

SPECULATIVE KNOWLEDGE DISTILLATION: BRIDGING THE TEACHER-STUDENT GAP THROUGH INTER-LEAVED SAMPLING

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent advances in knowledge distillation (KD) have enabled smaller student models to approach the performance of larger teacher models. However, popular methods such as supervised KD and on-policy KD, are adversely impacted by the knowledge gaps between teacher-student in practical scenarios. Supervised KD suffers from a distribution mismatch between training with a static dataset and inference over final student-generated outputs. Conversely, on-policy KD, which uses student-generated samples for training, can suffer from low-quality training examples with which teacher models are not familiar, resulting in inaccurate teacher feedback. To address these limitations, we introduce Speculative Knowledge Distillation (SKD), a novel approach that leverages cooperation between student and teacher models to generate high-quality training data on-the-fly while aligning with the student’s inference-time distribution. In SKD, the student proposes tokens, and the teacher replaces poorly ranked ones based on its own distribution, transferring high-quality knowledge adaptively. We evaluate SKD on various text generation tasks, including translation, summarization, math, and instruction following, and show that SKD consistently outperforms existing KD methods across different domains, data sizes, and model initialization strategies.




Figure 1: SKD outperforms supervised and on-policy KD for our tested tasks: Assamese-to-English translation, dialogue summarization, and arithmetic reasoning. For teacher models, we employ supervised FT GEMMA-7B-IT and QWEN-7B-IT as teacher models (IT means instruction-tuned); for student models, we use GEMMA-2B and QWEN-0.5B (either instruction-tuned or supervised FT student models checkpoints depending on the best performance). Supervised KD is trained on ground-truth outputs, while on-policy KD uses self-generated data. All models use greedy decoding for evaluation.

1 INTRODUCTION

Text generation capabilities of large language models (LLMs) have seen continuous improvement, largely driven by scaling up the number of parameters and the amount of training data (Kaplan et al., 2020; Anil et al., 2023; Dubey et al., 2024; OpenAI et al., 2024). However, the substantial inference-time costs and memory footprint associated with LLMs present significant challenges for practical deployment (Agarwal et al., 2024). Therefore, compressing LLMs while maintaining their performance is crucial for real-time practical applications.

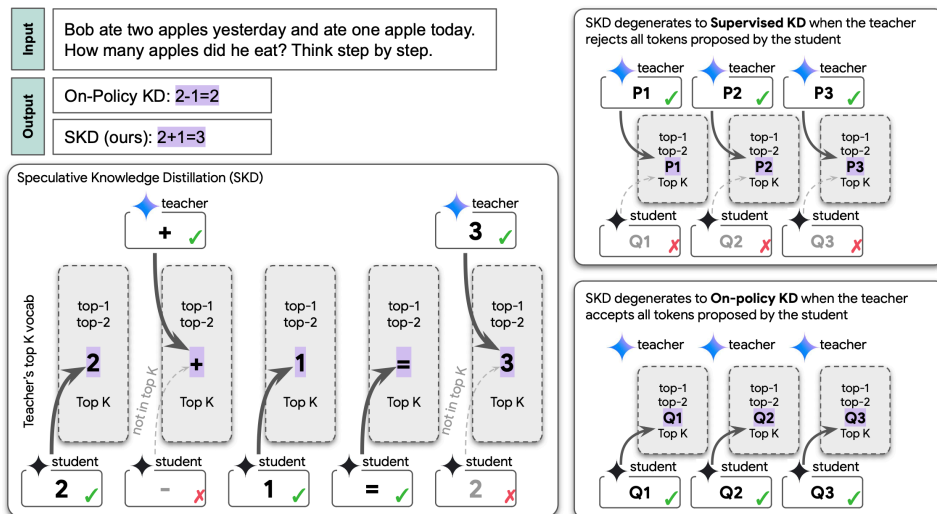


Figure 2: Overview of Speculative Knowledge Distillation (SKD) on an arithmetic reasoning task. Left: SKD addresses the limitations of on-policy knowledge distillation (KD) by filtering out low-quality student samples and replacing them with teacher generated tokens. **Right:** illustration of how SKD generalizes to both supervised KD (replacing with all teacher tokens) and on-policy (accepting all student tokens).

Knowledge Distillation (KD) (Hinton et al., 2015) is a widely used method to compress LLMs by transferring knowledge from a larger teacher model to a smaller student model. Traditional KD approaches, such as supervised KD (Sanh et al., 2020) and SeqKD (Kim & Rush, 2016b), rely on a static dataset of outputs to train the student model. However, this fixed dataset can lead to a distribution mismatch between the training data and the student’s generated samples at inference time, hindering the student’s learning. On-policy KD (Agarwal et al., 2024) attempts to address this train-inference mismatch by training the student on its self-generated samples, with the teacher provides feedback through its token-level probabilities. However, the student may generate low-quality samples that are out-of-distribution (OOD) for the teacher, especially early in the training, leading to inaccurate teacher feedback. Additionally, the performance of on-policy is highly sensitive to the student’s initialization, as illustrated in Figure 4.

In this work, we introduce **Speculative Knowledge Distillation (SKD)** that uses interleaved sampling between teacher and student to overcome the challenges faced by supervised and on-policy distillation (Figure 2). Similar to on-policy KD, SKD utilizes student-generated samples to address the train-inference mismatch. However, to mitigate the issue of low-quality student samples, SKD filters out intermediate tokens that the teacher is unlikely to produce and instead re-samples them from the teacher, drawing inspiration from speculative decoding (Leviathan et al., 2023). The design of SKD is supported by theoretical results in imitation learning (Ross & Bagnell, 2010; Ross et al., 2011), which suggest that gradually shifting the sample generation from the teacher to the student is an effective approach. Early in training, when the student model is likely to propose many out-of-distribution tokens, SKD functions similarly to supervised KD, replacing low-quality tokens with those from the teacher. As training progresses and the student’s sample quality improves, SKD behaves more like on-policy KD, accepting many of the student’s proposed tokens.

SKD significantly outperforms supervised and on-policy distillation for LLMs, as shown in Figure 1. Specifically, when distilling Gemma-7B to Gemma-2B, SKD achieves substantial gains compared to supervised fine-tuning in low-resource machine translation (41.8%), summarization (230%), and arithmetic reasoning (160%). Furthermore, training SKD on instruction-following data leads to impressive results on held-out MATH (Hendrycks et al., 2021) (198%) and GSM_{plus} (Li et al., 2024) (360%) testing suites. To provide a comprehensive understanding of SKD’s capabilities, we conduct a systematic evaluation under varying conditions and offer practical guidelines for tokens acceptance criteria from the teacher model, considering factors such as task type and model initialization. Our key contributions are:

- **A novel KD approach** that utilizes on-policy student samples *likely* to be generated by the teacher, mitigating the issues of low-quality samples in on-policy KD and dynamically switching between supervised and on-policy KD.

- **Superior performance** across various text generation tasks, including summarization, translation, and reasoning tasks, in both task-specific and task-agnostic scenarios.
- **Robustness and adaptability** across different model initialization, data sizes, and model families.

2 PRELIMINARIES

Language Models. An auto-regressive language model, denoted as M , predicts the probability distribution of the next token in a sequence. Conditioned on an input sequence X and the generated prefix $y_{<i} = \{y_1, y_2, \dots, y_{i-1}\}$ consisting of $(i - 1)$ tokens from a vocabulary V , M calculates the probability of the i -th token, $p_{y_i} = M(y_i|X, y_{<i})$. The auto-regressive language model M computes a logit l at the i -th token position. To obtain the probability p_{y_i} , we apply the softmax function with temperature t : $p_{y_i} = \frac{\exp(l_i/t)}{\sum_{j=1}^{|V|} \exp(l_j/t)}$. The temperature t introduces randomness into the text generation process. A higher temperature leads to more diverse sample output, while a lower temperature results in more focused and deterministic generation. In practical implementations, the vocabulary V is often truncated using top-k sampling (Radford et al., 2019), which improves text generation compared to searching over the entire vocabulary distribution. We draw inspiration from top-k sampling to develop our acceptance criteria for student-generated tokens. We include background of divergence metrics in the Appendix A.

Speculative Decoding (Leviathan et al., 2022; Chen et al., 2023) accelerates the decoding process by employing a smaller draft model to generate token candidates, which are then verified by larger approximation models. Our interleaved sampling is inspired by speculative decoding but introduces key differences. While speculative decoding ensures that final sampled tokens adhere to the larger model’s distribution, our method directly assesses the feasibility of student-proposed tokens within the teacher’s top K tokens (See Sec 3.2).

3 SPECULATIVE KNOWLEDGE DISTILLATION FOR LANGUAGE MODELS

Problem Setup. Given a teacher LM M_t and a student LM M_s , with differing model capacities, our goal is to train M_s to mimic the behavior of M_t on a specific task T . We assume that M_s has learnable parameters θ_s , while M_t ’s parameters θ_t are fixed. For task T , we have a set of prompts $\{x\}$ and corresponding output sequences $\{y\}$, which can be generated by either model or provided as ground truth. To measure the divergence between the token-level distributions of M_t and M_s for a given input-output pair (x, y) , we define the following metric:

$$D(M_t||M_s)(y|x) = \frac{1}{L_y} \sum_{i=1}^{L_y} D(M_t(\cdot|y_{<i}, x)||M_s(\cdot|y_{<i}, x)) \quad (1)$$

where L_y is the length of the output sequence y , and i indicates decoding steps. Our training objective is to minimize $D(M_t||M_s)(y|x)$ such that M_s can effectively imitate M_t on task T .

3.1 BASELINE KNOWLEDGE DISTILLATION APPROACHES

Supervised FT is a training method where a student policy is trained on a fixed dataset of input-output pairs (x, y) . The objective is to minimize the negative log-likelihood of the student’s predictions: $L_{SFT} = -\log p_s(y|x)$. It’s also a common initialization strategy for more advanced knowledge distillation techniques like Supervised KD, On-Policy KD, and SKD.

SeqKD (Kim & Rush, 2016a) is a supervised fine-tuning method that trains on sequences generated by a teacher model using maximum likelihood.

Supervised KD (Sanh et al., 2020) trains a student model to mimic the token-level probability distribution of a teacher model over a fixed ground truth dataset. This is achieved by minimizing the distance between the student’s and teacher’s predictions, as defined in Equation 1. The ground truth labels, denoted by y , are obtained from the fixed dataset.

On-Policy KD (Agarwal et al., 2024) addresses the training-inference mismatch in student models by sampling target tokens directly from the student’s own output. This approach calculates both the teacher and student’s token probabilities and employs Equation 1 to train the student model on its self-generated data.

Algorithm 1 Speculative knowledge distillation

Require: Student LLM M_s , Teacher LLM M_t , Prompt dataset $\{x_j\}_{j=1}^N$, Decoding length α .

- 1: **Hyperparameters:** Top K for token acceptance, Divergence metric D ,
- 2: **for** $step := 1$ to N **do** ▷ We assume batch size := 1 to simplify the illustration
- 3: **for** $i := 1$ to α **do**
- 4: $y_i \sim M_s(\cdot|y_{<i}, x_j)$ ▷ Sample y_i from student M_s
- 5: **if** $y_i \notin \text{top}_K(M_t(\cdot|y_{<i}, x_j))$ **then** ▷ If y_i is not within top K token of teacher M_t
- 6: $y_i \sim M_t(\cdot|y_{<i}, x_j)$ ▷ We replace student token to teacher’s resampled token
- 7: **end if**
- 8: **end for**
- 9: Apply gradient descent to minimize the minibatch loss $D(M_t||M_s)(y|x)$ in equation 1
- 10: **end for**

ImitKD (Lin et al., 2020) is a hybrid approach that can be viewed as naively combining supervised and on-policy KD. During training, it randomly selects samples from either the ground truth dataset or the student model’s own output with equal probability. This mixed dataset is then used to train the student model, optimizing the distance metric defined in Equation 1.

3.2 SPECULATIVE KNOWLEDGE DISTILLATION

Speculative knowledge distillation (SKD) draws two important inspirations from imitation learning (Ross et al., 2011; Ross & Bagnell, 2010): 1) Directly taking samples from a fixed dataset or sampling from teacher can cause the state distribution to be unrealistically good. As a result, student model never learns to correct previous mistakes, leading to poor performance at test time. 2) Directly drawing samples from student can result in low quality samples that can be out-of-distribution for the teacher, resulting in inaccurate feedback when assigning logits to the student. SKD addresses above limitations by using major components from speculative sampling (Leviathan et al., 2023): 1) student model proposes on-policy sample candidate, and 2) the teacher model evaluates these proposed samples concurrently and accepts only those that it deems likely to generate. After speculative sampling, SKD calculates teacher and student’s token probabilities and employ equation 1 to train the student model (L9 in Algorithm 1).

Interleaved On-Policy Sampling Given a prompt x , student model M_s will generate a sequence of tokens, $(y_i, \dots, y_{i+\gamma})$, following its own distribution $M_s(\cdot|y_{<i}, x)$ (L4 in Algorithm 1). The teacher model M_t will evaluate this sequence, accepting or rejecting each token based on pre-defined criteria, which resembles speculative decoding (Leviathan et al., 2023). If a token y_i is rejected, the subsequent tokens (y_{i+1}, \dots) will be discarded, and y_i will be resampled from M_t ’s distribution $M_t(\cdot|y_{<i}, x)$ (L5-7 in Algorithm 1). The student model will then continue to generate tokens on-policy based on $M_s(\cdot|y_{<i}, x, y_i)$ for the remaining γ tokens. This sample generation process is particularly advantageous for auto-regressive models, where early errors can propagate and affect future predictions (Ross & Bagnell, 2010). By focusing on on-policy samples that are likely to be generated by the teacher model, we can effectively utilize limited student capacity to mimic teacher.

Token Acceptance Criteria Using the default acceptance criteria in speculative decoding corresponds to sampling from teacher’s distribution, which would degenerate to supervised KD. Instead, we are interested in generating student samples that are likely under the teacher. To do so, we draw inspiration from top-k sampling (Radford et al., 2019) that draws samples using only the k highest probability tokens from the entire vocabulary V at each generation step. Following this intuition, we define our acceptance criteria based whether the student token y_i falls into the top K tokens of teacher’s distribution $M_t(\cdot|y_{<i}, x)$. Similar to speculative decoding, the teacher M_t can evaluate sequence $(y_i, \dots, y_{i+\gamma})$ proposed by student in parallel and return to student with the token position that discarded. We replace the discarded student token by re-sampling a token from the teacher’s distribution. Note that in Algorithm 1, we simplify the SKD presentation by setting $\gamma = 1$, but in practice use $\gamma = 5$ for efficient implementation.

Sample Transition During SKD Training. In the early stages of training, student models often propose low-quality samples that are frequently rejected by the teacher. This initial behavior resembles supervised KD, where the student is explicitly corrected and guided towards the teacher’s generated samples. As the student’s training progresses, the quality of its proposed samples improves, leading to a more on-policy interaction with the teacher. SKD dynamically adjusts the balance be-

tween supervised and on-policy KD based on the distribution gap between the teacher and student models. This adaptive approach allows the student to efficiently learn from the teacher’s samples. Moreover, SKD is a flexible framework that can degenerate both supervised and on-policy KD as special cases. In [Figure 2](#), when the student’s samples are consistently rejected or accepted by the teacher, SKD effectively degenerates to supervised or on-policy KD, respectively. This demonstrates the generality and versatility of SKD as a unified approach to knowledge distillation.

4 EXPERIMENTAL SETUP

Student and Teacher Models. We conducted experiments using two model families, GEMMA1 ([Team et al., 2024](#)) and QWEN2 ([Yang et al., 2024a](#)). We use supervised fine-tuned (SFT) GEMMA-7B-IT (IT indicates instruction-tuned) and QWEN-7B-IT as the teacher models, while GEMMA-2B-IT and QWEN-0.5B-IT were employed as student models. We further leveraged SFTed student models to compare KDs under different model initializations.

Hyperparameters. For all fine-tuning process, we use learning rate $1e-5$, warmup ratio 0.1 and 0.1 drop out rate for all fine-tuning processes. All SFT checkpoint is trained for three epoches and we select the checkpoint with the lowest validation loss. Details about SFT training dataset can be found in [Section 4.1](#). Without parameter search, we fixed acceptance criteria K to be 25 throughout main experiment and ablation ($K=25$ is a common choice for top-k sampling ([Radford et al., 2019](#))). See [Appendix B](#) for detailed study). We include training details and our hyper-parameter choice of baselines and SKD in the [Appendix Section C](#).

Baselines. We conduct SKD over widely accepted KD approaches listed in [Section 3.1](#): Supervised Fine-tuning, Supervised KD, on-policy KD and ImitKD. We compared SeqKD to supervised fine-tuning (SFT) using ground truth data and found that SFT consistently outperformed SeqKD. Therefore, we did not include SeqKD results.

4.1 DATASET AND EVALUATION.

We evaluated SKD and baseline methods across three generation tasks: low-resource translation, dialogue summarization, and arithmetic reasoning. Task-specific KD often operates with training data around 1K samples ([Quteineh et al., 2022](#)). Thus, we randomly draw 1K samples from our task data, and further test low-data regime with 100 samples. Additionally, we performed experiments using 1K and 10K on a task-agnostic instruction-following task in the math domain. It’s important to note that different baseline KD methods have access to different training data. Supervised KD, SFT, and ImitKD have access to both (x, y) pairs from the ground truth dataset, while on-policy and SKD only see the prompt x and self-generate the target label during training. Since our assumption that the teacher model should be an expert in a specific task, we used full training data for each task to supervise the FT teacher model. For SFTed student initialization, we supervise the FT student model with our constructed training set x, y (around 1000 instances). For reproducibility, all baselines and SKD use greedy decoding for evaluation. We include all results of supervised FT teacher model, supervised FT student model and instruction tuned student models in the [Appendix D](#).

Low Resource Translation. We utilize the Flores-200 ([NLLB Team, 2022](#)) Assamese-to-English translation dataset for low-resource translation. Due to the scarcity of Assamese-to-English training data, we employ the Flores-200 development set (997 instances) as our training set. Additionally, we split the Flores-200 testing set (1012 instances) into a development set (500 instances) and a testing set (512 instances) for evaluation. We use SOTA learned metric COMET ([Rei et al., 2022](#)) to evaluate the translation quality. We use 997 instances from our constructed training set to SFTed both teacher and student models.

Dialogue Summarization. For dialogue summarization, we utilize the DialogSum dataset ([Chen et al., 2021](#)). We randomly sample 1K instances from the DialogSum training set to create our input-output pairs (x, y) . For evaluation, we employ the dataset’s development set (500 instances) and test set (1500 instances). To assess summary quality, we use the ROUGE-L metric ([Lin, 2004](#)). We Supervised FT our teacher models on 10K instances from the DialogSum training set, and our student models on the 1K instances from our constructed training set.

Arithmetic Reasoning. For arithmetic reasoning, we utilize the GSM8K dataset ([Cobbe et al., 2021](#)). We split the GSM8K training set into a development set (473 instances) and a training set (7K

instances). We randomly sample 1K instances from the training set to create our input-output pairs (x, y) . For evaluation, we employ the development set and the GSM8K test set (1319 instances). We assess the quality of the generated output using answer accuracy. We Supervised FT our teacher models on 7K instances from the GSM8K training set, and our student models on the 1K instances from our constructed training set.

Math Instruction Following. For math instruction following, we utilize the UltraInteract dataset (Yuan et al., 2024). We randomly sample 11K instances from the UltraInteract training set, dividing them into a training set (10K instances) and a development set (1K instances). For evaluation, we employ the GSM_{plus} (Li et al., 2024), Math (Hendrycks et al., 2021), Asdiv (Miao et al., 2020) and SVAMP (Patel et al., 2021) test sets as held-out sets. We Supervised FT our teacher models on 10K instances from the UltraInteract training set, and our student models on the 1K instances from our constructed training set. We also tested the setting with 10K training instances for student model.

We explored two types of knowledge distillation: task-specific (Quteineh et al., 2022) and task-agnostic (Agarwal et al., 2024). In task-specific distillation, training and inference data share the same format, and the task is known during training. In task-agnostic distillation, only the domain is known, while the specific task varies during deployment. We categorized distillation of translation, dialogue summarization, and arithmetic reasoning as task-specific and distillation of math instruction following as task-agnostic. For both math instruction following and arithmetic reasoning, we use chain-of-thought (Wei et al., 2023) prompts to generate answer. All prompt formats can be found in Appendix E. Details about parsing and evaluation for all tasks can be found in the Appendix F.

5 RESULTS

We validate SKD’s effectiveness through comprehensive studies, including task-specific (Section 5.1) and task-agnostic distillation (Section 5.2), different model initializations (Section 5.3), low-data regimes (Section 5.4), and an ablation study on two-staged training (supervised KD followed by on-policy KD) (Section 5.5).

Table 1: SKD outperforms baseline KDs at in-domain tasks: low resource Assamese-to-English translation, dialogue summarization, arithmetic reasoning. We use GEMMA-2B and QWEN-0.5B instruction tuned model (w/o SFT) as students and SFTed GEMMA-7B-IT and QWEN2-7B-IT as teachers. $K=25$ for SKD.

Baselines	Supervised FT GEMMA-7B to GEMMA-2B			Supervised FT QWEN2-7B to QWEN2-0.5B		
	Translation	Summarization	GSM8K	Translation	Summarization	GSM8K
COMET	72.5	31.7	18.7	57.6	29.2	31.8
ROUGE-L	73.3	34.1	22.5	57.5	31.3	35.0
Acc	36.1	34.1	25.3	53.0	31.2	33.9
ImitKD	74.8	34.9	26.2	55.9	30.9	33.2
SKD	75.3	35.0	29.1	57.1	31.7	36.6

5.1 TASK-SPECIFIC DISTILLATION

In this section, we investigated task-specific distillation, where the training and inference prompt-response sets share the same format, and the task is known during training. We study two model families, supervised FT GEMMA-7B to instruction tuned GEMMA-2B (GEMMA-2B-IT) and supervised FT QWEN2-7B to instruction tuned QWEN2-0.5B (QWEN2-0.5B-IT). As shown in Table 1, the performance of supervised KD and on-policy KD is highly task-dependent, likely due to the model’s initial familiarity with each task (see Sec 5.3 for details). SKD, trained on GEMMA-2B-IT, consistently outperformed all baseline KD approaches across three in-domain text generation tasks.

When initialized with QWEN2-0.5B-IT, SKD outperformed all baselines on summarization and GSM8k. We found that on-policy KD consistently underperformed supervised KD and SKD for translation and summarization. This is primarily attributed to the low-quality samples generated by the student model during training. Although ImitKD (partially on-policy) can achieve competitive performance with a stronger student like GEMMA-2B due to its mixed data strategy, it falls behind SKD when using the QWEN-0.5B student due to the lower quality of on-policy samples (See exam-

ple in Sec 5.3). Our results demonstrate SKD’s superiority in five out of six test settings across two model families, highlighting its optimized performance for task-specific distillation.

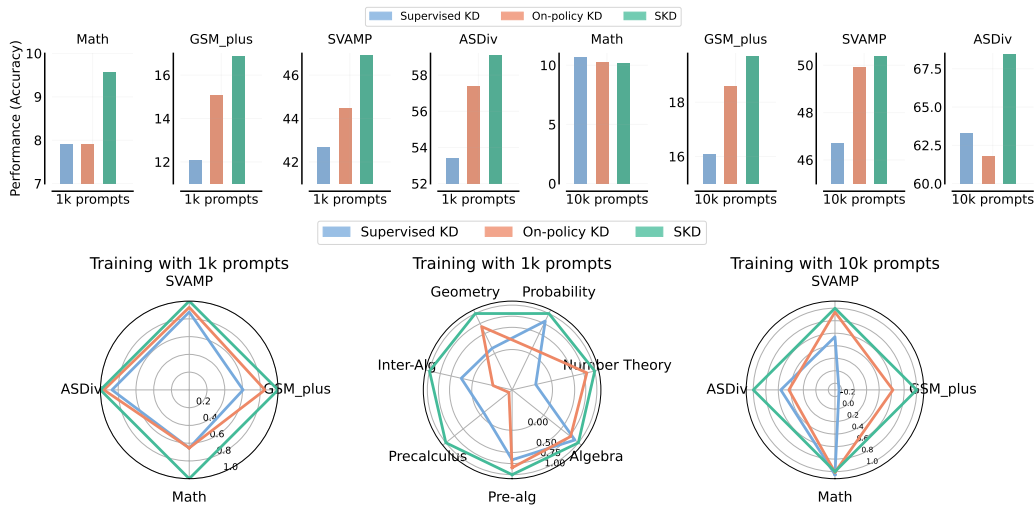


Figure 3: SKD outperforms baseline KDs at four held-out testing sets (Math, GSM_{plus} , SVAMP and ASDiv). We consider two data size setting, 1k and 10k respectively and report testing accuracy in the first row. In the second row, we calculate the performance gain of SKD over SFT and report the relative performance gains of baselines compared to SKD. As SKD outperforms all baselines in most cases, these values typically range from 0 to 1. In some instances, supervised KD may be outperformed by SFT, resulting in negative values. In the middle figure at the second row, we show that SKD outperforms all baselines under 7 math concepts at math testing set. SKD is performed under $K=25$.

5.2 TASK-AGNOSTIC DISTILLATION

To explore task-agnostic instruction following, where the exact nature of the task is unknown during training and can vary during deployment, we trained and evaluated SKD and baselines on a math domain instruction following task. Our training set included diverse math instruction following data, such as MathQA (Amini et al., 2019), math reasoning (Mishra et al., 2022), and tabular processing (Lu et al., 2023). We tested each KD approach’s performance on four held-out testing sets, ranging from simple math word problems (ASDiv and SVAMP) to grade-level math problems (GSM_{plus}) and competitive math problems (Math).

As illustrated in Figure 3, SKD initialized with GEMMA-2B-IT consistently outperformed supervised KD and on-policy KD across these four held-out sets when trained on 1k prompts. We also calculated the performance gain of SKD over SFT and reported baseline performance gains relative to SKD. Results showed that under 1k prompts, supervised KD and on-policy could approach close performance on simple tasks like SVAMP and ASDiv but still had a 20-40% performance gap compared to SKD (Second row on the left of Figure 3). The middle figure demonstrates SKD’s superiority over baselines on the Math task under 7 different concepts. Since the training set was task-agnostic instruction following data, it could be trained on a larger scale in practical scenarios. On the right side of Figure 3, we show that SKD can outperform or match all baseline methods with 10k prompts. Our results demonstrate SKD’s generalization capability under task-agnostic distillation. Detailed table result can also be found at Appendix D.

5.3 DIFFERENT MODEL INITIALIZATION

In this section, we explore how different student model initialization impact the effectiveness of various KD approaches. As shown in Figure 4, on-policy KD initialized with an instruction-tuned student model can surprisingly degrade performance over time. This raises the question of whether a teacher’s token-wise probability distribution always provides beneficial guidance. Our findings suggest that if the initial student sample quality is insufficient, the student model may become trapped in a sub-optimal state, unable to effectively learn from the teacher’s supervision. On-policy KD’s low quality sample can be found in Sec 5.3 and Appendix G.

Table 3: SKD outperforms other KD approaches with IT and SFTed student initialization across three text generation tasks at 100 data points. We use supervised SFT GEMMA-7B-IT as teacher and instruction tuned GEMMA-2B and supervised FT GEMMA-2B-IT as students. SKD is performed under $K=25$.

Baselines	KD initiates with IT Gemma-2B			KD initiates with SFTed Gemma-2B		
	Translation	Summarization	GSM8K	Translation	Summarization	GSM8K
	COMET	ROUGE-L	Acc	COMET	ROUGE-L	Acc
SFT	59.8	23.0	8.11	59.8	23.0	8.11
Supervised KD	65.8	31.7	17.0	60.9	32.0	14.4
On-policy KD	33.7	32.2	12.1	63.4	32.4	16.1
SKD	67.0	32.3	18.8	63.8	32.6	16.7

This is attributed to unavoidable overfitting during SFT with only 100 data points. In Appendix Section H, we illustrate how overfitting at the SFT stage can lead to suboptimal performance in post-training KDs. This finding suggests that end-to-end SKD, capable of training from any base model and bypassing the SFT stage, offers a more optimal solution than two-stage on-policy KD (SFT+KD) in extreme low-data scenarios.




Figure 5: Ablation study. We conduct an ablation study by training the student model with supervised KD for the first half and self-generated samples for the second half. Our results demonstrate that this mixed training strategy can mostly be outperformed by either supervised KD or on-policy KD across three tasks and under two model initialization. Moreover, it is significantly outperformed by our SKD. We excluded on-policy KD from the leftmost figure due to its extremely low COMET score.

5.5 TWO-STAGED TRAINING

Given that SKD transitions from supervised KD-like behavior at the beginning of training to on-policy KD towards the end, we investigated whether a simple approach of training with supervised KD for the first half of iterations and self-generated samples for the second half could achieve SKD-like performance. We conducted an ablation study by training the student model with supervised KD first, followed by on-policy KD (see Appendix Section I for details). As shown in Figure 5, SKD significantly outperforms this mixed training strategy. Surprisingly, this heuristic baseline was often outperformed by either supervised KD or on-policy KD across three tasks and two model initializations. This is likely due to the data discrepancy encountered during the training process – naively mixing samples from the teacher and student in two different stages lacks natural confidence measurement of the tokens; therefore, simply combining supervised and on-policy KD can even degrade the performance of either approach. This underscores the importance of our adaptive SKD, which implicitly covers intermediate transitions between two model behaviors.

5.6 SUMMARY AND ADDITIONAL RESULTS

SKD consistently outperforms baselines in both task-specific and task-agnostic distillation. Unlike on-policy methods, which are sensitive to student-generated sample quality, SKD maintains its advantage even with low-quality student model initialization. As shown in Appendix H, overfitting in SFT can hinder post-training KD’s performance. SKD’s end-to-end approach, bypassing SFT, provides a significant benefit, especially in low-data scenarios. Merging supervised and on-policy KD naively can lead to performance degradation due to data inconsistencies.

Speculative decoding, as described in (Leviathan et al., 2023), is a crucial technique for accelerating LLM inference. The quality of the draft model used in this process directly influences the speed-up

486 achieved. As demonstrated in Appendix J, our SKD-trained draft model significantly outperforms
487 the instruction-tuned base model in terms of token acceptance rate (71% to 85% higher), leading to
488 a 1.2X speed-up in speculative decoding.
489

490 6 RELATED WORK

491 **Knowledge Distillation** (KD) (Hinton et al., 2015) is a widely adopted technique to transfer knowl-
492 edge from a large, complex teacher model to a smaller, more efficient student model. This approach
493 has been successfully applied to autoregressive models (Sanh et al., 2020) and language models in
494 particular (Yang et al., 2024b). There are two primary categories of KD methods: hard label and
495 soft label. Hard label methods train the student model to mimic the teacher’s input-output pairs,
496 similar to supervised fine-tuning (Kim & Rush, 2016a; Taori et al., 2023; Chiang et al., 2023; Xu
497 et al., 2023). Various data augmentation techniques, such as fine-tuning on rationales (Hsieh et al.,
498 2023) or instruction-following formats (Wang et al., 2023), have been explored within this frame-
499 work. Soft label methods, on the other hand, compute logits for both the student and teacher models
500 and minimize a distance metric, such as KL-divergence, between their output distributions (Hinton
501 et al., 2015; Sanh et al., 2020). Samples for computing logits can be generated by either model or
502 both. Our proposed SKD falls under the category of soft label KD. To the best of our knowledge,
503 this work is the first to introduce interleaved teacher-student sampling for soft-label KD.

504 **Soft-label Knowledge Distillation** can be broadly categorized into supervised KD and on-policy
505 KD based on the source of training data. In supervised KD (Sanh et al., 2020), the training data is
506 fixed, leading to a discrepancy between training and inference time for the student model (Agarwal
507 et al., 2024). MiniLLM (Gu et al., 2024), for instance, frames this as a reinforcement learning
508 problem, optimizing reverse KL divergence over the fixed dataset. Lin et al. (2020) explored
509 different distance metrics, including JSD and total variation distance. On the other hand, on-policy
510 KD, pioneered by ImitKD (Lin et al., 2020), leverages samples generated by the student model
511 itself. This approach addresses the training-inference mismatch by directly sampling target tokens
512 from the student’s own output. Inspired by on-policy imitating learning (Ross et al., 2011), Agarwal
513 et al. (2024) introduced on-policy KD for language models, calculating both teacher and student
514 token probabilities to train the student on its self-generated data. DistillM (Ko et al., 2024) further
515 enhanced this by incorporating a mix of on-policy samples and ground truth. As discussed in §3.2,
516 both supervised KD and on-policy KD are special cases of SKD. SKD outperforms both baselines
517 and surpasses the performance from mixing ground-truth and on-policy samples.

518 **Speculative Decoding** (Leviathan et al., 2022; Chen et al., 2023) inspired our interleaved sampling
519 approach in SKD. However, instead of sampling from teacher’s distribution, SKD assesses the feasi-
520 bility of student-proposed tokens within the teacher’s top K tokens. Interleaved sampling allows us
521 to obtain high quality on-policy samples. Zhou et al. (2024); Liu et al. (2024) enhanced speculative
522 decoding alignment with various KD methods. Building on these, we benchmarked SKD’s speedup
523 against KD baselines, finding it superior in both speed and token acceptance.

524 7 CONCLUSION

526 We propose Speculative Knowledge Distillation (SKD), a novel method that addresses the limita-
527 tions of existing KD approaches. SKD leverages interleaved sampling to mitigate the discrepancy
528 between training and inference samples and remove low-quality student-generated data. Our ex-
529 periments consistently show SKD’s superiority in both task-specific and task-agnostic distillations.
530 SKD is robust to various model families, initialization, and dataset sizes. We also demonstrate its
531 superiority in token-level mixing over sequence-level mixing. Furthermore, SKD-trained student
532 models can accelerate speculative decoding, enabling new applications.
533
534
535
536
537
538
539

8 REPRODUCIBILITY STATEMENT

We discussed our model implementation, hyper-parameters, baseline details in Section 4. We include information about our dataset and evaluation metrics in Section 4.1. We discuss training and hyper-parameter details in Appendix C. We include examples of our prompt and response for all tasks in Appendix E. We include additional information about our evaluation metrics in Appendix F. All data, model and evaluation metrics are open sourced and available for research community. Upon the camera ready of the paper, we will include a link to a publicly accessible code repository, including all implementations about baseline KDs and SKD, running instructions and hyper-parameter settings. All results reported in this paper can be fully reproduced by the community.

REFERENCES

- Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes, 2024. URL <https://arxiv.org/abs/2306.13649>.
- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. MathQA: Towards interpretable math word problem solving with operation-based formalisms. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2357–2367, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1245. URL <https://aclanthology.org/N19-1245>.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Keane, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Mousaleem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. Palm 2 technical report, 2023. URL <https://arxiv.org/abs/2305.10403>.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. Accelerating large language model decoding with speculative sampling, 2023. URL <https://arxiv.org/abs/2302.01318>.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. DialogSum: A real-life scenario dialogue summarization dataset. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 5062–5074, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.449. URL <https://aclanthology.org/2021.findings-acl.449>.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An

- 594 open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- 595
- 596
- 597 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
- 598 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
- 599 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,
- 600 2021.
- 601 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
- 602 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony
- 603 Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark,
- 604 Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere,
- 605 Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris
- 606 Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong,
- 607 Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny
- 608 Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino,
- 609 Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael
- 610 Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Ander-
- 611 son, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah
- 612 Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan
- 613 Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Ma-
- 614 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy
- 615 Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak,
- 616 Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Al-
- 617 wala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini,
- 618 Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der
- 619 Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo,
- 620 Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Man-
- 621 nat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova,
- 622 Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal,
- 623 Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur
- 624 Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhar-
- 625 gava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong,
- 626 Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic,
- 627 Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sum-
- 628 baly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa,
- 629 Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang,
- 630 Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende,
- 631 Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney
- 632 Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom,
- 633 Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta,
- 634 Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petro-
- 635 vic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang,
- 636 Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur,
- 637 Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre
- 638 Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha
- 639 Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay
- 640 Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda
- 641 Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew
- 642 Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita
- 643 Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh
- 644 Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De
- 645 Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Bran-
- 646 don Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina
- 647 Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai,
- Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li,
- Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana
- Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil,
- Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Ar-
- caute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco

- 648 Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella
649 Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory
650 Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang,
651 Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Gold-
652 man, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman,
653 James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer
654 Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe
655 Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie
656 Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun
657 Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal
658 Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva,
659 Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian
660 Khabsa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson,
661 Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Ke-
662 neally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel
663 Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mo-
664 hammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navy-
665 ata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong,
666 Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli,
667 Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux,
668 Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao,
669 Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li,
670 Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Sa-
671 tadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lind-
672 say, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang
673 Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen
674 Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho,
675 Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser,
676 Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Tim-
677 othy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan,
678 Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu
679 Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Con-
680 stable, Xiaocheng Tang, Xiaofang Wang, Xiaojuan Wu, Xiaolan Wang, Xide Xia, Xilun Wu,
681 Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi,
682 Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef
683 Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. The llama 3 herd of models, 2024.
684 URL <https://arxiv.org/abs/2407.21783>.
- 684 Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis,
685 Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. Results of WMT22 metrics
686 shared task: Stop using BLEU – neural metrics are better and more robust. In Philipp Koehn,
687 Loïc Barrault, Ondřej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian
688 Federmann, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grund-
689 kiewicz, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Tom Kocmi,
690 André Martins, Makoto Morishita, Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, Mat-
691 teo Negri, Aurélie Névél, Mariana Neves, Martin Popel, Marco Turchi, and Marcos Zampieri
692 (eds.), *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pp. 46–68, Abu
693 Dhabi, United Arab Emirates (Hybrid), December 2022. Association for Computational Linguis-
694 tics. URL <https://aclanthology.org/2022.wmt-1.2>.
- 695 Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. Minillm: Knowledge distillation of large lan-
696 guage models, 2024. URL <https://arxiv.org/abs/2306.08543>.
- 697 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
698 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021.
699 URL <https://arxiv.org/abs/2103.03874>.
- 700 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
701 URL <https://arxiv.org/abs/1503.02531>.

- 702 Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner,
703 Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. Distilling step-by-step! outperforming larger
704 language models with less training data and smaller model sizes. In Anna Rogers, Jordan Boyd-
705 Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics:
706 ACL 2023*, pp. 8003–8017, Toronto, Canada, July 2023. Association for Computational Linguistics.
707 doi: 10.18653/v1/2023.findings-acl.507. URL [https://aclanthology.org/2023.
708 findings-acl.507](https://aclanthology.org/2023.findings-acl.507).
- 709 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child,
710 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language
711 models, 2020. URL <https://arxiv.org/abs/2001.08361>.
- 712 Yoon Kim and Alexander M. Rush. Sequence-level knowledge distillation. In Jian Su, Kevin
713 Duh, and Xavier Carreras (eds.), *Proceedings of the 2016 Conference on Empirical Methods in
714 Natural Language Processing*, pp. 1317–1327, Austin, Texas, November 2016a. Association for
715 Computational Linguistics. doi: 10.18653/v1/D16-1139. URL [https://aclanthology.
716 org/D16-1139](https://aclanthology.org/D16-1139).
- 717 Yoon Kim and Alexander M Rush. Sequence-level knowledge distillation. *arXiv preprint
718 arXiv:1606.07947*, 2016b.
- 719 Jongwoo Ko, Sungnyun Kim, Tianyi Chen, and Se-Young Yun. Distillm: Towards streamlined dis-
720 tillation for large language models, 2024. URL <https://arxiv.org/abs/2402.03898>.
- 721 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
722 decoding, 2022. URL <https://arxiv.org/abs/2211.17192>.
- 723 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
724 decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023.
- 725 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
726 decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023.
- 727 Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. Gsm-plus: A comprehensive
728 benchmark for evaluating the robustness of llms as mathematical problem solvers, 2024. URL
729 <https://arxiv.org/abs/2402.19255>.
- 730 Alexander Lin, Jeremy Wohlwend, Howard Chen, and Tao Lei. Autoregressive knowledge dis-
731 tillation through imitation learning. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu
732 (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Process-
733 ing (EMNLP)*, pp. 6121–6133, Online, November 2020. Association for Computational Linguistics.
734 doi: 10.18653/v1/2020.emnlp-main.494. URL [https://aclanthology.org/2020.
735 emnlp-main.494](https://aclanthology.org/2020.emnlp-main.494).
- 736 Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization
737 Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.
738 URL <https://aclanthology.org/W04-1013>.
- 739 Xiaoxuan Liu, Lanxiang Hu, Peter Bailis, Alvin Cheung, Zhijie Deng, Ion Stoica, and Hao Zhang.
740 Online speculative decoding, 2024. URL <https://arxiv.org/abs/2310.07177>.
- 741 Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter
742 Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured
743 mathematical reasoning. In *International Conference on Learning Representations (ICLR)*, 2023.
- 744 Shen-yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing
745 English math word problem solvers. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel
746 Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational
747 Linguistics*, pp. 975–984, Online, July 2020. Association for Computational Linguistics. doi: 10.
748 18653/v1/2020.acl-main.92. URL <https://aclanthology.org/2020.acl-main.92>.
- 749 Swaroop Mishra, Arindam Mitra, Neeraj Varshney, Bhavdeep Sachdeva, Peter Clark, Chitta Baral,
750 and Ashwin Kalyan. NumGLUE: A suite of fundamental yet challenging mathematical reason-
751 ing tasks. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings
752 of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long
753 and Short Papers)*, pp. 1000–1010, Toronto, Canada, July 2022. Association for Computational
754 Linguistics. doi: 10.18653/v1/2022.acl-long.100. URL <https://aclanthology.org/2022.acl-long.100>.
- 755

- 756 *Papers*), pp. 3505–3523, Dublin, Ireland, May 2022. Association for Computational Linguistics.
 757 doi: 10.18653/v1/2022.acl-long.246. URL [https://aclanthology.org/2022.
 758 acl-long.246](https://aclanthology.org/2022.acl-long.246).
 759
- 760 James Cross Onur Çelebi Maha Elbayad Kenneth Heafield Kevin Heffernan Elahe Kalbassi Jan-
 761 ice Lam Daniel Licht Jean Maillard Anna Sun Skyler Wang Guillaume Wenzek Al Youngblood
 762 Bapi Akula Loic Barrault Gabriel Mejia Gonzalez Prangthip Hansanti John Hoffman Sema-
 763 Jarrett Kaushik Ram Sadagopan Dirk Rowe Shannon Spruit Chau Tran Pierre Andrews Necip
 764 Fazil Ayan Shruti Bhosale Sergey Edunov Angela Fan Cynthia Gao Vedanuj Goswami Fran-
 765 cisco Guzmán Philipp Koehn Alexandre Mourachko Christophe Ropers Safiyyah Saleem Holger
 766 Schwenk Jeff Wang NLLB Team, Marta R. Costa-jussà. No language left behind: Scaling human-
 767 centered machine translation. 2022.
- 768 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Floren-
 769 cia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red
 770 Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Moham-
 771 mad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher
 772 Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brock-
 773 man, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann,
 774 Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis,
 775 Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey
 776 Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux,
 777 Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila
 778 Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix,
 779 Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gib-
 780 son, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan
 781 Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hal-
 782 lacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan
 783 Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu,
 784 Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun
 785 Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Ka-
 786 mali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook
 787 Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel
 788 Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen
 789 Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel
 790 Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez,
 791 Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv
 792 Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney,
 793 Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick,
 794 Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel
 795 Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Ra-
 796 jeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe,
 797 Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel
 798 Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe
 799 de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny,
 800 Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl,
 801 Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra
 802 Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders,
 803 Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Sel-
 804 sam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor,
 805 Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky,
 806 Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang,
 807 Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Pre-
 808 ston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vi-
 809 jayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan
 Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng,
 Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Work-
 man, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming
 Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao

- 810 Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL
811 <https://arxiv.org/abs/2303.08774>.
812
- 813 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic
814 evaluation of machine translation. In Pierre Isabelle, Eugene Charniak, and Dekang Lin (eds.),
815 *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp.
816 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL <https://aclanthology.org/P02-1040>.
817
- 818 Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP models really able to solve simple
819 math word problems? In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek
820 Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou
821 (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association
822 for Computational Linguistics: Human Language Technologies*, pp. 2080–2094, Online, June
823 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.168. URL
824 <https://aclanthology.org/2021.naacl-main.168>.
- 825 Husam Quteineh, Spyridon Samothrakis, and Richard Sutcliffe. Enhancing task-specific distillation
826 in small data regimes through language generation. In Nicoletta Calzolari, Chu-Ren Huang,
827 Hansaem Kim, James Pustejovsky, Leo Wanner, Key-Sun Choi, Pum-Mo Ryu, Hsin-Hsi Chen,
828 Lucia Donatelli, Heng Ji, Sadao Kurohashi, Patrizia Paggio, Nianwen Xue, Seokhwan Kim,
829 Younggyun Hahm, Zhong He, Tony Kyungil Lee, Enrico Santus, Francis Bond, and Seung-Hoon
830 Na (eds.), *Proceedings of the 29th International Conference on Computational Linguistics*, pp.
831 5955–5965, Gyeongju, Republic of Korea, October 2022. International Committee on Computational
832 Linguistics. URL <https://aclanthology.org/2022.coling-1.520>.
- 833 Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever.
834 Language models are unsupervised multitask learners. 2019. URL [https://api.
835 semanticscholar.org/CorpusID:160025533](https://api.semanticscholar.org/CorpusID:160025533).
- 836 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
837 Finn. Direct preference optimization: Your language model is secretly a reward model. In *Thirty-
838 seventh Conference on Neural Information Processing Systems*, 2023. URL [https://arxiv.
839 org/abs/2305.18290](https://arxiv.org/abs/2305.18290).
- 840 Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. COMET: A neural framework for MT
841 evaluation. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the
842 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2685–
843 2702, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/
844 2020.emnlp-main.213. URL <https://aclanthology.org/2020.emnlp-main.213>.
- 845 Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya
846 Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. COMET-22: Unbabel-IST
847 2022 submission for the metrics shared task. In Philipp Koehn, Loïc Barrault, Ondřej Bojar,
848 Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel,
849 Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Paco Guzman, Barry
850 Haddow, Matthias Huck, Antonio Jimeno Yepes, Tom Kocmi, André Martins, Makoto Morishita,
851 Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, Matteo Negri, Aurélie Névéol,
852 Mariana Neves, Martin Popel, Marco Turchi, and Marcos Zampieri (eds.), *Proceedings of the
853 Seventh Conference on Machine Translation (WMT)*, pp. 578–585, Abu Dhabi, United Arab
854 Emirates (Hybrid), December 2022. Association for Computational Linguistics. URL [https://
855 aclanthology.org/2022.wmt-1.52](https://aclanthology.org/2022.wmt-1.52).
- 856 Stephane Ross and Drew Bagnell. Efficient reductions for imitation learning. In Yee Whye
857 Teh and Mike Titterton (eds.), *Proceedings of the Thirteenth International Conference on
858 Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*,
859 pp. 661–668, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL [https://
860 proceedings.mlr.press/v9/ross10a.html](https://proceedings.mlr.press/v9/ross10a.html).
- 861 Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured
862 prediction to no-regret online learning. In *Proceedings of the fourteenth international conference
863 on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference
Proceedings, 2011.

- 864 Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version
865 of bert: smaller, faster, cheaper and lighter, 2020. URL [https://arxiv.org/abs/1910.
866 01108](https://arxiv.org/abs/1910.01108).
867
- 868 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy
869 Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model.
870 https://github.com/tatsu-lab/stanford_alpaca, 2023.
- 871 Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya
872 Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard
873 Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex
874 Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, An-
875 tonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo,
876 Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric
877 Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Hen-
878 ryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski,
879 Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu,
880 Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee,
881 Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev,
882 Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko
883 Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo
884 Pandya, Siamak Shakeri, Soham De, Ted Klimentko, Tom Hennigan, Vlad Feinberg, Wojciech
885 Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh
886 Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin
887 Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah
888 Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. Gemma: Open models based on
889 gemini research and technology, 2024. URL <https://arxiv.org/abs/2403.08295>.
- 890 Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and
891 Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions,
892 2023. URL <https://arxiv.org/abs/2212.10560>.
- 893 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc
894 Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models,
895 2023. URL <https://arxiv.org/abs/2201.11903>.
- 896 Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and
897 Lei Li. INSTRUCTSCORE: Towards explainable text generation evaluation with automatic feed-
898 back. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference
899 on Empirical Methods in Natural Language Processing*, pp. 5967–5994, Singapore, December
900 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.365. URL
901 <https://aclanthology.org/2023.emnlp-main.365>.
902
- 903 An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,
904 Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang,
905 Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jin-
906 gren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin
907 Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao,
908 Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wen-
909 bin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng
910 Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu,
911 Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024a. URL
912 <https://arxiv.org/abs/2407.10671>.
- 913 Chuanpeng Yang, Wang Lu, Yao Zhu, Yidong Wang, Qian Chen, Chenlong Gao, Bingjie Yan, and
914 Yiqiang Chen. Survey on knowledge distillation for large language models: Methods, evaluation,
915 and application, 2024b. URL <https://www.arxiv.org/abs/2407.01885>.
- 916 Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Jia Deng, Boji Shan, Huimin
917 Chen, Ruobing Xie, Yankai Lin, Zhenghao Liu, Bowen Zhou, Hao Peng, Zhiyuan Liu, and
Maosong Sun. Advancing llm reasoning generalists with preference trees, 2024.

918 Yongchao Zhou, Kaifeng Lyu, Ankit Singh Rawat, Aditya Krishna Menon, Afshin Rostamizadeh,
919 Sanjiv Kumar, Jean-François Kagy, and Rishabh Agarwal. Distillspec: Improving speculative
920 decoding via knowledge distillation, 2024. URL <https://arxiv.org/abs/2310.08461>.
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

A DIVERGENCE METRICS.

As previously discussed, for each token position, the auto-regressive language model (LM) M can generate a probability distribution, $p(\cdot|X, y_{<k})$. Assuming the teacher and student models share the same vocabulary V , we can directly employ divergence metrics, like KL divergence, to quantify the distance between their respective probability distributions. For token position k , let $P(C)$ and $Q(C)$ represent the vocabulary distributions of the teacher and student, respectively. The KL divergence at this position is calculated as $D_{KL}(P||Q) = \sum_{c \in C} P(c)(\log \frac{P(c)}{Q(c)})$. By using this divergence metric as our objective function, we can effectively conduct knowledge distillation between the teacher and student models.

B CAN SKD OUTPERFORM BASELINES UNDER DIFFERENT ACCEPTANCE CRITERIA?

In this section, we tested a wide range of K values and showed SKD’s performance across three text generation tasks. In Figure 6, we show that while extremely low or high K values can not yield the most optimal performance, a broad range of K values between 5 to 50 can consistently outperform supervised KD and on-policy KD across all three text generation tasks. The overall performance trend confirms our hypothesis that as K increases, SKD degenerates to on-policy KD like behavior, while decreasing K degenerates SKD to supervised KD like behavior. Consequently, neither the highest nor lowest values of K are optimal.

We maintained a fixed K of 25 (a common choice for top-K sampling) in our primary experiments to assess SKD’s general effectiveness without extensive parameter tuning. Our results in this section suggest that further task-specific optimization of K could potentially yield even better performance.




Figure 6: We test a wide range of K values and show SKD’s performance across three text generation tasks. We initialize student model with GEMMA-2B-IT. We use SKD with $K=25$ (a common choice for top-k sampling) for our experiments in prior sections. In this Figure, our results demonstrate that while extremely low or high K values can not yield the most optimal performance, a broad range of K values between 5 and 50 consistently outperforms supervised KD and on-policy KD across all three text generation tasks.

C TRAINING AND HYPERPARAMETER DETAILS FOR BASELINES AND SKD

For both GEMMA-2B-IT and QWEN2-0.5B-IT models, we set the learning rate to $1e-5$. We disabled dropout during sampling, as in direct preference learning (Rafailov et al., 2023). Maximum input and output lengths were set task-specifically: (256, 256, 1024, 1024) and (256, 512, 128, 1024) for translation, arithmetic reasoning, summarization, and math instruction, respectively. All baseline models used a batch size of 8 and a gradient accumulation step of 1. Student model is trained over 375, 375, 375 and 1225 training steps for summarization, GSM, math and translation tasks respectively. We train 7500 steps for full data training at Math. For extreme low data setting (with prompt size 100), we trained student model over 125 steps for summarization, GSM and translation tasks. We report checkpoint results with the lowest validation loss.

Initial experiments with top-p sampling (temperature=1, top-p=1) led to significantly degraded generation quality compared to greedy decoding (temperature=0). This was particularly evident in math instruction task, where accuracy halved. To balance exploration and performance, we gradually reduced temperature and top-p until top-p sampling matched or outperformed greedy decoding. We

1026 ultimately settled on temperature=0.5 and top-p=0.5. This setting effectively limits the answer space
1027 while allowing for sufficient exploration. To ensure fair comparison, we kept the student temperature
1028 and top-p parameters identical for on-policy KD and SKD.

1029 Two additional hyperparameters, teacher model top-p and temperature, can influence SKD’s per-
1030 formance. When a student token is rejected by the teacher, these parameters affect the resampled
1031 teacher token. Similar to prior work (Agarwal et al., 2024), we found that lower teacher temperatures
1032 improve distillation. We conducted a grid search over top-p {0.5, 1} and temperature {0.2, 0.3, 0.4}
1033 for translation, arithmetic reasoning, and summarization on respective validation sets. While the
1034 optimal temperature varied slightly across tasks, 0.2 generally performed well. Although a full vo-
1035 cabulary distribution (top-p=1) typically yielded the best results, reducing top-p to 0.5 improved
1036 ROUGE-L from 0.5 to 0.9 across different temperatures at Summarization task. Therefore, we
1037 set temperature to 0.2 for all tasks and top-p to 0.5 for summarization and 1 for other tasks. All
1038 hyperparameter settings are consistent between GEMMA-2B and QWEN2-0.5B-IT model.

1039 In Section B, we explore a wide range of K values and test SKD’s performance across three text
1040 generation tasks. Our results demonstrate that a wide range of K values can outperform baseline
1041 knowledge distillation methods, such as On-policy KD and Supervised KD. To ensure consistent
1042 evaluation, we selected a constant K value of 25 (a common choice for Top-K sampling) for all tasks
1043 and ablation studies. However, future research could explore more dynamic K selection strategies
1044 to potentially further enhance SKD performance.

1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

Table 4: In this Table, we show performance of GEMMA-2B-IT (student), supervised FT GEMMA-2B-IT (student) and supervised FT GEMMA-7B-IT (teacher) at both task specific distillation and task agnostic distillation.

Baselines	Task specific distillation			Task agnostic distillation			
	MT	Summ	GSM8K	Math	GSM _{plus}	SVAMP	ASDiv
	COMET	ROUGE-L	Acc	Acc	Acc	Acc	Acc
GEMMA-2B-IT	44.5	11.7	0	0.68	0.5	1.6	1.8
Supervised FT GEMMA-2B-IT	72.5	31.7	18.7	4.82	4.70	12.8	15.2
Supervised FT GEMMA-7B-IT	79.2	36.8	60.0	19.5	40.5	66.6	76.1

Table 5: In this Table, we show performance of QWEN-0.5B-IT (student), supervised FT QWEN-0.5B-IT (student) and supervised FT QWEN-7B-IT (teacher) at both task specific distillation and task agnostic distillation.

Baselines	Task specific distillation			Task agnostic distillation			
	MT	Summ	GSM8K	Math	GSM _{plus}	SVAMP	ASDiv
	COMET	ROUGE-L	Acc	Acc	Acc	Acc	Acc
Qwen-0.5B-it	44.9	15.2	0	7.0	3.9	1.7	3.9
Supervised FT QWEN-0.5B-IT	57.6	29.2	31.8	2.6	3.9	10.5	14.1
Supervised FT QWEN-7B-IT	79.4	38.1	73.8	46.0	60.4	84.8	90.0

Table 6: This table presents baseline and SKD performance across four held-out test sets. The teacher model is a supervised FT GEMMA-7B, while the student model is an instruction-tuned GEMMA-2B. The left side of the table shows results for all baselines and SKD using 1,000 prompts from our constructed training set. The right side displays results using 10,000 prompts.

Baselines	KD initiates with GEMMA-2B-IT (1k prompts)				KD initiates with GEMMA-2B-IT (10k prompts)			
	Math	GSM _{plus}	SVAMP	ASDiv	Math	GSM _{plus}	SVAMP	ASDiv
	Acc	Acc	Acc	Acc	Acc	Acc	Acc	Acc
SFT	4.82	4.70	12.8	15.2	10.9	16.8	42.3	56.8
Supervised KD	7.92	12.1	42.7	53.4	10.7	16.1	46.7	63.3
On-policy KD	7.92	15.1	44.5	57.4	10.3	18.6	49.9	61.8
SKD	9.56	16.9	46.9	59.1	10.2	19.7	50.4	68.4

D ADDITIONAL RESULTS

Tables 4 and 5 show the performance of GEMMA-2B-IT (student), QWEN-0.5B-IT (students) and supervised FT teacher models (GEMMA-7B-IT and QWEN-7B-IT) under task-specific and task-agnostic distillation. Table 6 presents task-agnostic instruction tuning results for baseline and SKD models across four held-out test sets, using both low-data (1k prompts) and full-data (10k prompts) training regimes. These results align with those shown in Figure 3.

E TASK INPUT AND OUTPUT FORMAT

We include input and output formats for four text generation tasks that we studied in this paper: math instruction following (Table 7), arithmetic reasoning (Table 8), dialogue summarization (Table 9) and translation (Table 10).

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Prompt for model:

Solve the following math problem step-by-step.

Simplify your answer as much as possible. Present your final answer as boxed {Your Answer}.

For a positive integer n , let $H_n = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n}$. Compute $\sum_{n=1}^{\infty} \frac{1}{(n+1)H_n H_{n+1}}$.

Output for model:

Step 1: First, let's understand what the problem is asking for. We are given the expression $H_n = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n}$ for a positive integer n , and we need to compute the sum $\sum_{n=1}^{\infty} \frac{1}{(n+1)H_n H_{n+1}}$.

Step 2: To begin solving this problem, we can start by manipulating $\frac{1}{(n+1)H_n H_{n+1}}$ to see if we can simplify it in any way. By doing this, we can rewrite $\frac{1}{(n+1)H_n H_{n+1}}$ as $\frac{H_{n+1} - H_n}{H_n H_{n+1}}$.

Step 3: This allows us to express the given sum as a telescoping series, where each term of the sum cancels out most of the preceding and following terms.

Step 4: Taking the sum $\sum_{n=1}^{\infty} \frac{1}{(n+1)H_n H_{n+1}}$ in parts, we can break it down term by term to simplify the expression and find the pattern of the telescoping series.

Step 5: By expressing each term of the sum as a difference involving the harmonic numbers H_n , we can observe how the terms cancel out, and this leads us to the result of the sum.

Step 6: To validate and confirm the result, we need to acknowledge the fact that $H_n \rightarrow \infty$ as $n \rightarrow \infty$, and we can prove this by showing that the sum of the harmonic series diverges.

Step 7: This proof involves demonstrating that the sum of the harmonic series increases without bound as n increases, which verifies the behavior of H_n as $n \rightarrow \infty$.

Step 8: Therefore, the final answer to the given problem is $\boxed{1}$, and it is backed by the property of harmonic numbers tending to infinity as the terms increase.

Answer:

$\boxed{1}$

Table 7: This outlines the input and output format for math instruction following task. The output includes chain-of-thought steps leading to the computed answer, which is presented in a boxed format. We evaluate these answers by comparing their numerical equivalence to the ground truth.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Prompt for model:

Q: American carmakers produce 5 650 000 cars each year. Then they distribute it to 5 car suppliers. The first supplier receives 1 000 000 cars each. The second supplier receives 500 000 cars more while the third car supplier receives the same number as the first and second suppliers combined. The fourth and the fifth supplier receive an equal number of cars. How many cars do the fourth and fourth suppliers each receive?

Output for model:

The second supplier receives $1\,000\,000 + 500\,000 = 1\,500\,000$ cars.

The third supplier receives $1\,000\,000 + 1\,500\,000 = 2\,500\,000$ cars.

So the first to third suppliers receive a total of $1\,000\,000 + 1\,500\,000 + 2\,500\,000 = 5\,000\,000$ cars.

Thus, $5\,650\,000 - 5\,000\,000 = 650\,000$ cars are left to be divided by the two other suppliers.

Therefore, the fourth and fifth suppliers receive $650\,000/2 = 325\,000$ cars each.

325,000

Table 8: This outlines the input and output format for arithmetic reasoning task. The output includes chain-of-thought steps leading to the computed answer, which is presented after "####". We evaluate these answers by comparing their numerical equivalence to the ground truth.

Prompt for model:

#Person1#: Hi, Mr. Smith. I'm Doctor Hawkins. Why are you here today?

#Person2#: I found it would be a good idea to get a check-up.

#Person1#: Yes, well, you haven't had one for 5 years. You should have one every year.

#Person2#: I know. I figure as long as there is nothing wrong, why go see the doctor?

#Person1#: Well, the best way to avoid serious illnesses is to find out about them early. So try to come at least once a year for your own good.

#Person2#: Ok.

#Person1#: Let me see here. Your eyes and ears look fine. Take a deep breath, please. Do you smoke, Mr. Smith?

#Person2#: Yes.

#Person1#: Smoking is the leading cause of lung cancer and heart disease, you know. You really should quit.

#Person2#: I've tried hundreds of times, but I just can't seem to kick the habit.

#Person1#: Well, we have classes and some medications that might help. I'll give you more information before you leave.

#Person2#: Ok, thanks doctor.

Output for model:

Mr. Smith's getting a check-up, and Doctor Hawkins advises him to have one every year. Hawkins'll give some information about their classes and medications to help Mr. Smith quit smoking.

Table 9: This outlines the input and output format for dialogue summarization task. The output is the summary of above dialogue. We evaluate summary by computing its ROUGE-L score to the ground truth.

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

Prompt for model:

Translate Assamese sentence into English. This is the Assamese sentence: [Assamese Texts]. Here is the English translation:

Output for model:

It was to last for the next 40 years and would be fought for real, by proxy armies, on battlefields from Africa to Asia, in Afghanistan, Cuba and many other places.

Table 10: This outlines the input and output format for translation task. The output is the English translation. We evaluate translation by computing its COMET score to the reference and source.

F EVALUATION METRICS

As in (Freitag et al., 2022), we adopted the learned evaluation metric COMET (Rei et al., 2020) for assessing our translation outputs, instead of the n-gram-based BLEU (Papineni et al., 2002). Specifically, we used the comet-da-22 checkpoint (Rei et al., 2022), a widely recognized metric that exhibits strong correlation with human judgments at WMT22 (Freitag et al., 2022). COMET is a learned metric trained on a large-scale dataset of human judgments for evaluating translation quality. It takes source, translation, and reference as input, producing a quality score between 0 and 1. A COMET score improvement of 0.5 to 1.0 indicates a substantial improvement.

Tables 7 and 8 illustrate the structured output format used to derive final answers from math instruction tasks and GSM-8k. We employed the evaluation code from (Yuan et al., 2024) to assess the mathematical equivalence of parsed answers and ground truths. This involved converting mathematical symbols like π into their numerical values for more accurate numerical comparisons, rather than relying solely on string matching. Implementation details can be found in the original paper (Yuan et al., 2024). We reported accuracy for both tasks.

We evaluated summarization output against references using ROUGE-L (Lin, 2004). ROUGE-L ranges from 0 to 100, with a 1-point increase indicating substantial improvement.

Prompt for model:

Translate Assamese sentence into English. This is the Assamese sentence: [Assamese Texts]. Here is the English translation:

Reference translation:

Italian is also the everyday language used by most of those who work in the state while Latin is often used in religious ceremonies.

On-policy student sampled output (COMET: 53):

The Estonian language is also used in the Estonian-speaking countries as a national language, and is used in the Estonian-speaking countries as a medium of instruction in the Estonian-speaking countries.

SKD student sampled output (COMET: 80):

The Latin language is still used in the country to a large extent, though it is not the official language.

Table 11: We sample training data from SKD and on-policy methods. Our results show that solely sampling from the student model can lead to hallucinated translations and limited training signals (On-policy output). Conversely, SKD generates fluent, semantically similar outputs while staying close to the student’s distribution (SKD output). Student model is supervised FT QWEN-0.5B.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

Prompt for model:

Q: Libby is building an igloo in her backyard using bricks of snow. She builds her igloo in rows, using a total of 10 rows of bricks of snow. The bottom half of the igloo has 12 bricks of snow in each row while the top half of the igloo has 8 bricks of snow in each row. How many bricks of snow did Libby use for her igloo? A:

Correct Answer:

In the bottom half of the igloo, there are $10 \text{ rows} / 2 = 5$ rows of bricks of snow.

So in the bottom half of the igloo, there are $5 \text{ rows} * 12 \text{ bricks of snow in each row} = 60$ bricks of snow.

In the top half of the igloo, there are $5 \text{ rows} * 8 \text{ bricks of snow in each row} = 5*8=40$ bricks of snow.

There is therefore a total of $60 + 40 = 100$ bricks of snow in the igloo. ##### 100

On-policy student sampled output:

$\$35 - \$160 = \$125$

Libby has $\$160 - \$125 = \$35$

Libby has $160 - 35 = \$125$

Libby has $160 - 35 = \$125...$

SKD student sampled output:

The bottom half of the igloo has $12 \times 10 = 120$ bricks of snow.

The top half of the igloo has $8 \times 10 = 80$ bricks of snow.

Libby used $120 + 80 = 200$ bricks of snow. ##### 200

Table 13: We sample training data from SKD and on-policy methods using the GEMMA-2B-IT model. These samples are collected during the early stages of training. The on-policy sample exhibit repetitive outputs and errors in arithmetic calculations. While our SKD model also produces incorrect answers, incorporating teacher’s corrections enables it to perform some intermediate steps correctly and maintain accurate arithmetic calculations.

Prompt for model:

Q: John Smith buys 3 cakes for \$12 each and splits the cost with his brother. How much did he pay? A:

Correct Answer:

He spent $3 * 12 = 36$ total So he paid $36 / 2 = 18$ ##### 18

On-policy student sampled output:

Each brother pays $\$12 / 3 = 4$ each

So he pays $\$12 - \$4 = 8$ ##### 8

SKD student sampled output:

The 3 cakes cost John $3 \times \$12 = \36

So he paid $\$36 / 2 = \18 each ##### 18

Table 14: We sample training data from SKD and on-policy methods using the GEMMA-2B-IT MODEL. These samples are collected during the late stages of training. We observe that teacher model can perform some intermediate correction on the reasoning step, ultimately leading to correct answer.

H OVER-FITTING AT SFT STAGE CAN CAUSE POST-TRAINING KDS TO SUB-OPTIMAL PERFORMANCE

Figure 7 demonstrates that overfitting during the SFT stage can hinder optimal performance in post-training KD. **End-to-end KD methods, such as SKD, which can train from arbitrary base model and bypass the SFT stage, offer significant benefits.** To investigate this, we selected three checkpoints from the supervised FT stage of the student model, corresponding to training steps 64, 128, and 192. Based on validation, the checkpoint at step 128 represents the converged model, while step 192 exhibits signs of overfitting (Figure a). Figure (b) shows that the converged checkpoint (step 128) outperforms the overfitted checkpoint (step 192). Moreover, we found that training from an earlier checkpoint (step 64) can yield better results than the overfitted checkpoint (step 192). This observation is consistent for both on-policy and SKD. However, SKD is an end-to-end training pipeline. Our findings confirm that overfitting during the SFT stage can lead to suboptimal performance in post-training KD. As shown in Section 5.4, training with extremely limited data can inevitably lead to overfitting during the SFT stage, further highlighting the superiority of our end-to-end SKD pipeline compared to two-stage KD approaches like on-policy KD.




Figure 7: In this Figure, we demonstrate that over-fitting during the SFT stage can hinder optimal performance in post-training knowledge distillation. **End-to-end KD methods, such as SKD, which can train directly from arbitrary base model and bypass the SFT stage, offer significant benefits.** Figure (a) presents three checkpoints from the SFT stage, corresponding to training steps 64, 128, and 192. Based on validation loss, the checkpoint at step 128 represents the converged model, while step 192 shows signs of overfitting. Figure (b) demonstrates that the converged checkpoint (step 128) outperforms the overfitted checkpoint (step 192). This finding is consistent for both on-policy KD and SKD.

I TRAINING DETAILS OF MIXED TRAINING OF SUPERVISED KD AND ON-POLICY KD

The parameter setting for this baseline is identical for the setup in Appendix C. We start to train student model with supervised KD samples for 188 steps (roughly 1.5 epoches), then we perform on-policy KD for another 187 steps. Total training step is 375. Initial learning rate is $1e-5$ and linear decay is applied during training process.

J SPEEDUP TO SPECULATIVE DECODING

Motivated by (Zhou et al., 2024), we further investigate the speedup achieved by each baseline model when used as a draft model to approximate a teacher model during speculative decoding (Leviathan et al., 2022)]. As depicted in Figure 8, SKD-trained student models consistently exhibit the highest speedup and token acceptance ratio (Approximation model will accept or reject tokens depending on whether draft token is sampled under its probability distribution) across translation and summarization tasks. These models accelerate instruction-tuned GEMMA-2B by 1.21x and 1.17x in translation and summarization, respectively, while also improving token acceptance ratio by 1.71x

and 1.85x. These results demonstrate that SKD not only produces superior performance but also generates the closest distribution to the teacher model. This opens up a new application: utilizing SKD to train powerful draft models and accelerate the speculative decoding process.




Figure 8: We utilize student models trained with various baseline KD methods as draft models and a supervised fine-tuned GEMMA-7B-IT model as the approximation model for speculative decoding. Our results demonstrate that SKD-trained student models significantly outperform the others, achieving the highest speedup and token acceptance ratio. SKD accelerates GEMMA-2B-IT by 1.21x and 1.17x in translation and summarization, respectively, while also improving token acceptance ratio by 1.71x and 1.85x.