

# Importance Weighting Can Help Large Language Models Self-Improve

Chunyang Jiang, Chi-Min Chan, Wei Xue<sup>✉</sup>, Qifeng Liu, Yike Guo<sup>✉</sup>

Hong Kong University of Science and Technology  
{cjiangaq, cchanbc}@connect.ust.hk, {weixue, liuqifeng, yikeguo}@ust.hk

## Abstract

Large language models (LLMs) have shown remarkable capability in numerous tasks and applications. However, fine-tuning LLMs using high-quality datasets under external supervision remains prohibitively expensive. In response, LLM self-improvement approaches have been vibrantly developed recently. The typical paradigm of LLM self-improvement involves training LLM on self-generated data, part of which may be detrimental and should be filtered out due to the unstable data quality. While current works primarily employ filtering strategies based on answer correctness, in this paper, we demonstrate that filtering out correct but with high *distribution shift extent* (DSE) samples could also benefit the results of self-improvement. Given that the actual sample distribution is usually inaccessible, we propose a new metric called *DS weight* to approximate DSE, inspired by the *Importance Weighting* methods. Consequently, we integrate DS weight with self-consistency to comprehensively filter the self-generated samples and fine-tune the language model. Experiments show that with only a tiny valid set (up to 5% size of the training set) to compute DS weight, our approach can notably promote the reasoning ability of current LLM self-improvement methods. The resulting performance is on par with methods that rely on external supervision from pre-trained reward models.

## Introduction

Recently, Large Language Models (LLMs) have made impressive achievements on a large amount of NLP tasks and applications (Li et al. 2023a; OpenAI 2023; Yang et al. 2023; Li et al. 2023b). Moreover, new capabilities emerge in LLMs with the model size scaled to hundreds of billions of parameters, especially the general reasoning capabilities (Kojima et al. 2022). Relevant techniques like in-context few-shot learning (Brown et al. 2020a), Chain-of-Thought prompting (Wei et al. 2022), and self-consistency (Wang et al. 2023a) were further proposed to get better performance.

Despite the remarkable capabilities of LLMs pre-trained on the large corpus, fundamentally improving the model’s performance still necessitates fine-tuning on a great amount of high-quality supervised data (Huang et al. 2023a), which is usually costly. To alleviate this problem, many works are committed to investigating the *self-improvement* ability of

LLMs (Shinn et al. 2023; Madaan et al. 2023; Vernikos et al. 2024). Among them, fine-tuning the LLM on self-generated data appears as one of the most promising way (Gülçehre et al. 2023; Huang et al. 2023a; Wang et al. 2023b; Xu et al. 2023; Li et al. 2024). This formula typically includes generating reasoning thoughts and answers on unsupervised datasets, filtering data, and fine-tuning models on the self-generated data (Huang et al. 2023a). It is regarded as an attractive approach for LLMs to self-supervise by utilizing unlabeled data without external supervision.

The primary challenge of utilizing self-generated data is the variability in data quality. While high-quality samples can enhance the model’s reasoning abilities, there are low-quality samples that may detrimentally affect performance (Li and Qiu 2023). For example, an incorrectly generated answer could mislead the model. Therefore, a good filtering strategy is decisive for effective self-improvement. Many approaches have been proposed to address this issue. Inspired by Self-Consistency (Wang et al. 2023a), LMSI (Huang et al. 2023a) adopts majority voting to select the most *consistent* answer, under the assumption that consistency is positively related to the correctness. MoT (Li and Qiu 2023) further introduces uncertainty to the filtering strategy, by utilizing entropy to exclude high-uncertainty data points. Self-Alignment (Li et al. 2024) demonstrates that prompting the LLM to self-filter is also feasible.

However, present methods mostly emphasize assessing the correctness of generated samples, yet ignore the distribution shift problem. Specifically, the distribution of the LLM self-generated data may differ from that of real-world data, and fine-tuning models on samples with high *distribution shift extent* (DSE) may defect the resulting performance (Shumailov et al. 2023a). In this paper, we demonstrate that even self-generated samples with correct answers can possess high DSE, potentially degrading model performance. Consequently, filtering out high DSE samples is essential to further promote the efficacy of LLM self-improvement.

To exclude samples with high DSE, the primary question is how to estimate the DSE, since the actual data distribution is usually inaccessible. We note *Importance Weighting* (IW) (Sugiyama, Krauledat, and Müller 2007) as a well-known approach to address the traditional distribution shift problem (Sugiyama and Kawanabe 2012), where the key

idea is deriving importance weights based on the distribution ratio between test and training data, and using it to rebuild an unbiased training loss. IW usually contains two steps: *weight estimation* computes test-over-training density ratio and *weighted classification* utilizes the ratio to weight each data point and train the model (Fang et al. 2020).

Inspired by IW, we propose *Distribution Shift Weight* (DS weight) as a new metric to measure the DSE of self-generated samples. Based on this, we build an LLM self-improvement framework that incorporates both the correctness and DSE in its filtering strategy. Specifically, given a question-only dataset, we first let a pre-trained LLM generate multiple reasoning thoughts as well as answers. Then we create a tiny valid set comprising a few human-written demonstrations. With the pre-trained LLM and valid set, we leverage a simple approximation for importance weights to compute DS weight, as a measure of DSE, for each training data point. We subsequently combine the results from majority voting (for correctness) and DS weight (for DSE) to filter the dataset and fine-tune the LLM. We denote our framework as **Importance Weighting-based Self-Improvement (IWSI)**. Experiments show that the performance of IWSI largely surpasses baseline self-improvement methods and rivals the enhancements achieved with supervision from the pre-trained reward model.

Our contributions are threefold: (1) We propose a metric called DS weight to approximate the DSE of LLM self-generated data, with help from a tiny valid set. (2) Leveraging DS weight, we build a novel self-improvement framework called IWSI where the filtering strategy considers both the answer correctness and DSE. (3) We empirically examine the effectiveness of our proposed method, analyze the impact of high DSE samples on LLM self-improvement, and explore how DS weight interacts with other filtering criteria.

## Related Work

### LLM Self-Improvement

Fundamentally improving LLMs’ reasoning ability essentially requires fine-tuning on a large amount of high-quality supervised data. However, this methodology faces the threat that the stock of high-quality language data will be exhausted in some day (Villalobos et al. 2022). Self-improvement emerges as a promising approach to utilize the inherent knowledge to make supervision for self-training LLMs. While LLMs can easily generate extensive data, the data quality is not always guaranteed (Huang et al. 2023b) and training on unfiltered data may even cause performance degradation (Shumailov et al. 2023b). Therefore, an essential requirement in LLM self-improvement is data filtering.

Pioneering works (Wang et al. 2023b; Bai et al. 2022; Xu et al. 2023) use language models to generate diverse types of data such as feedback, instructions, and questions. They filter data by heuristic rules as well as manual inspection, which is challenging and costly. LMSI (Huang et al. 2023a) proposed a framework including generating data for a question-only dataset and using the majority voting (self-consistency) (Wang et al. 2023a) to select the most consistent answers, which is empirically proven to be effective

among various tasks. LMSI also demonstrates that the answer correctness is positively relevant to self-consistency. Along with this work, MoT (Li and Qiu 2023) proposes further filtering the consistent answers by entropy, which measures the answer uncertainty. Self-Alignment (Li et al. 2024) shows it is feasible to prompt the LLM self-filtering the generated data. To comprehensively evaluate the generated data, some works use external pre-trained LMs as the reward model to score the generated data, such as GENIE (Yehudai et al. 2024) and ReST (Gülçehre et al. 2023). With external supervision from the reward model, their filtering strategies are typically more considered.

### Importance Weighting

Importance weighting (IW) is a primary approach to mitigate the influence of distribution shift problem (Sugiyama and Kawanabe 2012). The typical IW process includes two steps: weight estimation and weighted classification. Weight estimation approximates the importance weights, which are subsequently used in the weighted classification stage to build a weighted training loss (Fang et al. 2023).

Traditional IW methods mainly estimate the importance weights by assessing the matching between training and test distribution in different ways, such as maximum mean discrepancy in a reproducing kernel Hilbert space (Huang et al. 2006), KL divergence (Sugiyama et al. 2007), and squared loss (Kanamori, Hido, and Sugiyama 2009). While these methods work well in linear models, their performances degrade largely in deep learning scenarios (Fang et al. 2020). To overcome this, DIW (Fang et al. 2020) proposes an end-to-end dynamic solution, which uses a deep network to predict the importance weights, and repeats weight estimation and weighted classification stages to iteratively converge on the optimal solution.

In this paper, we use some lemmas and empirical results in DIW to build the DS weight for estimating the DSE of self-generated data.

## Methodology

Fig. 1 shows the overview of IWSI. Given an unsupervised (question-only) dataset  $\mathcal{D}_q$ , we first use the pre-trained LLM  $\mathcal{M}_L$  to generate multiple candidate answers as well as the reasoning thoughts for each question, using CoT prompts (Wei et al. 2022). Following LMSI (Huang et al. 2023a), we adopt the majority voting to keep the most *consistent* answer and corresponding thoughts for each question, resulting in the consistency-filtered dataset  $\mathcal{D}_c$ . Then we calculate *DS weight* for every data point in  $\mathcal{D}_c$ , with the help of a tiny valid set  $\mathcal{D}_v$ . Lastly, we filter  $\mathcal{D}_c$  into  $\mathcal{D}_{ds}$  utilizing the DS weight and fine-tune the model  $\mathcal{M}_L$ . The following sections elaborate on different components of IWSI.

### Candidate Answers Generation and Self-Consistency Filtration

In this stage, we let the pre-trained LLM  $\mathcal{M}_L$  generate candidate answers as well as reasoning thoughts for an unsupervised dataset  $\mathcal{D}_q$  which only contains unlabeled questions. Given a question  $q_i \in \mathcal{D}_q$ , we concatenate Few-

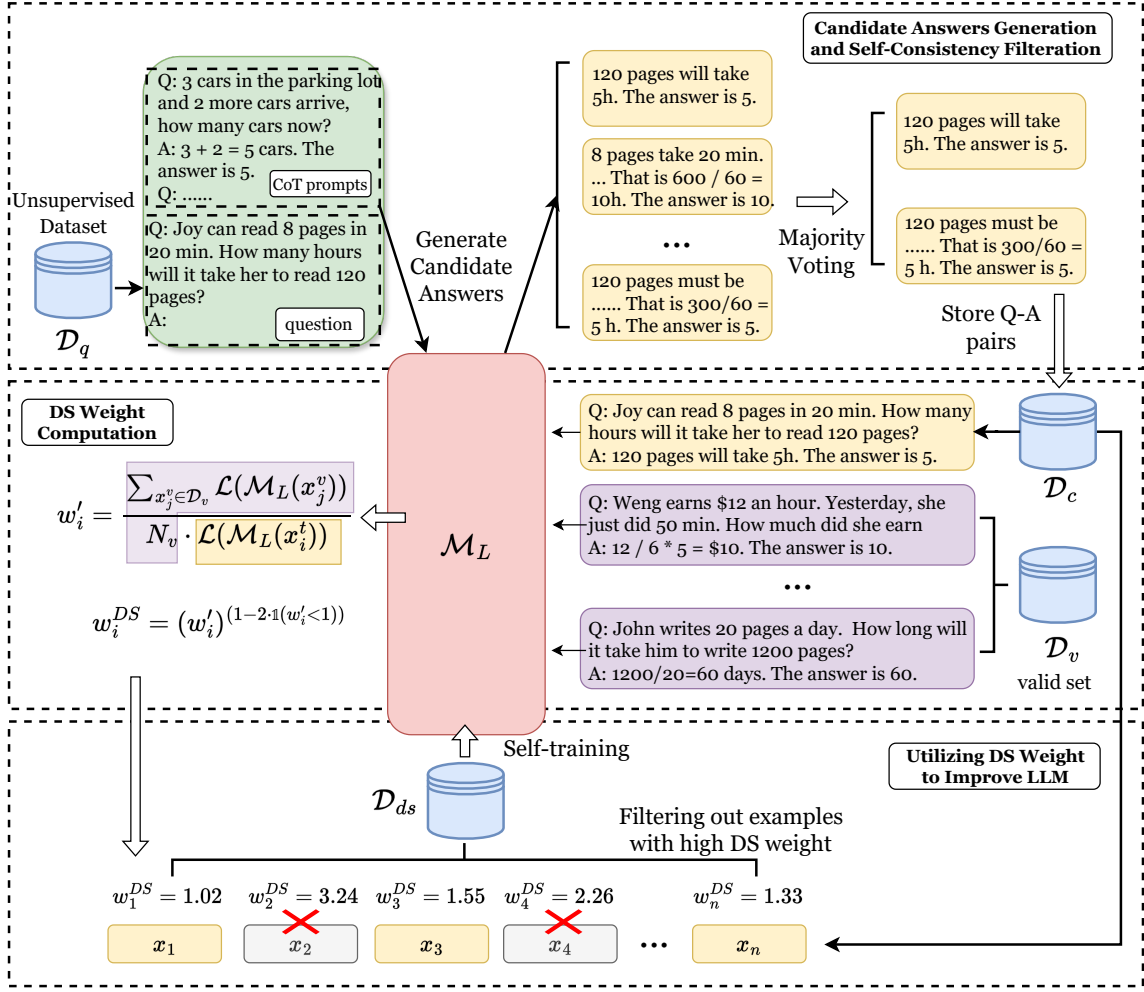


Figure 1: The overview of IWSI. Given the unsupervised dataset  $\mathcal{D}_q$ , the pre-trained LLM  $\mathcal{M}_L$  is first used to generate multiple candidate answers as well as the reasoning thoughts, prompted by CoT examples. Then IWSI uses majority voting to select the most consistent answer and corresponding thoughts, stored in  $\mathcal{D}_c$ . With the help of  $\mathcal{D}_v$ , IWSI calculates DS weight for each data point in  $\mathcal{D}_c$ . IWSI filters  $\mathcal{D}_c$  into  $\mathcal{D}_{ds}$  by keeping samples with the  $k$ %-lowest DS weight and lastly self-trains  $\mathcal{M}_L$ .

Shot-CoT (Wei et al. 2022) prompts with  $q_i$  to form the input text  $x_i$ . With temperature  $T > 0$ , we let  $\mathcal{M}_L$  sample  $m$  candidate answers  $[a_{i_1}, a_{i_2}, \dots, a_{i_m}]$  and their reasoning thoughts  $[r_{i_1}, r_{i_2}, \dots, r_{i_m}]$ . Then we select the most *consistent* answer  $\hat{a}_i$  by majority voting (Wang et al. 2023a),  $\hat{a}_i = \arg \max_{a_{i_j}} \sum_{k=1}^m \mathbb{1}(a_{i_j} = a_{i_k})$ , and keep the corresponding reasoning thoughts  $R_i = \{r_{i_j} | a_{i_j} = \hat{a}_i, 1 \leq j \leq m\}$ . By repeating over each question in  $\mathcal{D}_q$ , the consistency-filtered dataset  $\mathcal{D}_c$  is built.

### DS Weight Computation

To elaborate *DS Weight*, we first introduce some important preliminaries in the distribution shift problem and importance weighting methods.

*Distribution shift* problem denotes that the training data and test data are drawn from two different distributions  $p_{train}$  and  $p_{test}$ , and  $p_{train} \neq p_{test}$  (Sugiyama and Kawane 2012). A common assumption for distribution shift is

that there exists a function  $w^*(x)$ , holding that:

$$\mathbb{E}_{p_{test}(x)}[f(x)] = \mathbb{E}_{p_{train}(x)}[w^*(x) \cdot f(x)] \quad (1)$$

for any function  $f$  of  $x$  (Fang et al. 2020). Based on Eq. 1, importance weighting methods (Sugiyama, Krauledat, and Müller 2007; Sugiyama et al. 2007) deal with distribution shift in two steps: *weight estimation* finds a proper solution for  $w^*(x)$ ; *weighted classification* trains the model with a weighted loss derived by substituting  $f$  in Eq. 1 with the target loss function.

Obviously, it plays a decisive role in importance weighting that finding the appropriate importance weights  $\mathcal{W} = \{w_i\}^{N_t}$ , to approximate  $w^*(x)$  in Eq. 1. To simplify the question, DIW (Fang et al. 2020) provides an empirical surrogate goal with the help of a valid set:

$$\frac{1}{N_v} \sum_{j=1}^{N_v} \mathcal{L}(\mathcal{M}(x_j^v)) \approx \frac{1}{N_t} \sum_{i=1}^{N_t} w_i \cdot \mathcal{L}(\mathcal{M}(x_i^t)). \quad (2)$$

Here  $N_v$ ,  $N_t$ ,  $x^v$ , and  $x^t$  indicate the size of the valid set, the size of the training set, data in the valid set, and data in the training set.  $\mathcal{M}$  is the training model and  $\mathcal{L}$  represents the training loss.

While in DIW, Eq. 2 is used as a goal to train a deep model that predicts the desired  $\mathcal{W}$ , we use Eq. 2 to design a naive measurement for the distribution shift extent between training samples and valid set. Our intuition is that when the training data distribution is identical to the valid data distribution,  $w_i \equiv 1$  would be a proper solution to Eq. 2. Conversely, the larger the actual  $w_i$  differs from 1, the more different the training distribution and valid distribution are.

Based on this idea, we first design a naive estimation  $w'_i$  for  $x_i^t$  by regarding  $N_t$  as 1:

$$w'_i = \frac{\sum_{x_j^v \in \mathcal{D}_v} \mathcal{L}(\mathcal{M}_L(x_j^v))}{N_v \cdot \mathcal{L}(\mathcal{M}_L(x_i^t))} \quad (3)$$

where  $\mathcal{M}_L$  is the pre-trained LLM,  $\mathcal{L}$  denotes the *sft* loss (Brown et al. 2020b),  $\mathcal{D}_v$  is a tiny valid set and  $x_i^t$  is a self-generated training data point. Here we notice that the value range of  $w'_i$  is  $(0, +\infty)$  while the ideal value is 1, which creates asymmetry between the two deviation directions (lower than 1 and greater than 1) and makes filtering inconvenient. Therefore, to establish symmetry for both shift directions, we define *DS weight*  $w_i^{DS}$  as:

$$w_i^{DS} = \begin{cases} w'_i & \text{if } w'_i \geq 1 \\ \frac{1}{w'_i} & \text{if } w'_i < 1 \end{cases} \quad (4)$$

### Utilizing DS Weight to Improve LLM

With DS weight approximating DSE, we are able to further filter the self-generated data in  $\mathcal{D}_c$ , excluding data points that possibly possess higher DSE.

First, all data points are ranked with respect to their DS weight  $w_i^{DS}$ , and the  $k$ -percentile  $\sigma_{k\%}$  is selected, *s.t.*

$$\frac{\sum_i^{|\mathcal{D}_c|} \mathbb{1}(w_i^{DS} \leq \sigma_{k\%})}{|\mathcal{D}_c|} = k\% \quad (5)$$

where  $|\cdot|$  denotes the set size and  $w_i^{DS}$  is the corresponding DS weight of sample  $x_i$ . As a result, only samples whose  $w_i^{DS} \leq \sigma_{k\%}$  are kept to train the model  $\mathcal{M}_L$ . The training loss can be written as:

$$\mathcal{L}_F = \frac{1}{|\mathcal{D}_c| \cdot k\%} \sum_{x_i}^{\mathcal{D}_c} \mathbb{1}_{k\%}(x_i) \cdot \mathcal{L}(\mathcal{M}_L(x_i)) \quad (6)$$

where  $\mathbb{1}_{k\%}(x_i)$  equals to  $\mathbb{1}(w_i^{DS} \leq \sigma_{k\%})$  and  $\mathcal{L}$  represents the *sft* loss.

Another natural way to utilize DS weight is directly employing Eq. 3 to calculate a weighted loss, which is more analogous to the standard IW procedure. We also implement this variant in our work and denote it as *IWSI-w*. The weighted loss is:

$$\mathcal{L}_W = \frac{1}{|\mathcal{D}_c|} \sum_{x_i}^{\mathcal{D}_c} \text{Clip}(w'_i, C) \cdot \mathcal{L}(\mathcal{M}_L(x_i)) \quad (7)$$

where  $C$  is a constant. We clip  $w'_i$  to  $(0, C]$  for stabilizing the training process.

However, we found that *IWSI-w* is much less effective than *IWSI*. We believe this is mainly attributed to the inadequacy of Eq. 3. Empirical results and details are discussed in the experiment section.

## Experiment

### Setup

**Datasets** We conduct experiments on six datasets across three types of tasks: **Arithmetic Reasoning**: gsm8k (Cobbe et al. 2021) and SVAMP (Patel, Bhattamishra, and Goyal 2021). **Natural Language Inference**: Adversarial NLI subsets (Nie et al. 2020). ANLI-A1 and ANLI-A2 subsets are used. **Commonsense Reasoning**: OpenBookQA (Mihaylov et al. 2018) and StrategyQA (Geva et al. 2021).

For all datasets, only the questions are used to self-generate candidate answers. For gsm8k and SVAMP, we keep the original question format, which is the open-ended question. For the other four datasets, we unify the question format to the multiple choice question. The LLM must choose one option as its answer.

To build the valid set, we extract rationales from the original datasets apart from SVAMP, for which we manually write rationales. The size of valid sets varies among different datasets, but none of them exceeds 5% size of the corresponding training set. Appendix A provides more details about the split and statistics of all datasets.

**Baselines** The goal of our experiments is to verify whether incorporating DS weight into the filtering strategy in our proposed approach can help LLMs self-improve. Therefore, given the same base model, we compare *IWSI* with the fundamental self-improvement framework *LMSI* (Huang et al. 2023a), and some variants that we implement by adopting trendy filtering strategies designed for training LLMs on model-generated data.

*LMSI* (Huang et al. 2023a) is the first self-improvement framework that significantly improves LLMs' reasoning ability without any external supervision. The core idea of *LMSI* is adopting majority voting to select answers that are most likely correct, thus filtering the self-generated data.

*MoT* (Li and Qiu 2023) uses entropy to measure the uncertainty of the answers and further filters data. We combine this technique with *LMSI* and denote it as *Entropy-filter*.

*Self-Alignment* (Li et al. 2024) shows that LLM self-evaluation could be helpful in filtering strategy. We implement this idea with *LMSI* and denote it as *Self-filter*.

Works like *GENIE* (Yehudai et al. 2024) and *ReST* (Gülçehre et al. 2023) use pre-trained models to evaluate the self-generated samples. Intervened by external supervision, their filtering results are usually more comprehensive and meticulous. Following that, we also implement a variant of *LMSI* for reference, the *RM-filter*. *RM-filter* uses a pre-trained reward model to score the generated data, as *GENIE* (Yehudai et al. 2024) does.<sup>1</sup>

<sup>1</sup><https://huggingface.co/OpenAssistant/reward-model-deberta-v3-large-v2>

	<b>gsm8k</b>	<b>SVAMP</b>	<b>ANLI-A1</b>	<b>ANLI-A2</b>	<b>OpenBookQA</b>	<b>StrategyQA</b>	<b>Avg.</b>
base	7.0	14.7	16.4	14.6	31.8	48.3	22.1
LMSI	27.9	45.0	25.2	22.6	31.6	51.4	34.0
Entropy-filter	22.7	56.0	25.2	22.8	33.4	51.2	35.2
Self-filter	35.6	<b>62.7</b>	22.8	<b>25.6</b>	35.0	50.4	38.7
IWSI-w	37.0	43.3	21.8	21.8	31.8	49.2	34.2
<b>IWSI</b>	<b>37.6</b>	<b>62.7</b>	<b>27.2</b>	23.4	<b>37.0</b>	<b>54.6</b>	<b>40.4</b>
RM-filter	40.0	66.3	25.6	25.0	34.2	51.4	40.4

Table 1: Accuracy results on all datasets. Numbers in the table are the accuracy percent. The first part is the performance of the base model. The second part is the performance of three baseline self-improvement methods, our proposed method IWSI, and a variant IWSI-w. As RM-filter uses the external reward model, we list its performance separately at the bottom of the table.

**Implementation details** We select Llama3-8B as our base model (Touvron et al. 2023). For each question, we generate 15 candidates, with temperature  $T = 1.1$ . All training process is performed on eight RTX-4090 GPUs. The training batch size per device is set to 1 and the gradient accumulation steps is 4. We use LoRA (Hu et al. 2022) to do fine-tuning. We use AdamW (Loshchilov and Hutter 2019) optimizer and the learning rate is  $3e-4$ . Few-Shot-CoT prompts are only applied in generating candidate answers and the evaluation stage. CoT examples for each dataset, prompts used for Self-filter, and details about how to derive the answer from output texts are given in Appendix D. The source code and supplementary materials are available at <https://github.com/rubickkcibur/IWSI>.

## Main Results

The main comparison results are shown in Table 1. The evaluation metric is accuracy percent and all results are derived by greedy decoding. The top part is the performance of the base model. The middle part are self-improvement baselines and our proposed method IWSI. For reference, we list the performance of RM-filter at the bottom of the table. For fairness, we universally set the filtering percentage  $k = 80$  for IWSI, Entropy-filter, Self-filter, and RM-filter.

Among self-improvement methods (the middle part), IWSI is the only one that consistently outperforms LMSI, and it also achieves the best in almost all datasets. We further empirically demonstrate that the superiority of IWSI primarily stems from excluding self-generated samples with higher DSE, rather than merely from access to part of the information of the valid set (the mean loss value of valid samples). Details are in Appendix F.

For IWSI-w, the variant of IWSI that uses DS weight to compute weighted loss other than filtering data, it generally performs worse than IWSI, even though IWSI-w is more compliant with the standard importance weighting formula. The most possible reason is that unlike deep methods like DIW (Fang et al. 2020), which uses a deep neural network to learn the weights, our weight estimation (Eq. 3) is a pretty naive approach. While it largely reduces computational cost, it also omits the semantic similarity among training samples, potentially compromising efficacy. Therefore, the weighted loss in IWSI-w might make the training process difficult and

noisy. In contrast, IWSI only uses the weight as an indicator to rank the samples with respect to DSE, without directly incorporating the weight into the training loss, which makes the overall process more robust.

As for the RM-filter, we found that it does not always perform the best among all six datasets, even though it introduces external supervision by using a pre-trained reward model. As Table 1 shows, after incorporating both the answer correctness and DSE of samples, the overall performance of IWSI is comparable to that achieved with external supervision from a pre-trained reward model.

## Hyperparameter Study

We investigate the effect of varying the filtering threshold  $k$  and corresponding percentile  $\sigma_{k\%}$  (in Eq. 5). Fig. 2 shows the accuracy results on gsm8k, StrategyQA, and ANLI-A1. As the figure shows, either a too-large or too-small  $k$  value will make the performance degrade. When  $k$  is very large, more samples with high DSE will be kept, thus potentially harming the performance. If the  $k$  is pretty small, there will not be sufficient samples kept to support the model training. The optimal  $k$  value range varies across different tasks. In general, around 80% would be an appropriate choice.

Fig. 3 shows the varying  $k$ -percentile  $\sigma_{k\%}$  of DS weight. While  $\sigma_{k\%}$  of different datasets are similar when  $k$  is very small, the difference becomes larger as  $k$  increases. This phenomenon suggests that the boundary above which the DSE of samples can be regarded as "high" is relative according to different datasets.

## Valid Set Analysis

The valid set  $\mathcal{D}_v$  plays a crucial role in IWSI. It determines the calculation results of DS weight and subsequently steers the filtering strategy. Therefore, variation in the composition of the valid set can introduce randomness and thus potential instability. In this section, we take the gsm8k dataset as example to discuss the impact of valid set.

We employ the loss value distribution as the analytical tool and, for simplicity, we assume all distributions of different sample sets conform to the normal distribution. For example, the loss value distribution of valid set is denoted as  $\mathcal{N}_v(\mu_v, \sigma_v^2)$ , where  $\mu_v$  and  $\sigma_v$  are the *mean* and *standard deviation* respectively.




Figure 2: Accuracy results with varying  $k$  values.




Figure 3:  $\sigma_k\%$  (in Eq. 5) with varying  $k$  values.

Fig. 4 shows distributions of the valid set and self-generated samples before and after IWSI. Analogous to our intuition, the distributions differ significantly between valid set samples and self-generated samples before IWSI, and become much closer after IWSI, illustrating the effectiveness of IWSI in handling the distribution shift problem. Furthermore, we provide quantitative analyses and a case study in Appendix E for a better understanding of how the LLM generation was affected by IWSI.

The next question is *would the randomness of valid set composition cause great instability in IWSI*, since  $\sigma_v$  is apparently not small enough. The answer is “No” as long as there is an adequate valid set size  $N_v$ . Theoretically, in Eq. 3, it is only the *sample mean*, denoted as  $\bar{\mathcal{L}}_v$ , that matters.  $\bar{\mathcal{L}}_v$  is also subject to the normal distribution, with its standard deviation inversely proportional to the size  $N_v$ :

$$\begin{aligned} \bar{\mathcal{L}}_v &= \frac{\sum_{x_j^v \in \mathcal{D}_v} \mathcal{L}(\mathcal{M}_L(x_j^v))}{N_v} \\ \bar{\mathcal{L}}_v &\sim \mathcal{N}_{\bar{v}}(\mu_v, (\frac{\sigma_v}{N_v})^2) \end{aligned} \quad (8)$$

Eq. 8 implies that increasing  $N_v$  can scale down the variance of  $\bar{\mathcal{L}}_v$ , thus making the estimation more stable. More importantly, *it is completely irrelevant to the size of the training samples*. For instance, in gsm8k, if the valid set size is 100, the standard deviation of  $\bar{\mathcal{L}}_v$  is  $\bar{\sigma}_v = \frac{\sigma_v}{100} = 4.1 \times 10^{-3}$ , which is small enough to mitigate the interference of randomness.

$N_v$	Comp.1		Comp.2		Comp.3	
	$\bar{\mathcal{L}}_v$	acc	$\bar{\mathcal{L}}_v$	acc	$\bar{\mathcal{L}}_v$	acc
50	2.053	36.77	2.083	36.54	2.091	36.47
100	2.077	36.69	2.086	36.47	2.054	37.00

Table 2: Results of different valid sets on gsm8k.




Figure 4: Loss value distributions of the valid set samples, self-generated samples of base model (generated-base), and self-generated samples after IWSI (generated-IWSI), of gsm8k.  $\mu$  and  $\sigma$  denote the *mean* and *standard deviation*.

To empirically investigate the influence of different valid set compositions, we randomly constitute six subsets of the valid set of gsm8k and test IWSI with them. Table 2 shows the results.  $N_v$  denotes the valid set size.  $\bar{\mathcal{L}}_v$  is the sample mean of different composition. We use *acc* as the metric.

As we can see, the impact of different compositions on the accuracy results is quite minimal. We believe this is primarily attributed to the *double-robustness* of IWSI. First, the DS weight calculation is robust to the valid set composition, since it only uses the sample mean which varies very little. Furthermore, the filtering strategy is also robust to the DS weight, since the DS weight is used for ranking other than weighting. As a result, samples with extremely high DSE are probably always discarded even if DS weight changes.

### Orthogonality Analysis

In IWSI, two factors are considered in the filtering strategy, the answer correctness (represented by self-consistency) and the sample DSE (represented by DS weight). A natural question is what is the relationship between these two factors. Are they correlated to or independent of each other? To explore this question, we counted the percentage of samples with correct answers (using the ground truth labels) across different DS weight intervals, as Fig 5 shows. Along  $x$ -axis are the selected intervals:  $[1, 1.1)$ ,  $[1.1, 1.3)$ ,  $[1.3, 1.5)$ ,  $[1.5, 2)$ , and  $[2, \infty)$ . In each bar, the upper portion (yellow) indicates the ratio of correct answers, while the lower portion (blue) represents the ratio of wrong answers. For all datasets, we observe a general downward trend in the ratio of correct answers, as DS weight increases. The highest ratios of correct answers is found either in the  $[1, 1.1)$  interval (for gsm8k and ANLI-A1) or in the  $[1.1, 1.3)$  interval (for StrategyQA). However, both correct and wrong answers occupy a

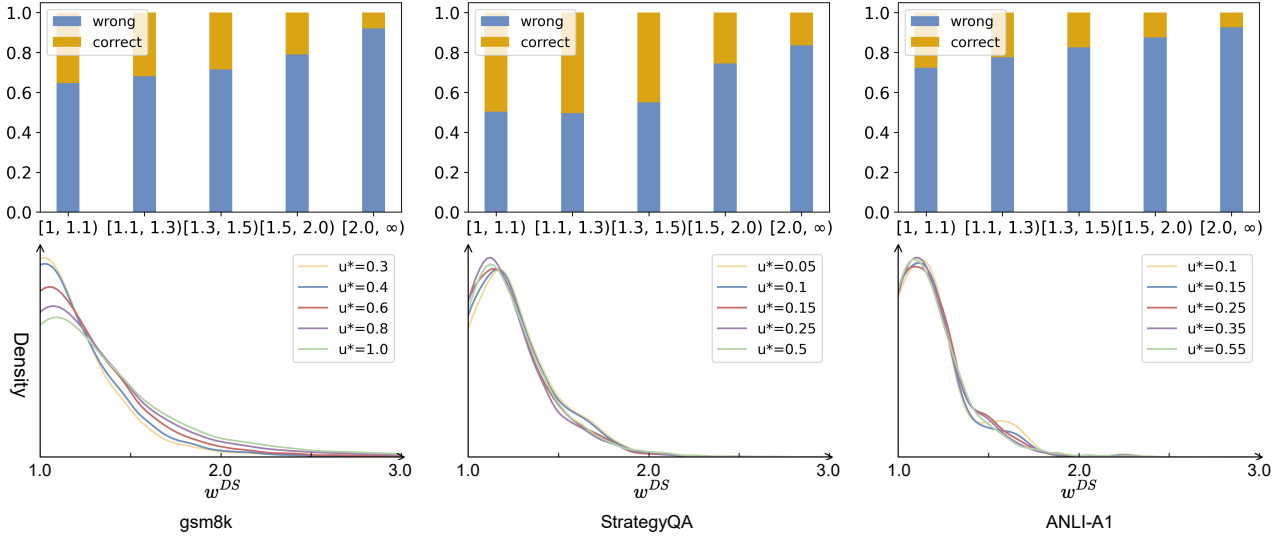


Figure 5: The first row shows relationship between answer correctness and DSE, where along  $x$ -axis are DSE intervals and the  $y$ -axis indicates the proportion of correct answers and wrong answers. The second row are the DS weight probability density function curves with varying uncertainty threshold  $u^*$  (Eq. 9).

portion that can not be ignored in every interval, suggesting a degree of independence between these two factors.

We delve deeper into the relationship between DSE and the answer uncertainty, which is first investigated by MoT (Li and Qiu 2023) regarding its impact on self-improvement. MoT also suggested using entropy to represent answer uncertainty. We briefly introduce the calculation: given a certain question  $q$ , the self-generated candidate answers  $[a_1, a_2, \dots, a_m]$ , and the most consistent answer  $\hat{a}$ , uncertainty  $u$  is computed in the following steps:

$$\begin{aligned}
 A^* &= \text{unique}(\{a_i\}^m) \\
 p(a_j^*) &= \sum_i^m \mathbb{1}(a_j^* = a_i) / m \\
 u &= - \sum_{a_j^*}^{A^*} p(a_j^*) \log p(a_j^*)
 \end{aligned} \tag{9}$$

where  $A^* = \{a_1^*, a_2^*, \dots\}$  is the unique answer set. The higher  $u$  is, the more uncertain the answer is. In extreme cases, if  $u = 0$ , all candidate answers are identical, and if each candidate answer has its unique value,  $u$  will reach the maximum  $\log m$ . For convenience, we normalize  $u$  with a divisor  $\log m$  and we denote the filter threshold as  $u^*$ .

We draw the probability density function (PDF) of DS weight for various uncertainty thresholds  $u^*$ . The second row of Fig. 5 shows the results. For arithmetic reasoning (gsm8k), as  $u^*$  increases, the peak of PDF falls and the PDF curve becomes flatter, indicating a growth in the proportion of samples with high DSE. Conversely, for common-sense reasoning (StrategyQA) and natural language inference (ANLI-A1), the relationship between uncertainty and DSE appears much weaker. The PDF curves are almost identical, with little variation at the peak, suggesting that DSE is

nearly orthogonal to the uncertainty.

### Perception of DSE

We conducted a case study on gsm8k to provide an intuitive perception about what a correct but with high DSE sample looks like. We compare the generated answers with the highest and lowest DSE for the same question. We found that cases with the highest DSE are usually notably absurd that we can easily tell them apart from human-written samples. We categorize these samples into 3 types:

- **Redundant samples.** Redundant samples include irrelevant or repeated information in the reasoning thoughts, making it confusing.
- **Jumping samples.** Jumping samples omit essential reasoning steps or even directly give the answer, making it less logically fluent.
- **Spurious samples.** The reasoning steps in a *spurious* sample (Guu et al. 2017; Jiang et al. 2023) are logically wrong. They get the correct answer just by coincidence.

We give more exact demonstrations in Appendix B.

### Conclusion

In this paper, we investigate the impact of sample DSE on LLM self-improvement. We propose *DS weight* to approximate the DSE inspired by importance weighting methods, and a novel framework IWSI where the filtering strategy comprehensively considers DSE and answer correctness. Empirical results demonstrate that the incorporation of DS weight significantly enhances the effectiveness of LLM self-improvement. Further analysis reveals that DSE is nearly orthogonal to other factors, suggesting a new direction to promote LLM self-improvement for the future work.

## Acknowledgements

This research was supported by Theme-based Research Scheme (T45-205/21-N) from Hong Kong RGC, and Generative AI Research and Development Centre from InnoHK. The corresponding authors are Wei Xue and Yike Guo.

## References

- Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; Chen, C.; Olsson, C.; Olah, C.; Hernandez, D.; Drain, D.; Ganguli, D.; Li, D.; Tran-Johnson, E.; Perez, E.; Kerr, J.; Mueller, J.; Ladish, J.; Landau, J.; Ndousse, K.; Lukosiute, K.; Lovitt, L.; Sellitto, M.; Elhage, N.; Schiefer, N.; Mercado, N.; DasSarma, N.; Lasenby, R.; Larson, R.; Ringer, S.; Johnston, S.; Kravec, S.; Showk, S. E.; Fort, S.; Lanham, T.; Telleen-Lawton, T.; Conerly, T.; Henighan, T.; Hume, T.; Bowman, S. R.; Hatfield-Dodds, Z.; Mann, B.; Amodei, D.; Joseph, N.; McCandlish, S.; Brown, T.; and Kaplan, J. 2022. Constitutional AI: Harmlessness from AI Feedback. *CoRR*, abs/2212.08073.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020a. Language Models are Few-Shot Learners. In *NeurIPS*.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020b. Language Models are Few-Shot Learners. In *NeurIPS*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *CoRR*, abs/2110.14168.
- Fang, T.; Lu, N.; Niu, G.; and Sugiyama, M. 2020. Rethinking Importance Weighting for Deep Learning under Distribution Shift. In *NeurIPS*.
- Fang, T.; Lu, N.; Niu, G.; and Sugiyama, M. 2023. Generalizing Importance Weighting to A Universal Solver for Distribution Shift Problems. In *NeurIPS*.
- Geva, M.; Khashabi, D.; Segal, E.; Khot, T.; Roth, D.; and Berant, J. 2021. Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. *Trans. Assoc. Comput. Linguistics*, 9: 346–361.
- Gülçehre, Ç.; Paine, T. L.; Srinivasan, S.; Konyushkova, K.; Weerts, L.; Sharma, A.; Siddhant, A.; Ahern, A.; Wang, M.; Gu, C.; Macherey, W.; Doucet, A.; Firat, O.; and de Freitas, N. 2023. Reinforced Self-Training (ReST) for Language Modeling. *CoRR*, abs/2308.08998.
- Guu, K.; Pasupat, P.; Liu, E. Z.; and Liang, P. 2017. From Language to Programs: Bridging Reinforcement Learning and Maximum Marginal Likelihood. In *ACL (1)*, 1051–1062. Association for Computational Linguistics.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*. OpenReview.net.
- Huang, J.; Gu, S.; Hou, L.; Wu, Y.; Wang, X.; Yu, H.; and Han, J. 2023a. Large Language Models Can Self-Improve. In *EMNLP*, 1051–1068. Association for Computational Linguistics.
- Huang, J.; Smola, A. J.; Gretton, A.; Borgwardt, K. M.; and Schölkopf, B. 2006. Correcting Sample Selection Bias by Unlabeled Data. In *NIPS*, 601–608. MIT Press.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; and Liu, T. 2023b. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *CoRR*, abs/2311.05232.
- Jiang, C.; Zhu, T.; Zhou, H.; Liu, C.; Deng, T.; Hu, C.; and Li, J. 2023. Path Spuriousness-aware Reinforcement Learning for Multi-Hop Knowledge Graph Reasoning. In *EACL*, 3173–3184. Association for Computational Linguistics.
- Kanamori, T.; Hido, S.; and Sugiyama, M. 2009. A Least-squares Approach to Direct Importance Estimation. *J. Mach. Learn. Res.*, 10: 1391–1445.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large Language Models are Zero-Shot Reasoners. In *NeurIPS*.
- Li, G.; Hammoud, H. A. A. K.; Itani, H.; Khizbullin, D.; and Ghanem, B. 2023a. CAMEL: Communicative Agents for "Mind" Exploration of Large Scale Language Model Society. *CoRR*, abs/2303.17760.
- Li, R.; Allal, L. B.; Zi, Y.; Muennighoff, N.; Kocetkov, D.; Mou, C.; Marone, M.; Akiki, C.; Li, J.; Chim, J.; Liu, Q.; Zheltonozhskii, E.; Zhuo, T. Y.; Wang, T.; Dehaene, O.; Davaadorj, M.; Lamy-Poirier, J.; Monteiro, J.; Shliazhko, O.; Gontier, N.; Meade, N.; Zebaze, A.; Yee, M.; Umaphathi, L. K.; Zhu, J.; Lipkin, B.; Oblukulov, M.; Wang, Z.; V, R. M.; Stillerman, J.; Patel, S. S.; Abulkhanov, D.; Zocca, M.; Dey, M.; Zhang, Z.; Moustafa-Fahmy, N.; Bhattacharyya, U.; Yu, W.; Singh, S.; Luccioni, S.; Villegas, P.; Kunakov, M.; Zhdanov, F.; Romero, M.; Lee, T.; Timor, N.; Ding, J.; Schlesinger, C.; Schoelkopf, H.; Ebert, J.; Dao, T.; Mishra, M.; Gu, A.; Robinson, J.; Anderson, C. J.; Dolan-Gavitt, B.; Contractor, D.; Reddy, S.; Fried, D.; Bahdanau, D.; Jernite, Y.; Ferrandis, C. M.; Hughes, S.; Wolf, T.; Guha, A.; von Werra, L.; and de Vries, H. 2023b. StarCoder: may the source be with you! *CoRR*, abs/2305.06161.
- Li, X.; and Qiu, X. 2023. MoT: Memory-of-Thought Enables ChatGPT to Self-Improve. In *EMNLP*, 6354–6374. Association for Computational Linguistics.
- Li, X.; Yu, P.; Zhou, C.; Schick, T.; Levy, O.; Zettlemoyer, L.; Weston, J. E.; and Lewis, M. 2024. Self-Alignment with Instruction Backtranslation. In *The Twelfth International Conference on Learning Representations*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR (Poster)*. OpenReview.net.

- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; Gupta, S.; Majumder, B. P.; Hermann, K.; Welleck, S.; Yazdanbakhsh, A.; and Clark, P. 2023. Self-Refine: Iterative Refinement with Self-Feedback. In *NeurIPS*.
- Mihaylov, T.; Clark, P.; Khot, T.; and Sabharwal, A. 2018. Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering. In *EMNLP*, 2381–2391. Association for Computational Linguistics.
- Nie, Y.; Williams, A.; Dinan, E.; Bansal, M.; Weston, J.; and Kiela, D. 2020. Adversarial NLI: A New Benchmark for Natural Language Understanding. In *ACL*, 4885–4901. Association for Computational Linguistics.
- OpenAI. 2023. GPT-4 Technical Report. *CoRR*, abs/2303.08774.
- Patel, A.; Bhattamishra, S.; and Goyal, N. 2021. Are NLP Models really able to Solve Simple Math Word Problems? In *NAACL-HLT*, 2080–2094. Association for Computational Linguistics.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023. Reflexion: language agents with verbal reinforcement learning. In *NeurIPS*.
- Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Gal, Y.; Papernot, N.; and Anderson, R. J. 2023a. The Curse of Recursion: Training on Generated Data Makes Models Forget. *CoRR*, abs/2305.17493.
- Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Gal, Y.; Papernot, N.; and Anderson, R. J. 2023b. The Curse of Recursion: Training on Generated Data Makes Models Forget. *CoRR*, abs/2305.17493.
- Sugiyama, M.; and Kawanabe, M. 2012. *Machine Learning in Non-Stationary Environments - Introduction to Covariate Shift Adaptation*. Adaptive computation and machine learning. MIT Press.
- Sugiyama, M.; Krauledat, M.; and Müller, K. 2007. Covariate Shift Adaptation by Importance Weighted Cross Validation. *J. Mach. Learn. Res.*, 8: 985–1005.
- Sugiyama, M.; Nakajima, S.; Kashima, H.; von Büna, P.; and Kawanabe, M. 2007. Direct Importance Estimation with Model Selection and Its Application to Covariate Shift Adaptation. In *NIPS*, 1433–1440. Curran Associates, Inc.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *CoRR*, abs/2302.13971.
- Vernikos, G.; Brazinskas, A.; Adámek, J.; Mallinson, J.; Severyn, A.; and Malmi, E. 2024. Small Language Models Improve Giants by Rewriting Their Outputs. In *EACL (1)*, 2703–2718. Association for Computational Linguistics.
- Villalobos, P.; Sevilla, J.; Heim, L.; Besiroglu, T.; Hobbhahn, M.; and Ho, A. 2022. Will we run out of data? An analysis of the limits of scaling datasets in Machine Learning. *CoRR*, abs/2211.04325.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q. V.; Chi, E. H.; Narang, S.; Chowdhery, A.; and Zhou, D. 2023a. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *ICLR*. OpenReview.net.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khashabi, D.; and Hajishirzi, H. 2023b. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In *ACL (1)*, 13484–13508. Association for Computational Linguistics.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E. H.; Le, Q. V.; and Zhou, D. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*.
- Xu, C.; Sun, Q.; Zheng, K.; Geng, X.; Zhao, P.; Feng, J.; Tao, C.; and Jiang, D. 2023. WizardLM: Empowering Large Language Models to Follow Complex Instructions. *CoRR*, abs/2304.12244.
- Yang, K.; Swope, A. M.; Gu, A.; Chalamala, R.; Song, P.; Yu, S.; Godil, S.; Prenger, R. J.; and Anandkumar, A. 2023. LeanDojo: Theorem Proving with Retrieval-Augmented Language Models. In *NeurIPS*.
- Yehudai, A.; Carmeli, B.; Mass, Y.; Arviv, O.; Mills, N.; Toledo, A.; Shnarch, E.; and Choshen, L. 2024. Genie: Achieving Human Parity in Content-Grounded Datasets Generation. *CoRR*, abs/2401.14367.