed samples from RQ^α are used here to regulate model behaviors.

Documents. We use FineWeb (Penedo et al., 2024) as the main source for our documents. To further amplify the dataset’s benefits for complex reasoning capabilities, we additionally mix in documents that involve difficult reasoning tasks like mathematics and coding from PILE (Gao et al., 2020) and MathPILE (Wang et al., 2023b).

Users and Motivations. For each document, we prompt LLaMA-3-8B-Instruct to build a grounded situation with users and motivations. To generate more reasonable and grounded situations, samples from RQ^α are used as demonstrations in this process:

$$(u', m') = LLaMA(P_g, d', demo) \quad (2)$$

where the outputs u' are m' are generated user and motivation, the inputs P , d' , and $demo$ are prompt for grounding, documents in the above corpus, and demonstrations from RQ^α respectively.

New Instructions. Finally we ask the model to play the role of the user and utter the most possible instructions when placed in the above grounding situation:

$$i' = LLaMA(P_i, d', u', m', demos) \quad (3)$$

where i' is the generated instruction, P_i is the prompt for instruction generation.

After generating new instructions, we first score them with methods in Section 3 and discard all the instructions whose scores are below 3 (distribution

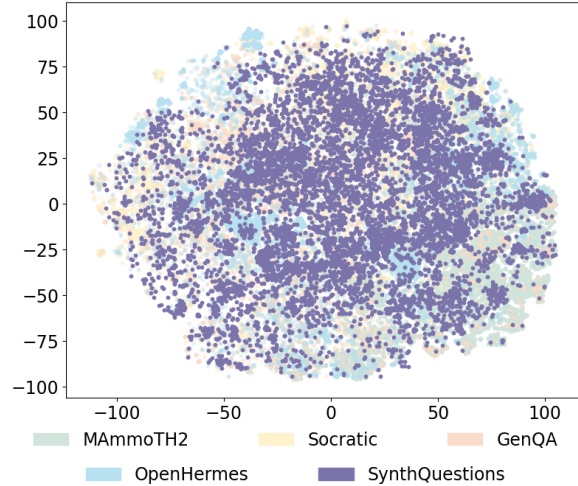


Figure 5: t-SNE plot of SYNTHQUESTIONS along with MAMmoTH2, SocraticChat, GenQA and OpenHermes. SYNTHQUESTIONS covers more area than other datasets.

of the full dataset is shown in Appendix E). We choose the threshold 3 as we notice that Specificity, Problem-Solving and Technical Accuracy are three more fundamental requirements for a valid instruction, and instructions failing these three requirements tend to be unclear or ambiguous. As this is a heuristic setting, we set the threshold 3 instead of directly filtering out instructions that do not meet these requirements to be more tolerant. To assure the diversity of the dataset, we conduct topic modeling with BERTopic (Grootendorst, 2022) following (Li* et al., 2024) and construct a 1M-size final dataset by including instructions with the highest scores in each topic. We refer to the final dataset as **SYNTHQUESTIONS**.

5 Data Analysis

In this section we demonstrate the diversity and complexity of SYNTHQUESTIONS through qualitative and quantitative evaluations.

5.1 Basic Statistics

Table 9 in Appendix D shows the basic statistics of SYNTHQUESTIONS along with other instruction datasets(tokenization is done with Tiktoken[†]). As can be seen in the table, our method generates data with the most average turn lengths, indicating the complexity of SYNTHQUESTIONS. Besides, SYNTHQUESTIONS ranks top in the lexical diversity calculated with MTL algorithm (McCarthy and Jarvis, 2010) among all synthesized datasets.

[†]<https://github.com/openai/tiktoken>

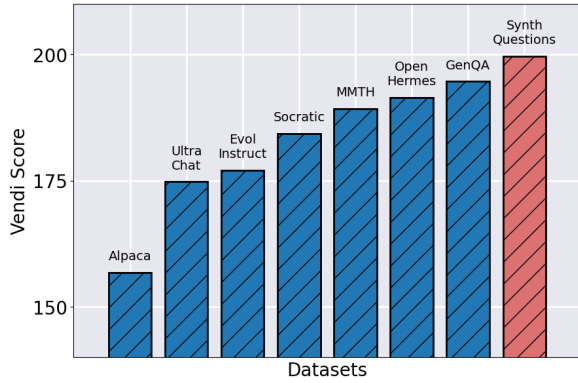


Figure 6: Comparison of diversity (vendi score) of synthesized datasets.

5.2 Diversity

Apart from lexical diversity in Table 9, we assess and visualize the semantic diversity of instructions in SYNTHQUESTIONS along with other synthetic datasets. We sample 10,000 instructions from each dataset above and calculate sentence embeddings with `all-mpnet-base-v2` model[‡]. Then t-SNE is applied to project semantic embeddings into a 2D space. Figure 5 display the t-SNE plots of the five most recent datasets (more datasets are visualized in Appendix E), where data points from SYNTHQUESTIONS occupy the most extensive area. This implies that dataset synthesized with our method covers more diverse topics or subjects.

Additionally, we provide Vendi Score as a supplemental quantitative metric of data diversity. Vendi Score (Friedman and Dieng, 2023) is a diversity evaluation metric designed for machine learning dataset, which increases linearly with the number of unique modes in the dataset. We compute Vendi Score on the previously sampled 10,000 data points for each dataset. Results are presented in Figure 6, where Vendi Score of SYNTHQUESTIONS ranks top among all synthesized datasets.

5.3 Complexity

We randomly sample 10,000 unfiltered instructions from all the synthetic datasets in Table 9 and score them with criteria from Arena Hard. Results in Figure 7 display the violin plots of score distribution of the 5 most recent dataset (full results in Appendix E), where the score distribution of SYNTHQUESTIONS is notably condensed to the higher end of spectrum compared to other datasets. Such

[‡]<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>




Figure 7: Comparison of complexity (Arena Hard Score) of all synthetic data.

observation signifies the efficiency in our framework to produce complicated instructions.

5.4 Safety

We analyze SYNTHQUESTIONS with LLaMA-Guard-3-8B, detecting potential harmful content in the dataset (results shown in Appendix F). Among the 1M dataset, we detected 4.32% data points with potential safety risks. Detailed results are shown below, where most of problematic data points (3.60%) lie in the "Specialized Advice" category, which we think are the ones requiring users to carefully discern the responses, instead of being directly harmful. For all the other categories, the potentially risky data points are less than 0.2%. When open-sourcing we will remove contents with potential harm.

6 Experiments

In this section, we first verify that data synthesized with our method can improve model's instruction following and reasoning performance effectively. Then we show the scaling curve of our synthesized data, demonstrating how model performance changes as data scale increases. Finally through ablation experiments we demonstrate the potential of our data to be further elicited by preference optimization and the necessity of each module in our method. Instruction tuning and preference optimization are respectively conducted with Megatron-LM[§] and Huggingface TRL[¶]. Full training details are shown in Appendix C.

[§]<https://github.com/NVIDIA/Megatron-LM>
[¶]<https://huggingface.co/docs/trl/index>

Model Groups	Models	Data Scale	Arena Hard	Alpaca Eval 2.0		
			WR(%)	LC(%)	WR(%)	SD
Larger or Proprietary Models	HumpBack-LLaMA2-70B	-	-	16.25	10.12	0.94
	GPT-3.5-Turbo-0301	-	18.1	18.09	9.62	0.89
	GPT-3.5-Turbo-1106	-	18.9	19.30	9.18	0.89
7B-8B Models w/ Proprietary Data	Mistral-7B-Instruct-v0.3	-	16.7 [†]	20.61	16.69	1.11
	Qwen-2-7B-Insturct [†]	-	23.5	21.86	19.62	1.15
	LLaMA-3-8B-Instruct	>10M	20.6	22.92	22.57	1.26
LLaMA3-8B w/ Open-source Data	OpenHermes2.5	1M	4.4	9.94	6.27	0.73
	GenQA	10M	3.0	9.05	7.11	0.82
	MAMmoTH2*	10M	16.6	<u>18.5</u>	-	-
Ours	SYNTHQUESTIONS	1M	<u>15.4</u>	18.87	19.15	1.15

Table 1: Performance of models on Alpaca Eval 2.0 and Arena Hard benchmarks. Among models with open-source data, the best performance is bolded and the second best performance is underlined. Results marked with † are evaluated by us. *Apart from synthetic data, MAMmoTH2 is further fine-tuned with external math and code datasets, which may explain its high performance on Arena Hard.

Model Groups	Models	IFEVAL	MMLU	ARC-C	GPQA	GSM8K	MATH
7B-8B Models w/ Proprietary Data	Mistral-7B-Instruct-v0.3	54.65	61.84	63.57	27.8 [†]	43.37	12.94 [†]
	Qwen-2-7B-Insturct	56.79	70.5	59.73	25.3	82.3	49.6
	LLaMA-3-8B-Instruct	74.08	68.5	82.4	34.6	80.6	29.1
LLaMA3-8B w/ Open-source Data	OpenHermes2.5	-	<u>65.7</u>	61.86	-	67.02	-
	GenQA	-	63.45	58.53	-	43.13	-
	MAMmoTH2	43.94 [†]	64.2	82.2	35.2	<u>70.4</u>	35.8
Ours	SYNTHQUESTIONS	57.05	65.79	<u>63.92</u>	<u>30.3</u>	70.53	<u>22.71</u>

Table 2: Performance of models on different close-ended knowledge and reasoning benchmarks. Notations and marks are the same with the above figure. Some unreported results are not reproduced due to high expense.

6.1 Main Results

Due to page limit, we omit the evaluation results of datasets that are relatively old and less competitive (e.g. UltraChat, ShareGPT and so on). Complete results for them can be found in Appendix G.

Alignment Benchmarks. To verify the effectiveness of our method and the quality of synthesized data, we train LLaMA-3-8B on SYNTHQUESTIONS. For evaluation, we select two prevailing alignment benchmarks, Alpaca Eval 2.0 (Dubois et al., 2024) and Arena Hard (Li* et al., 2024), which leverage gpt-4-1106-preview as the judge and are highly consistent with human annotation.

Results are shown in Table 1. Among all models trained on open-source datasets, model trained on SYNTHQUESTIONS shows the best performance on Alpaca Eval 2.0 and only falls behind MAMmoTH2 on Arena Hard. It is especially worth mentioning that on Alpaca Eval 2.0, model trained with SYNTHQUESTIONS outperforms ones trained with MAMmoTH2 and HumpBack, which are two synthesizing methods also utilizing web documents. Note that MAMmoTH2 is trained with 10M data and fur-

ther fine-tuned with open-source math and code datasets, which may explain its high win rate on Arena Hard. Models trained with SYNTHQUESTIONS is also comparable to latest models trained with proprietary data and reinforcement learning, showing a better win rate on Alpaca Eval 2.0 than Mistral-7B-Instruct-v0.3. Above experimental results prove the effectiveness of SYNTHQUESTIONS on improving models capabilities.

Closed-form Benchmarks. To demonstrate the robustness of our dataset, we evaluate the models on several closed-form benchmarks. We present the results in Table 2, where model trained on SYNTHQUESTIONS ranks either first or second on all benchmarks among models trained with open-source data.

Preference Optimization. We further investigate the potential of our synthesized dataset by applying DPO. We randomly sample 100K instructions from different datasets, generate 5 responses(T=0.8) with LLaMA-3-70B-Instruct, and label the preferences with ArmoRM-LLama3-8B-v0.1 (Wang et al.,

Models	Alpaca Eval (WR)	Arena Hard
LLaMA-3-8B-Instruct	22.56	20.6
LLaMA-3-70B-Instruct	33.18	44.5
SYNTHQUESTIONS	19.15	15.4
+DPO (MAmmoTH2)	28.46	15.6
+DPO (GenQA)	28.52	17.4
+DPO (OpenHermes)	<u>28.94</u>	<u>19.6</u>
+DPO (Ours)	33.81	24.8

Table 3: Performance of models with DPO on different datasets. Our DPO model even outperforms LLaMA-3-70B-Instruct on Alpaca Eval 2.0 Win Rate.

2024). We set the response with the highest score as the chosen, and the one with the lowest score as the rejected. We train our the SFT model and the performance of resulted model are presented in 3. Model trained on SYNTHQUESTIONS not only outperforms all recent synthesized datasets, but also beats LLaMA-3-8B-Instruct. Our model even surpasses the data generator LLaMA-3-70B-Instruct on Alpaca Eval 2.0 Win Rate. However, the performance on Arena Hard falls far behind, which may indicate that solving more difficult tasks still calls for an increase in model scales.

6.2 Scaling Effect

We investigate the impact of instruction scale to model performance. We randomly draw $2^i \cdot 100K$ subsets from the total 1M data and train LLaMA-3-8B models on them. Performance of the models are evaluated with Alpaca Eval 2.0, results of which are displayed in Figure 8. As the scale of the train set increases, model performance consistently improves on the benchmark, which corroborates the positive impact of expanding the scale of data on enhancing model performance, while also demonstrating the potential of our approach to continuously improve model capabilities by synthesizing more instructions. We also provides results of Qwen-2.5-7B trained on subsets with different scales in Appendix H, which shows similar trends.

6.3 Performance on Model of Other Structures and Sizes

To verify the generalizability of SYNTHQUESTIONS, we train another two models, Qwen2.5-7B and Qwen2.5-14B on a 100k-subset randomly drawn from different synthesized datasets including OpenHermes2.5, GenQA, MAmmoTH2 and SYNTHQUESTIONS. We evaluate the model perfor-

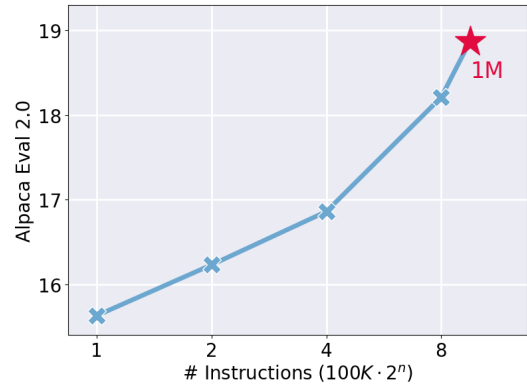


Figure 8: Performance of models trained with subsets in different scales.

mance on Alpaca Eval 2.0 and closed-form knowledge and reasoning benchmarks. As is shown in Table 4, models trained with SYNTHQUESTIONS ranks first on most of the benchmarks with a non-trivial advantage. Especially, in MATH dataset, model trained with SYNTHQUESTIONS outperforms other models by a large margin. We also notice that SYNTHQUESTIONS does not perform as strong on GPQA, which may indicate that it still requires enhancement in subjects like biology, physics, and chemistry. Nevertheless, present results can already demonstrate the effectiveness and generalizability of our method.

6.4 Ablations

In this section, we verify the necessity of each module by ablate certain parts of our synthesizing framework. Figure 5 presents the main results of our ablation study.

Attributed Grounding. We test the effect of our core idea, i.e. attributed grounding, by directly generating instructions without attributing them to documents, users or motivations. We leverage instructions from REALQUESTIONS as demonstrations and prompt LLaMA-3-70B-Instruct to generate new ones. When selecting demonstrations, we apply two strategies: randomly sampling and semantics-based selecting. For the latter strategy we randomly sample an instruction from REALQUESTIONS and then search for K nearest instructions from the whole unfiltered human instruction set (mentioned in Section 3). Detailed prompts are shown in Appendix J.1. We collect 100K instructions for each strategy and train models on them.

As shown in the table, data generated with both

Model	Dataset	AE (LC)	AE (WR)	IFEVAL	MMLU	ARC-C	GPQA	GSM8K	MATH
Qwen2.5-7B	OpenHermes2.5	16.09	9.85	48.08	68.86	82.17	32.32	78.62	29.8
	GenQA	11.33	6.23	42.21	67.49	84.64	33.33	72.63	29.74
	MAmmoTH2	11.38	6.72	41.61	69.68	86.95	25.76	68.02	35.48
	SYNTHQUESTIONS	17.25	16.03	48.68	70.2	87.88	28.12	79.91	41.22
Qwen2.5-14B	OpenHermes2.5	24.89	13.21	51.02	74.37	89.16	32.32	85.14	36.82
	GenQA	12.06	6.31	42.03	72.98	89.93	38.38	77.48	37.66
	MAmmoTH2	16.96	8.30	44.18	77.2	90.44	32.83	79.38	40.48
	SYNTHQUESTIONS	24.33	22.22	58.03	77.16	90.61	34.34	87.49	44.96

Table 4: Performance of Qwen2.5-7B and Qwen2.5-14B trained on different 100K datasets. AE denotes Alpaca Eval 2.0

Training Set	AE	GSM8K
SYNTHQUESTIONS-100K	15.63	58.30
- Attributed Grounding		
w/ KNN	10.85	45.26
w/o KNN	10.64	43.77
- Math/Code Docs	15.50	50.34

Table 5: Results of ablation study on grounding and additional documents. AE denotes Alpaca Eval 2.0.

strategy behave similarly, bringing little improvement to model performance compared to SYNTHQUESTIONS-100K. Such degeneration demonstrate the effectiveness of attributing process in our method, underlying the critical role of grounding.

Math and Code Documents. Apart from FineWeb, we add documents involving more difficult tasks like math or code when synthesizing new instructions. To study the effect of these documents, we randomly sample 100K instructions purely generated from FineWeb and train a model on it. Results in Table 5 show that though removing additional documents does not cause significant performance degradation on Alpaca Eval 2.0, the accuracy on GSM8K drops severely. Such phenomenon reveals that incorporating instructions covering more challenging tasks or fields are especially beneficial for model performance on reasoning and knowledge tasks.

7 Conclusion

In this paper, we propose a two-step instruction synthesizing framework aimed at generating better grounded instruction data. Our framework first attributes human instructions to documents, users and motivations, and then reversely generate grounded instructions from existing web documents through simulating the natural appearance of human instructions. With our synthesizing framework, we construct SYNTHQUESTIONS, a 1-million synthesized

instruction dataset. We fine-tuned LLaMA-3-8B models on our synthesized data and experiments shown that SYNTHQUESTIONS can enhance model capabilities effectively, achieving comparable performance with models trained with 10 times more data and preference training. Apart from decent performance, study about the scaling effect of SYNTHQUESTIONS demonstrates the potential of our method to further improve model capabilities by synthesizing larger scales of data.

Acknowledgment

This research is supported by Artificial Intelligence-National Science and Technology Major Project 2023ZD0121200 and National Natural Science Foundation of China under Grant 62222212.

Limitation

The main limitations of this work fall in two aspects. Firstly, while scaling curve shows the potential to further improve model performance, we do not test data scale larger than 1M. Secondly, while it is not the main topic of this work, a more thorough study about the optimal selection and distribution of web corpora used for synthesizing can be conducted.

Potential Risks

The dataset has not been assessed in terms of hallucination, which may lead language models to output false or unfaithful contents.

References

Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language](#). *Preprint*, arXiv:2004.10151.

- John Seely Brown, Allan Collins, and Paul Duguid. 1989. [Situated cognition and the culture of learning](#). *Educational Researcher*, 18(1):32–42.
- Jiuhai Chen, Rifaa Qadri, Yuxin Wen, Neel Jain, John Kirchenbauer, Tianyi Zhou, and Tom Goldstein. 2024. [Genqa: Generating millions of instructions from a handful of prompts](#). *Preprint*, arXiv:2406.10323.
- 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](#).
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. [Chatbot arena: An open platform for evaluating llms by human preference](#). *Preprint*, arXiv:2403.04132.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. [Free dolly: Introducing the world’s first truly open instruction-tuned llm](#).
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Zhi Zheng, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. [Enhancing chat language models by scaling high-quality instructional conversations](#). *arXiv preprint arXiv:2305.14233*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gouget, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenxin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baeviski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi,

- Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khan-delwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsim-poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Her-moso, Mo Metanat, Mohammad Rastegari, Mun-ish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pa-van Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Mah-eswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lind-say, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agar-wal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiao-jian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tat-sunori B. Hashimoto. 2024. [Length-controlled alpacaeval: A simple way to debias automatic evaluators](#). *Preprint*, arXiv:2404.04475.
- Nicholas Epley, Boaz Keysar, Leaf Van Boven, and Thomas Gilovich. 2004. [Perspective taking as ego-centric anchoring and adjustment](#). *Journal of personality and social psychology*, 87:327–39.
- Dan Friedman and Adji Bousso Dieng. 2023. [The vendi score: A diversity evaluation metric for machine learning](#). *Preprint*, arXiv:2210.02410.
- Leo Gao, Stella Biderman, Sid Black, Laurence Gold-ing, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The pile: An 800gb dataset of diverse text for language modeling](#). *Preprint*, arXiv:2101.00027.
- Team GLM, :, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadao Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Jingyu Sun, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). *Preprint*, arXiv:2406.12793.
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models](#). *Preprint*, arXiv:2203.15556.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *Preprint*, arXiv:2001.08361.
- Chuyi Kong, Yaxin Fan, Xiang Wan, Feng Jiang, and Benyou Wang. 2024. [Platolm: Teaching llms in multi-round dialogue via a user simulator](#). *Preprint*, arXiv:2308.11534.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. [Openassistant conversations – democratizing large language model alignment](#). *Preprint*, arXiv:2304.07327.
- Tianle Li*, Wei-Lin Chiang*, Evan Frick, Lisa Dunlap, Banghua Zhu, Joseph E. Gonzalez, and Ion Stoica. 2024. [From live data to high-quality benchmarks: The arena-hard pipeline](#).
- Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason Weston, and Mike Lewis. 2023. [Self-alignment with instruction back-translation](#). *arXiv preprint arXiv:2308.06259*.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2023. [#instag: Instruction tagging for analyzing supervised fine-tuning of large language models](#). *Preprint*, arXiv:2308.07074.
- Philip M McCarthy and Scott Jarvis. 2010. Mtd, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. [Cross-task generalization via natural language crowdsourcing instructions](#). *Preprint*, arXiv:2104.08773.
- Jinjie Ni, Fuzhao Xue, Kabir Jain, Mahir Hitesh Shah, Zangwei Zheng, and Yang You. 2023. [Instruction in the wild: A user-based instruction dataset](#). <https://github.com/XueFuzhao/InstructionWild>.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameesh Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perialman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav

- Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeff 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 Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben alal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). *Preprint*, arXiv:2406.17557.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Teknium. 2023. [Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants](#).
- Lev Semenovich Vygotsky. 1978. *Mind in society: The development of higher psychological processes*, volume 86. Harvard university press.
- Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. 2024. Interpretable preferences via multi-objective reward modeling and mixture-of-experts. In *EMNLP*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023a. [Self-instruct: Aligning language models with self-generated instructions](#). *Preprint*, arXiv:2212.10560.
- Zengzhi Wang, Rui Xia, and Pengfei Liu. 2023b. [Generative ai for math: Part i – mathpile: A billion-token-scale pretraining corpus for math](#). *Preprint*, arXiv:2312.17120.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#). *Preprint*, arXiv:2109.01652.
- Orion Weller, Nicholas Lourie, Matt Gardner, and Matthew E. Peters. 2020. [Learning from task descriptions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1361–1375, Online. Association for Computational Linguistics.
- Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou, Yongdong Zhang, and Zhendong Mao. 2023a. [Expertprompting: Instructing large language models to be distinguished experts](#). *Preprint*, arXiv:2305.14688.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023b. [Wizardlm: Empowering large language models to follow complex instructions](#). *Preprint*, arXiv:2304.12244.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Xiang Yue, Tuney Zheng, Ge Zhang, and Wenhu Chen. 2024. [Mammoth2: Scaling instructions from the web](#). *Preprint*, arXiv:2405.03548.
- Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. [Wildchat: Lm chatgpt interaction logs in the wild](#). *Preprint*, arXiv:2405.01470.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric P. Xing, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. 2024. [Lmsys-chat-1m: A large-scale real-world llm conversation dataset](#). *Preprint*, arXiv:2309.11998.
- 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](#). *Preprint*, arXiv:2306.05685.

A Dataset and Model Licenses

Here are the licenses of datasets and models used in this paper:

- **CC**: Chatbot Arena Conversations
- **CC-BY-SA-3.0**: Databricks-Dolly-15K
- **Apache-2.0**: OpenAssistant, ShareGPT
- **ODC-by**: WildChat
- **Others**: InstructionWild (Non-commercial Use), LLaMA-3 family (customized license), LMSYS-Chat-1M (Unknown),

B Performance on REALQUESTIONS

Below is the performance of LLaMA-3-8B trained on different human instruction datasets, where REALQUESTIONS ranks first.

Training Set	Alpaca Eval 2.0
OASST	4.51
Chatbot-Arena Convs.	5.17
UltraChat	6.20
ShareGPT	9.13
WildChat	14.62
REALQUESTIONS	16.77

Table 6: Performance of models trained on different datasets evaluated with Alpaca Eval 2.0.

C Training Details

Table below is the hyper-parameters used in finetuning models. For SFT, we use Megatron-LM and 8*8*Nvidia H100 GPUs. For DPO, we use Huggingface TRL and 8*Nvidia H100 GPU.

Parameters	Values
Epoch	3
Learning Rate	$2e^{-5}$
Global Batch Size	128
Gradient Accumulation	1
Gradient Checkpointing	False
Precision	BF16
Max Length	8192
Warmup Ratio	0.06
Weight Decay	0
Learning Rate Scheduler	Cosine

Table 7: Hyper-parameters of SFT.

Parameters	Values
Epoch	1
Learning Rate	$0.7e^{-6}$
Global Batch Size	128
Gradient Accumulation	8
Gradient Checkpointing	True
Precision	BF16
Max Length	8192
Warmup Ratio	0.1
Weight Decay	0
Learning Rate Scheduler	Cosine

Table 8: Hyper-parameters of DPO.

D Statistics of Datasets

Below is the basic statistics of common instructional datasets, where SYNTHQUESTIONS ranks top in tokens per turn and lexical diversity.

Source	Dataset Name	#Convs	#Turns	#Tokens/T	Lex-Div
Human	ShareGPT (Chiang et al., 2023)	62K	3.42	335	59.86
	Chatbot Arena Convs (Chiang et al., 2024)	33K	1.19	224	58.44
	InstructionWild (Ni et al., 2023)	110K	1	76	80.18
	OpenAssistant (Köpf et al., 2023)	48K	1.73	211	69.64
	WildChat (Zhao et al., 2024)	652K	2.52	519	86.58
	Databricks-Dolly (Conover et al., 2023)	15K	1	174	76.62
	LMSYS-Chat-1M (Zheng et al., 2024)	1M	2.01	248	59.94
Synthetic	Alpaca (Taori et al., 2023)	52K	1	74	64.08
	UltraChat (Ding et al., 2023)	208K	3.16	364	73.90
	Evol Instruct (Xu et al., 2023b)	143K	1	475	60.19
	SocraticChat (Kong et al., 2024)	50K	5.28	345	65.67
	OpenHermes (Teknium, 2023)	1M	1	346	60.10
	MAMmoTH2 (Yue et al., 2024)	10M	1	331	53.00
	GenQA (Chen et al., 2024)	11M	1.69	167	58.75
	SYNTHQUESTIONS	1M	1	802	77.19

Table 9: Statistics of different datasets. #Tokens/T represents number of tokens per turn. Lexical diversity is calculated with 10,000 samples randomly drew from the full dataset.

E Supplemental Visualization of Dataset Diversity and Complexity

Figure 9 presents the full visualization of dataset diversity and complexity of different synthesized datasets.




Figure 9: Left: t-SNE plot of SYNTHQUESTIONS along with Alpaca, EvolInstruct and UltraChat. The t-SNE plot of SYNTHQUESTIONS covers more area than other datasets. Right: Violin plots of the Arena Scores of all synthetic datasets.

Score	0	1	2	3	4	5	6	7
Percentage	0.77%	0.0011%	0.021%	0.22%	2.36%	12.93%	38.83%	44.88%

Table 10: The score distribution of the complete unfiltered dataset generated with our method.

F Dataset Safety

Table 11 shows the detection results from LLaMA-Guard-3-8B, which demonstrates that there are very few harmful content in our dataset.

Category	Specialized Advice	Intellectual Property	Defamation	Elections	Non-Violent Crimes	Sexual Content	Child Exploitation	Privacy	Violent Crimes	Hate	Indiscriminate Weapons	Self-Harm	Sex Crimes
Percentage	3.60%	0.158%	0.151%	0.098%	0.068%	0.059%	0.036%	0.034%	0.028%	0.027%	0.022%	0.013%	0.011%

Table 11: The percentages of detected potential harmful content in SYNTHQUESTIONS.

G Full Evaluation Results

Table 12 and Table 13 shows the full results of models on alignment benchmarks and closed-form benchmarks.

Model Groups	Models	Arena Hard		Alpaca Eval 2.0	
		WR(%)	LC(%)	WR(%)	SD
Larger or Proprietary Models	HumpBack-LLaMA2-70B	-	16.25	10.12	0.94
	GPT-3.5-Turbo-0301	18.1	18.09	9.62	0.89
	GPT-3.5-Turbo-1106	18.9	19.30	9.18	0.89
7B-8B Models w/ Proprietary Data	Mistral-7B-Instruct-v0.3	16.7 [†]	20.61	16.69	1.11
	Qwen-2-7B-Instruct [†]	23.5	21.86	19.62	1.15
	LLaMA-3-8B-Instruct	20.6	22.92	22.57	1.26
LLaMA3-8B w/ Open-source Data	UltraChat	3.6	8.29	5.44	0.71
	Evol Instruct	5.1	8.52	6.25	0.76
	ShareGPT*	4.8	9.13	6.55	0.79
	Tulu V2 Mix	8.7	9.91	7.94	0.86
	OpenHermes	4.4	9.94	6.27	0.73
	SocraticChat*	10.5	13.85	9.40	0.87
	WildChat	8.7	14.62	10.58	0.92
	GenQA	3.0	9.05	7.11	0.82
	MAmmoTH2	16.6	<u>18.5</u>	-	-
Ours	SYNTHQUESTIONS	<u>15.4</u>	18.87	19.15	1.15

Table 12: Performance of models trained on SYNTHQUESTIONS on Alpaca Eval 2.0 and Arena Hard. Among models with open-source data, the best performance is bolded and the second best performance is underlined. *marks results that are not reported and evaluated by us.

H Results of Qwen Model

Table 14 is the results of Qwen-2.5-7B trained on 100K, 200K and 400K subsets of SYNTHQUESTIONS. Due to increase in base model ability, model trained on only 400K data achieves performance near to LLaMA-3-8B with the full 1M dataset. Model performance also shows a positive relation with dataset scale when tested on Qwen model. Due to limited compute resource, we will complete this table with 800K and 1M dataset later.

Model Groups	Models	IFEVAL	MMLU	ARC-C	GPQA	GSM8K	MATH
7B-8B Models w/ Proprietary Data	Mistral-7B-Instruct-v0.3	54.65	61.84	63.57	27.8 [†]	43.37	12.94 [†]
	Qwen-2-7B-Instruct	56.79	70.5	59.73	25.3	82.3	49.6
	LLaMA-3-8B-Instruct	74.08	68.5	82.4	34.6	80.6	29.1
LLaMA3-8B w/ Open-source Data	UltraChat	-	65.23	62.12	-	50.57	-
	Evol Instruct	-	65.62	60.75	-	42.91	-
	ShareGPT	-	<u>66.03</u>	58.45	-	48.67	-
	Tulu V2 Mix	-	66.34	59.22	-	58.07	-
	OpenHermes	-	65.7	61.86	-	67.02	-
	WildChat	-	65.95	59.22	-	48.75	-
	GenQA	-	63.45	58.53	-	43.13	-
	MAmmoTH2	43.94 [†]	64.2	82.2	35.2	<u>70.4</u>	35.8
Ours	SYNTHQUESTIONS	57.05	65.79	<u>63.92</u>	<u>30.3</u>	70.53	<u>22.71</u>

Table 13: Performance of models on different close-ended knowledge and reasoning benchmarks. Among 7B-8B scales, the best performance is bolded and the second best performance is underlined. Results marked with † are evaluated by us.

Models	Alpaca Eval 2.0		
	LC(%)	WR(%)	SD
LLaMA-3-8B			
SYNTHQUESTIONS-1M	18.87	19.15	1.15
Qwen-2.5-7B			
SYNTHQUESTIONS-100K	17.25	16.03	1.10
SYNTHQUESTIONS-200K	17.89	14.64	1.06
SYNTHQUESTIONS-400K	18.19	16.04	1.11

Table 14: Performance of models trained on SYNTHQUESTIONS on Alpaca Eval 2.0.

I Criteria of Arena Hard

Criteria from Arena Hard

1. **Specificity:** Does the instruction ask for a specific output?
2. **Domain Knowledge:** Does the instruction cover one or more specific domains?
3. **Complexity:** Does the instruction have multiple levels of reasoning, components, or variables?
4. **Problem-Solving:** Does the instruction directly involve the AI to demonstrate active problem-solving skills?
5. **Creativity:** Does the instruction involve a level of creativity in approaching the problem?
6. **Technical Accuracy:** Does the instruction require technical accuracy in the response?
7. **Real-world Application:** Does the instruction relate to real-world applications?

J Prompts and Demonstrations

J.1 Prompt and Demonstrations Used in Attributing Step

Prompts Used in Attributing

[SYSTEM PROMPT]

Given a document and a query to an AI assistant.

1. You should link the document and the user query with a practical scene, considering user identity and motivation.
2. Decompose the query regarding ability, knowledge, output and extra information:

- Ability: The fundamental skills or capabilities required to address the problem.
- Knowledge: The relevant domain or subject matter related to the query.
- Output: The expected type of response or result.
- Extra information: Specific details or context from the scenario that ground the query in a real-world context (e.g., specific numbers, codes, or quotes from the document).

Here are some examples:

<example>

<document>

Etsy has become a leading online marketplace home to around 7.5 million active retailers, who recently generated \$1.7 billion worth of revenue in a single year alone. The Etsy marketplace is excellent for selling everything from handmade creations, like home decor and digital art, to vintage products.

But is it really possible to make money on Etsy, even as a beginner? The truth is that whether or not you'll be able to make money on Etsy will largely depend on how much time you're willing to invest in learning what it takes to become a successful Etsy seller.

Fortunately, starting your own Etsy store comes with plenty of beginner-friendly benefits. The online marketplace doesn't charge mandatory monthly fees and offers plenty of great resources to help you master the easy-to-use navigation platform.

We'll walk you through everything you need to know to set up your own Etsy store and start selling in no time. You'll also get the inside scoop on what differentiates successful shops from the competition.

Main takeaways from this article:

- Setting up an Etsy store is easy - the hard part is figuring out how to make money on Etsy. We'll walk you through what you need to know to become a successful seller.
- Clothing and textiles, jewelry, personalized items, homeware, and art & collectibles are among the top-selling product categories on Etsy.
- By launching your own Etsy shop, you can sell to an established audience, minimize payment processing hassles, avoid making significant upfront investments, and adopt a multichannel selling approach.
- Working with a print on demand partner can eliminate the need to worry about purchasing supplies, keeping up with inventory, or dealing with shipping.
- Using high-quality products and providing excellent customer service are vital components that can set your shop apart.
- While setting up your Etsy shop, business licenses, taxes and fees, and shipping costs are some requirements you must take

care of.

- Promoting your Etsy store through online marketing can greatly increase your odds of making money on Etsy.

</document>

<query>

Act as a an online business expert and tell me how I can use the information of the best selling products of my etsy store and use it to make more money, like listing in another website or something.

</query>

<scene>

The user might be a vendor who wants to increase the sales of his etsy store. He wants to advertise the best-selling products in his store, but has no idea where and how he can achieve this.

However, he does not need suggestions that are too general without detailed and actionable guidance. He wants to seek concrete suggestions from a business expert.

</scene>

<query_compositions>

Ability: Summarizing, Planning and Guiding.

Knowledge: Business, Online store, Advertising.

Extra Information: Etsy store

Output: A business plan or a concrete suggestion list.

</query_compositions>

</example>

<example>

<document>

Making money in stocks is usually a long-term game: Very few people make tons of money in stocks overnight. Here's how to sustainably grow your wealth with stocks.

How to make money in stocks

You can make money in stocks by opening an investing account and then buying stocks or stock-based funds, using the "buy and hold" strategy, investing in dividend-paying stocks and checking out new industries.

Open an investment account

Pick stock funds instead of individual stocks

Stay invested with the "buy and hold" strategy

Check out dividend-paying stocks

Explore new industries

</document>

<query>

You are an investment advisor, you will provide me with ideas of investments. You have \$100, and your only goal is to turn that into as much money as possible in the shortest time possible, without doing anything illegal. I will do everything you say and keep you updated on our current cash total. No manual labor.

```
</query>
<scene>
The user might be a high-school student who wants to make some
quick money to pay for his/her hobbies, but has not much
principle in pocket. The fastest way to make money is without
doubt investments, so he seeks investments that do not take
much principal but can earn money quickly without breaking the
laws. When asking the AI assistant for suggestions, he takes
$100 for an example to illustrate that he does not has much
money.
```

```
</scene>
```

```
<query_compositions>
```

```
Ability: Summarizing, Planning and Guiding.
```

```
Knowledge: Investment, Low cost investment, Business, Law.
```

```
Extra Information: $100
```

```
Output: An investment plan or suggestions
```

```
</query_compositions>
```

```
</example>
```

```
<example>
```

```
<document>
```

```
Have you ever considered the power of a one-page website?
```

```
Modern website designs lean towards minimalism; prioritizing
user experience with clean layouts, intuitive navigation, and
mobile-first thinking. Less is often more!
```

```
While multi-page website architecture emphasizes structure and
organization, the single-page website concept is all about
simplicity and focus. It places all the vital information about
your business or project on a single, scrollable page.
```

```
This can be very effective especially when you need to lead
visitors to a singular action without overwhelming them with
multiple pages.
```

```
In this blog post, you are going to learn how to create an
effective one-page website on WordPress.com that conveys its
core message and steers visitors to a specific action or
understanding.
```

```
Ready to get started?
```

```
</document>
```

```
<query>
```

```
Create a one-page website for a web development company named
Open Agency.
```

```
</query>
```

```
<scene>
```

```
The user might be a developer from a newly started web
development company named Open Agency. The company needs a one-
```

page website to introduce themselves, but they have not hired experts for advertising yet. As a result, the task of constructing the website is assigned to this developer. Unfortunately, he has no idea how to create such a one-page website, so he turns to an AI assistant for help with the query.

</scene>

<query_compositions>

Ability: Coding.

Knowledge: Web development, Advertising, Website creation.

Extra Information: None

Output: A brief code snippet for a one-page website.

</query_compositions>

</example>

<example>

<document>

(some codes ...)

The error log shown is:

```
torch.Size([2, 12, 12])
```

RuntimeError

Traceback (most recent

call last)

<ipython-input-22-d2f43f09fd01> in <module>()

```
74     status = 1 #F
```

```
75     while(status == 1): #G
```

```
---> 76         qval = model(state1) #H
```

```
77         qval_ = qval.data.numpy()
```

```
78         if (random.random() < epsilon): #I
```

3 frames

/usr/local/lib/python3.7/dist-packages/torch/nn/modules/linear.

py in forward(self, input)

```
101
```

```
102     def forward(self, input: Tensor) -> Tensor:
```

```
--> 103         return F.linear(input, self.weight, self.bias)
```

```
104
```

```
105     def extra_repr(self) -> str:
```

RuntimeError: mat1 and mat2 shapes cannot be multiplied (128x4 and 128x64)

mat1 should be the output of the convolutional network after it is flattened, and mat2 is the linear network following it.

Appreciate any help. Thanks!

</document>

<query>

I'm initializing my observation as np.zeros((111,)) and state representation is as follows: 109 Laser scan points, yaw and distance to goal total 111. I don't know why I'm getting the following error: [ERROR] [1684308219.676930, 2100.420000]: bad

```

callback: <bound method EvaderNode.scan_callback of <__main__.
EvaderNode object at 0x7f77a26aaca0>>
Traceback (most recent call last):
  File "/opt/ros/noetic/lib/python3/dist-packages/rospy/topics.
py", line 750, in _invoke_callback
    cb(msg)
  File "/home/cse4568/catkin_ws/src/pa2/src/evader_2.py", line
636, in scan_callback
    self.agent.train(32) # Set the batch size here
  File "/home/cse4568/catkin_ws/src/pa2/src/DQN.py", line 64, in
train
    target = reward + self.gamma * torch.max(self.q_target(torch
.tensor([next_state], dtype=torch.float32)))
  File "/home/cse4568/.local/lib/python3.8/site-packages/torch/
nn/modules/module.py", line 1110, in _call_impl
    return forward_call(*input, **kwargs)
  File "/home/cse4568/catkin_ws/src/pa2/src/DQN.py", line 27, in
forward
    return self.model(x)
  File "/home/cse4568/.local/lib/python3.8/site-packages/torch/
nn/modules/module.py", line 1110, in _call_impl
    return forward_call(*input, **kwargs)
  File "/home/cse4568/.local/lib/python3.8/site-packages/torch/
nn/modules/container.py", line 141, in forward
    input = module(input)
  File "/home/cse4568/.local/lib/python3.8/site-packages/torch/
nn/modules/module.py", line 1110, in _call_impl
    return forward_call(*input, **kwargs)
  File "/home/cse4568/.local/lib/python3.8/site-packages/torch/
nn/modules/linear.py", line 103, in forward
    return F.linear(input, self.weight, self.bias)
RuntimeError: mat1 and mat2 shapes cannot be multiplied (1x113
and 111x128)

```

And everytime it runs I'm getting different mat1 values. Find where I made the mistake and fix the code. You are welcome to make all the necessary changes and modifications to make it the best DQN implementation for my Autonomous robot navigation in maze like env. I already implemented the Evader node. You can modify the DQN to make it fit for the Evader:
(some codes ...)

</query>

<scene>

The user might be a student studying reinforcement learning, who is developing an algorithm based on DQN model. However, he is faced with an error "mat1 and mat2 shapes cannot be multiplied" in his code. He is not familiar with pytorch, so he copied his error log and codes to ask the assistant to debug for him.

</scene>

<query_compositions>

Ability: Coding, Debugging.

```
Knowledge: Python, PyTorch, Deep Learning.
Extra Information: A code snippet copied from the document (
Traceback...).
Output: The corrected code or suggestions on how to fix the bug.
</query_compositions>
</example>
```

[USER PROMPT]

Now imagine a practical scene which link the user query and the document. Describe such a scene with one brief paragraph, containing the user identity and the motivation. Then also decompose the query regarding ability, knowledge, extra information and output.

Remember you are not responding the query. Only output with the following JSON format without any additional explanation or chat :

```
{
  "scene": "xxx",
  "query_compositions": {
    "ability": "xxx",
    "knowledge": "xxx",
    "extra_information": "xxx",
    "output": "xxx"
  }
}
```

```
## Document
{document}
```

```
## Query
{query}
```

```
## Scene
```

J.2 Prompts Used in Synthesizing Step

Prompts Used in Synthesizing

[SYSTEM PROMPT]

You will be shown a document, you should imagine a scene where a user with a certain identity comes up with some query compositions and a query related to the document. Here are some examples:

```
{demos}
```

[USER PROMPT]

Now you should

1. Envision a real-world scenario based on the provided document

. Describe this scenario in one paragraph, detailing the logical steps from the document's content to a query directed at an AI assistant.

2. Then list the compositions of a query that could emerge from this scenario, including:

- Ability: The fundamental skills or capabilities required to address the problem.
- Knowledge: The relevant domain or subject matter related to the query.
- Output: The expected type of response or result.
- Extra information: Specific details or context from the scenario that ground the query in a real-world context (e.g ., specific numbers, codes, or quotes from the document).

3. Finally formulate a user query based on the scenario and query compositions you have identified. Ensure:

- Maximize the ability that is needed to solve the query. Avoid simple copying or extracting tasks.
- The query should be practical, complex and requires advanced skills. It should be challenging for the most capable AI.
- The query should be self-contained and answerable without additional resources.
- You must copy excerpts from the document into the query if extra information from the document is needed.
- As the AI assistant does not have search engine access, **avoid** creating queries that rely on external search engines.

When constructing query compositions and the final query, consider the following requirements:

- > Specificity: The query should ask for a specific output;
- > Domain Knowledge: The query should cover one or more specific domains;
- > Complexity: The query should have multiple levels of reasoning , compositions, or variables;
- > Problem-Solving: The query should directly involve the AI to demonstrate active problem-solving skills;
- > Creativity: The query should involve a level of creativity in approaching the problem;
- > Technical Accuracy: The query should require technical accuracy in the response;
- > Real-world Application: The query should relate to real-world applications.

Output the scene and query in JSON format. Before generating scene, query_composition and query, you should include your thought on how you design the real-world scenario and the query, so that each of the above requirements is satisfied.

Document


```

{document}

## Output Format
{{
  "thought": "xxx"
  "scene": "xxx",
  "query_compositions": {{
    "ability": "xxx",
    "knowledge": "xxx",
    "extra_information": "xxx",
    "output": "xxx"
  }},
  "query": "xxx"
}}

## Your Output

```

J.3 Prompts Used for Filtering Instructions

Prompts Used for Filtering Instructions

```

## Role
Prompt Evaluator

## Task
You will be given a prompt written for large language models,
and you should evaluate the prompt according to the provided
criteria.

## Evaluation Criteria
1. Specificity: Does the prompt ask for a specific output?
2. Domain Knowledge: Does the prompt cover one or more specific
domains?
3. Complexity: Does the prompt have multiple levels of reasoning
, compositions, or variables?
4. Problem-Solving: Does the prompt directly involve the AI to
demonstrate active problem-solving skills?
5. Creativity: Does the prompt involve a level of creativity in
approaching the problem?
6. Technical Accuracy: Does the prompt require technical
accuracy in the response?
7. Real-world Application: Does the prompt relate to real-world
applications?

## Rules
1. You should evaluate based on each aspects of the criteria
independently. First analyze the prompt according to each aspect
and then assign it with a score.
2. If a prompt satisfies one aspect, you should score it as 1.
Otherwise you should score it as 0.

```

3. Output your results with JSON dictionary format.

Output Sample

```
{
  "specificity": {"analysis": "analysis about specificity", "
score": n},
  "domain_knowledge": {"analysis": "analysis about domain
knowledge", "score": n},
  "complexity": {"analysis": "analysis about complexity", "
score": n},
  "problem_solving": {"analysis": "analysis about problem
solving", "score": n},
  "creativity": {"analysis": "analysis about creativity", "
score": n},
  "technical_accuracy": {"analysis": "analysis about technical
accuracy", "score": n},
  "real_world_application": {"analysis": "analysis about real-
world application", "score": n}
}
Here is the prompt to evaluate:
{prompt}
```

K Other Cases

Case 1

[Document]

You hear a lot about how creatives and account people don't get along, that couldn't be farther from the truth at this place. We eat lunch together, we make fun of each other and we laugh together. Soon we will be playing ping pong together. We are one big family. To my office family, thank you for making me feel so welcome in such a short amount of time. I can't wait to see what 2013 has in store for us.

[User]

A new employee in a creative agency.

[Motivation]

The user is trying to understand the company culture and dynamics. They want to know how to foster and maintain a positive and collaborative atmosphere in their own team.

[Instruction]

What are some actionable steps to foster a positive and collaborative team culture, like being 'one big family', in a creative agency, and how can we maintain it over time?

Case 2

[Document]

Your health your way Kids may not think to pay much attention to their health. As long as they can move around when they want and have a regular source of food their health and well being takes back seat to other more entertaining thoughts. But with their peers increasingly becoming

overweight and obese those concerns are starting to be pushed to the front of the line whether kids want to address the problem or not. Living healthy not only decreases chance of health risks but can also make a child's live much more well rounded and enjoyable. Living healthy isn't hard, it's all about making the right choices and creating good habits. In order to kick a healthy lifestyle into high gear kids have to get away from their video games and televisions and get moving.

[User]

A health-conscious parent.

[Motivation]

The user is browsing the internet for ways to encourage their child to adopt a healthy lifestyle. They come across an article emphasizing the importance of living healthy, especially for kids, and want to create a personalized plan for their child to make healthy choices and develop good habits.

[Instruction]

Design a personalized health plan for a 10-year-old child who spends most of their time playing video games and has a sedentary lifestyle. The plan should include a daily routine, dietary recommendations, and fun physical activities that can help them develop healthy habits and reduce the risk of obesity. Consider the child's age, interests, and current lifestyle when creating the plan. Provide a step-by-step guide on how to implement the plan and make adjustments as needed.

Case 3

[Document]

One of the many challenges that face senior citizens is keeping their home clean. A clean house can do wonders for your mental health and well-being, and no one should have to sacrifice the joy of a clean and presentable home because chores have become too difficult to complete. Chrisel's Affordable Care is your choice for in-home care and cleaning services in the Suffield, CT area. Often, keeping up with a loved one's house cleaning is left up to the family. Our job is to take the stress off you or your loved one's plate, so let us take care of those responsibilities for you. Chrisel's offers house cleaning packages that fit every budget. Contact us today to find out how we can help you. You don't want to let just anyone into your home. You trust Chrisel's Affordable Care with your loved ones, so why not trust us to keep their home neat and tidy? We're insured and bonded, and we offer free consultations. Every family has different needs. Chrisel's Affordable Care offers a basic cleaning service that covers your most immediate cleaning ones. This basic package includes:

[User]

A family caregiver.

[Motivation]

The user is exploring in-home care and cleaning services for their elderly loved one in the Suffield, CT area. They come across Chrisel's Affordable Care and want to determine the best cleaning package for their loved one's needs and budget.

[Instruction]

Considering the basic cleaning service package offered by Chrisel's Affordable Care, how can I create a customized cleaning package that meets my elderly loved one's specific needs and stays within their budget? Please provide a cost-benefit analysis of the recommended package and explain why it's the most suitable option.

Case 4

[Document]

The easiest way is to create a storyboard that change opacity of two images with blend. Then you can call the storyboard.begin() to run the storyboard you want for example from the tapped or click event.

[User]

A mobile app developer.

[Motivation]

The user is tasked with creating an engaging user interface for a new app. They want to implement an animation that transitions between two images with a blend effect. The developer has experience with storyboard animations but needs help with implementing this specific effect.

[Instruction]

Create a storyboard animation that transitions between two images with a blend effect, where the opacity of the images changes, and provide a code snippet or step-by-step guide on how to implement it in a mobile app.