

DavIR: Data Selection via Implicit Reward for Large Language Models

Haotian Zhou^{†,1}, Tingkai Liu^{†,2}, Qianli Ma¹, Yufeng Zhang¹, Jianbo Yuan¹
Pengfei Liu^{3,*}, Yang You^{4,*}, Hongxia Yang^{1,*}

[†] Equal Contribution

¹ByteDance, Inc.

²NeuroAI Scholar, Cold Spring Harbor Laboratory

³Generative Artificial Intelligence Research Lab, Shanghai Jiao Tong University

⁴School of Computing, National University of Singapore

Abstract

We introduce DavIR, a model-based data selection method for post-training Large Language Models. DavIR generalizes Reducible Holdout Loss to core-set selection problem of causal language modeling, and quantifies the “learnability” of a given datum with respect to a pre-trained LLM based on relative reduction in loss during fine-tuning, a metric we show to be closely related to the implicit reward model described in Direct Preference Optimization (DPO). We show that 6% of Alpaca dataset selected with DavIR can steer both the LLaMA and Gemma model family to produce superior performance compared to the same models trained on the full 52K dataset. We also show that Alpaca dataset compressed with DavIR can be combined with GSM8K dataset to effectively balance open-domain freeform QA and mathematical reasoning capabilities. Finally, we apply the DavIR objective to DPO and develop a normalized DavIR-DPO objective which improves alignment performance of Zephyr-7B-SFT model by 8% (relative) on AlpacaEval, compared against training on vanilla DPO objective.

1 Introduction

Large Language Models (LLMs) (Brown et al., 2020; Chowdhery et al., 2022; Touvron et al., 2023; Ouyang et al., 2022) have sparked a revolution in the field of Natural Language Processing (NLP), with far reaching impacts in domains such as law (Cui et al., 2023), medical (Singhal et al., 2022) and finance (Wu et al., 2023).

A critical step in the current paradigm of post-training LLMs is Supervised/Instruction Fine-tuning (SFT/IFT), which enables pre-trained models to exhibit strong instruction-following capabilities (Chung et al., 2022; Ouyang et al., 2022; Touvron et al., 2023; Wang et al., 2022; Zheng et al.,

2023). Selecting the most effective training data during this stage is particularly important since effective steering of LLM during SFT could be achieved by just a few thousand carefully curated data (Zhou et al., 2023). Previous approaches to selecting SFT training data focused on data quality and diversity (Ji et al., 2023; Zhou et al., 2023; Chen et al., 2023b,a; Li et al., 2023a), guided by the intuition of encouraging LLMs to output accurate and reliable information while maintaining generalization capabilities to a wide range of tasks and scenarios.

However, by focusing on the quality and diversity of the data, existing methods are *data-centric*, and are agnostic to the capabilities of the pre-trained model upon which fine-tuning occurs. Instead, following the “Superficial Alignment Hypothesis” (Zhou et al., 2023) which postulates that fine-tuning process unlocks the capabilities of pre-trained LLMs, we seek a model-centric data selection algorithm that chooses data that:

1. Quantifies the degree to which a model “learns” a data before and after training;
2. Does not require querying closed-source teacher models which may lead to security concerns;
3. Is theoretically grounded in the implicit reward function of the underlying LLM (see Section. 3.1).

We note that the previously proposed Reducible Holdout Loss (Mindermann et al., 2022) admits a simplification that satisfy the three requirements above (Rafailov et al., 2023). However, when applying RHO-like objectives to language modeling tasks, we observed a significant challenge: the RHO metric is highly correlated with the sequence length of the input data. This correlation introduces an undesirable bias in the data selection process, reducing the core-set selection to an approximation of length-based filtering. We show that this issue is inherent to the *sequential* nature of language

*Corresponding authors

modeling in state-of-the-art LLMs, and cannot be resolved by normalizing the total cross entropy loss by number of tokens in a datum (document).

Instead, a subtle yet crucial change in normalization of the RHO objective - normalizing with reference model loss instead of number of tokens - dramatically reduced the length dependency of the object. We term this modified RHO objective, and the consequent data selection method, **DavIR (Data Selection via Implicit Reward)**.

We demonstrate the effectiveness of DavIR across model families (LLaMA (Touvron et al., 2023), Gemma (Team et al., 2024a,b)) and across data domains (Alpaca (Taori et al., 2023), LIMA (Zhou et al., 2023), GSM8K (Cobbe et al., 2021)) and across benchmarks benchmarks (Self-Instruct (Wang et al., 2022), Vicuna (Zheng et al., 2023), Koala (Geng et al., 2023), OpenAssistant (Köpf et al., 2023), Helpful Base (Bai et al., 2022), GSM8K (Cobbe et al., 2021)). We show that DavIR outperforms *all* (to the best of author’s knowledge) state-of-the-art core-set selection methods across benchmarks.

Finally, as the introduction of normalization in the DavIR objective led to a deviation from the implicit reward model given by the vanilla DPO objective, we propose **DavIR-DPO** that incorporates the normalization proposed in the current work. We show that DavIR-DPO metric is the least correlated with the difference in length of paired responses in UltraFeedback dataset (Cui et al., 2024), and training Zephyr-7B-SFT model using the DavIR-DPO metric led to an 8% boost of length-controlled performance on AlpacaEval (Li et al., 2023b; Dubois et al., 2024) as compared to when trained using the vanilla DPO objective.

2 Background and Related Works

Supervised (Instruction) Fine-tuning of LLM

In training LLMs, Supervised (Instruction) Fine-tuning (SFT/IFT) plays a pivotal role during post-training by fine-tune LLMs with a small amount of data to enable instruction-following and multi-round dialogue. Two predominant methods of collecting SFT training data are 1) distillation from teacher models (e.g. Self-Instruct (Wang et al., 2022), Alpaca (Taori et al., 2023), Evol-Instruct (Xu et al., 2023)) and 2) manual annotation (e.g. InstructGPT (Ouyang et al., 2022), Vicuna (Zheng et al., 2023), LIMA (Zhou et al., 2023)).




Figure 1: **DavIR outperforms full data fine-tuning and data selection based on teacher LLM across model scales.** Performance comparison of 7B and 13B parameter models fine-tuned with data selected using DavIR (3,000 items), the full Alpaca dataset (52K), and data filtered using ChatGPT (9,229 items). “G” represents evaluation using GPT-4, and “H” represents human evaluation. The statistical significance of performance gain of DavIR over training on full dataset and other core-set selection methods are established in subsequent sections.

Implicit Reward in Direct Preference Optimization

First proposed in (Rafailov et al., 2023), Direct Preference Optimization (DPO) emerged as a post-training method following supervised-finetuning. DPO simplifies the RLHF pipeline by directly optimizing a language model using preference data, eliminating the need for explicit reward modeling and reinforcement learning. The simplicity of DPO and its effectiveness has led to wide adoption across models. A large body of follow-up works have been proposed that modify the DPO objective to improve robustness (Azar et al., 2023; Ji et al., 2024; Chowdhury et al., 2024; Chen et al., 2024; Wu et al., 2024), address issues of data scarcity (Liu et al., 2024; Jung et al., 2024) or provide stronger control over likelihood of producing winning and losing responses (D’Oosterlinck et al., 2024; Melnyk et al., 2024). Recent works have also began to explore length-dependencies of the DPO objective (Zhou et al., 2024; Park et al., 2024).

Core-set Selection for LLM

Core-set selection and dataset pruning has a long and rich history in ML research (Har-Peled and Kushal, 2005; Paul et al., 2021), where the goal is to find small subsets of training data which gives similar or superior performance as compared to training on the full dataset. A wide range of metrics have been explored for core-set selection, including model loss (e.g. RETRIEVE (Killamsetty et al., 2021), RHO (Mindermann et al., 2022)), gradient (e.g. CRAIG (Mirzasoleiman et al., 2020)), influence function (e.g. (Yang et al., 2022)) and clustering. (Birodkar et al., 2019; Sorscher et al., 2022).

Within the scope of LLMs, prior arts have primarily focused on data selection during pre-training, such as DoReMi (Xie et al., 2023a), RHO, DRO (Oren et al., 2019), and DSIR (Xie et al., 2023b). For post-training LLMs, recent works have focused on selecting training data based on quality, based on either 1) human annotation (e.g. LIMA (Zhou et al., 2023)), 2) LLM (e.g. AlpaGasus (Chen et al., 2023b)) or 3) validation loss on evaluation dataset (e.g. Instruction Mining (Cao et al., 2023)).

3 DavIR: Data Selection via Implicit Reward

3.1 DavIR in Supervised Fine-Tuning

As an ever increasing number of post-training datasets are developed for LLMs, it is important for practitioners to select a compute-permissible subset of the training data that achieves similar, or *better*, performance than the full available training corpus.

As such, the task that DavIR is set out to solve is one of core-set selection for post-training LLM: given a base model π_{base} , and a collection of training data $D_{\text{full}} = \{(x_i, y_i)\}_i$ (where (x_i, y_i) represents the prompt/response pair that constitutes a training datum), find a *minimal* subset of the training dataset $D_{\text{train}} \subset D_{\text{full}}$, $|D_{\text{train}}| \ll |D_{\text{full}}|$ such that the model trained on D_{train} achieves comparable, or **better**, performance than that trained on D_{full} .

At the core of the DavIR algorithm is the concept of “learnability” in post-training LLMs. We are motivated by the “Superficial Alignment Hypothesis” (Zhou et al., 2023) which suggests the post-training stage of LLMs involves using a small number of carefully selected training samples to steer a pre-trained LLMs to align with desired response patterns. In particular, this suggests that the training samples ought to be tightly coupled with the underlying capabilities of the base LLM model, or that such samples need to be “learnable” by the base model.

A simple and intuitive quantification of “learnability” is subtracting the evaluation loss of the base model π_{base} from that of the reference model π_{ref} (π_{base} trained on all of D_{full}):

$$\begin{aligned} S_{\text{RHO-LM}}(x, y) &= \mathcal{L}_{\text{base}}(y|x) - \mathcal{L}_{\text{ref}}(y|x) \\ &= [-\log \pi_{\text{base}}(y|x)] - [-\log \pi_{\text{ref}}(y|x)]. \end{aligned} \quad (1)$$

This approach is akin to that of Reducible Holdout Loss (RHO) (Mindermann et al., 2022) and we

refer to this vanilla generalization of RHO to causal language modeling as simply RHO-LM.

We note that the RHO-LM metric in Equation (1) is closely related to the implicit reward function in the Direct Preference Optimization (Rafailov et al., 2023) procedure. As shown in (Rafailov et al., 2023), under mild conditions, reward functions $r(x, y)$ consistent with the Bradley-Terry (BT) preference model (Bradley and Terry, 1952) can be equivalently written as:

$$\begin{aligned} r(x, y) &= \beta \log \frac{\pi(y|x)}{\pi_{\text{base}}(y|x)} \\ &= \beta \cdot [\mathcal{L}_{\text{base}}(x, y) - \mathcal{L}(x, y)] \end{aligned} \quad (2)$$

for some language model $\pi(y|x)$ obtained by training via Reinforcement Learning with Human Feedback procedure (using Proximal Policy Optimization) from a base model $\pi_{\text{base}}(y|x)$ using the said reward function $r(x, y)$ till optimality.

In other words, the reference model $\pi_{\text{ref}}(y|x)$ in RHO-LM can be obtained from $\pi_{\text{base}}(y|x)$ via reward maximization of the *implicit* reward function $r(x, y)$ in Equation (2). As such, selecting data via RHO-LM using the score function can be viewed as choosing data with maximum reward given by this implicit reward model.

However, we found empirically that the vanilla RHO-LM metric in Equation (1) is highly correlated with sequence length of the training data (see also Appendix B), an issue that persists despite aggregating the token-level losses via the averaging operation. This is inherently due to the sequential nature of language modeling, where increasing sequence length introduces additional contexts that constraints the distributions of all (subsequent) tokens. The effect of this length dependency is not to be under-estimated, as Table 1 shows that correlation between length and average (across tokens) cross-entropy loss as well as entropy of predictive probabilities could be as high as -0.9 (on a scale of [-1, 1]). Consequently, the RHO objective, which subtracts these length-dependent objectives, is also prone to be highly correlated with sequence length (see Table 2) - an issue that only applies to language modeling and was therefore overlooked by the original RHO work which dealt with image classification or NLP tasks with single classification objective (i.e. grammatical correctness in CoLA (Warstadt et al., 2018) and sentiment analysis in SST-2 (Socher et al., 2013)).

Fortunately, we found that a simple, yet highly

Dataset	Model	Pearson		Spearman	
		Entropy	Loss	Entropy	Loss
Alpaca	albert-base-v2	-0.67	-0.76	-0.90	-0.97
	bert-base-uncased	-0.64	-0.79	-0.88	-0.95
	gemma-2-2b	-0.66	-0.61	-0.80	-0.76
	gemma-2b	-0.68	-0.56	-0.83	-0.69
GSM8K	albert-base-v2	-0.78	-0.85	-0.83	-0.91
	bert-base-uncased	-0.73	-0.57	-0.77	-0.90
	gemma-2-2b	-0.75	-0.72	-0.79	-0.74
	gemma-2b	-0.77	-0.66	-0.80	-0.68
MBPP	albert-base-v2	-0.69	-0.83	-0.64	-0.85
	bert-base-uncased	-0.76	-0.87	-0.80	-0.93
	gemma-2-2b	-0.85	-0.82	-0.90	-0.85
	gemma-2b	-0.85	-0.82	-0.91	-0.83

Table 1: **Language modeling objectives are highly correlated with sequence length.** Pearson correlation and Spearman rank correlation of entropy and loss with respect to number of tokens in a given document. Note that all correlations are negative, indicating that token-level entropy/loss decrease as corresponding context length increases.

effective, normalization technique could dramatically mitigate the length-dependency of the RHO-LM metric, resulting in the normalized score function, which we term DavIR:

$$S_{\text{DavIR}}(x_i, y_i) = \frac{\mathcal{L}_{\text{base}}(x_i, y_i) - \mathcal{L}_{\text{ref}}(x_i, y_i)}{\mathcal{L}_{\text{base}}(x_i, y_i)} \quad (3)$$

Note that the denominator in the normalization could be either the base or the reference losses without impacting the ordering of the data via the DavIR metric S_{DavIR} (see Appendix. C for a simple proof). The reduction in both spearman and pearson correlation is shown in Table. 2.

Given the DavIR score function, the DavIR algorithm for supervised fine-tuning data selection is simply given as Algorithm 1.

Algorithm 1 DavIR for Supervised Fine-tuning

- 1: $\pi_{\text{ref}}(y|x) \leftarrow \pi_{\text{base}}(y|x)$ trained on D_{full}
 - 2: **for each** $(x_i, y_i) \in D_{\text{full}}$ **do**
 - 3: $\mathcal{L}_{\text{base}}(x_i, y_i) \leftarrow -\log \pi_{\text{base}}(y_i|x_i)$
 - 4: $\mathcal{L}_{\text{ref}}(x_i, y_i) \leftarrow -\log \pi_{\text{ref}}(y_i|x_i)$
 - 5: Compute $S_{\text{DavIR}}(x_i, y_i)$ as in Equation (3)
 - 6: **end for**
 - 7: Re-train π_{base} on top- $k_{D_{\text{full}}}$ $S_{\text{DavIR}}(x_i, y_i)$
-

As we later demonstrate, while vanilla RHO-LM is effective in selecting a subset of the training data, it far under-performs the length-regularized DavIR algorithm in the downstream tasks performances across multiple datasets and models (see Figure. 2).

In fact, as demonstrated in Table 6, DavIR is able to outperform all (to the best knowledge of the authors) existing core-set selection techniques on post-training LLMs.

Finally, we note that both RHO-LM and DavIR score functions capture the essence of training on data that are “learnable, worth learning, and not yet learnt” (Mindermann et al., 2022), while dramatically reducing the length-dependencies of the original RHO objective. By focusing on the same next-token-prediction objective as training LLM, and omitting confounding factors such as additional small proxy models (Xie et al., 2023a) or hold-out dataset (Mindermann et al., 2022), DavIR provides exact single datum-level measurement of “learnability” that tightly couples with the underlying capabilities of the pre-trained model.

3.2 DavIR in Direct Preference Optimization

The performance gain of DavIR over the vanilla RHO-LM motivated use to revisited the DPO training objective and the underlying BT preference model. In particular, we propose a simple generalization of the DPO objective with normalization from the reference model loss. In particular, inspired by the formula in Equation. 3, we propose the following DavIR-DPO loss:

$$\begin{aligned} \mathcal{L}_{\text{DavIR-DPO}}(\pi_{\theta}; \pi_{\text{ref}}) &= -\mathbb{E} \left[\log \sigma \left(\beta \log \frac{\pi_{\theta}(y_w | x)}{\pi_{\text{ref}}(y_w | x)} / |\log \pi_{\text{ref}}(y_w | x)| \right) \right. \\ &\quad \left. - \beta \log \frac{\pi_{\theta}(y_l | x)}{\pi_{\text{ref}}(y_l | x)} / |\log \pi_{\text{ref}}(y_l | x)| \right]. \end{aligned} \quad (4)$$

We remark that concurrent research on regularizing the DPO loss by the length of the responses has been proposed (Park et al., 2024).

4 Experiments and Results

4.1 Experimental Setup

Training Dataset. Training datasets used in the current study are shown in Table. 4. Note that both Alpaca-4 and Alpaca-3.5 were proposed in (Taori et al., 2023), with the same prompts but different responses generated by GPT-4 and GPT-3.5-Turbo respectively.

Test Dataset and Evaluation Method. For open-domain freeform QA style evaluation of LLaMA models, our test set is an amalgamation of

		Alpaca		GSM8K		LIMA		MBPP	
		RHO-LM	DavIR	RHO-LM	DavIR	RHO-LM	DavIR	RHO-LM	DavIR
gemma-2b	spearman ↓	0.75	0.30	0.58	0.06	0.20	0.02	0.39	0.62
	pearson ↓	0.64	0.33	0.59	0.07	0.23	0.01	0.41	0.68
gemma-2-2b	spearman ↓	0.83	0.47	0.66	0.45	0.17	0.01	0.89	0.39
	pearson ↓	0.64	0.46	0.65	0.44	0.18	0.10	0.81	0.33

Table 2: **DavIR reduces length dependency from the RHO-LM objective.** Absolute Pearson correlation and Spearman rank correlation of entropy and loss with respect to number of tokens in a given document. See Appendix. B for more detail.

DPO Type	Reference	Resp. Len. Diff.
Vanilla	(Rafailov et al., 2023)	0.38
AOT	(Melnyk et al., 2024)	0.12
APO (Down)	(D’Oosterlinck et al., 2024)	0.36
APO (Zero)	(D’Oosterlinck et al., 2024)	0.39
EXO (Pair)	(Ji et al., 2024)	0.42
Hinge	(Liu et al., 2024)	0.39
IPO	(Azar et al., 2023)	-0.10
NCA	(Chen et al., 2024)	0.29
Robust	(Chowdhury et al., 2024)	0.38
SPPO (Hard)	(Wu et al., 2024)	-0.11
DavIR	Here	0.07

Table 3: **Pearson correlation of DPO objective against difference in response length for different flavors of DPO loss type in UltraFeedback (Cui et al., 2024)**

Name	Number	Source
LIMA (Zhou et al., 2023)	1K	Human
Alpaca-4 (Taori et al., 2023)	52K	GPT-4
Alpaca-3.5 (Taori et al., 2023)	52K	GPT-3.5-Turbo
GSM8K (Cobbe et al., 2021)	7.5K	Human
UltraFeedback (Cui et al., 2024)	61K	GPT-4

Table 4: **Training datasets used for experimental validation of DavIR.**

800 prompts from HH-RLHF, Koala, Self-Instruct, Open Assistant, and Vicuna, covering multiple aspects of daily use, such as generating, math, coding, and instruction-following. The model generated responses were evaluated either with GPT-4 (adjusted for positional bias in evaluation prompt) or human evaluator (blind ranking) as referee. Note that only 100 questions were randomly selected for human evaluation (20 questions per dataset). The performance of models trained with DavIR filtered dataset is compared against either 1) same based model trained with other data selection method as in Fastchat (Zheng et al., 2023), or 2) against frozen model (e.g. Text-Davinci-003) as in AlpacaEval (Li et al., 2023b). Experiments with Gemma and Zephyr models were evaluated using

AlpacaEval2.0 (Dubois et al., 2024).

Models and Baselines. For IFT/SFT experiments, we used LLaMA-7B, LLaMA-13B (Touvron et al., 2023), Gemma-2B (Team et al., 2024a) models as our base models π_{base} . DavIR is compared against a wide range of baseline data selection methods (See Table. 5): 1) full dataset, 2) random sampling, 3) ChatGPT-based data filtering (Chen et al., 2023b), 4) RHO-LM (Mindermann et al., 2022), 5) EL2N (Paul et al., 2023), 6) Forgetting score (Toneva et al., 2019), and 7) Influence function-based DataInf (Kwon et al., 2024). Note that, for ChatGPT-based data filtering approach, a specific version of ChatGPT API was prompted to assign an integer quality score (1-5) to each data point in the Alpaca-3.5 dataset. To avoid introducing additional variabilities due to changes of ChatGPT API, for comparison against ChatGPT-based data filtering, we did not generate new data by querying ChatGPT, but instead directly used the 9k subset of Alpaca-3.5 reported in (Chen et al., 2023b). The comparison with which is shown in Figure. 1 where ChatGPT-based data filtering is referred to as “GPT-Select”. DPO experiments were conducted using Zephyr-7B-SFT (Tunstall et al.) and compared against training with the vanilla DPO objective.

4.2 DavIR in SFT

4.2.1 Impact of Length Normalization in DavIR

We first demonstrate the effect of normalization in the DavIR objective as compared to the RHO-LM objective. As shown in Figure. 2, while LLaMA models trained on subset of data selected using the RHO-LM objective can achieve comparable performance to that of the model trained on the full dataset, DavIR outperforms the full dataset baseline by a wide margin.

Name	Model Dep.	Val. Set	Method/Metric	Domain	Metric Selection
EL2N	Loss	No	ℓ_2 Loss between logits and 1-hot label	Image Class.	Larger
Forgetting	Loss	No	# epochs were data loss increase	Image Class.	Larger
DataInf	Gradient	Yes	Approx. influence function	Text Generation	Most Negative
AlpaGasus	ChatGPT	No	ChatGPT scores 0-5	Text Generation	Larger
RHO	Loss	Yes	Eq. (1)	Image Class./NLP Class.	Larger
DavIR	Loss	No	Eq. (3)	Text Generation	Larger

Table 5: **Baseline core-set selection methods.**




Figure 2: **Models fine-tuned with data selected by DavIR surpass the full dataset on Alpaca3.5.** This figure shows the win score comparison between models trained with different sizes of datasets and the full dataset, as well as the improvement brought by using the normalization method. We select the model fine-tuned on the full dataset as the baseline. Win Score is computed as $1 + (N_{win} - N_{lose})/N_{total}$, with 1 being equal performance.

4.2.2 16x Compression in Freeform Chat Dataset

As show in Figure. 1, both LLaMA-7B and LLaMA-13B model can be effectively fine-tuned with a 3K subset sampled from the 52K Alpaca dataset using DavIR. A natural question to ask is whether this is a result of simply reducing redundancy in the training dataset, which could also be achieved by simply randomly sampling the dataset. To address this question, we compared performance of DavIR against random sampling and fine-tuning on full Alpaca-4 dataset using Text-Davinci-003 as a frozen baseline model. We show in Figure. 3 that the number of training data, when randomly sampled, improve model performance logarithmically, dramatically under-performing the proposed method.

To further demonstrate DavIR’s effectiveness against other baseline methods, we compared DavIR’s performance scaling across number of training data (selected from Alpaca dataset) against 4 other core-set selection methods (EL2N, Forgetting Score, DataInf, RHO), evaluated on Gemma-




Figure 3: **DavIR significantly out perform random sampling.** Using Text-Davinci-003 as the frozen baseline model, we show that performance of random selection of the Alpaca-4 dataset scales logarithmically with number of training data, significantly under-performing DavIR. Note that the x-axis is log-scale. Win Rate is computed as N_{win}/N_{total} , where N_{win}, N_{total} are number of win and total number of test data.

2B model using AlpacaEval. As shown in Table. 6, DavIR is the only method that consistently outperforms full dataset baseline across number of training samples. Even in the low data regime (less than 5K selected from the 52K dataset) when DavIR is not the best performing method, its performance gap with the best performing method is small. We’d also like to emphasize that computing the DavIR score requires only computing validation losses. In contrast, the second best performing algorithm (DataInf) requires significantly more compute as it requires computing Influence Functions via gradient and approximated Hessian. To show that the performance gain of DavIR compared against other methods is *statistically* significant, we estimated the 95% confidence interval of the AlpacaEval score via bootstrap sampling (shown in Table. 12). Since the bootstrap sampled distribution of AlpacaEval score is highly Gaussian, we performed t-test between the sample distributions of AlpacaEval scores, will DavIR beating almost all baseline methods across number of data with very low p-values, providing conclusive evidence of the effectiveness of DavIR. Refer to Appendix. E

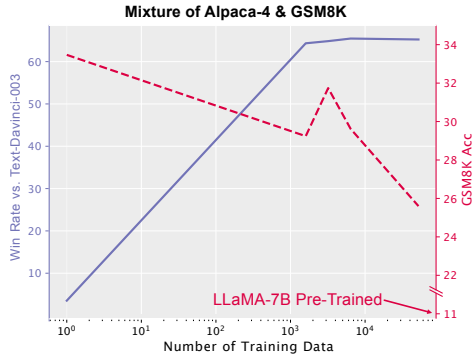


Figure 4: **Data mixing with LLaMA-7B and DavIR.** The x-axis represents the number of selected Alpaca-4 data points, plotted on a logarithmic scale.

for more details on statistical analysis.

Finally, to demonstrate the effectiveness of DavIR on other dataset, we performed data subset selection on the LIMA dataset. We show that even for a carefully curated dataset (1K), DavIR is still able to achieve a 3x compression (300) while achieving comparable performance to training on the full dataset (win score 1.01).

4.2.3 Balancing Open-Domain QA and Mathematical Reasoning

A key application of core-set selection in LLM in production is data flywheel scenario, where a constant stream of additional training data for multiple domains need to be filtered and combined to produce the best and most-balanced model for a wide range of downstream tasks. To that end, we evaluated the performance of LLaMA-7B model trained on a combination of the full GSM8K dataset (mathematical reasoning) and DavIR-filtered Alpaca-4 subset (freeform QA). As shown in Figure. 4, while we do observe the issue of “alignment tax” (Casper et al., 2023) where increasing Alpaca-4 data size caused a slight decrease in GSM8K accuracy, DavIR offers the flexibility for LLM developers to control the balance of open-domain QA with mathematical reasoning capabilities. In particular, the addition of 3.2K Alpaca-4 data (using 16.7% of total training data) boosts open-domain QA performance from <10% win-rate to >60% win-rate, at the cost of 2% reduction of GSM8K accuracy as compared to the model trained solely on GSM8K training set.

4.2.4 Generalization between Models

As shown above, DavIR is highly effective across both model sizes (LLaMA-7B/-13B) as well as

model families (LLaMA, Gemma). However, as DavIR was fundamentally motivated by the hypothesis that the best post-training data must be model-dependent, we sought to examine the data selected by different models. Comparing the best data subset selected using LLaMA-7B and LLaMA-13B models, we observe that only 516 of the top 800 data are scored highly by both models (see Table. 9 in Appendix. D), with the largest difference stemming from mathematical reasoning-related prompts as shown in Table. 10 in Appendix. D. Given that both models belong to the LLaMA family of model, and were presumably trained using similar training recipes (architecture, hyperparameters, datasets), we expect the discrepancy between data subsets selected by different models to widen when comparing between models of different families. This provides support for our intuition that the effectiveness in steering pre-trained model is highly dependent on the capabilities of the pre-trained model itself.

To probe the model-dependency of the post-training data selection from a different perspective, we explored a relaxed version of the DavIR algorithm. In this relaxed version, the base, reference model and re-trained models are allowed to be from different pre-trained models. For example, instead of using the LLaMA-7B base model throughout the data selection and re-training process, we experimented with computing DavIR score between LLaMA-7B base model and LLaMA-13B model trained on all of Alpaca dataset, and re-trained LLaMA-7B base model on the selected subset (and similarly for other combinations of base, reference and re-trained models). As shown in Table. 7, any misalignment in the models used in DavIR algorithm resulted in a decrease in performance, providing further support to the model-dependency of the optimal post-training data subset.

4.3 DavIR in DPO

We trained Zephyr-7B-SFT (Tunstall et al.) on UltraFeedback (Cui et al., 2024) paired preference dataset using both vanilla DPO objective as well as the DavIR-DPO objective in Equation. 4. Both models are evaluated on AlpacaEval against the text-davinci-003 model. We present the result in Table. 8 show that Zephyr trained using the DavIR-DPO objective outperforms that using the vanilla DPO objective, especially when evaluated using length-controlled metric (Dubois et al., 2024). Note that Zephyr model class was chosen for the DPO

Method	Selection	3000	4000	5000	6000	7000	8000	9000	10000
Random	N/A	10.6	12.7	15.9	16.1	17.0	16.3	16.7	17.6
EL2N	Lowest	10.3	13.5	17.6	17.6	19.0	18.4	18.7	18.5
	Highest	10.0	11.1	11.3	11.8	12.4	13.3	13.9	14.3
Forgetting	Lowest	10.4	12.8	16.7	16.6	17.1	17.0	17.7	16.7
	Highest	9.5	12.4	13.4	15.5	16.7	17.8	18.9	18.2
DataInf	Lowest	10.3	12.2	13.9	15.0	15.0	14.4	15.3	15.4
	Highest	10.3	13.7	15.9	18.6	18.7	19.8	18.5	18.8
RHO	Lowest	9.6	12.1	13.7	15.3	15.7	16.1	16.2	16.6
	Highest	9.9	12.3	14.5	15.8	15.3	15.5	16.5	16.7
DavIR	Lowest	10.6	12.7	14.2	14.3	13.1	12.0	13.2	13.6
	Highest	10.8	13.1	17.1	20.2	20.2	20.4	20.2	19.7
Full	N/A	<i>18.3</i>	<i>18.3</i>	<i>18.3</i>	<i>18.3</i>	<i>18.3</i>	<i>18.3</i>	<i>18.3</i>	<i>18.3</i>

Table 6: **Comparing DavIR to baselines for post-training Gemma-2B model on Alpaca dataset.** Performance reported here is the AlpacaEval win-rate against GPT-4. For completeness, we included the performances of Gemma-2B trained on Alpaca subsets selected by choosing both lowest and highest metric values (EL2N, Forgetting, DataInf, RHO, DavIR). Note that comparison against AlpaGasus is shown in Fig. 1, and is omitted here because we only have access to the full 9K+ subset reported by the authors of (Chen et al., 2023b) which limits our ability to perform ablation across number of data points. Refer to Appendix. E for statistical analysis of the performance comparison between DavIR and other baseline methods.

π_{base}	π_{ref}	$\pi_{retrain}$	Against	Data	Win Score
13B	13B + Alp-4	13B	13B + Alp-4	8K	1.2
7B	7B + Alp-4	7B	7B + Alp-4	8K	1.125
13B	13B + Alp-4	7B	7B + Alp-4	8K	0.825
7B	13B + Alp-4	7B	7B + Alp-4	8K	1.01

Table 7: **DavIR performs best when base, reference and re-trained models share the same pre-trained backbone.** For simplicity, 7B/13B refer to LLaMA-7B and LLaMA-13B respectively, $\pi_{retrain}$ refers to the pre-trained model that is trained with the DavIR selected data subset. Win Score is computed against models fine-tuning on D_{full} as shown in the ‘‘Against’’ column. Note that the first two rows correspond to experiments where there is no model mismatch and the bottom two rows correspond to the relaxed DavIR algorithm with model mismatch.

experiments, as opposed to LLaMA and Gemma as in the SFT experiments, for it differentiates between pretrained, instruction fine-tuned and DPO post-trained models, thus helping us isolate the effect of length-normalization at the DPO training stage.

5 Conclusion

We introduce DavIR, a model-based data selection method for LLM fine-tuning that focuses on ‘‘learnability’’ of data points given a base pre-trained model. We show that DavIR is closely related to, and is a generalization of, the Implicit Reward Model concept proposed in Direct Preference Optimization. By comparing DavIR to a wide range

	Length-Controlled Win-Rate	Win-Rate
Zephyr DPO	57.23	82.31
Zephyr DavIR-DPO	61.83	82.96

Table 8: **Comparing Zephyr trained on Davir-DPO vs. vanilla DPO objective.**

of data selection baselines, we demonstrate its effectiveness across models, data domain and data mixtures. Finally, we show that, by incorporating the proposed normalization back to the DPO objective, we are able to improve DPO performance after the supervised fine-tuning stage of LLM training.

6 Limitations & Discussions

Integration of DavIR to Data Flywheel As briefly discussed above, data compression techniques such as DavIR serve a critical, albeit incomplete, role in the data flywheel of training LLMs. In particular, DavIR does not take into account other aspects of data selection such as quality and diversity. In practice, DavIR needs to be used in conjunction with methods such as weighted sampling and prompt classification to ensure that the core-set selection is performed in a manner that does not artificially bias the distribution of the selected data. In this work, we provided a simple example of data mixture between Alpaca and GSM8K which hints at the importance of using DavIR in a manner that is conscious of data diversity. In production, how-

ever, more careful design of the data pipeline is required, for which DavIR could serve as the data compression module.

Application of DavIR to Reasoning Tasks A key limitation of DavIR is that its effectiveness varies based on the application domain of the training dataset. In particular, when we applied DavIR to compressing the GSM8K training dataset alone for LLaMA models, we did not observe a clear performance gain with subset of the training data. In fact, as shown in Figure. 5, we observed almost linear scaling of number of GSM8K training sample and the GSM8K evaluation accuracy, suggesting that the GSM8K dataset was *in-compressible* with LLaMA-7B using DavIR. We hypothesize that this could be caused by LLaMA-7B having insufficient underlying mathematical reasoning capabilities, leading to very large training data requirement. However, we could not rule out the possibility that perhaps Cross-Entropy Loss is a poor metric for how well data related to mathematical reasoning has been learnt by a given model, thereby rendering the normalized score metric unable to capture “learnability” of such data. We leave explorations of alternate metrics to Cross-Entropy Loss for future works.

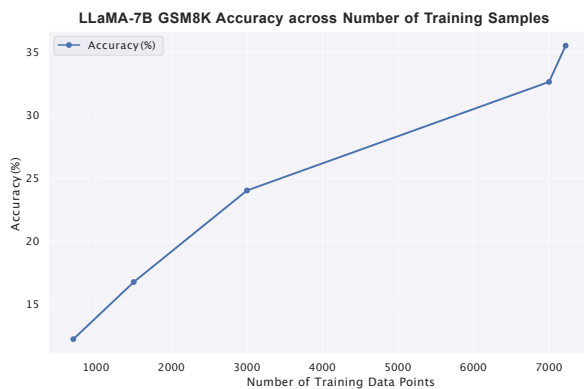


Figure 5: GSM8K training data was in-compressible with LLaMA-7B.

7 Ethics Statement

This research aims to provide a model-based algorithm for core-set selection of LLM alignment training data. Experimental validation in the current work leverages previously published datasets, and are employed in accordance with their intended use cases. While these datasets are widely used, we acknowledge that we cannot fully ascertain the extent to which they may contain discriminatory,

biased, or sensitive material.

Responsible Usage: Data selection via DavIR is based purely on model’s perceived degree of understanding, and makes no assumption about safety of the original training data. As such, caution must be exercised when deploying DavIR in production to ensure that necessary safety practices are adopted both before and after using DavIR for subset selection.

AI Assistant Usage: Claude-3.5-Sonnet and GPT-4o were used for grammatical correction in the current manuscript.

References

- Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. 2023. [A general theoretical paradigm to understand learning from human preferences](#). *Preprint*, arXiv:2310.12036.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *Preprint*, arXiv:2204.05862.
- Vighnesh Birodkar, Hossein Mobahi, and Samy Bengio. 2019. Semantic redundancies in image-classification datasets: The 10% you don’t need. *arXiv preprint arXiv:1901.11409*.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yihan Cao, Yanbin Kang, and Lichao Sun. 2023. [Instruction mining: High-quality instruction data selection for large language models](#). *Preprint*, arXiv:2307.06290.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David

- Lindner, Pedro Freire, et al. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Hao Chen, Yiming Zhang, Qi Zhang, Hantao Yang, Xiaomeng Hu, Xuetao Ma, Yifan Yanggong, and Junbo Zhao. 2023a. Maybe only 0.5% data is needed: A preliminary exploration of low training data instruction tuning. *arXiv preprint arXiv:2305.09246*.
- Huayu Chen, Guande He, Lifan Yuan, Ganqu Cui, Hang Su, and Jun Zhu. 2024. [Noise contrastive alignment of language models with explicit rewards](#). *Preprint*, arXiv:2402.05369.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Sriniwasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. 2023b. [Alpagasus: Training a better alpaca with fewer data](#). *Preprint*, arXiv:2307.08701.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Sayak Ray Chowdhury, Anush Kini, and Nagarajan Natarajan. 2024. [Provably robust dpo: Aligning language models with noisy feedback](#). *Preprint*, arXiv:2403.00409.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- 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.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2024. [Ultrafeedback: Boosting language models with scaled ai feedback](#). *Preprint*, arXiv:2310.01377.
- Jiaxi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. 2023. [Chatlaw: Open-source legal large language model with integrated external knowledge bases](#). *Preprint*, arXiv:2306.16092.
- Karel D’Oosterlinck, Winnie Xu, Chris Develder, Thomas Demeester, Amanpreet Singh, Christopher Potts, Douwe Kiela, and Shikib Mehri. 2024. [Anchored preference optimization and contrastive revisions: Addressing underspecification in alignment](#). *Preprint*, arXiv:2408.06266.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B. Hashimoto. 2024. [Length-controlled alpacaeval: A simple way to debias automatic evaluators](#). *Preprint*, arXiv:2404.04475.
- Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. [Koala: A dialogue model for academic research](#). Blog post.
- Sariel Har-Peled and Akash Kushal. 2005. Smaller coresets for k-median and k-means clustering. In *Proceedings of the twenty-first annual symposium on Computational geometry*, pages 126–134.
- Haozhe Ji, Cheng Lu, Yilin Niu, Pei Ke, Hongning Wang, Jun Zhu, Jie Tang, and Minlie Huang. 2024. [Towards efficient exact optimization of language model alignment](#). *Preprint*, arXiv:2402.00856.
- Yunjie Ji, Yong Deng, Yan Gong, Yiping Peng, Qiang Niu, Lei Zhang, Baochang Ma, and Xiangang Li. 2023. Exploring the impact of instruction data scaling on large language models: An empirical study on real-world use cases. *arXiv preprint arXiv:2303.14742*.
- Seungjae Jung, Gunsoo Han, Daniel Wontae Nam, and Kyoung-Woon On. 2024. [Binary classifier optimization for large language model alignment](#). *Preprint*, arXiv:2404.04656.
- Krishnateja Killamsetty, Xujiang Zhao, Feng Chen, and Rishabh Iyer. 2021. Retrieve: Coreset selection for efficient and robust semi-supervised learning. *Advances in Neural Information Processing Systems*, 34:14488–14501.
- Yongchan Kwon, Eric Wu, Kevin Wu, and James Zou. 2024. [Datainf: Efficiently estimating data influence in lora-tuned llms and diffusion models](#). *Preprint*, arXiv:2310.00902.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. [Openassistant conversations – democratizing large language model alignment](#). *Preprint*, arXiv:2304.07327.
- Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Luke Zettlemoyer, Omer Levy, Jason Weston, and Mike Lewis. 2023a. [Self-alignment with instruction back-translation](#). *Preprint*, arXiv:2308.06259.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023b. Alpacaeval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.

- Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J. Liu, and Jialu Liu. 2024. [Statistical rejection sampling improves preference optimization](#). *Preprint*, arXiv:2309.06657.
- Igor Melnyk, Youssef Mroueh, Brian Belgodere, Mattia Rigotti, Apoorva Nitsure, Mikhail Yurochkin, Kristjan Greenewald, Jiri Navratil, and Jerret Ross. 2024. [Distributional preference alignment of llms via optimal transport](#). *Preprint*, arXiv:2406.05882.
- Sören Mindermann, Jan Brauner, Muhammed Razzak, Mrinank Sharma, Andreas Kirsch, Winnie Xu, Benedikt Höltgen, Aidan N. Gomez, Adrien Morisot, Sebastian Farquhar, and Yarin Gal. 2022. [Prioritized training on points that are learnable, worth learning, and not yet learnt](#). *Preprint*, arXiv:2206.07137.
- Baharan Mirzasoleiman, Jeff Bilmes, and Jure Leskovec. 2020. Coresets for data-efficient training of machine learning models. In *International Conference on Machine Learning*, pages 6950–6960. PMLR.
- Yonatan Oren, Shiori Sagawa, Tatsunori B. Hashimoto, and Percy Liang. 2019. [Distributionally robust language modeling](#). *Preprint*, arXiv:1909.02060.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Ryan Park, Rafael Rafailov, Stefano Ermon, and Chelsea Finn. 2024. Disentangling length from quality in direct preference optimization. *arXiv preprint arXiv:2403.19159*.
- Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. 2021. Deep learning on a data diet: Finding important examples early in training. *Advances in Neural Information Processing Systems*, 34:20596–20607.
- Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. 2023. [Deep learning on a data diet: Finding important examples early in training](#). *Preprint*, arXiv:2107.07075.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). *Preprint*, arXiv:2305.18290.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2022. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. 2022. Beyond neural scaling laws: beating power law scaling via data pruning. *Advances in Neural Information Processing Systems*, 35:19523–19536.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikuła, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024a. [Gemma: Open models based on gemini research and technology](#). *Preprint*, arXiv:2403.08295.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan,

- Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshv, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonnell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iversen, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024b. [Gemma 2: Improving open language models at a practical size](#). *Preprint*, arXiv:2408.00118.
- Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. 2019. [An empirical study of example forgetting during deep neural network learning](#). *Preprint*, arXiv:1812.05159.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Shengyi Huang, Kashif Rasul, Alvaro Bartolome, Alexander M. Rush, and Thomas Wolf. [The Alignment Handbook](#).
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Han-naneh Hajishirzi. 2022. [Self-instruct: Aligning language model with self generated instructions](#). *arXiv preprint arXiv:2212.10560*.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2018. [Neural network acceptability judgments](#). *arXiv preprint arXiv:1805.12471*.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kam-badur, David Rosenberg, and Gideon Mann. 2023. [Bloomberggpt: A large language model for finance](#). *Preprint*, arXiv:2303.17564.
- Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. 2024. [Self-play preference optimization for language model alignment](#). *Preprint*, arXiv:2405.00675.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V. Le, Tengyu Ma, and Adams Wei Yu. 2023a. [Doremi: Optimizing data mixtures speeds up language model pretraining](#). *Preprint*, arXiv:2305.10429.
- Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy Liang. 2023b. [Data selection for language models via importance resampling](#). *Preprint*, arXiv:2302.03169.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. [Wizardlm: Empowering large language models to follow complex instructions](#). *arXiv preprint arXiv:2304.12244*.
- Shuo Yang, Zeke Xie, Hanyu Peng, Min Xu, Mingming Sun, and Ping Li. 2022. [Dataset pruning: Reducing training data by examining generalization influence](#). *arXiv preprint arXiv:2205.09329*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206*.

Zhanhui Zhou, Jie Liu, Jing Shao, Xiangyu Yue, Chao Yang, Wanli Ouyang, and Yu Qiao. 2024. [Beyond one-preference-fits-all alignment: Multi-objective direct preference optimization](#). *Preprint*, arXiv:2310.03708.

A Details on Human Evaluation of Model Performance

For open-domain freeform QA style evaluation of LLaMA models (an amalgamation of 5 test datasets), we used a combination of LLM (GPT-4) and human as referee. For human evaluation, 20 questions per test dataset were randomly selected as prompts, resulting in total 100 prompts. One human annotator (unpaid, college educated, age 20-25, proficient in English) was provided side-by-side comparison of two responses generated by two models for each question and asked to determine whether the response on the “left” is better/same/worse (win/tie/lose) than the response on the “right”. The annotator is blind to the identity of the model for which the responses were generated, and the responses were randomly ordered.

B Effect of Normalization for Score Function

In Figure. 6, we compare the sequence length of the Alpaca dataset ranked by either un-normalized and normalized score functions. It is apparent that without normalization, data with the highest scores (low ranking) correspond to data with very short sequence lengths. In contrast, the introduction of normalization completely removes the




Figure 6: Effect of normalization on score function and sequence length. The relationship between the sequence length and ranking for both with and without normalizing score function is shown for the (left) top-6400 subset and (right) full 52K Alpaca-4. We observe that, without normalization, data with highest score values have noticeably short sequence length, which is resolved by normalization.

C Choice of Denominator in Normalized Score Function Does Not Impact Ranking

Proposition 1. *Choosing either $\mathcal{L}_{ref}(x, y)$ or $\mathcal{L}_{base}(x, y)$ as the denominator for normalization does not affect the ranking of the learnability score.*

Specifically, if

$$\frac{\mathcal{L}_{base}(x_1, y_1) - \mathcal{L}_{ref}(x_1, y_1)}{\mathcal{L}_{base}(x_1, y_1)} > \frac{\mathcal{L}_{base}(x_2, y_2) - \mathcal{L}_{ref}(x_2, y_2)}{\mathcal{L}_{base}(x_2, y_2)} \quad (5)$$

then it also holds that

$$\frac{\mathcal{L}_{base}(x_1, y_1) - \mathcal{L}_{ref}(x_1, y_1)}{\mathcal{L}_{ref}(x_1, y_1)} > \frac{\mathcal{L}_{base}(x_2, y_2) - \mathcal{L}_{ref}(x_2, y_2)}{\mathcal{L}_{ref}(x_2, y_2)} \quad (6)$$

Proof. Assume that

$$\frac{\mathcal{L}_{base}(x_1, y_1) - \mathcal{L}_{ref}(x_1, y_1)}{\mathcal{L}_{base}(x_1, y_1)} > \frac{\mathcal{L}_{base}(x_2, y_2) - \mathcal{L}_{ref}(x_2, y_2)}{\mathcal{L}_{base}(x_2, y_2)} \quad (7)$$

which can be rewritten as

$$1 - \frac{\mathcal{L}_{ref}(x_1, y_1)}{\mathcal{L}_{base}(x_1, y_1)} > 1 - \frac{\mathcal{L}_{ref}(x_2, y_2)}{\mathcal{L}_{base}(x_2, y_2)}. \quad (8)$$

This implies

$$\frac{\mathcal{L}_{ref}(x_1, y_1)}{\mathcal{L}_{base}(x_1, y_1)} < \frac{\mathcal{L}_{ref}(x_2, y_2)}{\mathcal{L}_{base}(x_2, y_2)}. \quad (9)$$

Taking the reciprocal of both sides, we get:

$$\frac{\mathcal{L}_{base}(x_1, y_1)}{\mathcal{L}_{ref}(x_1, y_1)} > \frac{\mathcal{L}_{base}(x_2, y_2)}{\mathcal{L}_{ref}(x_2, y_2)}. \quad (10)$$

Subtracting one from both sides, we obtain:

$$\frac{\mathcal{L}_{base}(x_1, y_1) - \mathcal{L}_{ref}(x_1, y_1)}{\mathcal{L}_{ref}(x_1, y_1)} > \frac{\mathcal{L}_{base}(x_2, y_2) - \mathcal{L}_{ref}(x_2, y_2)}{\mathcal{L}_{ref}(x_2, y_2)}. \quad (11)$$

Thus, we have shown that normalizing by either \mathcal{L}_{base} or \mathcal{L}_{ref} does not affect the ranking of the learnability score. \square

D Analysis of Data Selected via DavIR

To explore what types of data are required by the model during the SFT process, we conducted a further analysis of the data selected by the 7B and 13B models. In Table 9, we show the number and percentage overlap of the data points selected by both LLaMA-7B and LLaMA-13B models.

No. Data	800	1,600	3,200	6,400	9,600	26,000	39,000
No. Overlap	516	1,111	2,399	5,123	8,033	24,358	38,075
Percent Overlap	64.5%	69.4%	74.9%	80.0%	83.7%	93.7%	97.6%

Table 9: The number and percentage overlap of data points selected by LLaMA-7B and LLaMA-13B.

Constituency Parsing via Benepar. We used the Benepar to parse constituency of the top 800 data points selected by the LLaMA-7B and LLaMA-13B models, which have 516 overlapping data points as shown in Table. 9. Benepar decomposes natural language statements into hierarchical representation of constituency, from which we visualize the top two level verb predicate and noun objects.

Upon close examination of Figure. 7a and Figure. 7b, we observe that, comparing the more powerful LLaMA-13B against LLaMA-7B, fewer creation tasks (e.g. write, generate, create) and more interpretative tasks (e.g. explain, describe) data were selected, with slighter more diverse long tail tasks.

Prompt Category Classification with LLM. Constituency parsing via Benepar, while helpful, does not effectively convey the semantics of the training data. Instead, we employed GPT-4 as classifier for a more precise semantically-oriented task classification.

In particular, we first classified Alpaca’s seed instructions into 7 primary categories. The categories are then used to further classify the first 800 data entries selected by models 7B and 13B. As shown in Table. 10, data in the "Problem Solving and Math" category had the most significant change between the two base models, increasing by 76.9% from 7B to 13B. We hypothesize that this could be due to the substantial difference in mathematical and reasoning capabilities between the 7B and 13B models, thereby increasing the learnability of these SFT data for the 13B model.

Category	7B	13B	Δ ($\Delta\%$)
Programming and Coding	60	56	-4(-6.6%)
Planning and Organization	63	57	-6(-9.5%)
Knowledge and Information Extraction	275	296	+21(+7.6%)
Language and Text Processing	53	45	-8(-15.1%)
Creative Writing and Entertainment	311	286	-25(-8.0%)
Problem Solving and Math	26	46	+20(+76.9%)
Recommendations and Suggestions	9	8	-1(-11.1%)
Others	3	6	-


Table 10: Comparison of first 800 data selected by LLaMA-7B and LLaMA-13B model based on category.

Finally, in Table. 11, we provide examples of data in Alpaca-4 with the highest normalized scores as computed using LLaMA-7B base model.


E Statistical Analysis of AlpacaEval

We first note that both the data selection procedure of DavIR and the model inference and evaluation procedures are deterministic (greedy decoding). Therefore, to establish a statistical significance analysis of the comparison between DavIR’s performance to other methods, we performed bootstrap estimation with 1000 samples from the 805 questions of the AlpacaEval dataset, giving us the 95% confidence interval shown in Table. 12.

As the bootstrap estimations of the performances on AlpacaEval is highly Gaussian for all experiments with similar variance, we performed t-test between DavIR and all other baseline method. The p-values, shown in Table. 13 below show that the performance gain of DavIR over other methods is indeed significant across number of samples (with p-values very close to 0).



(a) Top 800 data points selected by 7B model



(b) Top 800 data points selected by 13B model

Figure 7: Comparison of top 800 data points selected by different models

Category	Score	Instruction	Input	Output
Knowledge	12.4	Identify the odd one out: Twitter, Instagram, Telegram.		Telegram.
Knowledge	9.80	Convert the following sentence into the present continuous tense.	He reads books.	He is reading books.
Knowledge	8.78	Identify the most suitable adverb for the following sentence.	She works.	She works diligently.
Knowledge	8.21	Analyze the tone of the following sentences.	I am so upset with you right now.	The tone of the sentences is angry and accusatory.
Transcription	7.38	Rewrite the following sentence using active voice.	The news report was read by the captain.	The captain read the news report.
Knowledge	7.59	How did Julius Caesar die?		Julius Caesar was assassinated by a group of up to 60 conspirators, led by Gaius Cassius Longinus and Marcus Junius Brutus, in the Senate House on the Ides of March (15 March) of 44 BC.
Mathematics	6.09	Calculate the total surface area of a cube with a side length of 5 cm.		Total surface area of the cube: 150 cm ² .
Knowledge	5.96	Provide one example of a cultural practice.		An example of a cultural practice is the Japanese tea ceremony.
Mathematics	4.96	Use the given data to calculate the median.	[2, 3, 7, 8, 10].	The median of the given data is 7.
Knowledge	5.15	What is the product of 6 and 2?		The product of 6 and 2 is 12.
Knowledge	5.11	Give two examples of a liquid.		Two examples of liquids are water and oil.
Correction	5.12	Evaluate this sentence for spelling and grammar mistakes.	He finnished his meal and left the resturant.	He finished his meal and left the restaurant.
Knowledge	4.95	Arrange the given numbers in ascending order.	2, 4, 0, 8, 3.	0, 2, 3, 4, 8.
Knowledge	4.66	Given a set of numbers, find the maximum value.	Set: {10, 3, 25, 6, 16}.	The maximum value is 25.
Mathematics	4.53	Convert from Celsius to Fahrenheit.	Temperature in Celsius: 15.	Temperature in Fahrenheit: 59.
Knowledge	4.16	Transform the following sentence into the passive voice.	I bought a book.	A book was bought by me.
Mathematics	4.12	Convert the given equation into an algebraic expression.	$3x + 5y = 9$.	$3x + 5y - 9 = 0$.
Knowledge	4.10	What are the three primary colors?		The three primary colors are red, blue, and yellow.
None	3.93	Render a 3D model of a house.		<nooutput>This type of instruction cannot be fulfilled by a GPT model.

Table 11: Example data with highest DavIR scores selected with LLaMA-7B base model.

Method	Selection	3000	4000	5000	6000	7000	8000	9000	10000
Random	N/A	10.7(8.6~12.8)	12.7(10.7~14.7)	15.9(13.6~18.3)	16.1(13.8~18.5)	17.0(14.8~19.3)	16.4(14.1~18.8)	16.7(14.5~18.9)	17.6(15.1~20.2)
EL2N	Lowest	10.4(8.5~12.2)	13.5(11.3~15.7)	17.6(15.4~19.8)	17.6(15.4~20.1)	19.0(16.5~21.6)	18.5(16.2~20.9)	18.7(16.5~21.3)	18.4(16.0~21.0)
	Highest	10.0(8.2~11.9)	11.1(9.3~13.2)	11.2(9.3~13.2)	11.8(9.7~13.9)	12.4(10.3~14.7)	13.2(11.2~15.4)	13.9(11.8~16.2)	14.3(12.3~16.7)
Forgetting	Lowest	10.4(8.6~12.3)	12.8(10.7~15.0)	16.6(14.4~19.1)	16.5(14.2~18.9)	17.1(14.8~19.5)	17.0(14.7~19.4)	17.7(15.6~20.1)	16.7(14.6~19.1)
	Highest	9.5(7.7~11.4)	12.4(10.5~14.6)	13.4(11.3~15.5)	15.5(13.2~17.9)	16.7(14.6~18.9)	17.8(15.6~20.3)	18.9(16.6~21.3)	18.2(15.7~20.5)
DataInf	Lowest	10.2(8.4~12.1)	12.3(10.3~14.4)	13.9(11.9~16.0)	15.1(12.7~17.4)	15.0(13.0~17.4)	14.3(12.1~16.6)	15.3(13.2~17.5)	15.4(13.3~17.7)
	Highest	10.3(8.4~12.2)	13.8(11.6~15.9)	15.9(13.7~18.2)	18.6(16.2~21.2)	18.8(16.2~21.2)	19.8(17.4~22.6)	18.5(16.1~20.9)	18.9(16.4~21.3)
RHO	Lowest	9.6(7.9~11.6)	12.1(10.0~14.3)	13.8(11.7~16.0)	15.3(13.1~17.5)	15.7(13.5~18.0)	16.1(13.8~18.4)	16.2(13.9~18.4)	16.6(14.2~19.0)
	Highest	9.9(8.1~11.9)	12.2(10.3~14.5)	14.6(12.3~16.8)	15.8(13.5~18.0)	15.4(13.1~17.6)	15.5(13.2~17.9)	16.5(14.2~18.9)	16.7(14.4~19.0)
DavIR	Lowest	10.5(8.7~12.5)	12.7(10.7~14.8)	14.2(12.0~16.4)	14.3(12.1~16.5)	13.1(11.1~15.3)	12.0(9.9~13.8)	13.2(11.1~15.4)	13.6(11.4~15.7)
	Highest	10.8(8.8~12.9)	13.1(11.0~15.2)	17.1(14.8~19.6)	20.2(17.7~22.7)	20.2(17.6~22.7)	20.3(17.9~22.9)	20.2(17.8~22.5)	19.7(17.3~22.1)
Full	N/A	18.3(15.9~20.6)	18.3(15.9~20.6)	18.3(15.9~20.6)	18.3(15.9~20.6)	18.3(15.9~20.6)	18.3(15.9~20.6)	18.3(15.9~20.6)	18.3(15.9~20.6)

Table 12: DavIR comparison with baselines with 95% Confidence Interval.

Method	Selection	3000	4000	5000	6000	7000	8000	9000	10000
Random	N/A	5.3E-03	2.4E-16	1.3E-93	0.0E+00	0.0E+00	0.0E+00	0.0E+00	8.8E-229
EL2N	Lowest	4.4E-26	-	-	3.1E-307	1.3E-78	1.2E-187	2.8E-134	5.7E-99
	Highest	6.0E-65	6.4E-276	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00
Forgetting	Lowest	3.9E-22	6.2E-08	5.2E-17	0.0E+00	0.0E+00	0.0E+00	1.5E-321	0.0E+00
	Highest	1.7E-160	1.8E-43	0.0E+00	0.0E+00	0.0E+00	0.0E+00	3.9E-108	1.1E-146
DataInf	Lowest	2.9E-38	2.4E-63	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00
	Highest	2.6E-29	-	-	4.4E-136	4.4E-116	1.1E-17	7.9E-180	3.8E-46
RHO	Lowest	3.5E-141	1.8E-84	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00
	Highest	6.6E-86	5.2E-67	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00
DavIR	Lowest	1.1E-12	1.5E-16	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00	0.0E+00
Full	N/A	-	-	-	3.9E-189	6.1E-193	1.8E-224	3.2E-209	6.2E-124

Table 13: p-values of t-test comparing DavIR and all other selection methods presented in Table. 6 and Table. 12. Note that since the hypothesis is that DavIR out-performs other methods, only results where DavIR out-performs the baseline methods have corresponding p-values.