

---

# LLM Routing with Benchmark Datasets

---

Tal Shnitzer<sup>\*1</sup>, Anthony Ou<sup>\*2</sup>, Mírian Silva<sup>3</sup>, Kate Soule<sup>3</sup>, Yuekai Sun<sup>4</sup>,  
Justin Solomon<sup>2</sup>, Neil Thompson<sup>2</sup>, Mikhail Yurochkin<sup>3</sup>

<sup>1</sup>Broad Institute, <sup>2</sup>MIT, <sup>3</sup>MIT-IBM Watson AI Lab, <sup>4</sup>University of Michigan

## Abstract

There is a rapidly growing number of open-source Large Language Models (LLMs) and benchmark datasets to compare them. While some models dominate these benchmarks, no single model typically achieves the best accuracy in all tasks and use cases. In this work, we address the challenge of selecting the best LLM out of a collection of models for new tasks. We propose a new formulation for the problem, in which benchmark datasets are repurposed to learn a “router” model for this LLM selection, and we show that this problem can be reduced to a collection of binary classification tasks. We demonstrate the utility and limitations of learning model routers from various benchmark datasets. The extended version of the paper is available at this [arXiv link](#).

## 1 Introduction

The popularity of LLMs and their many potential uses prompted the development of comprehensive benchmarks, i.e., datasets representing different tasks and domains to compare LLMs. For example, HELM [1] consists of 42 scenarios covering a variety of uses, MMLU [2] is a multiple-choice question answering with 57 tasks organized by topics, among others [3, 4]. While there always will be an LLM that is the best *on average* across benchmarks, there is unlikely ever to be a model that is strictly the best *on each* of the hundreds of datasets comprising various benchmarks.

In this paper, we study the problem of identifying the best LLM for a new task. To learn about the strengths and weaknesses of candidate LLM, we use benchmark datasets that give insights into the performance across tasks and domains. We cast the learning of model strengths as a binary supervised learning task, where the features are input embeddings of samples across tasks, and the labels are whether the model “did well” on the corresponding inputs, e.g., generated correct class label, answered a question correctly, or followed input instructions sufficiently well. See Figure 1 for an illustration. We emphasize that when training the proposed “correctness” predictor, it is important to consider their ability to generalize out-of-distribution (OOD) data since our goal is to estimate

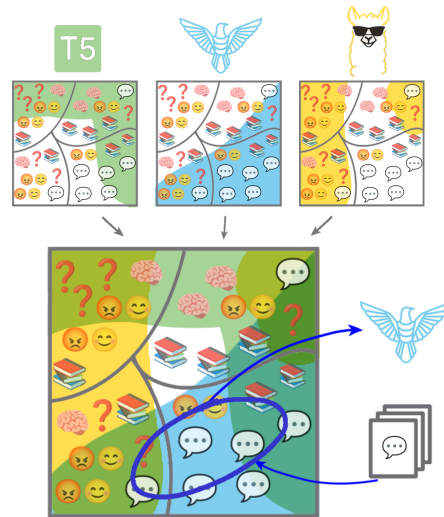


Figure 1: We learn the strengths of candidate LLMs (T5, Falcon, Llama) on various tasks (emojis inside boxes: QA, reasoning, etc.) and domains (4 sections within each box: finance, legal, etc.) from benchmark datasets. We achieve this by training a binary classifier per LLM (decision boundaries marked with colors in the upper part of the figure). For a new task (paper stack), we score each LLM with these binary classifiers and recommend an LLM (here Falcon) to the user.

---

\*Equal contributions

LLM performance on new tasks that were not seen during training. Training predictors given data from multiple domains that need to be generalized to unseen domains is an active area of research in ML literature [5–7]. We provide extended related work discussion in Appendix A.

## 2 Learning from Benchmarks

We introduce notations describing a typical NLP benchmark. Let  $\{x_1^d, \dots, x_{n_d}^d\}_{d=1}^D$  be a collection of inputs across  $D$  tasks. Each input text  $x_i^d$  corresponds to a reference answer  $r_i^d$ , i.e., an ideal generation for this input. There is a task-dependent metric  $F_d(x, o, r)$  that measures how well a response  $o$  for an input  $x$  corresponds to the reference  $r$ . To test an LLM $_m$ ,  $m \in \{1, \dots, M\}$ , on the benchmark, for each task  $d = 1, \dots, D$ , its responses are generated  $\{o_{im}^d = \text{LLM}_m(x_i^d)\}_{i=1}^{n_d}$  and compared to the corresponding references to obtain performance metrics  $\{f_{im}^d = F_d(x_i^d, o_{im}^d, r_i^d)\}_{i=1}^{n_d}$ .<sup>1</sup> At this point, most of the benchmark studies will take a (weighted) average of the performance metrics and report a single score for every LLM to rank them in performance. Instead, we reuse these evaluation results to formulate a supervised learning problem to understand better the strengths and weaknesses of various LLMs based on their performance on data points and tasks.

Our goal is to learn a routing function  $g_m(x)$  for each LLM,  $m = 1, \dots, M$ , that can predict performance of this LLM,  $\{f_{im}^d\}_{i=1}^{n_d}$ . We denote the ‘‘correctness’’ of model  $m$  on an input  $x$  by  $y(x, m) \in \{0, 1\}$ . Correctness is evaluated as follows: generate a response  $o_{im}^d$  with LLM  $m$  on input  $x_i^d$ , compare it to the corresponding reference  $r_i^d$ , and output 1 if the model’s response is good enough, i.e.,  $f_{im}^d > \eta_d$ , and 0 otherwise, where  $\eta_d$  is some threshold that can be task and/or metric specific. For tasks like classification or multiple-choice QA,  $y(x_i^d, m) = f_{im}^d$ , while for various evaluation metrics used in summarization and instruction following tasks [8–10], the notion of correctness can help to account for the heterogeneity of popular metrics and task difficulty levels.

For each LLM,  $m = 1, \dots, M$ , we train a predictor of its correctness with (multi-task) empirical risk minimization:  $\min_{g_m} \sum_{d=1}^D \sum_{i=1}^{n_d} \ell(g_m(x_i^d), y(x_i^d, m))$ , where we choose  $\ell$  to be a binary cross-entropy loss and  $g_m$  is any standard probabilistic classifier, i.e.,  $g_m(x)$  estimates  $P(y(x, m) = 1|x)$ .

We embed all inputs with a sentence transformer [11] and use a  $k$ -nearest neighbors classifiers [12] as correctness predictors  $\{g_m\}_{m=1}^M$ . We choose this approach for learning correctness predictors to emphasize the utility of learning from benchmarks even with a basic method and instead focus on the question specific to our problem that has not been studied in prior works on OOD generalization: *Can we improve the quality of LLM routing with an imperfect correctness predictor?*

## 3 LLM routing with (imperfect) correctness predictors

The goal of LLM routing is to identify an LLM that will have the highest frequency of being correct on a new task  $d'$ , given the inputs  $\{x_i^{d'}\}_{i=1}^{n_{d'}}$  from this task:  $\arg \max_m \tilde{S}(m, d')$ , where  $\tilde{S}(m, d') = \frac{1}{n_{d'}} \sum_{i=1}^{n_{d'}} y(x_i^{d'}, m)$ . Here,  $\tilde{S}(m, d')$  is the ‘‘oracle’’ score that we want to estimate. The most intuitive estimator is simply using the correctness predictor

$$S_1(m, d') = \frac{1}{n_{d'}} \sum_{i=1}^{n_{d'}} g_m(x_i^{d'}), \quad (1)$$

but prior work has shown that accurately estimating  $P(y|x)$ , i.e., calibration, is challenging on OOD data [13]. Meanwhile,  $g_m$  may still produce accurate predictions after thresholding the predicted probability even if the class probabilities are not estimated well, which is often the case with neural networks [14]. This motivates another score:

$$S_2(m, d') = \frac{1}{n_{d'}} \sum_{i=1}^{n_{d'}} \bar{g}_m(x_i^{d'}), \quad \text{where } \bar{g}_m(x_i^{d'}) = \mathbb{I}(g_m(x_i^{d'}) > t), \quad (2)$$

where  $t \in (0, 1)$  is some threshold, e.g.,  $t = 0.5$ ,  $\mathbb{I}$  is an indicator function, and  $\bar{g}_m(x) \in \{0, 1\}$  can be interpreted as the prediction of  $g_m$  on  $x$ . This score, however, does not take into account the potential ‘‘imperfection’’ of  $g_m$ , i.e., lower accuracy on OOD data from task  $d'$ . To address this issue, we model the out-of-distribution confidence of the predictions  $\bar{g}_m$ .

**A simple OOD confidence model.** We model LLM correctness as follows:

$$y(x, m)|x, d' = \begin{cases} \bar{g}_m(x) & \text{with probability } p(d', m) \\ 1 - \bar{g}_m(x) & \text{with probability } 1 - p(d', m), \end{cases} \quad (3)$$

<sup>1</sup>W.l.o.g., we can consider the same LLM with a different prompting strategy as a different LLM.

i.e.,  $p(d', m) \in [0, 1]$  is the probability that  $\bar{g}_m$  is the correct prediction on a data point from task  $d'$ . The above model can be condensed as follows:

$$y(x, m)|x, d' \sim \text{Bern}(\bar{g}_m(x)p(d', m) + (1 - \bar{g}_m(x))(1 - p(d', m))). \quad (4)$$

In this simplistic (and approximate) model, we assume that  $p(d', m)$  does not depend on the input  $x$  after conditioning on the task  $d'$ . We treat the problem of estimating  $p(d', m)$  as a supervised learning task, taking advantage of the task partition. Specifically, we assign a task descriptor  $u(d) \in \mathbb{R}_+$  to every task that measures the distance of the data from task  $d$  to the other available tasks combined. Then we collect the values of  $p(d, m)$ , i.e., the accuracy of  $\bar{g}_m$  on  $d$ , and fit a non-parametric regression model to predict  $p(d, m)$  from  $u(d)$ . At test time, we compute  $u(d')$  for a new task  $d'$  based on the inputs  $\{x_i^{d'}\}_{i=1}^{n_{d'}}$  and predict  $p(d', m)$  using the fitted regression model (we use Gaussian kernel smoother). We provide details in Appendix B.

Finally, given the model of LLM correctness 4,  $\tilde{S}(m, d')$  is a random variable (corresponding to  $\tilde{S}(m, d')$ ) distributed as a (scaled) sum of two Bernoulli random variables. To arrive at our final score for LLM routing, we take its expected value:

$$S_3(m, d') = S_2(m, d')p(d', m) + (1 - S_2(m, d'))(1 - p(d', m)). \quad (5)$$

When selecting an LLM with  $S_3$ , we consider an alternative to the arg max criterion based on our correctness model 4, which defaults to the best model on average across benchmark datasets when we are not sufficiently confident that a candidate model will be better:

$$\begin{cases} m_3 & \text{if } P(\tilde{S}(m_3, d') > \tilde{S}(m^*, d')) > \eta \\ m^* & \text{otherwise,} \end{cases} \quad (6)$$

where  $m_3 = \arg \max_m S_3(m, d')$ , i.e., the best LLM for the new task according to  $S_3$ , and  $m^* = \arg \max_m \sum_{d=1}^D \tilde{S}(m, d)$ , i.e., the best LLM across the benchmark datasets. In the experiments, we set  $\eta = 0.6$ . We summarize our LLM routing procedures in Appendix B.

## 4 Model routing on HELM

**Data:** We select 29 datasets from the HELM benchmark [1] representing scenarios such as question answering (including a subset of MMLU [2]), text classification, language, knowledge, and reasoning, among others. We present additional information about these datasets in Table 5.

**Models:** We evaluate 18 open-source models ranging in size from 3B to 70B, including base and chat variations of Llama 2 in different sizes. All models are summarized in Table 6.

**Model routing:** The best model on average (BMA) across the 29 considered HELM datasets is llama-2-70b. Our goal is to show that learning model routers from benchmark data can simultaneously outperform BMA and reduce inference costs by recommending smaller LLMs for tasks where they can perform well. We compare models selected with the three scores,  $S_1$ ,  $S_2$ , and  $S_3$ , presented in Section 3 to the performance of llama-2-70b, i.e., the BMA. All correctness predictors  $g_m$ s are kNN classifiers with  $k = 5$ . We report results with neural networks in Appendix C.1.

We also report the performance of the best model according to the ‘‘oracle’’ score  $\tilde{S}$ , which is the upper bound on what can be achieved with model routing, and  $S_3$  with the true  $p(d', m)$ , i.e., the accuracy of (an imperfect)  $g_m$  on  $d'$ . Finally, we compare the scoring LLMs with the average log-likelihood (LL) (or negative perplexity) of the response they generate on the inputs from the task of interest. This last baseline requires producing generations with every LLM at test time to make a selection, while all of our scores only require generating with the chosen LLM.

	Acc.	Pearson	# Params
$S_1$ eq. 1	0.662	0.685	40.3B
$S_2$ eq. 2	0.676	0.636	44.3B
$S_3$ eq. 5, 6	0.727	0.492	49.8B
$S_3$ true $p$	<b>0.735</b>	<b>0.799</b>	<b>33.8B</b>
LL	0.684	0.714	—
BMA	0.688	—	70.0B
Oracle	0.773	—	29.1B

**Results:** We conduct 29 sets of experiments, each time selecting 28 of the datasets as the benchmark data for training the LLM routers and using the remaining task as the new task  $d'$  for evaluating the performance. In Table 1 we report averages across experiments for the performance of the selected

model (Acc.), Pearson rank correlations between model accuracies and model scores, and the number of parameters of the selected model (# Params). We present additional comparison metrics in Table 2.

First, we notice that accounting for imperfections of the correctness predictors (their average accuracy is 0.59) has clear benefits: when we have access to the true accuracy of correctness predictors, the corresponding score,  $S_3$  true  $p$ , noticeably outperforms all other scores. Our simple kernel smoothing estimator of this accuracy (MAE= 0.116) allows us to obtain a practical model routing score  $S_3$  that outperforms BMA (11ama-2-70b) while choosing smaller models for some of the tasks (as evident by the average number of parameters of the chosen models).  $S_2$  sacrifices some accuracy but chooses even smaller performant models. Overall, learning from benchmarks allows us to obtain LLM routers that can improve overall performance while utilizing smaller models where appropriate. Finally, we note that log-likelihood (LL) also performs well. However, routing with it requires passing each test input through *each* candidate LLM, which has 347B parameters in total.

**Reducing the OOD gap:** The average accuracy of correctness predictors across tasks and models for the experiments in Table 1 is 0.59. It is a fairly low accuracy for binary classification, which we attribute to the diversity of tasks in HELM leading to substantial distribution shifts when predicting the correctness of LLMs on held-out tasks. We investigate the quality of model routing when we reduce this OOD gap. A simple strategy to reduce this gap is to collect a small number of labeled in-distribution samples. This can be accomplished by asking a practitioner to provide reference answers ( $r_i^{d'}$ s) for a small number of inputs from their task to evaluate the correctness of candidate LLMs on these in-distribution inputs and use it to improve correctness predictors.

We simulate this scenario by moving  $\min(\alpha n^{d'}, 50)$  samples from the data from a new task  $d'$  to the data for training the correctness predictors. The upper limit of 50 samples is to maintain practical utility while accounting for varying dataset sizes (see Table 5). We conduct 29 sets of experiments, repeating each one 10 times to obtain standard deviations (randomness is due to random selection of data points from a new task for reducing the OOD gap). We summarize the average accuracy of models selected with various routing scores for varying  $\alpha$  in Figure 2 ( $\alpha = 0$  corresponds to Table 1). Results for Pearson correlation are in Figure 4(a).




Figure 2: Reducing the OOD gap.

We see that even a small number of in-distribution samples ( $\alpha = 0.05$ ) can reduce the OOD gap (corresponding average accuracy of correctness predictors is 0.65; see Figure 4(b)) and noticeably improves the model routing performance of all three of our scores. When the number of in-distribution samples further increases,  $S_1$  starts to outperform  $S_3$ . We attribute this observation to kNN being well-calibrated in-distribution, i.e., the correctness predictors provide reliable estimates of their own confidence  $P(y|x)$ , which are used by  $S_1$  in (1). Finally, we note a fairly large variance in the results due to random selection of the in-distribution training samples from  $d'$ , suggesting that active learning [15] can help to improve LLM routing further.

**Learning from more benchmarks:** We anticipate learning LLM routers from benchmarks to be the most effective when new tasks are similar to the benchmark tasks, thus reducing the OOD gap without any labeling burden for a practitioner. To empirically investigate this hypothesis, in Figure 3 we visualize the relation between the quality of model routing with  $S_3$ , measured with Pearson correlation between model scores and accuracies of candidate LLMs, and the distance  $u(d')$  from a new task  $d'$  to the training benchmark data. Here we aggregated results across different  $\alpha$  values from Figure 2. For smaller distance values the correlation is approaching 1, while for large distances it sometimes deteriorates. Results for other scores demonstrate a similar trend and are presented in Appendix C.3 along with additional

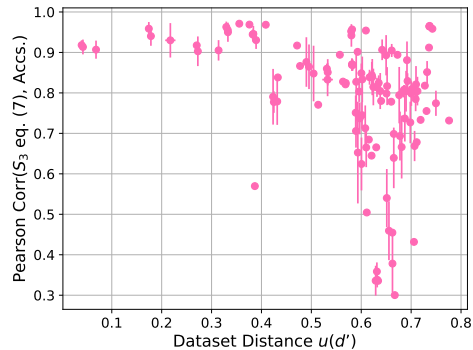


Figure 3: Correlation( $S_3$ , Accs.) and  $u(d')$ .

details. Our work demonstrates the potential of learning from benchmarks for LLM routing and investigates 3 model scores in the context of OOD generalization when routing LLMs for new tasks.

We present results on routing smaller LLMs on HELM in Appendix E, results on the MixInstruct benchmark [16] in Appendix D, and discussion of potential future work directions in Appendix F.

## Acknowledgments and Disclosure of Funding

The MIT FutureTech research group recognizes the generous support of the MIT-IBM Watson AI Lab, Boston Scientific and Good Ventures.

The MIT Geometric Data Processing group acknowledges the generous support of Army Research Office grants W911NF2010168 and W911NF2110293, of Air Force Office of Scientific Research award FA9550-19-1-031, of National Science Foundation grant CHS-1955697, from the CSAIL Systems that Learn program, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, from a gift from Adobe Systems, and from a Google Research Scholar award.

## References

- [1] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.
- [2] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- [3] Edward Beeching, Clémentine Fourrier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. Open llm leaderboard. [https://huggingface.co/spaces/HuggingFaceH4/open\\_llm\\_leaderboard](https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard), 2023.
- [4] Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, September 2021.
- [5] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *Computer Vision—ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14*, pages 443–450. Springer, 2016.
- [6] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- [7] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021.
- [8] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*, 2020.
- [9] Thibault Sellam, Dipanjan Das, and Ankur Parikh. BLEURT: Learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, 2020.
- [10] Weizhe Yuan, Graham Neubig, and Pengfei Liu. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277, 2021.
- [11] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2019.

- [12] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- [13] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019.
- [14] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017.
- [15] Burr Settles. Active learning literature survey. 2009.
- [16] Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14165–14178, Toronto, Canada, jul 2023. Association for Computational Linguistics.
- [17] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [18] Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. The UCI Machine Learning Repository, 2023.
- [19] Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, et al. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the ninth workshop on statistical machine translation*, pages 12–58, 2014.
- [20] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6391–6400, 2019.
- [21] Chunyuan Li, Haotian Liu, Liunian Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping Jin, Houdong Hu, Zicheng Liu, Yong Jae Lee, et al. Elevator: A benchmark and toolkit for evaluating language-augmented visual models. *Advances in Neural Information Processing Systems*, 35:9287–9301, 2022.
- [22] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [23] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32, 2019.
- [24] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- [25] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- [26] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- [27] Sebastian Raschka. Model evaluation, model selection, and algorithm selection in machine learning. *arXiv preprint arXiv:1811.12808*, 2018.
- [28] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

- [29] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [30] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021.
- [31] Hui Xu and Robert Tibshirani. Estimation of prediction error with known covariate shift. *arXiv preprint arXiv:2205.01849*, 2022.
- [32] Mayee Chen, Karan Goel, Nimit S Sohoni, Fait Poms, Kayvon Fatahalian, and Christopher Ré. Mandoline: Model evaluation under distribution shift. In *International conference on machine learning*, pages 1617–1629. PMLR, 2021.
- [33] Subha Maity, Mikhail Yurochkin, Moulinath Banerjee, and Yuekai Sun. Understanding new tasks through the lens of training data via exponential tilting. In *International Conference on Learning Representations*, 2023.
- [34] Yiding Jiang, Vaishnavh Nagarajan, Christina Baek, and J Zico Kolter. Assessing Generalization of SGD via Disagreement. In *International Conference on Learning Representations*, 2021.
- [35] Jiefeng Chen, Frederick Liu, Besim Avci, Xi Wu, Yingyu Liang, and Somesh Jha. Detecting errors and estimating accuracy on unlabeled data with self-training ensembles. *Advances in Neural Information Processing Systems*, 34:14980–14992, 2021.
- [36] Nathan Hoyen Ng, Neha Hulkund, Kyunghyun Cho, and Marzyeh Ghassemi. Predicting out-of-domain generalization with neighborhood invariance. *Transactions on Machine Learning Research*, 2023.
- [37] Devin Guillory, Vaishaal Shankar, Sayna Ebrahimi, Trevor Darrell, and Ludwig Schmidt. Predicting with confidence on unseen distributions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1134–1144, 2021.
- [38] Saurabh Garg, Sivaraman Balakrishnan, Zachary Chase Lipton, Behnam Neyshabur, and Hanie Sedghi. Leveraging unlabeled data to predict out-of-distribution performance. In *International Conference on Learning Representations*, 2022.
- [39] Yaodong Yu, Zitong Yang, Alexander Wei, Yi Ma, and Jacob Steinhardt. Predicting out-of-distribution error with the projection norm. In *International Conference on Machine Learning*, pages 25721–25746. PMLR, 2022.
- [40] Yixin Liu and Pengfei Liu. Simcls: A simple framework for contrastive learning of abstractive summarization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 1065–1072, 2021.
- [41] Mathieu Ravaut, Shafiq Joty, and Nancy Chen. Summareranker: A multi-task mixture-of-experts re-ranking framework for abstractive summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4504–4524, 2022.
- [42] Lingjiao Chen, Matei Zaharia, and James Zou. FrugalGPT: How to Use Large Language Models While Reducing Cost and Improving Performance. *arXiv preprint arXiv:2305.05176*, 2023.
- [43] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [44] LAION-AI. Open assistant, 2023.
- [45] 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. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, March 2023.

- [46] Julian Salazar, Davis Liang, Toan Q Nguyen, and Katrin Kirchhoff. Masked language model scoring. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, 2020.
- [47] Lingjiao Chen, Matei Zaharia, and James Y Zou. Frugalml: How to use ml prediction apis more accurately and cheaply. *Advances in neural information processing systems*, 33:10685–10696, 2020.



## A Related work

**Benchmarking** Comparing models or algorithms across various tasks is a standard practice in ML and AI literature. Prior to Foundation Models [17], it was typical to apply *the same learning algorithm* to train a model on each of the datasets and compare the performance against other learning algorithms. The UCI Machine Learning Repository [18] is one prominent example of such a collection of datasets often used to compare learning algorithms. With the emergence of Foundation Models, i.e., models with billions of parameters trained on massive datasets using large compute clusters, the paradigm changed to evaluating *the same model* (or a few-shot tuned version of it) on a variety of tasks [19–21]. In the context of Large Language Models, many benchmarks [22, 23, 2, 4, 24, 1, 3, 16] were proposed to help determine the most capable LLM. Benchmarks typically average the performance of models across tasks and provide a final ranking, discarding the rest of the information. In this work, we use the byproducts of benchmark evaluations, i.e., the per-sample performance of various LLMs across tasks, to learn about their individual strengths and identify the best LLM for a new task.

**Model selection** Selecting the best model, or model selection, is a classical topic in statistics and ML [25–27]. However, the typical problem setting is quite different: classical methods like cross-validation aim to estimate the population error of a model trained on samples from the population distribution. In other words, the goal is to find the best model for in-distribution test data, i.e., data sampled from the same distribution as the train data. The notion of “train” data is quite elusive for LLMs, as they are usually trained on massive datasets with trillions of tokens with a simple task of next token prediction [28, 29]. However, the tasks we evaluate them on are often more structured, e.g., classification and question-answering, and are specific to domains that may or may not be sufficiently represented in the train data. In addition, techniques like  $k$ -fold cross-validation require training the model multiple times, which is infeasible for LLMs.

**Out-of-distribution model selection** Recognizing the limitations of the model selection methods for in-distribution test data [30, 7], recent work has proposed a variety of methods to select models when deployed on data that may differ from the train data. These methods rely on ideas such as bootstrapping [31], reweighing [32, 33], agreement of models or ensembles [34–36], or aligning model accuracy in-distribution with a confidence threshold [37–39]. Most of these methods are nontrivial to extend to generation use-cases of LLMs; some require training multiple models, and some need well-defined in-distribution data related to the new task.

**Routing LLMs** Prior work on selecting LLMs primarily considers choosing one that produces the best generation for a given input. [40, 41, 16] train dedicated scoring or ranking models that can be applied to model generations. Unlike our work, these approaches require generating outputs with *every* candidate LLM to make a decision, which can be computationally prohibitive with a large pool of candidate LLMs. FrugalGPT [42] calls LLMs sequentially until a dedicated scoring model deems the generation acceptable. Prior works in this group require training data sufficiently representative of each of the tasks and domains of interest to train the corresponding ranking and scoring models. In this paper, instead, we use data from benchmarks to learn the strengths and weaknesses of LLMs across tasks and domains. The resulting model router requires generating outputs only with the chosen LLM at test time.

## B Correctness predictors and confidence estimation

We provide additional details on the correctness predictors used in our experiments, along with more details on the dataset distance and the Gaussian kernel smoother for estimating the accuracy of the correctness predictors on new tasks,  $p(d', m)$ s.

**Correctness predictors in our experiments** While any probabilistic classifier may fit our setting, in the experiments, we mainly used a simple kNN classifier applied in an embedded space. Recall that we have  $D$  benchmark datasets with inputs  $\{x_i^d\}_{i=1}^{n_d}$  for  $d = 1, \dots, D$ . To compute our correctness predictor based on the benchmark datasets, we first embed all their inputs. We denote the combined set of embedded inputs from the benchmark datasets as  $\mathcal{D} = \{\phi(x_1^d), \dots, \phi(x_{n_d}^d)\}_{d=1}^D$ , where  $\phi$  is a sentence transformer [11]. We use `a11-mpnet-base-v2` from Hugging Face in all experiments.

Given a sample  $x_i^{d'}$  from a new task  $d'$ , we embed it using the same  $\phi$  and define the classifier,  $g_m$ , for each model  $m$  by:

$$g_m(x_i^{d'}) = \frac{1}{k} \sum_{e \in \text{NN}(\phi(x_i^{d'}), k, \mathcal{D})} y(e, m), \quad (7)$$

where  $y(e, m) \in \{0, 1\}$  is the correctness of model  $m$  on the (embedded) input  $e$ , and  $\text{NN}(\phi(x_i^{d'}), k, \mathcal{D})$  is the set of  $k$  closest embedded neighbors from  $\mathcal{D}$  to the new embedded sample  $\phi(x_i^{d'})$ , according to the cosine distance. Then,  $\bar{g}_m$ , as defined in (2), is a binary kNN classifier. Finally, we compute the per-model correctness predictors,  $S_1(m, d')$  and  $S_2(m, d')$ , for the new task  $d'$ , according to (1) and (2), respectively.

Next, we describe a method for estimating the probability  $p(d', m)$  in our confidence model and the  $S_3(m, d')$  score, (5). This method comprises a dataset distance and a kernel smoother, defined as follows.

**Dataset distance** Our dataset distance  $u(d)$  is a one-sided variant of the Chamfer distance with extended neighborhood size. We define it formally below:

$$u(d) = \frac{1}{n_d} \sum_{i=1}^{n_d} nn(x_i^d, \mathcal{D}_{-d}), \quad (8)$$

where  $\mathcal{D}_{-d}$  is the set of (embedded) inputs from the  $D$  datasets excluding inputs from  $d$  (for a new task  $d'$ ,  $\mathcal{D}_{-d'} = \mathcal{D}$  since  $d'$  is not part of the  $D$  benchmark datasets we use for training LLM routers), and  $nn(x_i^d, \mathcal{D}_{-d})$  is the average distance from the input  $x_i^d$  to its closest  $\kappa$  neighbors in  $\mathcal{D}_{-d}$ :

$$nn(x, \mathcal{D}) = \frac{1}{\kappa} \sum_{e \in \text{NN}(\phi(x), \kappa, \mathcal{D})} \text{cosine}(\phi(x), e), \quad (9)$$

where  $\text{NN}(\phi(x), \kappa, \mathcal{D})$  is the set of  $\kappa$  closest embedded neighbors of  $\phi(x)$  in  $\mathcal{D}$  according to cosine distance. We set  $\kappa = 19$  for the dataset distance in all experiments.

**Kernel smoother** For each LLM  $m = 1, \dots, M$ , to obtain the corresponding kernel smoother estimate we iterate over the available benchmark datasets, each time holding one out and computing pairs  $(u(d), p(d, m))$  for held out dataset  $d$ , where  $p(d, m)$  is the accuracy of  $g_m$  on data from  $d$  after training on  $\mathcal{D}_{-d}$ . We repeat this process 10 times for 15 values of in-distribution mixing parameter  $\alpha$  (similar to the experimental setup in Figure 2 but using benchmark datasets  $d = 1, \dots, D$  instead of  $d'$ ) to obtain the training set of distance-accuracy pairs  $\{u_z, p_z(m)\}_{z=1}^Z$ . In the HELM experiments in Section 4,  $Z = 28 * 10 * 15 = 4200$  (28 is the number of datasets from HELM after holding one out as the new task for evaluating the performance).

For a new task  $d'$ , we compute  $u(d')$  using the inputs from this task and our benchmark datasets and estimate  $p(d', m)$  for each  $m$  with simple Gaussian kernel smoothing:

$$p(d', m) = \frac{\sum_{z=1}^Z p_z(m) \mathcal{K}(u(d'), u_z)}{\sum_{z=1}^Z \mathcal{K}(u(d'), u_z)}, \quad (10)$$

where  $\mathcal{K}(u(d'), u_z) = \exp\left(-\frac{(u(d') - u_z)^2}{2\sigma^2}\right)$ . We set  $\sigma = 0.09$  in all experiments, which is the value we found to perform well through some preliminary experimentation.

Finally, we note that the proposed confidence model, including the definitions of the dataset distance and kernel smoother, can be combined with any classifier  $g_m$ , and is not restricted to the kNN classifier used for the correctness predictor in our experiments.

**Additional notes regarding  $S_3$**  Recall that when selecting a model with  $S_3(m, d')$  we use an additional step described in (6) that facilitates the selection of the best model on average when we are not sufficiently confident in the model with the highest  $S_3(m, d')$  score. Probability expression,  $P(\tilde{S}(m_3, d') > \tilde{S}(m^*, d'))$ , required for this step is not available in closed form, as  $\tilde{S}$  is distributed as a (scaled) sum of two Bernoulli random variables, but it is straightforward to estimate via Monte Carlo sampling from the corresponding Bernoulli distributions.

When reporting correlations for  $S_3$  (e.g., Pearson and Spearman correlations in Table 1), we use  $S_3(m, d')$  as is, i.e., as defined in (5).

Table 2: LLM routing on HELM: Comparison of various model scores for LLM routing with the Oracle model selection and performance of the best model on average (BMA).

	Acc.	Ratio to Best	Pearson	Spearman	% BMA	# Params	Rank
$S_1$ eq. 1	0.662	0.855	0.685	0.465	0.17	40.3B	6.172
$S_2$ eq. 2	0.676	0.868	0.636	0.468	0.10	44.3B	5.897
$S_3$ eq. 5, 6	<u>0.694</u>	<u>0.898</u>	<u>0.727</u>	<u>0.492</u>	0.48	49.8B	<u>5.310</u>
$S_3$ true $p$	<b>0.735</b>	<b>0.944</b>	<b>0.799</b>	<b>0.596</b>	0.22	<b>33.8B</b>	<b>3.800</b>
LL	0.684	0.869	0.714	0.459	0.10	—	6.517
BMA	0.688	0.884	—	—	1.00	70.0B	6.069
Oracle	0.773	1.000	—	—	0.21	29.1B	1.000

Table 3: Average metrics for the experiments using MLP. Best results are highlighted with bold, and second best with an underline (excluding Oracle).

	Acc.	Ratio to Best	Pearson	Spearman	% BMA	# Params	Rank
$S_1$ eq. 1	<u>0.693</u>	<u>0.893</u>	0.750	<u>0.559</u>	0.19	57.6	<u>4.345</u>
$S_2$ eq. 2	0.683	0.878	0.675	0.547	0.45	57.9	5.041
$S_3$ eq. 5, 6	0.690	0.884	<u>0.755</u>	0.544	0.81	66.3	5.566
$S_3$ ATC	0.646	0.827	0.651	0.365	0.56	<b>48.7</b>	7.283
$S_3$ true $p$	<b>0.740</b>	<b>0.943</b>	<b>0.840</b>	<b>0.686</b>	0.43	<u>50.4</u>	<b>3.455</b>
LL	0.684	0.869	0.714	0.459	0.10	—	6.517
BMA	0.688	0.884	—	—	1.00	70.0	6.069
Oracle	0.773	1.000	—	—	0.21	29.1	1.000

## C Additional results for model routing on HELM

### C.1 Results with MLP as the correctness predictor

**MLP Classifier** We present LLM routing with more sophisticated correctness predictors. In this section, we use a small Multi-Layer Perceptron (MLP) as  $g_m$  to classify the embedded inputs. Each MLP comprises three fully connected layers (two hidden layers, 1500 units each) with ReLU activation and sigmoid output. These MLPs were trained with the Adam optimizer [43] with learning rate 0.01 to minimize the binary cross entropy loss as described in Section 2. Training occurred over 100 epochs with early stopping should the validation accuracy not increase for 10 consecutive epochs. We also explore the combination of  $S_3$  with ATC (instead of the kernel smoother), a prior OOD accuracy estimator proposed by [38].

**Results** We conducted the model routing experiment analogous to the experiment in Table 1 of the main paper. Results are reported in Table 3. We note several differences compared to results with the kNN classifier in Tables 1 and 2. First is the clear improvement in the  $S_1$  score; the MLP classifier performs better on OOD data. As before  $S_3$  true  $p$  outperforms the other scores indicating the potential value of using the proposed confidence model 4. However, in this experiment,  $S_3$  with the kernel smoother no longer improves upon  $S_1$ . We also report results with ATC as the estimator of the accuracy of correctness predictors, which performs noticeably worse than our kernel smoothing estimator. The corresponding MAE is 0.118 for the kernel smoother and 0.177 for the ATC, demonstrating the advantage of our estimator in this application. Finally, across all scores the selected model sizes have considerably increased; this suggests that the MLP classifier underestimates the performance of smaller models compared to kNN.

### C.2 Reducing the OOD gap

We present additional results for this experiment in Figure 4. (a) shows Pearson correlation improvement as we increase  $\alpha$ , similar to the trends in accuracy improvement in Figure 2; (b) demonstrates that the accuracy of correctness predictors  $g_m$ s improves as we increase the number of samples from  $d'$  used for training them, thus reducing the OOD gap; (c) shows the mean absolute error (MAE) of




Figure 4: Additional results for Reducing the OOD gap experiment in Figure 2.




Figure 5: LLM routing with  $\leq 13$ B parameter models compared to Llama 2 70B.

our kernel smoothing estimator of the accuracy of correctness predictors  $p(d', m)$  – the estimator does not improve as much with increased  $\alpha$ , thus  $S_3$  eventually becomes worse than  $S_1$  in terms of correlation and accuracy of the selected models.

### C.3 Dataset distance and Pearson correlation

The dataset distance  $u(d')$  is computed as in (8). As evident from (9), dataset distance will usually decrease for larger values of  $\alpha$  as inputs from  $d'$  are moved into  $\mathcal{D}$  (assuming that inputs from  $d'$  are on average closer to each other than they are to inputs from other tasks). In this experiment, this serves as a mechanism to study the performance of LLM routing on closer datasets, providing insights into the benefits of learning LLM routers on *more benchmarks* where it is more likely that dataset distance for a new task is small.

In Figure 6 we present relations between dataset distance  $u(d')$  and Pearson correlation between various model scores and accuracies of candidate LLMs. For results with  $S_3$  see Figure 3.

## D Model Routing on MixInstruct

We further demonstrate our approach in a different setting and task type, on the MixInstruct benchmark dataset [16]. The dataset is composed of instruction-following tasks, divided into train/validation/test sets of 100K/5K/5K samples, and includes evaluations of  $N = 11$  open-source LLMs using common




Figure 6: Correlation of scores and LLM accuracies on new tasks and corresponding data distances.

Table 4: Average metrics for per-instance LLM selection on the MixInstruct test set. Best results are highlighted with bold and second best with an underline (excluding Oracle).

	BERTScore $\uparrow$	BARTScore $\uparrow$	BLEURT $\uparrow$	MCPI
Random	66.36	-3.76	-0.77	-
LL	65.83	-4.12	-0.96	$N$
BMA: Open-Assisant	<u>74.68</u>	-3.45	-0.39	-
BMA: Vicuna	<u>69.60</u>	-3.44	-0.61	-
MLM-Scoring [46]	64.77	-4.03	-0.88	$N$
SimCLS [40]	73.14	<u>-3.22</u>	<u>-0.38</u>	$N$
SummaReranker [41]	71.60	-3.25	-0.41	$N$
PairRanker [16]	72.97	<b>-3.14</b>	<b>-0.37</b>	$N$
Ours	<b>74.75</b>	-3.40	<u>-0.38</u>	<b>2</b>
Oracle	77.67	-2.87	-0.15	$N$

metrics, e.g. BERTScore [8], BARTScore [10], and BLEURT [9]. In [16], this benchmark was used to compare different LLM ranking methods in per-instance model selection. We follow the same setting and apply our score  $S_1(m, d')$  to the test set, per-instance, where we use the 100K-sample train set as the benchmark data for training our LLM router. See Appendices B and D for details on the score computation and the experiment parameters, respectively. Due to the per-instance setting, and since the test set was constructed from in-distribution data, we focus on our simplest router model  $S_1, (1)$ .

We compare our approach with the scoring methods examined by [16], as well as scoring based on the average log-likelihood (LL) of the model responses to the inputs. Additionally, we present the metrics for the best models on average (BMA), Open-Assistant [44] and Vicuna [45]. We report the results of BERTScore, BARTScore and BLEURT in Table 4, along with the number of model calls per instance (MCPI), for  $N$  LLMs, performed during inference time. All compared methods require model generations for every point in the test set, by each of the examined LLMs, whereas our approach requires only one model generation and one call to some general embedding function. In addition, all methods, except for LL, require training auxiliary language models, whereas our approach is a simple kNN classifier on the embedded inputs. While our approach does not consistently outperform the compared methods, these results demonstrate the potential of using benchmark datasets for model routing with significantly better inference-time efficiency.

**Effect of benchmark dataset sparsity** To highlight the potential of our approach in this setting, we examine the effect of the reference benchmark data sparsity. We apply our method to different subsets of the test set,  $X_{\text{test}}$ , where the subsets are defined by limiting the maximal average distance of each test set point to the closest points from the reference (train) set, denoted by  $\text{NN}_{\text{train}}$ , i.e.  $X'_C = \left\{ x' \in X_{\text{test}} \mid \frac{1}{|\text{NN}_{\text{train}}(x')|} \sum_{x \in \text{NN}_{\text{train}}(x')} \text{dist}(x', x) < C \right\}$ , where  $C$  is the maximal average distance and  $X'_C$  is the resulting subset of the test set. Figure 7 presents the metric scores for the different subsets using our method, the oracle (best possible choices), and LL scoring. We also report the percentage of the test set that is used in each subset. This figure depicts that our predictor approaches the oracle metrics as the average distance to the reference points decreases. This suggests that adding more benchmark datasets to reduce the sparsity of the reference space may lead to better LLM selections with our approach.

**Correctness predictor and metrics** In the experiments on the MixInstruct dataset [16], we construct  $S_1(m, d')$  following the scoring approach described in Appendix B, where the MixInstruct train set was defined as the benchmark dataset. Then, instead of computing a per-dataset score for the entire test set, we compute the score for each test point, i.e.,  $S_1(m, x_i^{d'}) = g_m(x_i^{d'})$ , and select a model  $m$  per-point based on this score. The reported metrics in Table 4 and Figure 7 are averaged over the output evaluations of these per-point model selections. In our experiments, to compute  $g_m(x_i^{d'})$  we use the BERTScore metric on the closest train set points (as  $y(x, m)$  in 7). This was motivated by the conceptual relation between the implementation of our approach and the BERTScore, which relies on embedding space distances, and was validated empirically.

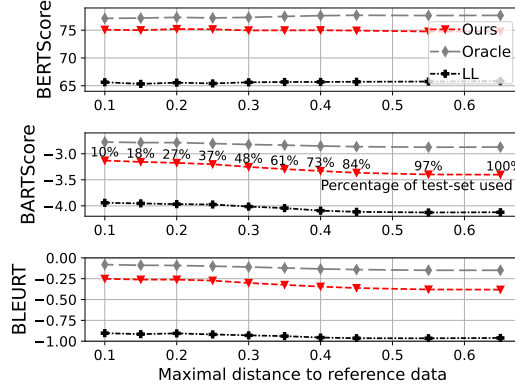


Figure 7: Average metrics on subsets of the MixInstruct test set, defined by limiting the maximal average distance between test instances and their closest neighbors in the reference (train) set.

**kNN parameter** We set  $k = 10$  for the kNN classifier, slightly higher than in the HELM experiments. This choice was motivated by the in-distribution properties of the test set in MixInstruct, which is constructed from different parts of the same datasets that comprise the train set. We note that the metrics did not significantly vary for different choices of  $k \in [5, 100]$ .

## E How useful are smaller LLMs?

While a given LLM may work best on average, these models tend to be the biggest and therefore most expensive to run. Practitioners can achieve gains in cost, compute, and latency if we can successfully predict whether a smaller LLM can be adequate for a given task. Identifying good smaller models for tasks of interest will also redefine the cost/benefit tradeoff behind automating certain tasks, potentially incentivizing the automation of new tasks that were previously cost-prohibitive to automate with larger LLMs.

To evaluate the potential of smaller LLMs we revisit our HELM experiment in Figure 2. In Figure 8, we perform LLM routing using *only models with  $\leq 13B$  parameters* and compare it to the performance of Llama 2 70B. Oracle’s performance demonstrates that it is conceptually possible to outperform a large model by routing smaller LLMs. Results with our scores  $S_1$  and  $S_2$  (see Figure 5 for breakdown by scores) demonstrate that it is also practically feasible to match the performance of the 70B model by combining learning from benchmarks with a small number ( $\alpha = 0.04$ , i.e., 2-40 samples) of labeled samples from a new task that a practitioner can provide to save on the inference costs in their LLM application.




Figure 8: LLM routing with  $\leq 13B$  parameter models compared to Llama 2 70B.

## F Future work

Our work demonstrates the potential of learning from benchmarks for LLM routing and investigates 3 model scores in the context of OOD generalization when routing LLMs for new tasks. We summarize potential next steps for improving the quality and efficacy of LLM routers.

The major challenge of LLM routing is OOD generalization of correctness predictors. Thus, using more benchmarks and modern methods for improving OOD generalization to learn correctness predictors is a promising next step. A practitioner can also provide labels for a few samples from their task, possibly guided by active learning techniques, to adapt or fine-tune correctness predictors. Even when reducing the OOD gap is too challenging, our score accounting for the (potentially low)

accuracy of correctness predictors demonstrated strong results when this accuracy,  $p(d', m)$ , is known for a new task, thus encouraging the development of methods for estimating it better.

We also anticipate that routing “expert” LLMs fine-tuned for a specific domain can improve the results. Regions of the sample space where such models are “correct” should mostly align with the domains of their expertise (recall Figure 1), making it easier to learn the corresponding correctness predictors and simplifying LLM routing when a new task is from a specific domain.

Our experiments in Figure 8 demonstrate the utility of LLM routing with *smaller* models, which can reduce costs and facilitate using LLMs in a broader set of domains. Thus, we want to explore modifications to our scores that will encourage the selection of smaller LLMs when their anticipated performance is comparable to the larger, more reliable models. Prior work on frugal API selection [47, 42] provides a good starting point to explore this direction.

Table 5: HELM dataset details

Dataset	Size (instances)	Type
RAFT-ADE Corpus V2	40	Binary Classification
RAFT-Banking 77	40	77 Class Classification
RAFT-NeurIPS Impact Statement Risks	40	Binary Classification
RAFT-One Stop English	40	3 Class Classification
RAFT-Overruling	40	Binary Classification
RAFT-Semiconductor Org Types	40	3 Class Classification
RAFT-Systematic Review Inclusion	40	Binary Classification
RAFT-TAI Safety Research	40	Binary Classification
RAFT-Terms of Service	40	Binary Classification
RAFT-Tweet Eval Hate	40	Binary Classification
RAFT-Twitter Complaints	40	Binary Classification
IMDB	1000	Binary Classification
Civil Comments-demographic=all	1000	Binary Classification
bAbI-QA-task=all	1000	Q&A: one word answers
BoolQ	1000	Binary Classification
Entity Matching-Dataset=Beer	182	Binary Classification
Entity Matching-Dataset=Dirty iTunes Amazon	218	Binary Classification
Entity Matching-Dataset=Abt Buy	1000	Binary Classification
Entity Data Imputation-Dataset=Restaurant	242	Q&A: one word answers
Entity Data Imputation-Dataset=Buy	182	Q&A: one word answers
BBQ-subject=all	1000	Multiple Choice Questions
Legal Support	1000	Multiple Choice Questions
LSAT QA-task=all	461	Multiple Choice Questions
MMLU-Subject=Abstract Algebra	111	Multiple Choice Questions
MMLU-Subject=College Chemistry	108	Multiple Choice Questions
MMLU-Subject=Computer Security	111	Multiple Choice Questions
MMLU-Subject=Econometrics	126	Multiple Choice Questions
MMLU-Subject=US foreign policy	111	Multiple Choice Questions
Truthful QA-task=mc single	654	Multiple Choice Questions
Total: 29 datasets	9946	



Table 6: Candidate LLMs

Name	Model Size, B	Average Accuracy on the 29 HELM tasks
codegen-16b-mono	16	0.451
dial-flan5-xl	3	0.454
falcon-40b	40	0.641
flan-t5-xl	3	0.650
flan-t5-xxl	11	0.658
flan-ul2	20	0.668
gpt-jt-6b-v1	6	0.576
gpt-neox-20b	20	0.492
mpt-7b-instruct	7	0.514
mt0-xxl	13	0.543
llama-2-13b	13	0.624
llama-2-13b-chat	13	0.623
llama-2-13b-chat-beam	13	0.603
llama-2-70b	70	<b>0.688</b>
llama-2-70b-chat	70	<u>0.687</u>
llama-2-7b	7	0.610
llama-2-7b-chat	7	0.605
starcoder	15	0.587
Total: 18 LLMs	347	