

Data Curation Alone Can Stabilize In-context Learning

Ting-Yun Chang and Robin Jia
University of Southern California
{tingyun, robinjia}@usc.edu

Abstract

In-context learning (ICL) enables large language models (LLMs) to perform new tasks by prompting them with a sequence of training examples. However, it is known that ICL is very sensitive to the choice of training examples: randomly sampling examples from a training set leads to high variance in performance. In this paper, we show that carefully curating a subset of training data greatly *stabilizes* ICL performance without any other changes to the ICL algorithm (e.g., prompt retrieval or calibration). We introduce two methods to choose training subsets—both score training examples individually, then select the highest-scoring ones. CONDACC scores a training example by its average dev-set ICL accuracy when combined with random training examples, while DATAMODELS learns linear regressors that estimate how the presence of each training example influences LLM outputs. Across five tasks and two LLMs, sampling from *stable* subsets selected by CONDACC and DATAMODELS improves average accuracy over sampling from the entire training set by 7.7% and 6.3%, respectively. Surprisingly, the *stable* subset examples are not especially diverse in content or low in perplexity, in contrast with other work suggesting that diversity and perplexity are important when prompting LLMs.

1 Introduction

In-context learning (ICL) is a new paradigm for few-shot learning with pretrained large language models (LLMs) without any parameter updates. In ICL, an LLM can perform a new task simply by conditioning on a prompt¹ consisting of a sequence of labeled training examples. First introduced by GPT-3 (Brown et al., 2020), ICL with LLMs has reached state-of-the-art few-shot performance across many tasks (Rae et al., 2021; Smith et al.,

¹Prior work has different definitions of *prompt*; in this paper, we fix the task templates and follow Rubin et al. (2021) to denote *prompt* as a sequence of training examples for ICL.

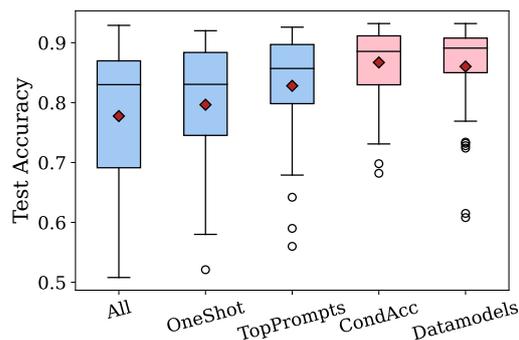


Figure 1: 4-shot ICL performance of GPTJ on SST2. Each boxplot summarizes the results of 50 sampled prompts. Compared with baselines (blue), our methods (pink) can greatly stabilize performance, having higher average accuracy (red diamonds) and lower variance.

2022; Thoppilan et al., 2022; Chowdhery et al., 2022). Compared with alternatives that use fine-tuning (Devlin et al., 2018; Schick and Schütze, 2020; Gao et al., 2020), ICL does not require task-specific training, which enables its use with very large language models, and it uses a unified model for all tasks, enabling easier deployment.

Despite its impressive few-shot performance, ICL often exhibits unintuitive behavior (Min et al., 2022). The standard ICL approach is to randomly sample a few examples from a training set to construct a prompt (Brown et al., 2020); however, prior work (Liu et al., 2021; Zhao et al., 2021; Lu et al., 2021) has found that ICL is very sensitive to the choice of training examples and their order in the prompt. ICL is also sensitive to small changes in prompt format (Chen et al., 2022).

In this paper, we show that carefully curating a smaller training dataset from a larger pool can make ICL much more stable. We define a training subset \mathcal{E} to be *stable* if randomly sampling a sequence of examples as a prompt from \mathcal{E} yields much higher average and worst-case accuracy than randomly sampling from the original training set. We propose two methods to identify such a sta-

ble subset. Our CONDACC method scores a training example by its average dev-set ICL accuracy when combined with random training examples; these scores are closely related to Data Shapley values (Ghorbani and Zou, 2019). Our DATAMODELS method fits a linear regressor that predicts the LLM’s output based on which example is present at each index in the prompt; we score a training example highly if the associated weights from the linear model indicate that its presence improves accuracy. For both methods, we then select training examples with the highest scores to form the stable subset. While some prior work improves ICL accuracy by retrieving a suitable prompt for each test example (Liu et al., 2021; Rubin et al., 2021; Su et al., 2022), we show that it is possible to achieve stably good accuracy with a randomly sampled prompt for all test examples, when given the “right” training (sub)set.

Our subset selection methods greatly improve performance across 5 classification datasets and 4 LLMs, with main experiments on GPTJ-6B (Wang and Komatsuzaki, 2021) and OPT-13B (Zhang et al., 2022a). On average, CONDACC and DATAMODELS outperform the baseline that uses the entire training set without selection (named ALL) by 7.7% and 6.3%, respectively, when comparing the average accuracy over multiple sampled prompts. In contrast, baselines that choose examples found in high-performing prompts (TOPPROMPTS) or examples that lead to high one-shot ICL accuracy (ONESHOT) to form the subsets do not perform as well (see Figure 1). Our stable subset examples generalize to out-of-distribution test data, and we can even find stable subsets for binary classification tasks that only contain examples of one label; these findings suggest that the stable subset examples help LLMs understand the overall task definition.

Finally, we study what makes stable subset examples special by analyzing sequence length, perplexity, and diversity in both raw text and embedding spaces. We find that these examples do not have abnormally long sequence lengths or high perplexities. In contrast with prior work optimizing diversity for prompt selection (Su et al., 2022; Ye et al., 2022), we find our stable subsets no more diverse than random subsets of the training data. In summary, we show that curating training data appropriately leads to more stable and accurate ICL performance; we hope future work can develop new strategies for writing such helpful ex-

amples. Code and data are publicly available at <https://github.com/terarachang/DataICL>.

2 Problem Setups

We use the original ICL formulation proposed by GPT-3, also known as the *direct* method, for all our experiments. Specifically, an LLM performs in-context learning on a new task based on a task-specific prompt \mathcal{Z} formed by concatenating K labeled training examples, i.e., $\mathcal{Z} = [z_1, \dots, z_K]$, where each z_j is a training example (x, y) consisting of an input x and label y . The LLM then makes predictions on a test input x_{test} conditioned on the prompt \mathcal{Z} followed by x_{test} , denoted by $\arg \max_{y \in \mathcal{C}} P(y | \mathcal{Z}, x_{\text{test}})$, where \mathcal{C} is the set of possible labels.

Given a training set D_{tr} , a dev set D_{dev} , a held-out test set D_{test} , and a predefined number of shots K , our goal is to select a stable training subset $\mathcal{E} \subset D_{\text{tr}}, |\mathcal{E}| > K$, such that randomly sampling a sequence of K examples from \mathcal{E} to form a prompt generally yields good performance on D_{test} .

We propose two setups, **Labeled** and **Unlabeled**. In **Labeled**, our goal is to study which training examples consistently lead to high ICL accuracy. We assume access to a large labeled $D_{\text{tr}}^L = \{(x_i, y_i)\}_{i=1}^{N_{\text{tr}}}$ of input-label pair, and a small labeled D_{dev} . **Unlabeled** is closer to the true few-shot learning setup (Perez et al., 2021), where we only have access to D_{dev} and a large unlabeled training set $\{x_i\}_{i=1}^{N_{\text{tr}}}$. We pair each input x_i with every possible label $y \in \mathcal{C}$ to create an unlabeled training set $D_{\text{tr}}^U = \{(x_i, \tilde{y}) | \tilde{y} \in \mathcal{C}\}_{i=1}^{N_{\text{tr}}}$. In both setups, our goal is to select a subset of E training examples from either D_{tr}^L or D_{tr}^U . Note that we may select examples with incorrect labels under **Unlabeled**. We only use the large labeled D_{test} for evaluating our selection methods.

3 Methods

We propose the following steps to identify a stable subset of E training examples:

1. Construct \mathcal{D}_{ICL} , a large set of M prompts, each consisting of K examples randomly sampled from D_{tr} . We then run ICL on the dev set M times given different prompts.
2. Estimate the value of each training example based on the results of Step 1. We propose two methods to do this: CONDACC (§3.1) and DATAMODELS (§3.2).

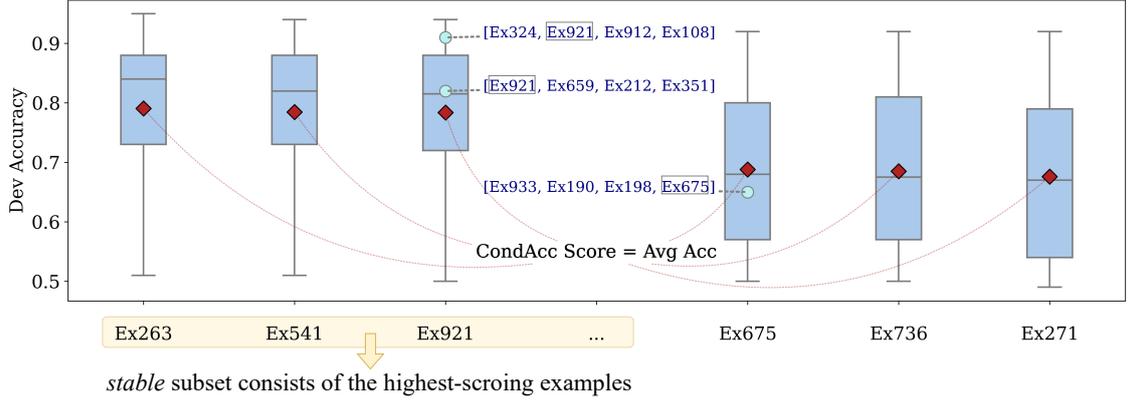


Figure 2: An overview of our CONDACC method, which scores training examples individually using its average accuracy (red diamonds) when combined with other random training examples. Each boxplot summarizes the dev-set accuracies conditioned on a training example appearing in the sampled prompts.

3. Select training examples with the highest importance scores per class to make up \mathcal{E} (§3.3).

3.1 CondAcc

We hypothesize that a good training example leads to higher accuracy, on average, when occurs in a prompt. Given a prompt \mathcal{Z} , we denote its dev set ICL accuracy as $\text{Acc}(\mathcal{Z})$. Thus, a simple way to score the i -th training example $(x_i, y_i) \in D_{\text{tr}}$ is to calculate the expected accuracy conditioned on this example appearing in a prompt \mathcal{Z} from \mathcal{D}_{ICL} :

$$s_{ca}(i) = \mathbb{E}_{\mathcal{Z} \sim \mathcal{D}_{\text{ICL}}} [\text{Acc}(\mathcal{Z}) \mid (x_i, y_i) \in \mathcal{Z}] \quad (1)$$

We ensure that each training example occurs in \mathcal{D}_{ICL} multiple times in different orders and with different examples.

In Appendix A.1, we show that Eq. 1 is similar to Data Shapley value (Ghorbani and Zou, 2019), where we define the valuation function of Data Shapley on subsets of K training examples since we focus on K -shot ICL.

3.2 Datamodels

The CONDACC method does not consider the order of training examples, which has a great impact on ICL performance. Also, it takes a simple average over the dev set, ignoring the LLM’s confidence in individual dev examples. Thus, we propose another data valuation method that leverages Datamodels (Ilyas et al., 2022) for ICL.

Ilyas et al. (2022) study a complex target model’s behavior in terms of the training data by replacing it with a linear, easy-to-analyze, proxy model called a *datamodel*. Specifically, given a subset of training data \mathcal{S} , a datamodel predicts the *outcome* of a

test input \bar{x} if training the target model on \mathcal{S} . The *outcome* $f(\bar{x}; \mathcal{S})$ is defined as the margin of the correct class, i.e., the logit for the correct class minus the highest incorrect logit. To train datamodels, Ilyas et al. (2022) first create a dataset consisting of subset-outcome pairs $(\mathcal{S}, f(\bar{x}; \mathcal{S}))$, requiring training the target model from scratch multiple times on different subsets to obtain the outcomes. In our work, the target model is a pretrained LLM, and is inference-only during ICL. Thus, our dataset collection only requires inference of the LLM multiple times with different prompts in \mathcal{D}_{ICL} .

In particular, given a prompt $\mathcal{Z} = [z_1, \dots, z_K]$, we run ICL on a dev example (\bar{x}, \bar{y}) and define an LLM’s *outcome* as $f(\bar{x}; \mathcal{Z}) = o(\bar{y} | \mathcal{Z}, \bar{x}) - \max_{y' \in \mathcal{C}, y' \neq \bar{y}} o(y' | \mathcal{Z}, \bar{x})$, where $o(y | \mathcal{Z}, \bar{x})$ is the output logit of the LLM on label y before softmax. We hypothesize that we can approximate an LLM’s *outcome* with a linear regressor taking two simple features: the existence of a training example and its index in \mathcal{Z} , which we consider to be the most important factors in ICL. Our linear datamodel is parameterized by weights $w \in \mathbb{R}^{N_{\text{tr}} \times K}$ and bias $b \in \mathbb{R}$; we use $w(i, j) \in \mathbb{R}$ to denote the weight for the i -th training example appearing at index j . For each dev example $(\bar{x}, \bar{y}) \in D_{\text{dev}}$, we train a datamodel $g_{w,b}$ to predict the LLM’s *outcome* of \bar{x} , $f(\bar{x}; \mathcal{Z}) \in \mathbb{R}$, with mean-squared error:

$$g_{w,b}(\bar{x}; \mathcal{Z}) = \sum_{j=1}^K w(\text{id}(z_j), j) + b, \quad (2)$$

$$\min_{w,b} \frac{1}{M} \sum_{n=1}^M (g_{w,b}(\bar{x}; \mathcal{Z}_n) - f(\bar{x}; \mathcal{Z}_n))^2, \quad (3)$$

where z_j is the training example at index j in the

prompt \mathcal{Z} , and $\text{id}(\cdot)$ maps z_j back to its example ID i in the training set, i.e., $z_j = (x_i, y_i)$.

By definition, $f(\bar{x}; \mathcal{Z}) > 0$ if the LLM is correct on \bar{x} . Hence, a positive $w(i, j)$ indicates that having the i -th training example at index j in the prompt encourages answering correctly. We hypothesize that a good training example should have a beneficial effect for many dev examples, regardless of its index in the prompts. Thus, we can aggregate the datamodels of all dev examples and marginalize over all possible orders to calculate the score of the i -th training example:

$$s_{dm}(i) = \sum_{(\bar{x}, \bar{y}) \in D_{\text{dev}}} \sum_{j=1}^K \mathbb{1}\{w_{\bar{x}}(i, j) > 0\}, \quad (4)$$

where $w_{\bar{x}}(i, j)$ is the weight value $w(i, j)$ of the datamodel for \bar{x} . Calculating the total number of positive weights empirically works better than averaging the weights of all the datamodels.

In Appendix A.3, we validate our hypothesis that we can linearize an LLM’s *outcomes* with two simple features. Table 6 shows that our datamodels can accurately approximate the outcomes of unseen prompts outside of \mathcal{D}_{ICL} across different tasks.

3.3 Select Training Examples

Now that we assign a score for each training example (Eq. 1, 4), let C be the number of classes and $E' = E/C$, we select the top- E' training examples of each class with the highest scores to form the stable subset $\mathcal{E} \subset D_{\text{tr}}, |\mathcal{E}| = E$.

4 Experiment

4.1 Setups

Tasks. We experiment on 5 classification tasks: SST-2 (Socher et al., 2013), BoolQ (Clark et al., 2019), Subj (Pang and Lee, 2004), Scicite (Cohan et al., 2019), and AGNews (Zhang et al., 2015). We set the stable training subset size $E = 20$ for all the tasks. For binary classification tasks, we set $K = 4$ and do not balance the classes in the prompts. Thus, the collected \mathcal{D}_{ICL} for a binary task covers all 2^4 label patterns, including prompts with all positive $([1, 1, 1, 1])$ and all negative $([0, 0, 0, 0])$ labels, allowing us to better understand the impact of label patterns on ICL. For multiclass tasks (Scicite and AGNews), we balance the classes, sampling a training example per class to form the prompts. Table 8 in the appendix summarizes our setups.

Data Splits. For all the tasks, we use class-balanced D_{tr} , D_{dev} , and D_{test} sampled from the original training set, as we do not have the gold labels of the original test set. We choose $|D_{\text{tr}}| = 1000$ to ensure a diverse range of training examples for subset selection, and $|D_{\text{test}}| = 1000$ for reliable evaluation. D_{dev} consists of 50 examples per class. All three sets are balanced, randomly sampled from the original training set, and do not overlap.

Models. We run our main experiments with two LLMs: GPTJ-6B (Wang and Komatsuzaki, 2021) and OPT-13B (Zhang et al., 2022a). More experiments on GPT-Neo-2.7B (Black et al., 2021) and OPT-6.7B can be found in Table 9 in the appendix, where our methods also work well.

4.2 Evaluation and Baselines

Recall that our goal is to select a training subset $\mathcal{E} \subset D_{\text{tr}}$ that is more stable than D_{tr} . To evaluate a selection method, we randomly sample 50 prompts from the selected subset \mathcal{E} , run ICL on the test set D_{test} , and report the average accuracy, standard deviation, and worst accuracy.

As shown in Zhao et al. (2021), when a prompt only contains examples of a single label, LLMs are prone to always predict that label on every test example. Thus, in our main experiments (§5.1), we ensure that every sampled prompt contains at least one example from every class. In §5.2, we separately investigate performance when the prompt contains only one label of binary tasks. We split the selected subset \mathcal{E} into two subsets, \mathcal{E}_0 and \mathcal{E}_1 , where \mathcal{E}_0 only contains negative training examples and \mathcal{E}_1 only contains positive examples. We then sample 50 prompts from \mathcal{E}_0 and \mathcal{E}_1 , respectively.

Besides the two proposed selection methods, CONDACC and DATAMODELS, we design 5 baseline methods: ALL, CALIB, RANDOM, ONESHOT, and TOPPROMPTS. ALL uses the entire training set as \mathcal{E} . CALIB uses the same prompts as ALL, but with calibration to prevent LLMs biased toward certain labels, where we closely follow the implementation of Zhao et al. (2021). RANDOM randomly selects a balanced training subset of $E = 20$ examples. ONESHOT first runs ICL with $K = 1$, using each training example alone as the prompt, and then scores the example by the corresponding ICL accuracy on D_{dev} ; these scores are used in the same way as our main methods to select examples (§3.3). ONESHOT tests if we can extrapolate ICL performance from $K = 1$ to $K > 1$. TOP-

	SST-2		BoolQ		Subj		Scicite		AGNews		Avg. Tasks
	Avg std	Min	Avg std	Min	Avg std	Min	Avg std	Min	Avg std	Min	
<i>GPTJ-6B</i>											
ALL	77.8 _{11.2}	50.8	61.0 _{3.8}	49.7	59.8 _{8.3}	50.1	43.8 _{7.2}	33.6	83.5 _{3.8}	70.4	65.2
+ CALIB	75.5 _{9.5}	53.6	61.2 _{3.9}	50.4	70.4 _{7.7}	55.7	35.4 _{2.6}	32.8	85.2 _{2.7}	78.0	65.5
RANDOM	74.6 _{11.4}	50.3	60.0 _{4.3}	49.5	59.9 _{10.4}	50.1	46.4 _{6.9}	35.5	82.5 _{4.7}	67.1	64.7
ONESHOT	79.6 _{10.5}	52.1	63.8 _{2.7}	56.4	63.3 _{10.1}	50.1	44.8 _{5.9}	33.8	83.3 _{3.4}	71.9	67.0
TOPPROMPTS-5	82.8 _{8.6}	56.0	62.3 _{3.0}	54.3	65.5 _{9.7}	50.1	50.4 _{6.0}	36.9	84.4 _{3.3}	74.3	69.1
TOPPROMPTS-10	78.5 _{9.3}	52.4	61.2 _{4.0}	51.1	65.1 _{10.7}	50.1	49.4 _{5.5}	36.2	85.4 _{2.4}	76.3	67.9
CONDACC	86.7 _{5.9}	68.2	65.1 _{1.6}	61.1	70.5 _{10.4}	50.2	52.3 _{4.4}	42.0	87.3 _{2.6}	70.5	72.4
DATAMODELS	86.0 _{7.5}	60.8	65.2 _{0.9}	63.4	69.4 _{10.7}	50.4	54.5 _{3.8}	43.9	86.9 _{1.4}	82.8	72.4
UN-ALL	71.0 _{11.9}	50.0	60.8 _{3.5}	49.6	60.1 _{8.8}	50.1	42.0 _{7.0}	33.5	75.1 _{9.9}	46.5	61.8
UN-ONESHOT	81.9 _{6.3}	68.5	62.6 _{3.3}	55.6	61.0 _{8.7}	50.1	43.5 _{6.7}	33.4	78.1 _{4.2}	69.8	65.4
UN-TOPPROMPTS-5	80.1 _{10.5}	56.8	61.2 _{3.3}	51.9	60.7 _{10.0}	50.1	48.7 _{6.9}	33.0	76.4 _{9.8}	53.0	65.4
UN-CONDACC	85.3 _{6.8}	60.5	63.7 _{2.2}	56.0	66.0 _{10.6}	50.1	54.2 _{3.4}	45.9	87.1 _{1.1}	84.6	71.3
<i>OPT-13B</i>											
ALL	68.5 _{14.0}	50.0	65.2 _{5.6}	49.7	60.9 _{10.2}	49.8	42.8 _{3.6}	35.0	81.6 _{5.9}	64.2	63.8
+ CALIB	84.7 _{6.8}	51.7	65.5 _{4.9}	51.8	63.7 _{8.9}	47.9	35.5 _{1.8}	31.2	81.8 _{4.1}	70.7	66.2
RANDOM	67.7 _{14.1}	50.0	64.7 _{6.4}	49.3	61.2 _{9.5}	49.9	41.2 _{4.6}	33.3	78.0 _{7.5}	61.4	62.6
ONESHOT	75.6 _{13.1}	50.7	68.3 _{2.3}	62.7	60.5 _{9.9}	49.9	41.9 _{3.8}	33.4	84.2 _{2.9}	73.1	66.1
TOPPROMPTS-5	69.6 _{14.7}	50.0	63.5 _{6.3}	51.0	67.4 _{12.7}	50.0	45.9 _{4.3}	36.0	83.9 _{3.1}	74.0	66.1
TOPPROMPTS-10	72.9 _{15.6}	50.0	65.5 _{5.2}	50.4	68.5 _{13.4}	49.9	44.6 _{3.9}	36.7	84.4 _{3.5}	70.9	67.2
CONDACC	83.6 _{9.1}	56.1	69.4 _{2.1}	62.8	70.6 _{11.9}	50.0	49.4 _{3.3}	41.1	87.0 _{1.0}	83.6	72.0
DATAMODELS	81.3 _{10.3}	60.3	69.3 _{3.8}	57.3	63.0 _{9.4}	50.1	46.3 _{3.9}	37.4	85.7 _{1.7}	81.8	69.1
UN-ALL	61.6 _{13.6}	50.0	64.8 _{5.3}	49.3	55.8 _{8.9}	35.6	41.9 _{3.6}	35.7	67.3 _{17.2}	26.4	58.3
UN-ONESHOT	74.8 _{15.6}	50.0	68.0 _{2.5}	59.8	54.8 _{6.2}	47.1	41.5 _{4.1}	33.7	82.3 _{4.5}	64.9	64.3
UN-TOPPROMPTS-5	70.5 _{17.0}	50.0	66.2 _{3.4}	54.6	63.4 _{12.3}	48.3	45.7 _{4.7}	33.6	81.8 _{6.9}	51.8	65.5
UN-CONDACC	80.3 _{12.8}	50.0	69.0 _{2.6}	61.5	63.7 _{11.7}	49.9	48.1 _{4.0}	39.2	84.6 _{3.1}	72.5	69.2

Table 1: Main results with different selection methods. The last column average accuracies across all tasks. Overall, the proposed methods CONDACC and DATAMODELS perform the best. Our method under the unlabeled setup (UN-CONDACC) even outperforms the ALL baseline that uses gold labels.

PROMPTS-5 and TOPPROMPTS-10 select the union of the examples from the top- $\{5, 10\}$ prompts in \mathcal{D}_{ICL} with the highest dev set accuracy, where the subsets contain at most $K \times 5$ and $K \times 10$ examples, respectively. Finally, we apply baselines and our CONDACC method to the unlabeled setup (§2), named with the UN- prefix.

5 Results

5.1 Main Results

The proposed methods outperform all baselines. Table 1 shows the test set accuracy with different training subset selection methods. Among methods without calibration, our CONDACC and DATAMODELS methods are the most *stable*, achieving substantially higher average and worst-case accuracy across all tasks, and lower variances on most tasks. Compared with CALIB, our methods perform better in 8/10 setups. Overall, CONDACC and DATAMODELS outperform the no-selection baseline ALL by 7.7% and 6.3% on average, respectively.

TOPPROMPTS is the strongest baseline. Within the baselines, ALL and RANDOM have similar performance. Applying calibration improves the worst accuracy on most tasks and the average accuracy on some tasks, but is not always beneficial, especially on Scicite. ONESHOT outperforms ALL and RANDOM on SST-2 and BoolQ, but performs similarly on other tasks, indicating that we cannot extrapolate ICL behavior from $K = 1$ to $K = 3$ or $K = 4$. TOPPROMPTS-5 and TOPPROMPTS-10 are the strongest baselines, performing especially well on SST-2, Subj, and Scicite, showing that the training examples that compose the highest-accuracy prompts are more stable than others.

Our method works without training set labels. Randomly sampling prompts from the unlabeled training set (UN-ALL) underperforms sampling from the original labeled training set (ALL), especially on SST-2 and AGNews. This shows that gold labels do matter in ICL in general, in contrast with the findings of Min et al. (2022). However, apply-

ing our selection method to the unlabeled training set (UN-CONDACC) surprisingly outperforms not only UN-ALL but ALL (which uses correctly labeled examples), although some selected training examples actually have the wrong labels, implying that having gold-labeled prompts is not necessary for ICL. Other baselines, UN-ONESHOT and UN-TOPPROMPTS, outperform UN-ALL but substantially underperform our method. Overall, UN-CONDACC outperforms baselines UN-ALL and ALL by 10.2% and 5.7% on average, respectively.

Does UN-CONDACC benefit from gold labels?

We study the number of the stable training examples selected by UN-CONDACC that indeed have gold labels. In most of the tasks, the numbers are much higher than the expected numbers by majority guess, where BoolQ is the exception with half of the selected examples having wrong labels. We thus study if we can achieve even better results by correcting those selected examples that have wrong labels with their gold labels; the other correct examples in the subset remain unchanged. After the label correction on BoolQ, the average and worst accuracy drops by 1.9% and 4.5% respectively on GPTJ, 0.4% and 5.7% respectively on OPT. These results again suggest that on the one hand, ICL benefits from gold-labeled examples in most cases; on the other hand, some training examples with wrong labels can surprisingly achieve better performance. The full results are in Table 10 in the appendix.

Finally, we study the alternative that uses more shots in A.8. Table 12 shows that using 4 curated examples (CONDACC) can outperform $K = 24$, 16 randomly sampled ones in SST-2 and AGNews.

5.2 Single-Label Prompts

We now evaluate whether it is possible to achieve good accuracy while only using training examples of a single class in a prompt (See §4.2 for more details). Table 2 compares the results of different methods. First, the baselines ALL and TOPPROMPTS perform near chance in most cases, as the LLMs are biased by the prompts to predict the same label on every test example. In contrast, single-label prompts sampled from the subsets of CONDACC and DATAMODELS substantially outperform majority guessing across all setups. We conclude that the selected training examples are beneficial because they help LLMs understand the overall definition of the task. Thus, even when used in single-label prompts, they can still give LLMs

	SST-2		BoolQ	
	[0,0,0,0]	[1,1,1,1]	[0,0,0,0]	[1,1,1,1]
<i>GPTJ-6B</i>				
ALL	51.7 _{1.7}	52.6 _{3.1}	52.8 _{1.6}	50.7 _{1.0}
TOPPROMPTS	52.1 _{2.1}	56.0 _{3.7}	54.2 _{1.9}	51.8 _{1.6}
CONDACC	61.8 _{4.9}	60.3 _{2.8}	58.3 _{2.1}	55.5 _{1.7}
DATAMODELS	72.8 _{4.9}	68.4 _{4.9}	61.7 _{1.6}	56.9 _{2.2}
<i>OPT-13B</i>				
ALL	54.0 _{4.1}	73.3 _{7.4}	52.4 _{2.6}	51.2 _{1.7}
TOPPROMPTS	53.0 _{3.0}	76.5 _{6.9}	53.0 _{3.2}	51.4 _{2.1}
CONDACC	66.3 _{3.7}	81.0 _{4.6}	65.3 _{2.9}	53.7 _{1.9}
DATAMODELS	63.9 _{4.2}	84.5 _{2.6}	69.2 _{2.0}	60.1 _{2.2}

Table 2: Results of single-labeled prompts with different selection methods. Each prompt consists of 4 training examples with the same labels.

OOD Tasks	IMDb		BoolQ Cst.	
	Avg std	Min	Avg std	Min
<i>GPTJ-6B</i>				
ALL	86.5 _{5.7}	63.6	56.6 _{3.0}	50.1
TOPPROMPTS	87.2 _{5.2}	63.0	56.7 _{2.6}	49.9
CONDACC	90.5 _{1.8}	84.8	58.9 _{1.7}	54.6
DATAMODELS	91.6 _{1.5}	84.0	57.6 _{1.9}	54.0
<i>OPT-13B</i>				
ALL	79.2 _{12.1}	50.1	59.8 _{2.9}	51.6
TOPPROMPTS	80.5 _{14.0}	50.8	60.3 _{3.5}	51.0
CONDACC	83.5 _{10.8}	54.6	60.1 _{2.1}	56.7
DATAMODELS	84.1 _{9.3}	58.9	60.6 _{3.3}	54.3

Table 3: Accuracy of IMDb and BoolQ Contrast Set, where the prompts consist of the selected SST-2 and BoolQ training examples, respectively.

useful signal to perform the desired task.

5.3 Out-of-Distribution Tasks

We further evaluate on out-of-distribution (OOD) tasks, where there is a distribution shift between prompts and test data. Specifically, we apply our selection methods on a source task, sampling 50 prompts from the selected subsets of the source task as done in the main experiments, and then evaluate on test data of a target task. We use SST-2 and BoolQ as the source tasks, and IMDb (Maas et al., 2011) and BoolQ Contrast Set (Gardner et al., 2020) as our target tasks, respectively. Table 3 shows that our CONDACC and DATAMODELS methods still outperform baselines on OOD tasks, especially on IMDb, implying that instead of simply overfitting the source tasks, the selected stable examples are indeed task-level examples that can generalize to OOD test data.

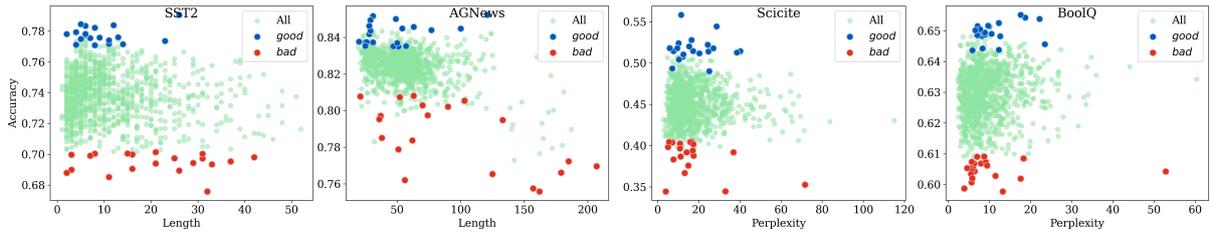


Figure 3: Accuracy versus sequence length (left) and accuracy versus perplexity (right). Each dot corresponds to a training example. Examples in *good* subsets are not outliers with abnormally long lengths or high perplexities.

6 Analysis

We further analyze what makes the selected training examples special along different dimensions: sequence length, perplexity, and diversity. We compare *good* with *bad* training examples, where we use our CONDACC method to identify a *bad* (resp., *good*) subset by selecting E' training examples with the *lowest* (resp., *highest*) importance scores in each class (§3.3). Please refer to Table 11 in the appendix for the full results of the *bad* subset.

6.1 Sequence Length and Perplexity

In Figure 3, we plot the accuracy against either sequence length or perplexity, where each dot corresponds to a training example. Here, the accuracy (Y-axis) is the importance score s_{ca} assigned to each training example by CONDACC in Eq. 1, i.e., the average dev-set accuracy when that example is combined with random other training examples in a random order.

Sequence Length. The first two subfigures show that while the *bad* examples (red dots) span across different sequence lengths, the *good* examples (blue dots) do not cover the tail distribution of long sequences. We observe little correlation between accuracy and sequence length across different tasks and LLMs (see more in Figure 6), except for a slightly negative correlation when the sequence length is very long, suggesting that abnormally long training examples can hurt ICL performance.

Perplexity. We calculate the perplexity of the inputs of training examples with respect to the same LLMs we run ICL on. Figure 3 shows that *good* examples are not outliers that have high perplexity. This could suggest future work filter out examples that have extraordinarily high perplexity in the training set before running ICL, and could be combined with active learning for ICL (Zhang et al., 2022b; Su et al., 2022), as we only need the

unlabeled inputs to calculate perplexity. However, we observe no correlation between accuracy and perplexity across all the tasks and LLMs (Figure 7), implying that using perplexity alone is not enough for identifying good training examples. Our findings are inconsistent with concurrent work (Gonen et al., 2022), which shows that lower prompt perplexity strongly correlates with better performance, probably because Gonen et al. (2022) focus on perplexities under different instructions, while we focus on the differences between training inputs.

6.2 Diversity

DIV-I and DIV-F. We measure the diversity of a training subset with DIV-I (Yuan et al., 2020) and DIV-F (Zhdanov, 2019) metrics, following prior work in active learning. DIV-I measures diversity in raw text, while DIV-F measures diversity in a feature space. For DIV-F, we use SentBERT (Reimers and Gurevych, 2019) to encode the inputs of training examples into sentence embeddings, following Su et al. (2022). We compare the selected *good* subset and *bad* subset with 5000 randomly sampled subsets $\subset D_{tr}$, each containing E' training examples per class. Figure 4 shows that *good* subsets (blue dots) sometimes have low diversity scores in both metrics, especially on BoolQ and AGNews. Overall, they do not seem to be more diverse than randomly sampled subsets. Our findings are different from Su et al. (2022), which finds that diverse training subsets are better for prompt retrieval. We hypothesize that diversity matters more when retrieving similar training examples for each test input, but is less important when using a single fixed prompt. On the other hand, the *bad* subsets have much higher DIV-I scores than random subsets across different tasks, because they often include examples with long sequence lengths (see Figure 3), covering more distinct unigrams. However, in the SentBERT feature space, the *bad* subsets are often less diverse than random subsets.

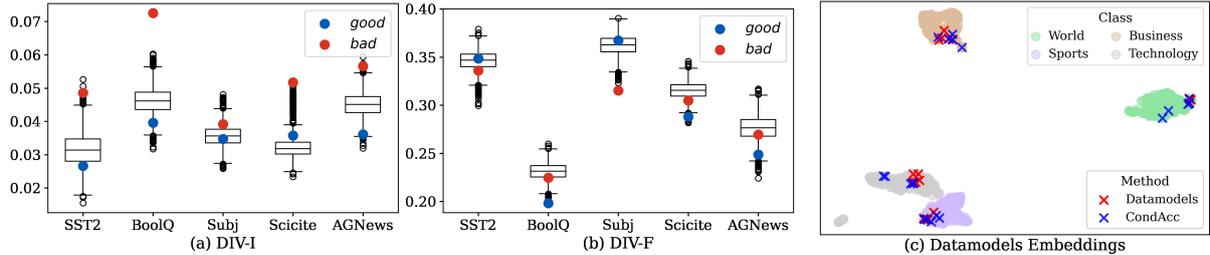


Figure 4: Different ways to visualize the diversity of examples. (a) and (b) compare the diversity of the *good* subset, *bad* subset, and randomly sampled subsets (boxplot). For both DIV-I and DIV-F, a higher number means a subset is more diverse. Overall, *good* subsets are no more diverse than random subsets. (c) visualizes the stable training examples selected by CONDACC and DATAMODELS methods in Datamodels embeddings space, where each dot is a training example in AGNews. Both methods choose tightly clustered examples instead of diverse ones.

Datamodels Embeddings. In §3.2, we learn a datamodel for each dev example in D_{dev} . Here, we concatenate the weights assigned on a training example learned by all the datamodels, creating an embedding $\in \mathbb{R}^{|D_{\text{dev}}| \times K}$ for each training example. We then project the embeddings of the entire training set to a two-dimensional space with UMAP (McInnes et al., 2018) for visualization. Figure 4 and Figure 8 in the appendix show that both CONDACC and DATAMODELS choose tightly clustered sets of examples in the embedding space, instead of diverse ones scattering over the training set. Moreover, the two methods actually select similar examples in the embedding space, although having very different scoring methods.

6.3 Do LLMs find the same stable examples?

We further study if the identified stable subset examples are transferable across different LLMs. Given two LLMs, we calculate the Pearson correlation coefficient between their example scores assigned by CONDACC (§3.1) and the actual number of overlapped examples in the stable subsets of the two LLMs. Table 4 shows the mixed results: some pairs have moderate correlation, especially when both LLMs are from the OPT family; however, we find little correlation between many pairs and there are only a few overlapped examples between the stable subsets identified by different LLMs.

Interestingly, when using the four stable SST2 examples shared by GPTJ-6B and OPT-13B as the prompt, evaluating all $4!$ permutations, we achieve very high test accuracy: 88.6 ± 3.7 on GPTJ, 89.8 ± 3.0 on OPT-6.7B, and 87.3 ± 5.2 on OPT-13B. This may indicate that there exist successful factors of training examples shared among LLMs. We include the four stable examples in

	Corr	Overlap
<i>GPTJ-6B vs OPT-13B</i>		
SST-2	0.15	4
AGNews	0.46	1
BoolQ	0.41	2
Subj	-0.03	2
Scicite	0.02	2
<i>GPTJ-6B vs OPT-6.7B</i>		
SST-2	0.27	3
AGNews	0.08	1
<i>OPT-6.7B vs OPT-13B</i>		
SST-2	0.76	3
AGNews	0.42	1

Table 4: Pearson correlation between the example scores (s_{ca}) of two LLMs and the number of overlapped examples in their stable subsets.

Examples
Review: <i>k-19 : the widownmaker is a great yarn.</i> Sentiment: positive
Review: <i>spiffy animated feature</i> Sentiment: positive
Review: <i>a plot cobbled together from largely flat and uncreative moments</i> Sentiment: negative
Review: <i>has the thrown-together feel of a summer-camp talent show : hastily written, underrehearsed, arbitrarily plotted and</i> Sentiment: negative

Table 5: The four stable subset examples shared by GPTJ-6B and OPT-13B in SST-2 dataset, where the last example is also selected by OPT-6.7B. All three models achieve high average accuracy and low variance across the $4!$ permutations of these four examples.

Table 5 and hope future work can discover what distinguishes them from other examples.

7 Discussion and Related Work

7.1 Prompt Retrieval

The performance of ICL greatly depends on the choice of the prompt. A common way to do prompt selection is to retrieve the top- K similar training examples for each test input (Liu et al., 2021; Rubin et al., 2021; Su et al., 2022), where the similarity is captured by sentence embeddings (Reimers and Gurevych, 2019; Robertson et al., 2009). Such instance-dependent prompt retrieval is critical for semantic parsing tasks (Rubin et al., 2021), as LLMs need to see relevant logical forms in the context to generate the appropriate predicates for the test example. In this paper, we focus on classification tasks and identify task-level training examples that work for *all* test examples, avoiding the retrieval process.

7.2 The influence of in-context labels

The proportions of labels appearing in context can greatly bias LLMs’ predictions (Zhao et al., 2021). However, with careful prompt selection, Zhang et al. (2022b) finds that LLMs can perform well without observing the entire label space of a classification task in the prompt. In §5.2, we identify a subset, instead of just one prompt, of single-label examples that perform well. On the other hand, the correctness of in-context labels may not matter as much, as Min et al. (2022) find that randomly flipping them barely hurts performance. Kim et al. (2022) re-examine the importance of gold labels, showing it varies largely across tasks and experimental settings; our unlabeled experiments also show varying importance of gold labels across different tasks. Our method also identifies some training examples with wrong labels that can yield surprisingly good performance.

7.3 Data Valuation

Given a learning algorithm trained on a training set to produce a predictor, data valuation quantifies the value of each training example to the predictor performance. Prior work includes influence functions (Koh and Liang, 2017), Data Shapley (Ghorbani and Zou, 2019), DVRL (Yoon et al., 2020), TracIn (Pruthi et al., 2020), and Datamodels (Ilyas et al., 2022). Our setup also aims to attribute the performance of a predictor (in our case, an LLM) to each training example (in our case, in-context examples). Our CONDACC method closely resembles Data Shapley, and we adapt Ilyas et al. (2022)

as our DATAMODELS method. However, the main difference is that we are doing K -shot ICL, where training examples are used as prompts, and there are no parameter updates to LLMs.

In concurrent and independent work, Nguyen and Wong (2023) also propose methods based on Data Shapley and Datamodels to study the influence of training examples. The main differences are: (1) we adapt Datamodels to consider the positions of training examples, while Nguyen and Wong (2023) follow the original implementation and study example ordering with influence scores. (2) They propose a Perplexity baseline that selects examples according to individual perplexity, while we use perplexity in analysis to study the correlation between example perplexity and their average performance. (3) Nguyen and Wong (2023) explore a larger number of shots in ICL ($K \in [10, 52]$), while we assume a few-shot setting and fix $K \in \{3, 4\}$. Our findings on good training examples are consistent with each other: both papers find little correlation between performance and potential factors such as example length and perplexity. In general, our work demonstrates the importance of data curation on in-context examples, even in the unlabeled and OOD scenarios, while Nguyen and Wong (2023) focus more on developing influence-based example selection frameworks. Taken together, two papers present a comprehensive view of data valuation for in-context learning.

8 Conclusion

We propose two methods to identify stable training subsets for in-context learning, achieving substantially higher average and worst-case accuracy across different setups. Our CONDACC method is intuitive and easy-to-implement, while our DATAMODELS method provides informative weights that enable further analyses. The success of our methods implies that when provided with the “right” training set (in our case, a subset of 20 examples to randomly sample prompts from), ICL could be far less sensitive to the choice of training examples and their orders in a prompt. Our analyses on stable subsets find that they do not contain outliers with especially long sequence lengths or high perplexities, and are also no more diverse than random subsets of the training data. We hope our work is a step towards developing guidelines for finding or writing better training examples.

Limitations

The main limitation of our work is the high memory and computation cost. As both our methods estimate the importance of training examples based on the prompt-performance statistics, we first need to run in-context learning on the dev set multiple times with different prompts in \mathcal{D}_{ICL} . Although ICL does not require any parameter updates, LLMs still require a large amount of memory footprint during inference, especially when the model size is large and the average sequence length is long. For each setup, our \mathcal{D}_{ICL} contains around 50,000 prompts (see Table 8) and 50 dev examples per class, so the most expensive setup (running OPT-13B on BoolQ) costs more than 500 GPU hours on an RTX A6000 GPU. Our preliminary study shows that the proposed methods need the statistics of at least 10,000 randomly sampled prompts to perform well. Future work may use search algorithms instead of random sampling during data collection to reduce the number of prompts in \mathcal{D}_{ICL} . We also release our collected data of every setup in <https://github.com/terarachang/DataICL> to support future studies on ICL.

In this paper, we only study classification tasks, for the sake of easy evaluation. Future work may study the influence of in-context examples in generation tasks under different evaluation metrics. Due to hardware constraints, we do not study LLMs of sizes over 13B, and we fix the number of shots and prompt templates for simplicity. In independent work, [Nguyen and Wong \(2023\)](#) complement these limitations of our paper, showing that similar approaches work well on larger models and a diverse number of shots for in-context example selection. Still, the influence of in-context examples for gigantic LLMs larger than 100B parameters has not been studied. Due to emergent abilities of LLMs ([Wei et al., 2022](#)), it is unclear whether our methods and findings would still apply when prompting these gigantic LLMs.

Acknowledgements

We thank Ameya Godbole, Johnny Wei, Wang Zhu, Albert Xu, Deqing Fu, Qinyuan Ye, the members of USC NLP, and our anonymous reviewers for valuable feedback on the paper. We thank Sameer Singh and Matt Gardner for helpful discussions. We thank Jesse Thomason for his support. This work was funded in part by gifts from Open Philanthropy and Google.

References

- Akari Asai, Mohammadreza Salehi, Matthew E. Peters, and Hannaneh Hajishirzi. 2022. Attentional mixtures of soft prompt tuning for parameter-efficient multi-task knowledge sharing. *ArXiv*, abs/2205.11961.
- Stephen H. Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Xiangru Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. 2022. [Promptsources: An integrated development environment and repository for natural language prompts](#).
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#).
- 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.
- Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown, and He He. 2022. On the relation between sensitivity and accuracy in in-context learning. *arXiv preprint arXiv:2209.07661*.
- 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*.
- 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*.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*.
- Arman Cohan, Waleed Ammar, Madeleine van Zuylen, and Field Cady. 2019. Structural scaffolds for citation intent classification in scientific publications. In *North American Chapter of the Association for Computational Linguistics*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference*

- of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, et al. 2020. Evaluating models’ local decision boundaries via contrast sets. *arXiv preprint arXiv:2004.02709*.
- Amirata Ghorbani and James Zou. 2019. Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*, pages 2242–2251. PMLR.
- Hila Gonen, Srini Iyer, Terra Blevins, Noah A Smith, and Luke Zettlemoyer. 2022. Demystifying prompts in language models via perplexity estimation. *arXiv preprint arXiv:2212.04037*.
- Andrew Ilyas, Sung Min Park, Logan Engstrom, Guillaume Leclerc, and Aleksander Madry. 2022. Data-models: Predicting predictions from training data. *arXiv preprint arXiv:2202.00622*.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Dániel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. Opt-impl: Scaling language model instruction meta learning through the lens of generalization. *arXiv preprint arXiv:2212.12017*.
- Junyeob Kim, Hyuhng Joon Kim, Hyunsoo Cho, Hwiyeol Jo, Sang-Woo Lee, Sang goo Lee, Kang Min Yoo, and Taek Kim. 2022. Ground-truth labels matter: A deeper look into input-label demonstrations. *ArXiv*, abs/2205.12685.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Annual Meeting of the Association for Computational Linguistics*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Sewon Min, Xinxin Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837*.
- Tai Nguyen and Eric Wong. 2023. In-context example selection with influences. *arXiv preprint arXiv:2302.11042*.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Annual Meeting of the Association for Computational Linguistics*.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. True few-shot learning with language models. *Advances in Neural Information Processing Systems*, 34:11054–11070.
- Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. Estimating training data influence by tracing gradient descent. *Advances in Neural Information Processing Systems*, 33:19920–19930.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. 2021. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *ArXiv*, abs/1908.10084.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2021. Learning to retrieve prompts for in-context learning. *arXiv preprint arXiv:2112.08633*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2021. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.

- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deep-speed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y 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.
- Hongjin Su, Jungo Kasai, Chen Henry Wu, Weijia Shi, Tianlu Wang, Jiayi Xin, Rui Zhang, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, et al. 2022. Selective annotation makes language models better few-shot learners. *arXiv preprint arXiv:2209.01975*.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Xi Ye, Srini Iyer, Asli Celikyilmaz, Ves Stoyanov, Greg Durrett, and Ramakanth Pasunuru. 2022. Complementary explanations for effective in-context learning. *ArXiv*, abs/2211.13892.
- Jinsung Yoon, Sercan Arik, and Tomas Pfister. 2020. Data valuation using reinforcement learning. In *International Conference on Machine Learning*, pages 10842–10851. PMLR.
- Michelle Yuan, Hsuan-Tien Lin, and Jordan Boyd-Graber. 2020. Cold-start active learning through self-supervised language modeling. *arXiv preprint arXiv:2010.09535*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022a. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Yiming Zhang, Shi Feng, and Chenhao Tan. 2022b. Active example selection for in-context learning. *ArXiv*, abs/2211.04486.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.
- Fedor Zhdanov. 2019. Diverse mini-batch active learning. *arXiv preprint arXiv:1901.05954*.

A Appendix

A.1 Connection with Data Shapley

Data Shapley has a valuation function $V(\mathcal{S})$ for any subset of training examples. We define $V(\mathcal{S})$ as the expected dev set ICL accuracy across all permutations of \mathcal{S} if $|\mathcal{S}| = K$, and $V(\mathcal{S}) = 0$ otherwise since we focus on K -shot learning. Then, the Data Shapley value $s_{shap}(i)$ for the i -th training example $e_i = (x_i, y_i)$ is:

$$\mathbb{E}_{\{z_1, \dots, z_{K-1}\} \sim (D_{\text{tr}} \setminus e_i)} [V(\{z_1, \dots, z_{K-1}, e_i\})] - \mathbb{E}_{\{z_1, \dots, z_K\} \sim (D_{\text{tr}} \setminus e_i)} [V(\{z_1, \dots, z_K\})].$$

We claim that $s_{shap}(i)$ is a monotonically increasing (in fact, affine) function of $s_{ca}(i)$, thus establishing a very tight connection between CONDACC and Data Shapley.

First, note that the first term is exactly equal to $s_{ca}(i)$, the expected conditional accuracy when e_i occurs in a prompt, since V returns the expected accuracy over all orders of $\{z_1, \dots, z_{K-1}, e_i\}$.

Similarly, the second term is the expected conditional accuracy when e_i does not occur in a prompt. Denote this quantity as $t(i)$, and let A denote the overall expected accuracy when randomly sampling a prompt. Since the probability of choosing a given example to be in a prompt is exactly $\frac{K}{N_{\text{tr}}}$, we have

$$\begin{aligned} A &= \frac{K}{N_{\text{tr}}} s_{ca}(i) + \frac{N_{\text{tr}} - K}{N_{\text{tr}}} t(i) \\ t(i) &= \frac{N_{\text{tr}}}{N_{\text{tr}} - K} \left(A - \frac{K}{N_{\text{tr}}} s_{ca}(i) \right) \\ &= \frac{N_{\text{tr}}}{N_{\text{tr}} - K} A - \frac{K}{N_{\text{tr}} - K} s_{ca}(i). \end{aligned}$$

Now we can rewrite $s_{shap}(i)$ as follows:

$$\begin{aligned} s_{shap}(i) &= s_{ca}(i) - t(i) \\ &= s_{ca}(i) - \left(\frac{N_{\text{tr}}}{N_{\text{tr}} - K} A - \frac{K}{N_{\text{tr}} - K} s_{ca}(i) \right) \\ &= \frac{N_{\text{tr}}}{N_{\text{tr}} - K} A + \frac{N_{\text{tr}}}{N_{\text{tr}} - K} s_{ca}(i). \end{aligned}$$

Since N_{tr} , K , and A are all constants that do not depend on i , and $N_{\text{tr}}, K > 0$, this establishes that $s_{shap}(i)$ is a monotonically increasing affine function of $s_{ca}(i)$.

A.2 Relation to Prompt Tuning

At test time, our setup is similar to Prompt Tuning (Lester et al., 2021), where a fixed prompt for

	GPTJ-6B		OPT-13B	
	L_1	Corr	L_1	Corr
SST2	0.133	0.962	0.264	0.930
Boolq	0.147	0.941	0.167	0.937
Subj	0.269	0.946	0.260	0.949
Scicite	0.088	0.937	0.151	0.938
AGNews	0.296	0.891	0.340	0.840

Table 6: Test results of datamodels. Our datamodels can accurately predict LLMs’ outcomes on unseen prompts, having high correlation and low L_1 distance to the ground-truth outcomes.

a task is prepended to the test inputs. In our case, a prompt is a sequence of training examples randomly drawn from the selected subset (fixed for all test examples). In Prompt Tuning, it is a set of continuous embeddings, called a *soft prompt*, learned through backpropagation.

At training time, however, Prompt Tuning needs to backpropagate through the LLM and thus is more memory-expensive and tends to suffer from training instability (Asai et al., 2022). In comparison, our method finds good prompts without accessing the LLM’s parameters, which is a realistic setup as many LLMs only provide API access.

A.3 Evaluating Datamodels

Recall that given a prompt \mathcal{Z} and a dev example (\bar{x}, \bar{y}) , a datamodel learns to approximate an LLM’s outcome of \bar{x} (§3.2). To evaluate how accurate our datamodels can approximate an LLM, we create a test set for datamodels \mathcal{D}_{DM} , consisting of 5000 pairs of newly sampled \mathcal{Z} and the LLM’s ground-truth outcomes of every dev example, where our sampling assures that every \mathcal{Z} is made up of an unseen combination of training examples. We evaluate the learned datamodels on \mathcal{D}_{DM} , calculating the correlation and L_1 distance between the predicted outcomes and the ground-truth outcomes. More specifically, each datamodel yields 5000 outcomes, which we calculate the Pearson correlation coefficient with the ground-truth outcomes of the associated dev example. We report the average correlation and L_1 distance over all datamodels in Table 6. We also randomly sample 10,000 outcomes across all datamodels to visualize the ground truths against predictions in Figure 5.

Overall, our datamodels can accurately approximate LLMs’ outcomes across different tasks and

Task	Example	Label Mapping
SST-2	Review: contains no wit , only labored gags Sentiment: negative	negative/positive
BoolQ	Exercise: read the text and answer the question by yes or no. Good Samaritan laws offer legal protection to people who give reasonable assistance... Question: do good samaritan laws protect those who help at an accident? yes	no/yes
Subj	Input: the tucks have a secret , they 're immortal . Type: objective	objective/subjective
Scicite	Is the following citation from a scientific paper describing a method, a result, or background? However, how frataxin interacts with the Fe-S cluster biosynthesis components... Answer: background	method/result/ background
MNLI	"yeah well you're a student right" Based on the previous passage, is it true that "Well you're a mechanics student right"? Yes, no, or maybe? maybe	yes/maybe/no
AGNews	Article: Wall St. Bears Claw Back Into the Black (Reuters) Reuters - Short-sellers... Answer: Business	World/Sports/ Business/Technology

Table 7: Our templates and label mappings in different tasks. For simplicity, all the label words we use consist of a single token, so we can easily get the probability of each label.

LLMs. As our datamodels only consider two simple features, the existence of a training example and its index in \mathcal{Z} , accurate test predictions may imply that these two features have a dominating effect on ICL.

A.4 Training Details of Datamodels

As the pattern of in-context labels (e.g., $[0,0,0,0]$, $[0,0,0,1]$, $[1,0,0,1]$) has a great impact on LLMs' predictions, for each label pattern, we train a set of $|D_{\text{dev}}|$ datamodels. Specifically, we apply two-phase training: in the first phase, we train on all data with shared weights. In the second phase, we first bucket prompts in \mathcal{D}_{ICL} by their label patterns. Initializing from the weights learned in the first phase, we then separately fine-tune a set of datamodels for each label pattern. For example, for a binary task with 4-shot learning, there are $2^4 = 16$ label patterns; thus, we have 16 sets of datamodels after the second-phase training, namely, $16 \times |D_{\text{dev}}|$ datamodels in total. We find that having two-phase training leads to more accurate predictions of LLMs' outcomes in appendix A.3. Thus, when assigning scores for training examples, we aggregate all sets of datamodels to obtain s_{dm} . When creating datamodel embeddings (§6), we use the unified weights learned by the first phase.

A.5 Implementation Details

We use PyTorch and Huggingface transformers to implement in-context learning on GPT-Neo-2.7B (Black et al., 2021), GPTJ-6B (Wang and Komatsuzaki, 2021), OPT-6.7B (Zhang et al., 2022a),

and OPT-13B. We run all our evaluations on a single RTX6000 GPU (48GB). Most of our experiments can also be run on an RTX3090 GPU (24GB), except that OPT-13B model requires a GPU with larger memory.

Our data collection on \mathcal{D}_{ICL} costs hundreds of GPU hours on an RTX6000. Once we finish the collection, our DATAMODELS method only takes about 5 seconds to train a datamodel on an i7-10700 CPU, and our CONDACC method does not involve any training, but simply calculates accuracy.

Table 7 shows our task templates and label mappings, where we closely follow Bach et al. (2022); Lu et al. (2021). Table 8 summarizes our experimental setups on different tasks.

A.6 More Experiments

Table 9 shows experiments on more LLMs and the MNLI task, where we evaluate on test data using 50 sampled prompts, as done in the main experiments (§4.2). Since ICL performs poorly on MNLI (majority: 33.3%) in both prior work (Su et al., 2022) and our results in Table 9, we do not experiment more on this task. Overall, our methods CONDACC and DATAMODELS substantially outperform other 4 baselines on all setups.

A.7 Why not Instruction-Finetuned LLMs?

The massive multitask learning in instruction-tuning leads to leakage in the datasets we evaluate on. For example, T0 (Sanh et al., 2021), FLAN-T5 (Chung et al., 2022), and OPT-IML (Iyer et al., 2022) are all trained on several tasks in our paper.

Task	N_{class}	Bal.	K	$ D_{\text{tr}}^L $	$ D_{\text{tr}}^U $	Permut ^L	Permut ^U	$ \mathcal{D}_{\text{ICL}}^L $	$ \mathcal{D}_{\text{ICL}}^U $
SST-2	2	N	4	1000	2000	$\binom{1000}{4} \times 4!$	$\binom{2000}{4} \times 4!$	100,000	50,000
BoolQ	2	N	4	1000	2000	$\binom{1000}{4} \times 4!$	$\binom{2000}{4} \times 4!$	100,000	50,000
Subj	2	N	4	1000	2000	$\binom{1000}{4} \times 4!$	$\binom{2000}{4} \times 4!$	100,000	50,000
Scicite	3	Y	3	999	2997	$(333)^3 \times 3!$	$(999)^3 \times 3!$	40,000	50,000
AGNews	4	Y	4	1000	4000	$(250)^4 \times 4!$	$(1000)^4 \times 3!$	40,000	50,000

Table 8: Setups on different tasks. (1) We balance the classes in the prompts (Bal.) for multiclass tasks. (2) To create the unlabeled training set D_{tr}^U , we pair each input with every possible label; therefore, D_{tr}^U is N_{class} times larger than the gold-labeled training set D_{tr}^L . (3) Permut^L and Permut^U denote the total number of possible permutations of training examples for K -shot ICL in the labeled and unlabeled setups, respectively. (4) $|\mathcal{D}_{\text{ICL}}^L|$ and $|\mathcal{D}_{\text{ICL}}^U|$ denote the number of the prompts we collect in the labeled and unlabeled setups, respectively.

	SST-2		AGNews		MNL	
	Avg	std	Min	Avg	std	Min
<i>GPT-Neo-2.7B</i>						
ALL	64.5	13.0	50.0	74.8	5.8	61.8
RANDOM	65.2	12.8	50.0	74.3	6.8	56.2
ONESHOT	66.1	12.8	50.0	78.3	4.4	64.9
TOPPROMPTS-5	64.1	12.7	50.0	79.9	3.4	71.1
CONDACC	76.5	10.5	52.4	82.3	2.2	77.4
DATAMODELS	72.6	14.2	50.4	83.5	1.4	80.0
<i>OPT-6.7B</i>						
ALL	76.8	11.8	52.4	67.9	15.8	26.0
RANDOM	72.6	12.7	50.2	66.1	14.7	27.6
ONESHOT	84.7	6.6	65.7	76.3	7.8	56.7
TOPPROMPTS-5	78.8	10.9	50.4	78.2	9.7	30.8
CONDACC	88.2	5.8	59.4	83.2	4.3	67.2
DATAMODELS	87.4	5.3	71.7	84.2	3.0	74.0
<i>GPTJ-6B</i>						
ALL	43.8	2.9	35.8			
RANDOM	44.1	2.7	33.8			
ONESHOT	41.7	4.3	33.5			
TOPPROMPTS-5	44.7	2.5	38.0			
CONDACC	45.6	1.8	40.0			
DATAMODELS	44.9	1.7	40.3			

Table 9: More results of GPT-Neo-2.7B and OPT-6.7B (left) and GPTJ-6B on MNL task (right). Overall, the proposed subset selection methods CONDACC and DATAMODELS significantly outperform other baselines.

	GPTJ-6B	OPT-13B	Majority
SST-2	20	19	10
BoolQ	11	10	10
Subj	20	20	10
Scicite	16	11	6.6
AGNews	18	13	5

Table 10: The number of gold-labeled training examples identified by UN-CONDACC in the unlabeled setup. The subset size $E = 20$ for all tasks.

	GPTJ-6B	OPT-13B	Majority
SST-2	66.0	12.3	55.3
BoolQ	50.5	3.4	55.6
Subj	51.8	4.2	50.6
Scicite	33.9	1.2	36.2
AGNews	60.8	9.4	54.6

Table 11: Results of the *bad* training subsets, which consist of examples of the lowest scores assigned by CONDACC method.

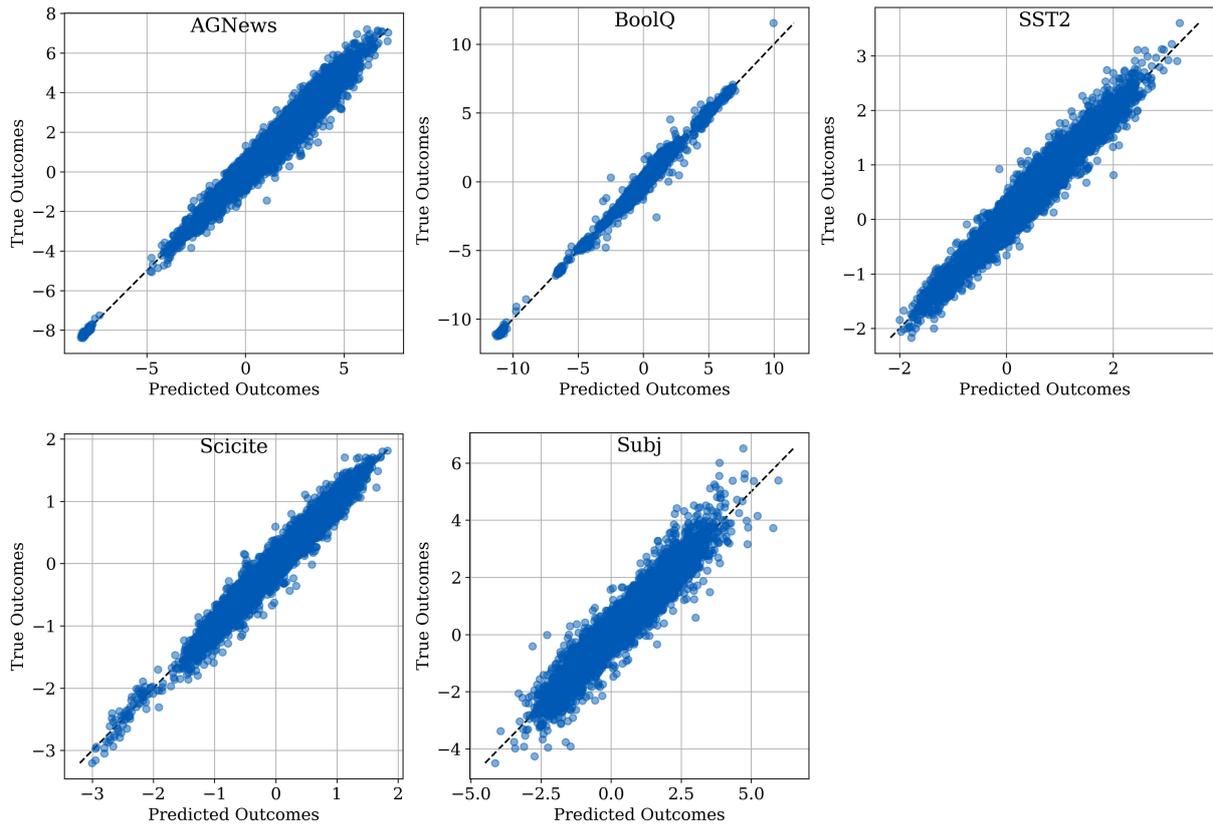
A.8 Larger Number of Shots

We further compare our methods with an alternative that takes as many labeled, balanced, in-context examples as the context window can fit, named MAXSHOT. Similar to the ALL baseline, MAXSHOT samples 50 prompts from the entire training set, each containing K training examples. The differences are that MAXSHOT uses a much larger K and balances the classes in the prompt for binary tasks as well. Here, our LLM is GPTJ-6B, which has a context window of 2048 tokens.

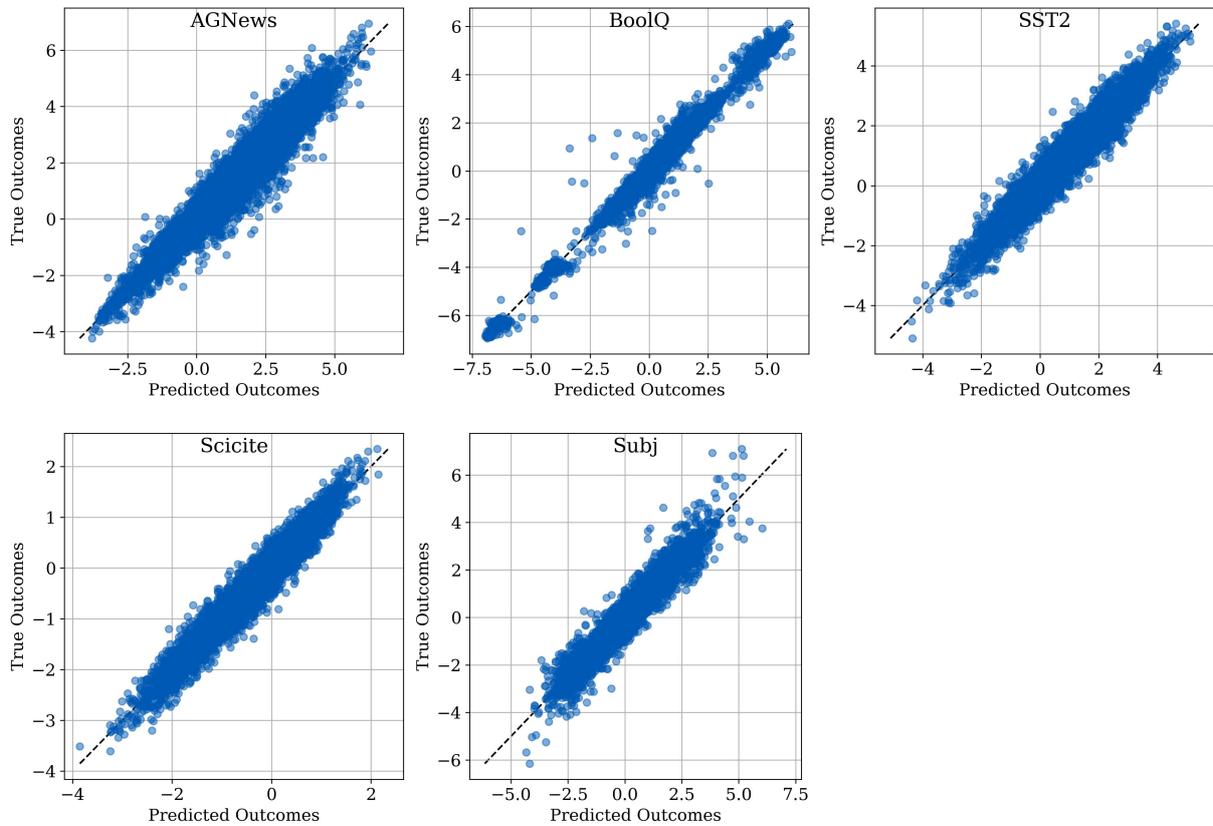
Table 12 shows the number of shots K for each task and compares the average test set accuracy of MAXSHOT with the ones of ALL and CONDACC in Table 1. The Δ_{All} column shows that using $K \in [8, 24]$ examples in the prompt substantially improves over only 3 or 4 examples in most tasks, except for AGNews. However, MAXSHOT only outperforms CONDACC on two out of five tasks (Δ_{Our} , blue). This shows the advantages of curated training examples over randomly sampled ones.

	K	Avg std	Δ_{All}	Δ_{Our}
SST-2	24	85.8 _{5.5}	8.0	-0.9
Subj	24	77.1 _{8.2}	17.2	6.6
BoolQ	8	63.6 _{1.8}	2.7	-1.5
Scicite	12	57.1 _{6.2}	13.2	4.8
AGNews	16	82.3 _{4.8}	-1.2	-5.0

Table 12: Test results of MAXSHOT on GPTJ, where Δ_{All} and Δ_{Our} show its improvements of average accuracy over ALL and CONDACC, respectively (both only use $K = 3, 4$ training examples).

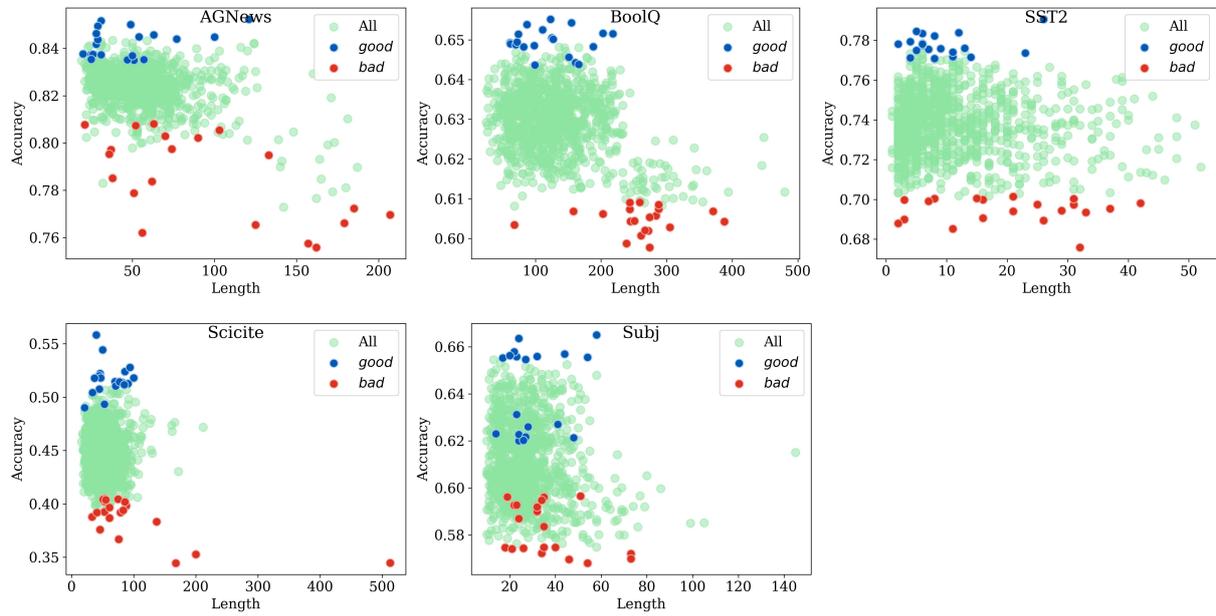


(a) GPTJ-6B

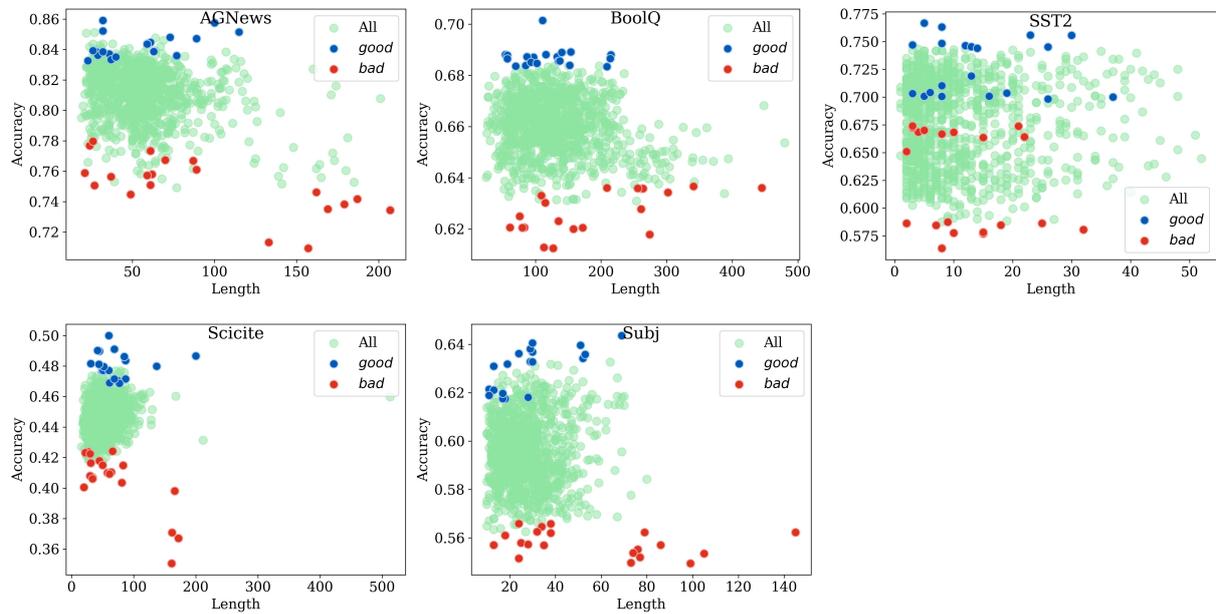


(b) OPT-13B

Figure 5: The ground-truth outcomes of an LLM versus predicted outcomes of datamodels on the test set of datamodels, which contains a set of newly sampled prompts with unseen combinations of training examples. The high correlations show that our datamodels can make accurate predictions.



(a) GPTJ-6B



(b) OPT-13B

Figure 6: The accuracy versus sequence length across different tasks. Each dot corresponds to a training example. Note that we select the top E' examples per class. As a class may have much lower average accuracy than others, the *good* (resp., *bad*) examples may not be the examples with the globally highest (resp., lowest) accuracy.

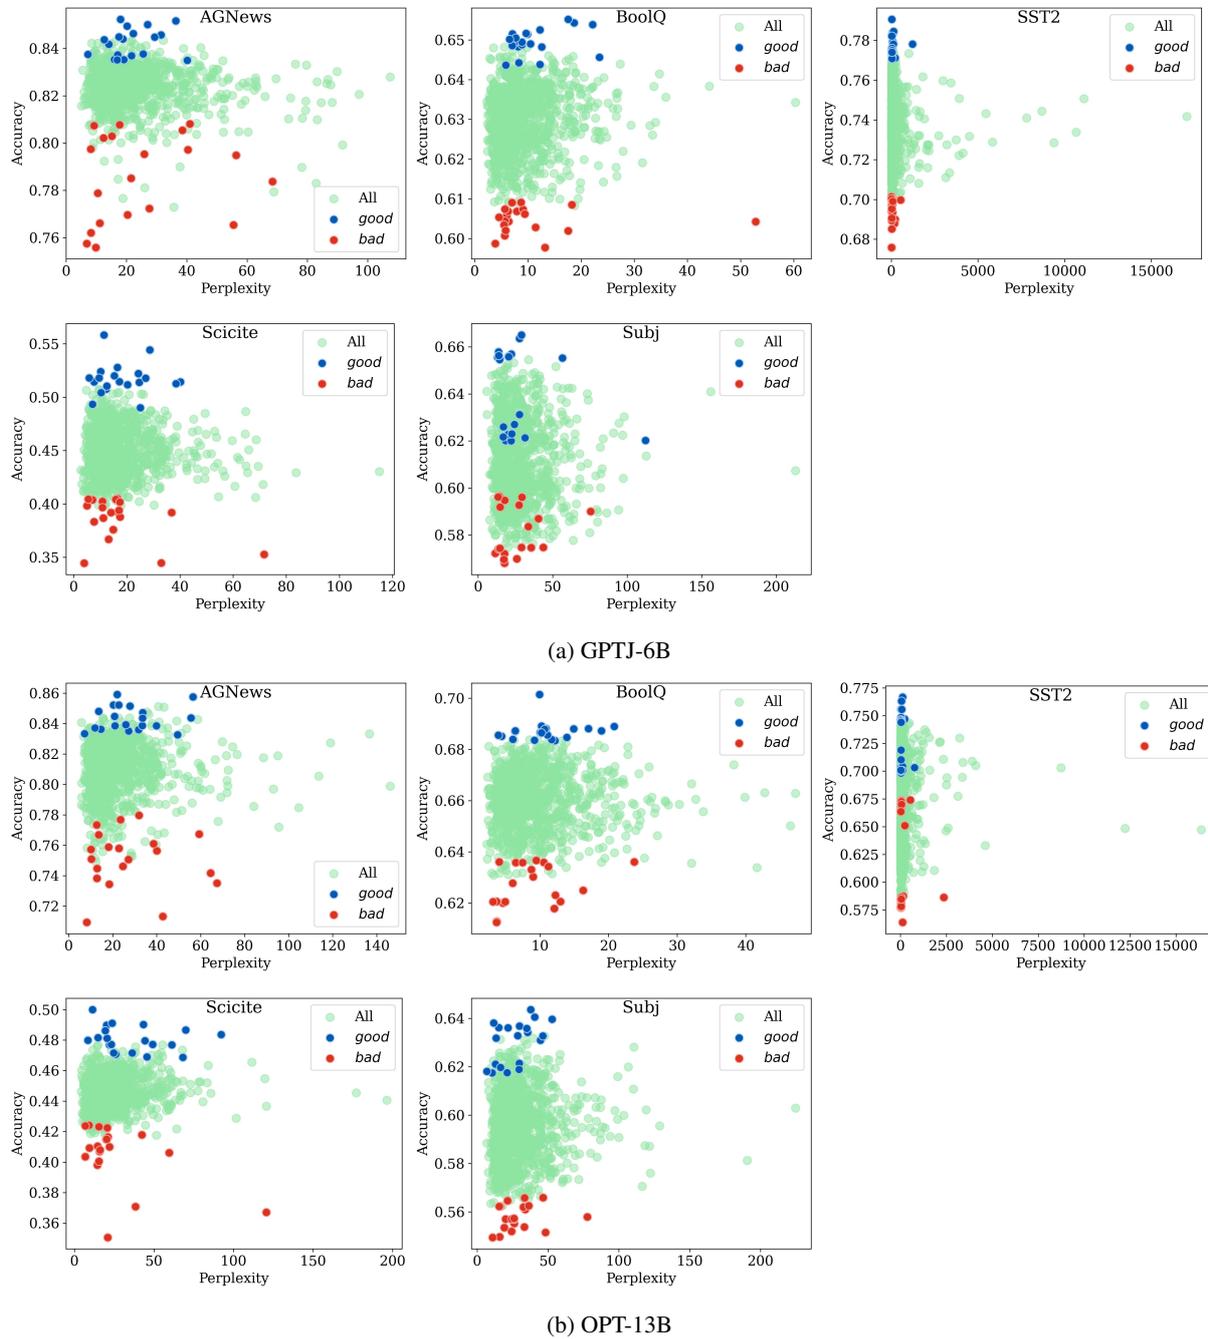


Figure 7: The accuracy versus perplexity across different tasks. Each dot corresponds to a training example. We do not observe any correlation between perplexity and accuracy. Note that we select the top E' examples per class. As a class may have much lower average accuracy than others, the *good* (resp., *bad*) examples may not be the examples with the globally highest (resp., lowest) accuracy.

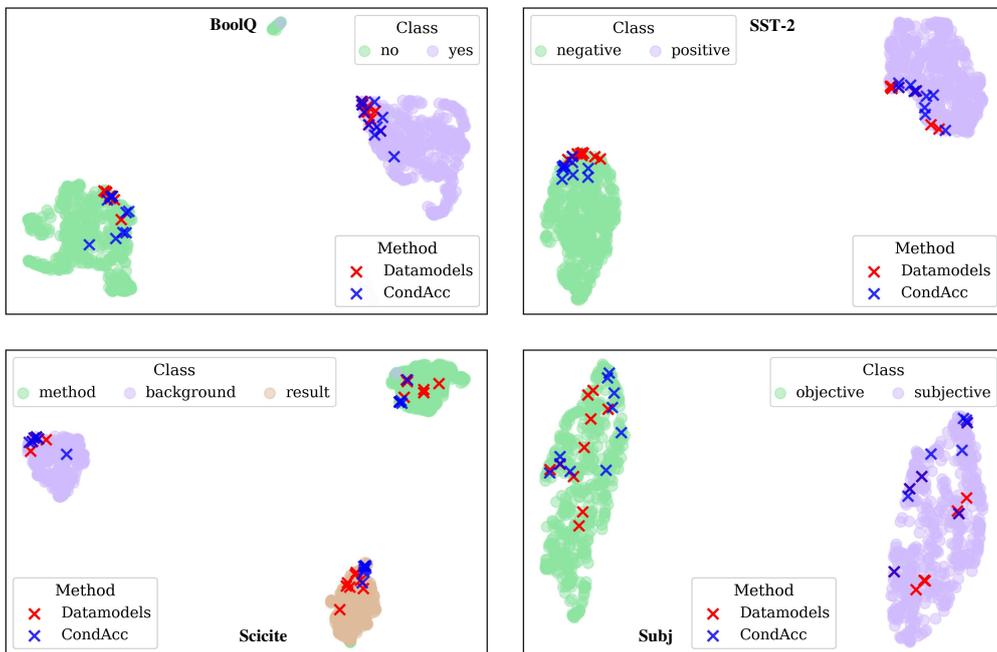


Figure 8: Visualizing the *good* training examples selected by CONDACC and DATAMODELS in datamodels embedding space across different tasks with GPTJ. Each dot is a training example, where datamodels spontaneously learn to encode class information in the embeddings.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations section
- A2. Did you discuss any potential risks of your work?
We use public available language models and commonly used datasets
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract and Sec1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Sec 3,4

- B1. Did you cite the creators of artifacts you used?
Sec 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
We use public available artifacts
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Sec 4 and Limitations
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Sec4, Table6

C Did you run computational experiments?

Sec5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix A.5

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Sec 4

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Sec 4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Not applicable. Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.