

Constructing Domain-Specific Evaluation Sets for LLM-as-a-judge

Ravi Raju, Swayambhoo Jain, Bo Li,
Jonathan Li, Urmish Thakker

SambaNova Systems

{ravi.raju, swayambhoo.jain}@sambanovasystems.com

Abstract

Large Language Models (LLMs) have revolutionized the landscape of machine learning, yet current benchmarks often fall short in capturing the diverse behavior of these models in real-world applications. A benchmark’s usefulness is determined by its ability to clearly differentiate between models of varying capabilities (separability) and closely align with human preferences. Existing frameworks like Alpaca-Eval 2.0 LC (Dubois et al., 2024a) and Arena-Hard v0.1 (Li et al., 2024a) are limited by their focus on general-purpose queries and lack of diversity across domains such as law, medicine, and multilingual contexts. In this paper, we address these limitations by introducing a novel data pipeline that curates diverse, domain-specific evaluation sets tailored for LLM-as-a-Judge frameworks. Our approach leverages a combination of manual curation, semi-supervised learning to generate clusters, and stratified sampling to ensure balanced representation across a wide range of domains and languages. The resulting evaluation set, which includes 1573 samples across 14 categories, demonstrates high separability (84%) across ten top-ranked models, and agreement (84%) with Chatbot Arena and (0.915) Spearman correlation. The agreement values are 9% better than Arena Hard and 20% better than AlpacaEval 2.0 LC, while the Spearman coefficient is 0.7 more than the next best benchmark, showcasing a significant improvement in the usefulness of the benchmark. We further provide an open-source evaluation tool that enables fine-grained analysis of model performance across user-defined categories, offering valuable insights for practitioners. This work contributes to the ongoing effort to enhance the transparency, diversity, and effectiveness of LLM evaluation methodologies.

	Live	Annotator	Diverse	Practitioner control
Chatbot Arena	✓	👤	✓	✗
Alpaca-Eval 2.0 LC	✗	👤	✗	✗
Arena-Hard v0.1	✓	👤	✗	✗
Ours	✓	👤	✓	✓

Figure 1: Compared to other benchmark frameworks our approach introduces a data pipeline that curates unlabeled data into categories that contain domains/capabilities that the practitioner cares about. It has the capability to be refreshed with new data and is diverse compared to alternatives.

1 Introduction

Large Language Models (LLMs) have dramatically changed the landscape of machine learning research and have been incorporated in products for the past few years. Along with their rise, a multitude of benchmarks and frameworks (Liang et al., 2023) have been proposed to assess the capabilities of LLMs which include knowledge tasks such as MMLU (Hendrycks et al., 2021a), reasoning tasks like GSM8k (Cobbe et al., 2021) and more standard NLP tasks (Zellers et al., 2019; Narayan et al., 2018). However, these benchmarks fail to capture the behavior that a user experiences in a chat/generative applications. Typically, human evaluations are seen as a gold standard to determine which LLM responses are preferable over others in a chat setting but is time-consuming and expensive to conduct (Chiang et al., 2024).

To address this shortcoming, Zheng *et al.* introduced the concept of LLM as a judge as an automatic evaluator alternative, which uses another LLM the judging of model completions to another LLM such as GPT-4 or GPT-4o (Zheng et al., 2023a; OpenAI et al., 2024). Alpaca-Eval is another benchmark designed under the paradigm of LLM as an evaluator where a target LLM’s com-




Figure 2: Alpaca-Eval category breakdown




Figure 3: Arena-Hard v0.1 category breakdown

pletions are compared against a reference LLM’s output (the default being GPT-4 Turbo) and assigned a winrate against the reference (Li et al., 2023). It has seen widespread adoption since it is cheap, fast, and mitigates length bias (Chiang et al., 2024). Similarly, Arena-Hard v0.1 is recent benchmark which focuses on distilling the Hard category of Chatbot Arena into a smaller evaluation set (Li et al., 2024a). They use a topic clustering pipeline to cluster prompts with OpenAI’s embedding model (text-embedding-3-small) (OpenAI, 2024b) and score each cluster based on difficulty, creativity, and reasoning ability with GPT-3.5 Turbo. They also introduce a notion of separability (how well can a benchmark differentiate between models) and agreement with human preferences (i.e. ChatBot Arena) as a measure of benchmark quality.

Unfortunately, there are still some limitations with the current open-source LLM-as-a-judge framework. Alpaca-Eval 2.0 LC is dominated by general chat queries/instructions and has few prompts in domains such as coding, medical, finance, law and mathematics as shown in Figure 2. Arena-Hard v0.1 addresses some of these deficiencies by upweighting coding and mathematics prompts and restricting the general chat queries to 30% in the evaluation set. However, both evaluation sets are strictly in English therefore not accessing the model’s multilingual capability and have a smaller number of prompts in more niche categories like law and medicine. As models are acquiring more capabilities across various data types such as charts/tables, domains and languages, it becomes crucial to determine how to evaluate each model’s ability in a scalable manner.

In this paper, we attempt to address challenges from Alpaca-Eval 2.0 LC and Arena-Hard v0.1 by introducing more diversity across domain knowledge and languages. To accomplish this, we introduce a simple data pipeline methodology to create a new evaluation set designed for these specific contexts. First, we source prompts from various open source datasets (shown in Table 4) to ensure our evaluation set has high data diversity. For the next step, we generate embeddings from a subsample of each of these datasets using an embedding model. To label the corresponding embeddings, we manually curate a seed set of prompts and label them to human-defined specific categories, generate those embeddings and train a k-NN classifier which we can use to classify the unlabeled data that we sampled. In order to make sure that no cluster/category dominates, we employ stratified sampling to ensure balanced representation across all domains and languages in the evaluation set. We further refine the quality of the prompts by manual curation and ensure that each category has a sufficient number of prompts to mitigate the inherent variability in LLM-as-a-judge and ultimately end up with 1573 samples in the evaluation set.

There are several advantages to our approach as shown in Figure 1. Similar to Arena-Hard v0.1, our approach is robust to contamination as we can periodically run our data pipeline on the same data to get new samples or potentially even a new data mixture. As mentioned earlier, our methodology allows introduction of new datasets which enables diversity rather than offered by Arena-Hard v0.1 and Alpaca-Eval. In addition, our evaluation set more closely mirrors Chatbot Arena rankings; Figure 4 shows a visual comparison of model rankings.




Figure 4: Visual comparison between our method, Arena-Hard v0.1, and Alpaca-Eval 2.0 LC on 10 models on separability of winrates. Our method has fewer overlaps of confidence intervals than the other baselines.

In particular, our evaluation set places Gemini-1.5-Flash (DeepMind, 2024) over Gemma2 27B Instruct (Team, 2024) which aligns with ChatBot Arena rankings whereas the others rank Gemma2 27B over Gemini-1.5-Flash. Moreover, since we use open source models for the entire pipeline, practitioners can mold the pipeline and generate evaluation sets to test domains and capabilities they care about.

After we have obtained the evaluation set, we execute the same procedure as LLM-as-a-Judge by generating the outputs completions from GPT-4o and using them as reference to construct a leaderboard from ten various open and closed-sourced models. With this labeling approach, we are able to breakdown the composition of prompts into various categories and report category win rates. We plan to release an evaluation tool which displays the category winrate for all models on the leaderboard and an explorer which displays both the target model as well as the reference model’s completions for a prompt and the reasoning given by the LLM judge. This analysis tool allows users to obtain fine-grained insights on where different models succeed and fail for their particular use-case.

Our main contributions can be summarized below:

- We introduce a new methodology that enables creation of a benchmark that tests for diverse skill sets of models. We open-source our evaluation infrastructure so practitioners can view how different models perform on separate tasks according to how they define

their categories. This fine-grained breakdown allows the practitioner to select models that work well for their particular use case.

- Our benchmark creation methodology encourages more diversity and transparency to the practitioner compared to other alternatives. In comparison to other baselines like Alpaca-Eval and Arena-Hard v0.1, our benchmark has 84% separability, 84% agreement with confidence interval (95%) with respect to Chatbot Arena rankings, 0.915 Spearman’s correlation coefficient with respect to Chatbot Arena rankings and 0.04 Brier Loss Score.
- We also analyze the aforementioned metrics on our evaluation set with 4 LLM judges: GPT-4o (OpenAI et al., 2024), GPT-4o-mini (OpenAI, 2024a), Llama 3.1 405B Instruct and Llama 3.1 70B Instruct (Dubey et al., 2024). Our overall findings suggest that while open-source models can be used to separate between model rankings, agreement with Chatbot Arena model rankings is roughly 10% (405B) and 20% (70B) than GPT-4o.

2 Related Work

At their core, benchmarks are tool to estimate LLM capabilities. There are many different flavors of benchmarks, spanning either across domains or various tasks. Some popular benchmarks include: Boolq (Clark et al., 2019), MMLU (Hendrycks et al., 2021a), GSM8k (Cobbe et al., 2021), MATH (Hendrycks et al., 2021b), XSUM (Narayan et al., 2018), Hellaswag (Zellers et al., 2019), and MGSM (Shi et al., 2022). An expanded framework of static benchmark is AutoBencher which automatically creates new benchmarks which finds holes in knowledge of current SOTA LLMs (Li et al., 2024b).

These types of benchmarks have ground-truth references and compare how closely the LLM’s completion aligns with those references. An inherent limitation with static benchmarks is that they are hosted on the internet and thus are susceptible to test leakage contamination (Sainz et al., 2023; Yang et al., 2023). The other style of benchmarking relies on constructing a human evaluation trials on a set of evaluation prompts. Due to the expensive nature of human evaluation, a recent, cheaper alternative is to use SOTA LLMs to evaluate model completions either through single score or pairwise

comparison with a reference answer, popularly referred to as LLM-as-a-Judge (Li et al., 2023; Zheng et al., 2023a; Li et al., 2024a; Dubois et al., 2024b; Verga et al., 2024).

This motivates the need for "live, refreshable" benchmarks so that the integrity of the benchmark can be maintained. LiveBench is a framework which sources data from arXiv papers, news articles, and datasets to periodically replace the stale prompts (White et al., 2024). Chatbot arena is an open platform that allows online users to send prompts to two different models and compare/contrast the models' response (Chiang et al., 2024). Users can then vote on which completion was superior. Other live benchmarks include DynaBench (Kiela et al., 2021), LiveCodeBench (Jain et al., 2024), and R2E (Jain et al.). Our work lies in the intersection between LLM-as-a-Judge and live benchmarks as our data pipeline enables periodic refreshing of the evaluation set from existing clusters. Furthermore, our data pipeline is fairly general as it can consume a variety of diverse datasets (relative to Arena-Hard v0.1 and Alpaca-Eval), consists of using open-source models, and is flexible enough to work on the user's desired data.

3 Methodology

In this section, we describe our approach to creating novel evaluation set using LLM-as-a judge. We enumerate the datasets that we source from to create our unlabeled corpus and subsequently describe our data pipeline for generating the evaluation set.

3.1 Data Sources

We use data sources from a variety of source to ensure we cover a variety of domains as well as languages. The domains we target can be broadly classified as the following: medical, law, finance, mathematics and coding. The languages we cover are standard but also more esoteric: Japanese (ja), Arabic (ar), Thai (th), Hungarian (hu), Russian (ru), Serbian (sr), Slovenian (sl), and Turkish (tr). Prompts that don't neatly fit into these groups fall into a catch-all general category. A complete list of all the data we use can be found in Table 4 in the Appendix.

3.2 Data pipeline

Our data pipeline can be divided into 3 distinct steps, as shown in Figure 5. We first take the data corpus and use an embedding model to generate

their corresponding embedding. Each embedding encapsulates some level of semantic understanding of its associated prompt, and nearby embeddings typically encode similar semantic information.

To generate the labels for the unlabeled data, we take inspiration from semi-supervised learning (Hady and Schwenker, 2013). We manually define a set of categories, curate a seed set of prompts which fall into those categories (assigning them distinct labels) and embed those prompts with the aforementioned embedding model. We train a k -NN model (Mucherino et al., 2009) on top of those embeddings and use the k -NN to label the larger unlabeled corpus.

The final step in our pipeline involves applying stratified sampling (Parsons, 2017) to each cluster. The reason for this last step is that we want our evaluation set to retain diversity of our larger data corpus rather than uniform random sampling. For each category, we sub-sample 100 prompts from the aggregate clusters and disregard clusters which have a lower count than the number of prompts we sampled. To obtain our final evaluation set, we manually curate the remaining prompts to ensure high quality, varied task capability and data diversity.

4 Experimental Setup

In this section, we discuss finer details about the data pipeline we mentioned in the prior section. experimental setup on a set of ten highly rated models¹ as well as defining the metrics which determine the quality of the benchmark.

4.1 Data pipeline details

For the data pipeline, we use semi-supervised learning via a k -NN classifier. We consider 13 categories comprising of domains: finance, law, medical, maths, coding and languages: Arabic, Russian, Serbian, Hungarian, Japanese, Thai and Slovenian. We follow usual supervised training and via hyperparameter sweep over validation set yield $k = 40$ as the best value of k .

To generate the embeddings of the unlabeled data collected, we use the e5-mistral-7b-instruct embedding model (Wang et al., 2024) for its strong performance on the Massive Text Embedding

¹gpt-4o-2024-05-13, claude-3-5-sonnet-20240620, claude-3-opus-20240229, gemini-1.5-flash-latest, google/gemma-2-27b-it, Meta-Llama-3-70B-Instruct, claude-3-sonnet-20240229, Qwen/Qwen2-72B-Instruct, Meta-Llama-3-8B-Instruct, Mixtral-8x7B-Instruct-v0.1

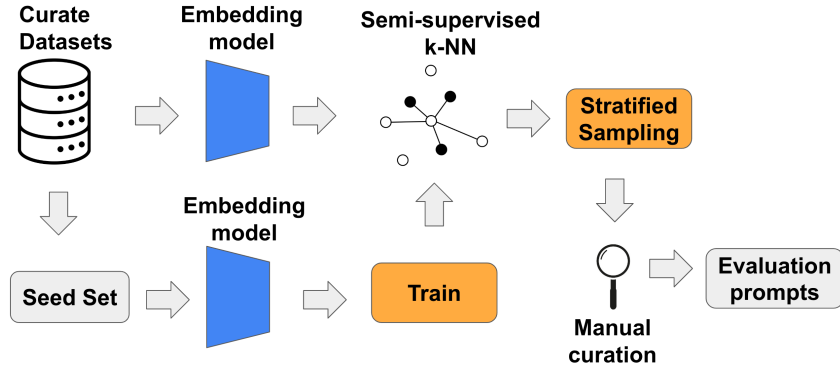


Figure 5: **Data pipeline:** After aggregating the prompts from datasets, we generate embeddings using a text embedding model. We set aside a set of prompts to use as a seed set for training the k-NN, label them into each category we care about, and generate their corresponding embeddings to train the k-NN with the embedding model. Subsequently, we classify the unlabeled data with our trained k-NN to create clusters of categories. We balance the clusters with stratified sampling and then manually curate the remaining prompts by removing overly long prompts (greater than 5000 words) and checking for low-quality content (nonsense prompts, NSFW etc.) to obtain the final evaluation set.

Benchmark (MTEB) Leaderboard (Muennighoff et al., 2022) and multilingual capability. If the k -NN encounters a sample which it is not familiar with or uncertain to label, we want those samples to be classified as general prompts. We use entropy of k -NN classifier probabilities of various categories for a given prompt as the measure of uncertainty. If entropy is too high, we bucket the sample into the default/general category (Settles, 2010). We set the entropy threshold to be 1.5 based on careful error analysis on the validation set.

After labeling with k -NN, we conducted stratified sampling within each cluster, selecting 100 samples for curation. We then filtered out excessively long prompts (longer than 5000 words) that could overwhelm the judge’s context window. Additionally, we reviewed the remaining prompts to eliminate those that were nonsensical or of low quality. During the evaluation, we observed that categories with a small number of examples had a significant impact on the category’s win rate. The inherent variability of the LLM-as-a-Judge evaluation, even with a fixed random seed and temperature set to 0.0, made it challenging to discern which model performed better in those categories. To mitigate this uncertainty, we ensured that any category with fewer than 90-100 examples was supplemented with additional data, enabling us to obtain meaningful and interpretable results. Our final evaluation set comprises 1573 examples.

4.2 LLM-as-a-Judge Details

We follow a similar scoring setup as Arena-Hard (Li et al., 2024a) and Alpaca-Eval (Dubois et al., 2024a) where we use GPT-4o as a judge model and GPT-4o as a reference model as well. For each model we want to test, we obtain the completions and ask GPT-4o to record which model response is better for the input prompt. In order to mitigate positional bias, we swap the completions between the model we are evaluating and the reference on a coin flip.

4.3 Obtaining Confidence Intervals

We follow the setup outlined in Li et al. (Li et al., 2024a; Chiang et al., 2024). We use the Bradley-Terry model in order to model the preference distribution between models on the leaderboard and the reference model (GPT-4o in our case). We aggregate preference pairs between models and perform 100 rounds of bootstrapping to obtain 95% confidence intervals for each model ranking.

We conduct the same analysis with annotations, denoting for each prompt which model response was preferred, from the Alpaca-Eval repo to obtain mean ELO rankings and 95% confidence intervals according to their leaderboard. Since similar artifacts (model preference comparisons) are not updated on Arena-Hard v0.1, we take the model winrates (ELO scores not listed) and 95% confidence intervals from their repo². For Chatbot Arena, we do the same thing and took model winrates/ELO

²7/26/2024

scores as well as the confidence intervals from the website³ as a source of ground truth.

4.4 Metrics

There are four different metrics we use to judge the efficacy of a benchmark. The first of these is Spearman’s correlation coefficient, which measures the rankings order between the two benchmarks. The other metrics are: separability, agreement with Confidence Interval (CI), and Brier Score. Separability refers to how well the benchmark can separate various models with high confidence. In particular, if on benchmark A model M1 has a higher ELO/winrate than model M2 and C_M refers to the confidence intervals of model M, S is a binary variable indicating if benchmark A is able to separate between model M1 and M2, $S = \mathbf{1}_{C_{M1} \cap C_{M2} = \emptyset}$. The separability is then calculated as a ratio over all possible model pairs. Agreement with CI measures how well benchmarks A and B confidently distinguish between two models with the same ordering. The Brier Score evaluates an LLM benchmark’s ability to predict the ranking of a pair of competing models, rewarding confidence in accurate predictions and penalizing confidence in incorrect ones. More details behind these metrics can be found in (Li et al., 2024a). Ultimately, we want our benchmark to align with Chatbot Arena as that is seen as an oracle for modeling human preferences.

5 Results

5.1 Separability, Agreement with CI (95%), Pair Brier Score

Our main results can be found in Table 1. With the exception of Chatbot Arena, our benchmark’s separability is 84.4% compared to other baselines like Arena-Hard v0.1 (80%) and Alpaca-Eval 2.0 LC (73.33%), which shows that our benchmark can better differentiate amongst different models.

One interesting datapoint regarding separability is Chatbot Arena’s score of 100% which may be attributed to a combination of two factors: 1) Chatbot Arena has more battles than any of the benchmarks listed in Table 1 and 2) Chatbot Arena includes battles between many different models rather than fixing a reference model like the other benchmarks. By providing the Bradley-Terry model bootstrapping process with more varied battles, Chatbot Arena is able to produce tighter confidence intervals, suggesting a future avenue for investigation

is whether confidence estimation should include multiple reference answers during judging to more closely simulate Chatbot Arena.

Our benchmark showed an 84.44% agreement with CI with respect to Chatbot Arena, which is higher than Arena-Hard v0.1’s 75.50% and Alpaca-Eval 2.0 LC’s 64.44%. This demonstrates that our benchmark has higher alignment with respect to Chatbot Arena which is supposed to be approximation of human preferences. In addition, our benchmark has a Spearman’s correlation coefficient of 0.915, indicating a strong correlation in rankings order compared to Alpaca-Eval 2.0 LC’s 0.2969. While our leaderboard ranking consists of 10 models, the pool of models we have included are the latest SOTA models that have been released so as to have the maximum amount of overlap possible. Finally, our benchmark scored a Brier score of 0.0417, which is lower than Alpaca-Eval 2.0 LC’s 0.0937, demonstrating better confidence in accurate predictions.

5.2 Diversity

Due to our data sources being quite diverse rather than simply just ChatBot Arena (Chiang et al., 2024), we are able to have more diversity in our evaluation set. To demonstrate this, we label Arena-Hard v0.1 with our kNN model using the entropy threshold to get a distribution of categories in that evaluation set. As shown in Figure 3, there is an over-representation of coding prompts, which comes from a byproduct of their data pipeline filtering for the hardest, highest quality which skews towards coding. Similarly, Alpaca-Eval’s prompt distribution shown in Figure 2 demonstrates that there is a large emphasis on general chat queries, along with some coding and math prompts while medical and law prompts are relatively underrepresented.

Our evaluation set breakdown in Figure 6(a) which covers more domains than the baseline, such languages like Arabic, Japanese, Hungarian and more. The close to equal distribution amongst the categories is likely due to the effect stratified sampling. We compare how our evaluation set category breakdown compared with LM-SYS Conversations (using our k-NN labeling approach) (Zheng et al., 2023b) in Figure 6(b), which is a snapshot of cleaned Chatbot Arena conversations from April to June 2023. In Figure 6(b), "Other" refers to the languages our k-NN classifier recognizes but groups them together collectively. We note that this distri-

³7/25/2024

	Chatbot Arena	Arena Hard v0.1	Alpaca-Eval 2.0 LC	Ours
Separability	100%	80%	73.33%	84.44%
Agreement with CI (95%)	N/A	75.50%	64.44%	84.44%
Spearman’s Correlation	N/A	0.187	0.2969	0.915
Brier Score	N/A	N/A	0.0937	0.0417

Table 1: Main results comparing the various benchmarks.

bution looks similar to Alpaca-Eval and the general category may contain additional languages not recognized by the classifier so it may have exceeded the entropy threshold.

5.3 Category Separability

Due to our unique ability to categorize the prompts, we can compute category separability for all the various categories in our evaluation set. Across 14 different categories, we do the same bootstrapping procedure on the category data to obtain the mean winrate/ELO and 95% CI, shown in Table 2. In general, there is a drop in separability when we look both at ELO ratings and winrate due to each category having a lower number of samples and thus larger CIs as a result.

The category-wise separability can act as an indicator which categories are superior at testing out the performance of models. Interestingly, across ELO and winrate rankings, Hungarian has the best separability of all categories, achieving 66.67% and 75.56% respectively. The medical category seems to be lowest separability around 55.56% and 68.89% respectively. The separability also indicates to use which categories we may need to add more samples to improve the confidence intervals.

5.4 Using different judges

We conduct an ablation of judge models on our evaluation, as we want to understand the effect of judge models on separability, Agreement with CI (95%) and Brier Score. We consider GPT-4o mini as one of the judges to be a small-closed source foil to GPT-4o. The other judges that we consider are open source models such as: Llama 3.1 405B instruct (using SambaNova’s developer API)⁴ and Llama 3.1 70B Instruct-Turbo⁵. We follow the same setup as gpt-4o with these other judge models.

Our results are shown in Table 3. In terms of separability, GPT-4o-mini and 405B get 82.2% and 70B get 84% separability, comparable to GPT-4o’s

⁴cloud.sambanova.ai


⁵<https://api.together.ai/models/meta-llama/Meta-Llama-3.1-70B-Instruct-Turbo>

Category	Ranking winrate	Ranking ELO
ar	73.33%	57.78%
ru	71.11%	55.56%
finance	75.56%	57.78%
sr	71.11%	53.33%
tr	73.33%	55.56%
general	77.78%	55.78%
hu	75.56%	66.67%
ja	71.11%	57.78%
medical	68.89%	55.56%
law	73.33%	51.11%
th	71.11%	57.78%
coding	73.33%	55.56%
sl	77.78%	53.33%
math	73.33%	55.56%


Table 2: Winrate and ELO separability for different categories

separability. 405B and GPT-4o-mini attain similar Agreement with CI (95%) close to 76% while 70B is almost 10 points lower; GPT-4o is the clear winner having the highest agreement with CI (95%). With the exception of 70B, all models get similar Brier Scores indicating that the Bradley-Terry models used to generate the rankings on confidence intervals for each judge are similarly confident. 70B’s high Brier score (relative to other judges), in addition to Agreement with CI, indicates that it poor judge than the other listed in Table 3.

The Spearman’s correlation coefficient (with respect to ChatBot Arena rankings) seems to indicate that GPT-4o-mini, Llama 3.1 405B, and 70B are poor judges getting a correlation of only 0.0787 vs. GPT-4o’s 0.915. Looking at Figure 7, it seems this aberration comes from both judges rating Claude Sonnet 3.5 over GPT-4o, Llama 3 70b over Claude Opus and Gemma2 27B over Gemini 1.5 Flash. Of course, Spearman’s correlation only measures correlation the final rank order of models with respect to ChatBot Arena and is a strictly weaker metric than Agreement with CI (95%). This finding seems to suggest while weaker closed-source models (like



(a) Our evaluation set category breakdown



(b) LMSys Conversations category breakdown

Figure 6: We look at the category breakdown on our evaluation set (Figure 6(a)) compared to LM-SYS Conversations (Figure 6(b)). We can clearly see that our evaluation set covers more languages and niche domains such as law and medical categories are a higher percentage of our evaluation set.

	GPT-4o	GPT-4o-mini	Llama 3.1 405B	Llama 3.1 70B
Separability	84.44%	82.22%	82.22%	84.44%
Agreement with CI (95%)	84.44%	76.77%	75.55%	66.66%
Spearman’s Correlation	0.915	0.0787	0.0787	0.0787
Brier Score	0.0417	0.062	0.0603	0.0955

Table 3: Comparing various judges on our evaluation set.

GPT-4o-mini) and open-source judge models seem to be able to separate other models based on capability, they still lack the preciseness that GPT-4o offers to align with rankings from Chatbot Arena.




Figure 7: Visual comparison of different judge’s separability on our benchmark.

6 Limitations/Future Work

There are certain limitations to our work. Currently, the categories we enumerate in our data pipeline is manually specified by humans and significant curation is done to ensure high quality prompts; for future work, we want to expand to using LLMs as category generators as well as quality checkers to automate the human effort out of this pipeline. Moreover, the diversity of prompts in the multilingual categories could be limited as we consider bucket all subdomains of a language into the same category. Sub-categorization of domains in non-English languages is left for future work.

For improving our leaderboard, we wish to add more models to be more representative of the entire spectrum of other leaderboards and future increasing the quality of the Bradley-Terry models we use to obtain the model’s confidence intervals. In order to improve category separability, we look to creating a methodology on figuring out the minimum number of samples required to improve separability.

The other aspect of future work relies to details regarding LLM-as-a-judge evaluation. Typically, the judge models are ablated but less explored is

the quality of the reference answer and whether one can use a weaker model instead of a stronger one to see if metrics are maintained. Current metrics define how separable a benchmark is and how much it aligns with human preferences but fails to account for the composition and diversity of the underlying data. For future work, we seek to quantify the diversity of each benchmark to understand how many capabilities/domains it spans.

7 Conclusion

We introduce a data pipeline that leverages via semi-supervised learning with a k-NN to enable practitioners to create benchmarks on their own data for targeted domains. Through evaluations of ten various closed and open-sourced models, we demonstrated that our benchmark achieves higher separability and agreement with CI with respect to Chatbot Arena, nearly 5 and 10 percentage points higher than the next best baseline, respectively. Our benchmark covers a wide variety of topics such as finance, medicine, legal and different languages absent in other LLM as a judge benchmarks. We hope that LLM developers can use our data pipeline to create their own benchmarks to evaluate their models for their particular use-case.

References

- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. [MathQA: Towards interpretable math word problem solving with operation-based formalisms](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2357–2367, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- 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.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [Boolq: Exploring the surprising difficulty of natural yes/no questions](#). *Preprint*, arXiv:1905.10044.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- DeepMind. 2024. [Gemini flash](#). Accessed: 2024-08-14.
- 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 Collet, 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 Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei 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 Grattafori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baeovski, 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 Leonhardt, 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, Changan 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, Dingkang 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, Han-

wen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, 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 Khandelwal, 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 Khabza, 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 Hermoso, Mo Metanat, Mohammad Rastegari, Munish 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, Pavan 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 Maheswari, 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 Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, 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, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiaojian 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 Tatsunori B. Hashimoto. 2024a. [Length-controlled alpacaeval: A simple way to debias automatic evaluators](#). *Preprint*, arXiv:2404.04475.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2024b. [Alpaca-farm: A simulation framework for methods that learn from human feedback](#). *Preprint*, arXiv:2305.14387.
- Gaurang Bharti. 2024. [finance-alpaca \(revision 51d16b6\)](#).
- Mohamed Farouk Abdel Hady and Friedhelm Schwenker. 2013. *Semi-supervised Learning*, pages 215–239. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. [Livecodebench: Holistic and contamination free evaluation of large language models for code](#). *Preprint*, arXiv:2403.07974.
- Naman Jain, Manish Shetty, Tianjun Zhang, King Han, Koushik Sen, and Ion Stoica. R2e: Turning any github repository into a programming agent environment. In *ICML 2024*.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577.
- Jon Durbin. 2024. [airoboros-gpt4-1.2](#).
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in nlp](#). *Preprint*, arXiv:2104.14337.
- Anastassia Kornilova and Vladimir Eidelman. 2019. [BillSum: A corpus for automatic summarization of US legislation](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 48–56, Hong Kong, China. Association for Computational Linguistics.
- Jonathan Li, Rohan Bhambhoria, and Xiaodan Zhu. 2022. [Parameter-efficient legal domain adaptation](#). In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 119–129, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E. Gonzalez, and Ion Stoica. 2024a. [From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline](#). *Preprint*, arXiv:2406.11939.
- Xiang Lisa Li, Evan Zheran Liu, Percy Liang, and Tatsunori Hashimoto. 2024b. [Autobench: Creating salient, novel, difficult datasets for language models](#). *Preprint*, arXiv:2407.08351.
- 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.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. [Holistic evaluation of language models](#). *Preprint*, arXiv:2211.09110.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Truthfulqa: Measuring how models mimic human falsehoods](#). *Preprint*, arXiv:2109.07958.
- Kelly Marchisio, Wei-Yin Ko, Alexandre Bérard, Théo Dehaze, and Sebastian Ruder. 2024. [Understanding and mitigating language confusion in llms](#). *Preprint*, arXiv:2406.20052.
- Antonio Mucherino, Petraq J. Papajorgji, and Panos M. Pardalos. 2009. *k-Nearest Neighbor Classification*, pages 83–106. Springer New York, New York, NY.
- Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. [Mteb: Massive text embedding benchmark](#). *arXiv preprint arXiv:2210.07316*.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). *Preprint*, arXiv:1808.08745.

OpenAI. 2024a. *Gpt-4o mini: Advancing cost-efficient intelligence*. Accessed: 2024-08-14.

OpenAI. 2024b. *New embedding models and api updates*. Accessed: 2024-08-14.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, 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, Rameez Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov,

Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. *Gpt-4 technical report*. *Preprint*, arXiv:2303.08774.

Van Parsons. 2017. *Stratified Sampling*.

Nazneen Rajani, Lewis Tunstall, Edward Beeching, Nathan Lambert, Alexander M. Rush, and Thomas Wolf. 2023. No robots. https://huggingface.co/datasets/HuggingFaceH4/no_robots.

Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, and Eneko Agirre. 2023. *NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10776–10787, Singapore. Association for Computational Linguistics.

Burr Settles. 2010. Active learning literature survey.

Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2022. *Language models are multilingual chain-of-thought reasoners*. *Preprint*, arXiv:2210.03057.

Shivalika Singh, Freddie Vargus, Daniel Dsouza, Börje F. Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura OMahony, Mike Zhang, Ramith Hettiarachchi, Joseph Wilson, Marina Machado, Luisa Souza Moura, Dominik Krzemiński, Hakimeh Fadaei, Irem Ergün, Ifeoma Okoh, Aisha Alaagib, Oshan Mudannayake, Zaid Alyafeai, Vu Minh Chien, Sebastian Ruder, Surya Guthikonda, Emad A. Alghamdi, Sebastian Gehrmann, Niklas Muennighoff,

Max Bartolo, Julia Kreutzer, Ahmet Üstün, Marzieh Fadaee, and Sara Hooker. 2024. [Aya dataset: An open-access collection for multilingual instruction tuning](#). *Preprint*, arXiv:2402.06619.

Gemma Team. 2024. [Gemma](#).

Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.

Pat Verga, Sebastian Hofstatter, Sophia Althammer, Yixuan Su, Aleksandra Piktus, Arkady Arkhangorodsky, Minjie Xu, Naomi White, and Patrick Lewis. 2024. [Replacing judges with juries: Evaluating llm generations with a panel of diverse models](#). *Preprint*, arXiv:2404.18796.

Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. [Improving text embeddings with large language models](#). *Preprint*, arXiv:2401.00368.

Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddhartha Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, and Micah Goldblum. 2024. [Livebench: A challenging, contamination-free llm benchmark](#). *Preprint*, arXiv:2406.19314.

Shuo Yang, Wei-Lin Chiang, Lianmin Zheng, Joseph E. Gonzalez, and Ion Stoica. 2023. [Rethinking benchmark and contamination for language models with rephrased samples](#). *Preprint*, arXiv:2311.04850.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) *Preprint*, arXiv:1905.07830.

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](#). *Preprint*, arXiv:2306.05685.

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. 2023b. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.

A Appendix

A.1 Data Sources

Table 4 includes various datasets across multiple domains such as medical, legal, financial, and multilingual categories. These sources were selected

to ensure a wide range of coverage, contributing to the diversity of the evaluation set. The datasets listed here were crucial for constructing the domain-specific evaluation sets, allowing for the thorough testing of models across different contexts and languages.

A.2 Judge Template

Below is our judge template that we used for our LLM-as-a-judge evaluation:

Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below. You should choose the assistant that follows the user’s instructions and answers the user’s question better, as well as answering in the desired language of the user. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of their responses. Begin your evaluation by comparing the two responses and provide a short explanation. Avoid any position biases and ensure that the order in which the responses were presented does not influence your decision. Do not allow the length of the responses to influence your evaluation. Do not favor certain names of the assistants. Be as objective as possible. Your evaluation should only focus on the correctness of the response. After providing your explanation, output your final verdict by strictly following this format: `[[A]]` if assistant A is better, `[[B]]` if assistant B is better, and `[[C]]` for a tie.

For the judge prompt, we used the default prompt from the MT-Bench work with one notable change (Zheng et al., 2023a). When we evaluated multilingual prompts with LLM-as-a-judge, the judge at times incorrectly awards wins to models which don’t necessarily follow instructions. Given the sentence "Please respond 'How does the economy work?' in Hungarian," two models might respond differently: 1) one provides a detailed English response with bulleted lists, while 2) the other responds concisely in Hungarian. The judge model will rate the model answering in the incorrect language higher, which is clearly not a measure of the model’s multilingual capability (Marchisio et al., 2024). In order to reduce these incorrect decisions, we modified the judge prompt to specifically penalize responses that respond to the prompt in the incorrect language.

In addition to issues with multilingual queries, we also note specifically for coding that GPT-4o seems to prefer models which provide detailed ex-

Datasets

LMSys Chatbot Arena (Chiang et al., 2024)
PubMedQA (Jin et al., 2019)
MathQA (Amini et al., 2019)
No Robots (Rajani et al., 2023)
Aya (Singh et al., 2024)
Legal reddit (Li et al., 2022)
Legal Summ. BillSum (Kornilova and Eidelman, 2019)
Airoboros-gpt4 (Jon Durbin, 2024)
Finance Advisor (Gaurang Bharti, 2024)
Finance Bier QA (Thakur et al., 2021)
MMLU (Hendrycks et al., 2021a)
TruthfulQA (Lin et al., 2022)
GSM8K (Cobbe et al., 2021)

Table 4: Dataset Sources used as input to the data pipeline in Figure 5.

planations to the code even if the code provided is of lower quality compared to a model which has better code quality but is not as verbose. This leads to scenarios where models that have chat but lower benchmark performance (e.g. HumanEval (Chen et al., 2021)) obtain higher winrate than models which are objectively better on coding prompts. To circumvent this issue, we explicitly prompt GPT-4o that it should focus on the correctness of the response as opposed to the style of the response.

A.3 Evaluation Tool

With the notion of self-defined categories and using the LLM-as-a-judge framework, we create an evaluation tool which loads an internal leaderboard from a csv file and breaks down the winrate into several categories the user defined. The UI shows the leaderboard in a dataframe and shows the winrates in set of bar plots across different categories. A screenshot of the tool can be seen in Figure 8.

There is also a feature which enables the user to view completions on the evaluation from both the model the user is interested in, the reference model, and the judge model to examine its reasoning. This tool enables the user to examine where the model they are developing is performing better than other competitors and areas where improvement is required.




Figure 8: A screenshot of our evaluation tool.