

# Stable On-Policy Distillation through Adaptive Target Reformulation

Ijun Jang Jewon Yeom Juan Yeo Hyunggu Lim Taesup Kim<sup>†</sup>  
Graduate School of Data Science, Seoul National University  
{ijun0824, jewon0908, juanyeo, sjrksjek, taesup.kim}@snu.ac.kr

## Abstract

Knowledge distillation (KD) is a widely adopted technique for transferring knowledge from large language models to smaller student models. However, conventional supervised KD often suffers from a distribution mismatch between training and inference. While on-policy KD approaches attempt to mitigate this issue by learning directly from student-generated outputs, they frequently encounter training instabilities because the distributional gap between the novice student and the expert teacher is often too wide to bridge directly. These challenges manifest as pathological gradients in forward KL objectives or diversity collapse in reverse KL regimes. To address these limitations, we propose *Veto*, an objective-level reformulation that constructs a geometric bridge in the logit space. Unlike prior methods that mix data samples, *Veto* creates an intermediate target distribution that promotes alignment between the teacher and the student. By introducing a tunable parameter  $\beta$ , *Veto* serves as an *Adaptive Gradient Veto* that stabilizes optimization by suppressing harmful gradients on low-confidence tokens, while simultaneously acting as a *Decisiveness Knob* to balance reward-driven performance with output diversity. Extensive experiments across various reasoning and generation tasks demonstrate that *Veto* consistently outperforms supervised fine-tuning and existing on-policy baselines.

## 1 Introduction

Knowledge distillation (KD) (Buciluă et al., 2006; Hinton et al., 2015) is widely used for transferring capabilities from proprietary models to efficient open-source counterparts and facilitating model self-improvement (Xu et al., 2024). Beyond model compression, KD has recently become an important mechanism for improving large language

models (LLMs) through self-training and alignment, where a strong teacher supervises a weaker or partially trained student. However, traditional supervised KD (Sanh et al., 2019; Kim and Rush, 2016) suffers from exposure bias (Ranzato et al., 2015; Bengio et al., 2015): the mismatch between teacher-provided trajectories and the student’s self-generated outputs leads to degraded performance in autoregressive tasks (Zhang et al., 2019; Arora et al., 2022), especially in long-horizon generation and reasoning settings where errors compound over time.

On-policy knowledge distillation addresses this limitation by aligning training with inference-time behavior, learning directly from student-generated outputs (Agarwal et al., 2024; Lin et al., 2020). This enables the student to receive feedback on its own predictions and adapt within regions of the probability space it is likely to visit at test time. Building on this idea, recent methods attempt to bridge the teacher-student gap using reverse KL divergences (Gu et al., 2023; Wen et al., 2023) or interleaved sampling strategies (Xu et al., 2025). While these approaches try to bridge the gap by mixing teacher and student tokens at the *data level*, they largely overlook a critical challenge: *the stability of the optimization objective itself*. Even with mixed data, forcing a novice student to match an expert’s sharp distribution creates a steep optimization cliff.

In practice, early-stage student policies are highly noisy and often assign near-zero probability to teacher-preferred tokens. As illustrated in Table 1, this leads to unreliable teacher feedback and pathological gradients. Standard forward KL objectives suffer from gradient explosion on such *ignorant tokens* (Agarwal et al., 2024), while reverse KL objectives, although numerically stable, lack explicit control over the trade-off between mode-seeking decisiveness and distributional diversity, frequently resulting in premature mode collapse.

<sup>†</sup>Corresponding author.




Figure 1: **Overview of stable on-policy knowledge distillation.** **Top (Forward KL):** Standard Forward KL triggers explosive gradients when students are *ignorant of teacher-preferred tokens* (low  $P_S$ ). Veto reformulates the logit-space target to suppress harmful updates to stabilize early optimization. **Bottom (Reverse KL):** Standard Reverse KL causes mode collapse and diversity loss. By tuning  $\beta$ , Veto balances generation decisiveness with distributional diversity, allowing the student to capture multiple teacher modes effectively.

These failure modes stem from the geometry of the divergence objective, rather than from model architecture or data generation strategy.

In this work, we propose *Veto*, an objective-level reformulation that stabilizes on-policy knowledge distillation by constructing a distributional bridge between the teacher and the student. Instead of mixing data samples, *Veto* mixes the distributions themselves in the logit space to create a geometric target that emphasizes regions of agreement. This intermediate target effectively “vetoes” harmful updates on low-confidence tokens, acting as a safe stepping stone for the student. We show that a single scalar parameter  $\beta$  plays a dual role: it acts as an *Adaptive Gradient Veto* that prevents gradient explosion in forward KL, and as a *Decisiveness Knob* that enables controlled entropy regularization and explicit diversity-decisiveness trade-offs in reverse KL. Extensive experiments across reasoning and generation tasks demonstrate that *Veto* consistently improves training stability and outperforms supervised fine-tuning and existing on-policy distillation baselines.

**Contributions** Our contributions are summarized as follows:

- We introduce *Veto*, an objective-level reformulation for on-policy knowledge distillation that bridges the teacher-student gap via logit-space interpolation, improving stability without modifying architectures.
- We show that a single parameterized target dis-

tribution unifies the treatment of forward and reverse KL objectives, mitigating gradient explosion and mode collapse.

- We provide theoretical analysis connecting *Veto* to adaptive gradient suppression in forward KL and entropy-regularized policy gradients in reverse KL.
- We empirically validate *Veto* on reasoning, code generation, and summarization tasks, demonstrating consistent gains over supervised and on-policy baselines.

## 2 Related Works

Knowledge Distillation (KD) (Buciluă et al., 2006; Hinton et al., 2015) has become a cornerstone technique for compressing Large Language Models (LLMs) by transferring knowledge from a high-capacity teacher to a more efficient student. Early approaches, often categorized as Supervised KD, typically train the student to mimic the teacher’s distribution over a fixed dataset (Sanh et al., 2019) or sequence-level outputs (Kim and Rush, 2016). Despite their effectiveness, these methods are prone to exposure bias (Ranzato et al., 2015; Bengio et al., 2015), which arises from the mismatch between ground-truth-conditioned training and self-generated outputs. This discrepancy often results in degraded performance in autoregressive generation tasks (Zhang et al., 2019; Chiang and Chen, 2021; Arora et al., 2022).

To mitigate this training-inference mismatch, recent research has shifted its focus to on-policy KD,




Figure 2: **Comparison of gradient magnitudes on ignorant tokens** ( $P_T > 0.1, P_S < 0.01$ ). While standard Forward KL triggers pathological gradient explosions exceeding  $10^7$ , Veto effectively suppresses these spikes within a stable range. This demonstrates *Veto*’s ability to ensure optimization stability by suppressing harmful updates during early training.

which aligns the training distribution with the student’s inference distribution. Inspired by imitation learning (Ross et al., 2011), ImitKD (Lin et al., 2020) initiated this line of work by training the student on a mixture of fixed and self-generated sequences. Building on this, GKD (Agarwal et al., 2024) introduced a generalized framework that enables distillation entirely on student-generated outputs, demonstrating that learning from self-generated mistakes can substantially improve performance.

Subsequent work has further refined on-policy distillation objectives. MiniLLM (Gu et al., 2023) and f-distill (Wen et al., 2023) leverage reverse KL and f-divergences to encourage mode-seeking behavior. Furthermore, SKD (Xu et al., 2025) introduces interleaved sampling to actively correct low-quality student generations, improving feedback quality. Despite these advances, existing approaches largely overlook the stability of the optimization objective itself, particularly during the early stages of training when student outputs are highly noisy. In this paper, we propose *Veto*, which simultaneously acts as a gradient stabilizer for forward KL and a tunable control for decisiveness in reverse KL, offering a unified solution that ensures stable optimization without the architectural overhead of previous methods.

### 3 Preliminaries

#### 3.1 On-Policy Knowledge Distillation

Knowledge distillation (KD) (Buciluă et al., 2006; Hinton et al., 2015) aims to transfer the predictive capabilities of a high-capacity teacher model

$P_T$  to a smaller student model  $P_S$ . In the context of Large Language Models (LLMs), conventional supervised KD typically minimizes the divergence between the student’s output and fixed teacher-generated trajectories. However, this approach often leads to exposure bias (Ranzato et al., 2015), as the student is not trained to handle the distribution shifts it encounters during autoregressive generation at inference time.

To mitigate this, on-policy KD (Agarwal et al., 2024) aligns the training distribution with the student’s inference distribution by sampling sequences  $y = (y_1, \dots, y_T)$  directly from the student’s own policy  $P_S(\cdot | x)$ . Formally, for a given prompt  $x$ , the objective is to minimize the expected token-level divergence along these self-generated trajectories:

$$\mathcal{L} = \mathbb{E}_{x \sim P_{\text{data}}, y \sim P_S(\cdot | x)} \left[ \sum_{t=1}^{|y|} D_{\text{KL}}(P_T(\cdot | x, y_{<t}) \| P_S(\cdot | x, y_{<t})) \right]$$

By receiving feedback on its own mistakes, the student learns to correct its behavior within the specific regions of the probability space it is likely to visit during test time (Agarwal et al., 2024; Xu et al., 2025).

#### 3.2 Asymmetric Properties of KL Divergence

The behavior of on-policy distillation is fundamentally governed by the direction of the Kullback-Leibler (KL) divergence. The forward KL divergence,  $D_{\text{KL}}(P_T \| P_S)$ , is known for its *zero-avoiding* property, which forces the student to assign non-zero probability to any token the teacher considers likely. In contrast, the reverse KL divergence,  $D_{\text{KL}}(P_S \| P_T)$ , is *mode-seeking*, encouraging the student to concentrate its probability mass on the primary modes of the teacher’s distribution. While these properties provide the theoretical basis for various distillation frameworks (Gu et al., 2023; Wen et al., 2023), they also introduce significant numerical challenges when applied to on-policy settings where the student’s early-stage outputs are highly unstable.

## 4 Methodology

#### 4.1 Stability Challenges in On-Policy KD

Despite the theoretical advantages of learning on-policy, the mismatch between the teacher’s expertise and the student’s initial ignorance often leads to optimization failures.

---

<Example 1>

**Question:** There are 350 trees in the park, 70% of which are pine trees. How many are not pine trees?

**Answer:** To find out how many trees are not pine trees, we need to subtract the number of pine trees from the total number of trees. → **NO ANSWER**

**Reference:** There are  $350 \times 70 / 100 = \langle 350 * 70 / 100 = 245 \rangle$  245 pine trees. Therefore,  $350 - 245 = \langle 350 - 245 = 105 \rangle$  105 trees are not pine trees. → **105**

---

<Example 2>

**Question:** The Donaldsons pay \$15 per hour for babysitting. The Merck family pays \$18 per hour and the Hille family pays \$20 per hour for babysitting. Layla babysat for the Donaldsons for 7 hours, the Merck family for 6 hours and the Hille family for 3 hours. How many dollars did Layla earn babysitting?

**Answer:** To find out how much Layla earned in total, we need to multiply the number of hours she worked by her hourly rate. → **NO ANSWER**

**Reference:** Donaldsons pay \$15 / hour =  $15 * 7 = \langle 15 * 7 = 105 \rangle$  105. The Mercks pay \$18 / hour =  $18 * 6 = \langle 18 * 6 = 108 \rangle$  108. The Hilles pay \$20 / hour =  $20 * 3 = \langle 20 * 3 = 60 \rangle$  60. The total amount Layla earned is  $105 + 108 + 60 = \langle 105 + 108 + 60 = 273 \rangle$  273. Layla earned \$273 for babysitting. → **273**

---

<Example 3>

**Question:** Bryan works as a social media account manager. He does marketing posts, advertisement posts, and customer outreach posts. His client has him spend four hours on customer outreach posts and half that time on advertisement posts each day. Bryan works eight hours a day. How much time in hours each day does he spend on marketing posts?

**Answer:** He spends  $4/2 = 2$  hours on advertisement posts. So he spends  $8 - 4 - 2 = 0$  hours on marketing posts. → **WRONG LOGIC**

**Reference:** Let M be the number of hours Bryan spends on marketing posts. He spends  $4 / 2 = \langle 4 / 2 = 2 \rangle$  2 hours on advertisement posts a day. He works  $M + 4 + 2 = 8$  hours a day. Thus, Bryan spends  $M = 8 - 4 - 2 = \langle 8 - 4 - 2 = 2 \rangle$  2 hours on marketing posts each day. → **2**

---

<Example 4>

**Question:** Drew is reseeding his lawn with grass seed. One bag of grass seed covers 250 square feet of lawn. His lawn is 22 feet from the house to the curb and 36 feet from side to side. He bought four bags of seed. How many extra square feet could the leftover grass seed cover after Drew reseeds his lawn?

**Answer:** The area of the lawn is  $22 * 36 = 792$  square feet. Four bags of grass seed cover  $4 * 250 = 1000$  square feet. So, the leftover grass seed can cover  $792 - 1000 = -208$  square feet. Since the result is negative, it means that the leftover grass seed cannot cover any more area. → **WRONG LOGIC**

**Reference:** Drew's lawn is  $22 * 36 = \langle 22 * 36 = 792 \rangle$  792 square feet. He bought enough grass seed to cover  $250 * 4 = \langle 250 * 4 = 1000 \rangle$  1000 square feet. Thus, the leftover grass seed can cover  $1000 - 792 = \langle 1000 - 792 = 208 \rangle$  208 square feet. → **208**

---

Table 1: **Low-quality outputs of On-policy KD.** Student produces low-quality samples, leading to unreliable or noisy teacher feedback. Optimization is hindered by pathological gradients stemming from noisy teacher feedback on low-quality, out-of-distribution student samples.

**Gradient Explosion in Forward KL** The forward KL gradient with respect to student parameters  $\theta$  is defined as  $\nabla_{\theta} \mathcal{L}_{\text{FWD}} \approx -\sum_y \frac{P_T(y)}{P_S(y)} \nabla_{\theta} P_S(y)$ . When the student is ignorant of teacher-preferred tokens ( $P_S(y) \rightarrow 0$  while  $P_T(y) > 0$ ), the ratio  $\frac{P_T(y)}{P_S(y)}$  diverges. As shown in Figure 2, this leads to pathological gradient explosions during early training stages, where student is ignorant of teacher-preferred tokens (Agarwal et al., 2024).

**Mode Collapse in Reverse KL** While reverse KL avoids divergence, it lacks a mechanism to explicitly control the intensity of mode-seeking behavior. As a result, reverse-KL based distillation can encourage the student to concentrate on

dominant teacher modes, potentially reducing distributional diversity.

#### 4.2 Veto: Bridging the Gap via Target Reformulation

Standard on-policy KD often suffers from a severe distribution mismatch: forcing an early-stage student to immediately match an expert teacher's complex distribution can lead to optimization instability. To address this, we propose *Veto*, an objective-level reformulation that acts as a **geometric bridge** between the teacher and the student.

Instead of targeting the teacher directly, we construct an intermediate target distribution  $Q$  in the logit space. Let  $z_T = \log P_T(\cdot | x)$  and  $z_S = \log P_S(\cdot | x)$  denote the logits. We define  $Q$

---

**Algorithm 1** Stable On-policy Distillation via *Veto* with Linear  $\beta$  Decay
 

---

**Require:** Student  $P_\theta$ , Teacher  $P_T$ , Dataset  $\mathcal{X}$ , Initial hyperparameter  $\beta$ , Total steps  $N$ , Learning rate  $\eta$

```

1: Initialize student parameters  $\theta$ 
2: for each training step  $i = 1 \dots N$  do
   $\triangleright$  Linear decay from  $\beta$  to 0
3:    $\beta \leftarrow \beta \cdot (1 - \frac{i}{N})$ 
   $\triangleright$  On-policy sampling
4:   Sample prompt  $x \sim \mathcal{X}$ 
5:   Sample trajectory  $y \sim P_\theta(\cdot|x)$ 
6:   for each token  $t = 1 \dots |y|$  do
     $\triangleright$  Get logits
7:      $z_T = \log P_T(\cdot|x, y_{<t})$ 
8:      $z_S = \log P_\theta(\cdot|x, y_{<t})$ 
     $\triangleright$  Compute reformulated target
9:      $Q \propto \exp(z_T + \beta \cdot z_S)$ 
10:    if using Forward KL regime then
     $\triangleright$  Adaptive gradient Vetoing
11:      $\mathcal{L}_t = D_{KL}(Q \parallel P_\theta)$ 
12:    else if using Reverse KL regime then
     $\triangleright$  Decisiveness control
13:      $\mathcal{L}_t = D_{KL}(P_\theta \parallel Q)$ 
14:    end if
15:  end for
   $\triangleright$  Optimization
16:   $\mathcal{L} = \sum_t \mathcal{L}_t$ 
17:   $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}$ 
18: end for
  
```

---

as:

$$\begin{aligned}
 Q(y|x) &\propto \exp(z_T(y|x) + \beta \cdot z_S(y|x)) \\
 &= \frac{1}{Z(x)} P_T(y|x) \cdot P_S(y|x)^\beta
 \end{aligned}$$

In this formulation,  $Q$  defines a **geometric bridge** by forming an intermediate distribution that geometrically interpolates the teacher and student distributions in probability space. This formulation represents a *Product of Experts (PoE)*. Intuitively, it works as a **consensus filter**: the target probability  $Q(y|x)$  is high only when the token is supported by *both* the teacher (quality) and the student (confidence). The parameter  $\beta \geq 0$  controls the position of this bridge, allowing the student’s current ignorance to effectively *Veto* overwhelming signals from the teacher.

### 4.3 Theoretical Analysis

We provide a unified analysis of *Veto*, explicitly showing how the reformulated target ensures stabil-

ity (Analysis I) and bridges KD with RL (Analysis II). Detailed proofs are in Appendix A.

**Analysis I: Stability via Adaptive Veto** In standard Forward KD, the loss explodes when the student assigns near-zero ( $P_S \rightarrow 0$ ) probability to a token the teacher prefers, as the term  $\log P_S$  diverges. *Veto* resolves this structurally.

**Theorem 1** (Adaptive Gradient *Veto*). *In standard Forward KD ( $\beta = 0$ ), gradients diverge as  $P_S(y) \rightarrow 0$ . However, in Veto ( $\beta > 0$ ), the target  $Q$  incorporates the student’s uncertainty, ensuring the loss converges to 0.*

*Proof Intuition.* As shown in Appendix A, substituting  $Q$  into the loss gradient, the diverging term becomes proportional to

$$\mathcal{L}(y) \approx P_S(y)^\beta \log P_S(y).$$

By L’Hôpital’s rule, the polynomial term  $P_S^\beta(y)$  decays to zero faster than the logarithmic term  $\log P_S(y)$  diverges, effectively acting as a gate that suppresses updates on ignorant tokens.

**Theorem 2** (Sharpening Effect). *The student converges to a sharpened version of the teacher with implicit temperature scaling  $T = 1 - \beta$ .*

*Derivation Logic.* At the optimal fixed point where  $P_S^* = Q$ , we have the relation  $P_S^* \propto P_T \cdot (P_S^*)^\beta$ . Rearranging terms yields  $(P_S^*)^{1-\beta} \propto P_T$ , which implies:

$$P_S^*(y|x) \propto P_T(y|x)^{\frac{1}{1-\beta}}$$

Since  $0 \leq \beta < 1$ , the exponent  $\frac{1}{1-\beta} > 1$ , naturally encouraging the student to be more decisive (sharper) than the teacher.

### Analysis II: Bridge to Reinforcement Learning

In the reverse regime, *Veto* acts as a *Decisiveness Knob* by bridging the gap between KD and RL.

**Theorem 3** (Bridge to REINFORCE). *The gradient of the Reverse GKD objective with respect to the student parameters  $\theta$  is equivalent to a policy-gradient update with scaled entropy regularization:*

$$\nabla_\theta \mathcal{L}_{\text{REV}} = \mathbb{E}_{y \sim P_S} \left[ \nabla_\theta \log P_S(y) \cdot A(y) \right],$$

where

$$A(y) = \underbrace{-\log P_T(y)}_{\text{Reward } r(y)} + \underbrace{(1 - \beta) \log P_S(y)}_{\text{Scaled Entropy Cost}}.$$

Method	GSM8K	HumanEval		DialogSum
	Accuracy	Pass@1	Pass@10	Win-rate
Teacher SFT	74.7	64.7	72.2	65.0
Student SFT	30.7	26.9	34.6	54.0
Supervised KD	33.4	26.8	34.5	54.3
SKD	33.6	24.8	34.8	53.6
On-policy KD	35.1	22.9	35.3	54.3
<b>Ours (Veto)</b>	<b>39.9</b>	<b>29.0</b>	<b>37.7</b>	<b>56.5</b>

Table 2: **Performance across three distinct domains.** We evaluate mathematical reasoning on GSM8K (Accuracy), code generation on HumanEval (Pass@1, Pass@10), and dialogue summarization on DialogSum (Win-rate). Pass@1 and Pass@10 quantify the proportion of problems for which at least one correct solution is obtained from 1 and 10 generated samples, respectively, whereas the Win-rate represents the proportion of model outputs that are preferred by GPT-4o-mini in pairwise evaluations. Our method outperforms baselines in all student categories. Note that we use  $\beta = 0.8$  for reasoning,  $\beta = 1.0$  for code tasks, and  $\beta = 0.3$  for summarization.

**The Spectrum of Decisiveness** This formulation reveals that  $\beta$  controls the trade-off between mode-covering and mode-seeking:

- $\beta = 0$  (**Standard Reverse KD**): The student tries to match the teacher’s distribution exactly, inheriting both its modes and its uncertainty.
- $0 < \beta < 1$  (**Geometric KD**): The "Middle Ground." The student seeks high-reward regions (like RL) while retaining a diversity budget proportional to  $(1 - \beta)$  to prevent premature collapse.
- $\beta \rightarrow 1$  (**Pure RL**): Zero entropy regularization. The objective becomes equivalent to REINFORCE, where the student collapses to the single highest-reward mode.

## 5 Experiments

### 5.1 Implementation Details

**Model details** We use Qwen2-0.5B-IT (Yang et al., 2024) as the student model and Qwen2-7B-IT (Yang et al., 2024) as the teacher model. We first fine-tune the teacher model in a supervised manner on task-specific data. The student model is then trained using outputs generated via on-policy KD, and optimized using our proposed objective with forward KL.

**Hyperparameters** All experiments use a learning rate of  $1e-5$ , a warmup ratio of 0.1, and a dropout rate of 0.1, with training conducted for three epochs with 2 H100 GPUs with 80GB of memory. The hyperparameter  $\beta$  was selected via grid search and scheduled with linear decay (see [Choice of  \$\beta\$  Scheduling Strategy](#)).

### 5.2 Mathematical Reasoning

We conduct experiments on the GSM8K (Cobbe et al., 2021) dataset. Following the setup used in the (Xu et al., 2025), we first fine-tune the teacher model in a supervised manner on 7K instances from the GSM8K training set. We then randomly sample 1K instances from the training set to construct our input–output pairs  $(x, y)$  for student training. For evaluation, we employ the GSM8K test set, which consists of 1,319 instances. The quality of generated outputs is assessed using answer accuracy.

As shown in Table 2, our method significantly outperforms supervised fine-tuning of the student model, yielding a 9.2% absolute improvement in accuracy (30.7%  $\rightarrow$  39.9%). Moreover, our approach surpasses prior baselines, including SKD and on-policy KD, by up to 6.3%, demonstrating the effectiveness of our objective-level reformulation in on-policy distillation settings.

### 5.3 Code Generation

For code generation, we use prompts from the WizardCoder (Luo et al., 2024) dataset. From this dataset, we randomly sample 10K instances for supervised fine-tuning of the teacher model, 1K instances for student training, and 1K instances for evaluation.

We evaluate code generation performance on the HumanEval (Chen, 2021) benchmark using pass@k metrics. As shown in Table 2, compared to on-policy KD, our method improves Pass@1 from 22.9 to 29.0 (+6.1) and Pass@10 from 35.3 to 37.7 (+2.4), indicating clear gains in code synthesis performance.

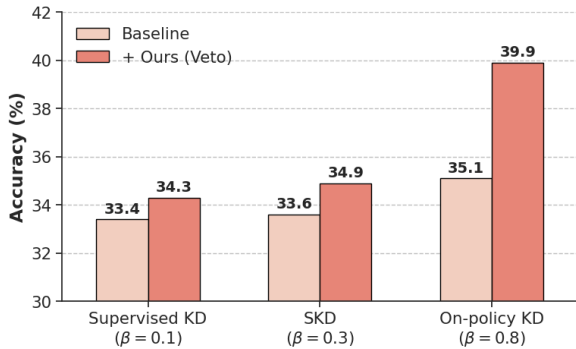


Figure 3: **Effect of different on-policy data generation strategies** on student performance for mathematical reasoning. Accuracy (%) is reported along with absolute improvements over the corresponding supervised student baseline for supervised KD, SKD, and on-policy KD under different values of  $\beta$ .

#### 5.4 Summarization

For dialogue summarization, we utilize the DialogSum (Chen et al., 2021) dataset. We randomly sample 1K instances from the DialogSum training set to construct input–output pairs  $(x, y)$  for student training. For evaluation, we use the official test set (1,500 instances). We first fine-tune the teacher model in a supervised manner on 10K instances sampled from the DialogSum training set, and then train the student model on the 1K sampled instances using our proposed distillation method. Summary quality is evaluated using win-rate (Zheng et al., 2023) based on pairwise comparisons judged by GPT-4o-mini.

As reported in Table 2, our method achieves the highest win-rate among all student models. In particular, it improves the win-rate from 54.3 to 56.5 (+2.2) compared to both supervised KD and on-policy KD, and outperforms SKD by 2.9 % (53.6%  $\rightarrow$  56.5%). These results indicate that our approach produces more faithful and higher-quality summaries, even when trained with limited student data.

#### 5.5 Ablation Study

For ablation study, we tested on GSM8k with Qwen2-0.5B-IT (Yang et al., 2024) as the student model and Qwen2-7B-IT (Yang et al., 2024) as the teacher model. Unless otherwise specified, the default objective for our experiments is conducted on on-policy outputs with *Veto* (forward KL) objective.

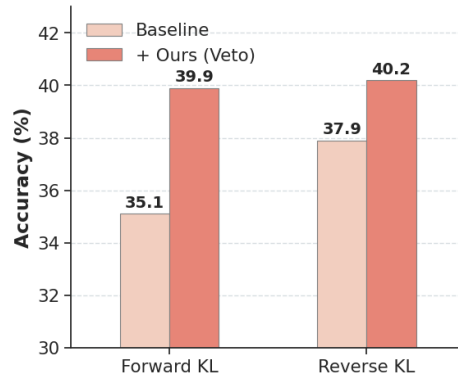


Figure 4: **Ablation comparing forward and reverse KL divergence objectives.** Accuracy (%) is reported for standard on-policy knowledge distillation using each KL loss and for the same settings augmented with the proposed method.

**Data Generation Setting** In Figure 3, we evaluate the effectiveness of our proposed method across various data generation strategies commonly used in knowledge distillation. Specifically, we apply our objective to supervised KD, speculative KD (SKD), and on-policy KD to assess its robustness. For supervised KD, which relies on a fixed teacher-generated dataset, our method achieves an accuracy of 34.3%, representing an absolute improvement of 0.9% over the standard baseline. Similarly, in the SKD setting, we observe a 1.3% performance gain. Most notably, the most significant improvement occurs in the on-policy setting, where our method reaches 39.9% accuracy, a +4.8% gain compared to previous on-policy distillation. These results demonstrate that our proposed objective consistently improves student performance and ensures optimization stability regardless of the specific data generation strategy utilized.

**Adaptability of KL Loss** We conduct an ablation study to evaluate the generalizability of *Veto* across different KL divergence objectives: Forward KL and Reverse KL. As summarized in Figure 4, our proposed method consistently enhances performance regardless of the underlying divergence used for distillation. By applying *Veto* to forward KL, we observe a significant absolute accuracy gain of 4.8% (35.1%  $\rightarrow$  39.9%). This improvement validates *Veto*’s role as an *Adaptive Gradient Veto*, which stabilizes early-stage optimization by effectively suppressing pathological gradients on ignorant tokens. When applied to Reverse KL, our method also improves performance by 2.3% (37.9%  $\rightarrow$  40.2%). In this context, *Veto* acts as a

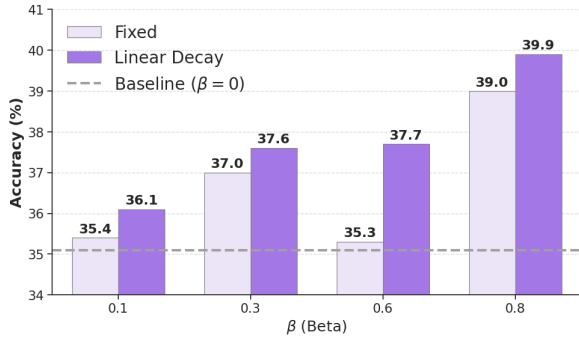


Figure 5: **Ablation comparing forward and reverse KL divergence objectives.** Accuracy (%) is reported for standard on-policy knowledge distillation using each KL loss and for the same settings augmented with the proposed method.

*Decisiveness Knob*, allowing for controlled entropy regularization and preventing the student from converging too prematurely to a single mode. These results demonstrate that *Veto* is objective-agnostic. It serves as a versatile, objective-level reformulation that provides a unified solution for ensuring optimization stability and enhancing generation quality in various on-policy distillation settings.

**Choice of  $\beta$  Scheduling Strategy.** To further investigate the dynamics of the parameter  $\beta$ , we explore various scheduling strategies. In our framework, a higher  $\beta$  prioritizes optimization stability by more aggressively suppressing pathological gradients on tokens where the student is ignorant. Conversely, a lower  $\beta$  allows the student to follow the teacher’s distribution more closely, which is beneficial once the student has gained sufficient proficiency. Based on this intuition, we implemented a linear decreasing schedule to ensure stability during the critical early stages of training. In this regime,  $\beta$  starts at a high initial value and linearly decays to 0 as the student’s distribution gradually aligns with that of the teacher. We compared this strategy with a constant schedule, where  $\beta$  remains fixed throughout the entire training process. As shown in Figure 5, the linear decreasing schedule consistently outperforms the constant scheduling across all tested values. Specifically, we found that starting from an initial  $\beta = 0.8$  yields the highest performance. This suggests that while a strong *Veto* gate is essential for stabilizing the early ignorant phase of the student, gradually relaxing this constraint allows the model to capture the finer nuances of the teacher’s distribution as training progresses.

Method	$\beta$	Accuracy
Teacher SFT	-	77.2
Student SFT	-	49.3
Supervised KD	-	47.6
SKD	-	47.1
On-policy KD	-	49.8
<b>Ours (<i>Veto</i>)</b>	<b>0.3</b>	<b>51.2</b>

Table 3: Mathematical reasoning accuracy on **GSM8K**. We distill Gemma2-2B-IT using SFTed Gemma2-9B-IT as a teacher. **Ours (*Veto*)** shows the best performance among student models.

**Generalizability to Various Models** To evaluate the generalizability of our proposed method, we extend our experiments to the Gemma2 model family (Team et al., 2024), employing Gemma2-9B-IT as the teacher and Gemma2-2B-IT as the student. We conducted evaluations on the GSM8K dataset, a representative benchmark for mathematical reasoning, with the decisiveness parameter fixed at  $\beta = 0.3$ . Our results show that *Veto* consistently yields significant performance improvements in this setting as well. This confirms that our objective-level reformulation is not limited to a specific model architecture but is robustly applicable across various state-of-the-art model families and complex reasoning tasks.

## 6 Conclusion

In this work, we proposed *Veto*, an objective-level reformulation for on-policy knowledge distillation (KD) that improves optimization stability. By constructing a geometric target distribution in logit space, *Veto* emphasizes agreement between the teacher and the student, effectively suppressing the pathological gradients typically encountered in forward KL objectives. Simultaneously, it serves as a *Decisiveness Knob* in reverse KL regimes, balancing reward-driven performance with distributional diversity to prevent premature mode collapse. Experiments across reasoning, code generation, and summarization tasks demonstrate that *Veto* consistently outperforms supervised fine-tuning and existing on-policy distillation baselines. These results suggest that *Veto* is an objective-agnostic and versatile framework capable of enhancing the efficiency and generation quality of smaller language models across diverse domains.

## Acknowledgment

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00345809, Research on AI Robustness Against Distribution Shift in Real-World Scenarios; and No. RS-2023-00222663, Center for Optimizing Hyperscale AI Models and Platforms).

## Limitations

Despite the performance gains and stability improvements achieved by Veto, several limitations remain for future investigation:

**Dependence on Teacher Quality** As a knowledge distillation framework, the effectiveness of Veto is inherently bounded by the teacher model’s capabilities. While our method encourages agreement through the geometric target, it cannot generate knowledge entirely absent from the teacher’s distribution.

**Hyperparameter Sensitivity** The parameter  $\beta$  plays a crucial role in balancing stability and diversity. Our experiments show that different tasks (e.g., reasoning vs. summarization) require different initial  $\beta$  values and scheduling strategies to achieve optimal results. Finding an automated way to determine the optimal  $\beta$  without extensive grid search remains an open challenge.

**Limited Data Scale** Our experiments were primarily conducted in low-resource settings, training student models on 1,000 sampled instances for each task. While Veto demonstrates strong performance in these settings, the scalability of the objective-level reformulation to massive datasets and its impact on large-scale pre-training or instruction tuning requires further empirical validation.

**Evaluation Benchmarks** While we validated Veto on reasoning, coding, and summarization, our evaluation did not cover the full spectrum of LLM capabilities, such as multi-turn dialogue, creative writing, or long-context understanding.

## Ethical Considerations

In this work, we introduce Veto, an objective-level reformulation for on-policy knowledge distillation aimed at improving the efficiency and stability of smaller language models. While our method demonstrates significant performance gains across various tasks, we acknowledge the following:

## Potential for Bias Inheritance and Mitigation

Knowledge distillation inherently carries the risk of a student model inheriting biases from its teacher. However, Veto’s unique mechanism—which emphasizes agreement between the teacher and student and suppresses updates on tokens where the student is uncertain—may naturally serve as a soft filter against noisy or unreliable teacher feedback. While this "Veto" effect potentially mitigates the transfer of extreme hallucinations or inconsistent biases compared to standard KD, it is not an explicit de-biasing tool. Therefore, we still recommend caution when deploying the model in high-stakes environments where subtle, systemic biases might persist.

**Use of AI Assistants** In accordance with the ACL Policy on AI Assistance, we acknowledge the use of ChatGPT<sup>1</sup> to assist with code debugging and writing polishing. All experimental designs, data analyses, and scientific claims presented in this work were verified by the authors.

**Dataset Usage and Licensing** Our experiments were conducted using publicly available benchmarks, including GSM8K, HumanEval, and DialogSum. We have ensured that our use of these datasets adheres to their respective licenses.

## References

- Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, and Olivier Bachem. 2024. On-policy distillation of language models: Learning from self-generated mistakes. In *The Twelfth International Conference on Learning Representations*.
- Kushal Arora, Layla El Asri, Hareesh Bahuleyan, and Jackie Chi Kit Cheung. 2022. Why exposure bias matters: An imitation learning perspective of error accumulation in language generation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 700–710.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. *Advances in neural information processing systems*, 28.
- Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541.

<sup>1</sup><https://chatgpt.com/>

- Mark Chen. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021. *DialogSum: A real-life scenario dialogue summarization dataset*. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5062–5074, Online. Association for Computational Linguistics.
- Ting-Rui Chiang and Yun-Nung Chen. 2021. Relating neural text degeneration to exposure bias. *arXiv preprint arXiv:2109.08705*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2023. Minillm: Knowledge distillation of large language models. *arXiv preprint arXiv:2306.08543*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 1317–1327.
- Alexander Lin, Jeremy Wohlwend, Howard Chen, and Tao Lei. 2020. Autoregressive knowledge distillation through imitation learning. *arXiv preprint arXiv:2009.07253*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2024. Wizardcoder: Empowering code large language models with evol-instruct.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Yuqiao Wen, Zichao Li, Wenyu Du, and Lili Mou. 2023. F-divergence minimization for sequence-level knowledge distillation. *arXiv preprint arXiv:2307.15190*.
- Wenda Xu, Rujun Han, Zifeng Wang, Long T Le, Dhruv Madeka, Lei Li, William Yang Wang, Rishabh Agarwal, Chen-Yu Lee, and Tomas Pfister. 2025. Speculative knowledge distillation: Bridging the teacher-student gap through interleaved sampling.
- Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. 2024. A survey on knowledge distillation of large language models. *arXiv preprint arXiv:2402.13116*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 43 others. 2024. *Qwen2 technical report*. *Preprint*, arXiv:2407.10671.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. Bridging the gap between training and inference for neural machine translation. *arXiv preprint arXiv:1906.02448*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.

## A Detailed Mathematical Proofs

### A.1 Proof of Theorem 1 (Adaptive Gradient Veto)

In standard forward KD ( $\beta = 0$ ), the loss is defined as  $\mathcal{L} \approx -c \log P_S(y)$ . As  $P_S(y) \rightarrow 0$ ,  $\log P_S(y) \rightarrow -\infty$ , causing the loss and its gradient to diverge.

In Veto ( $\beta > 0$ ), substituting the PoE target  $Q \propto P_T P_S^\beta$  into the loss gives:

$$\mathcal{L} \approx -\tilde{c} P_S(y)^\beta \log P_S(y) \quad (1)$$

To evaluate the limit as  $P_S(y) \rightarrow 0^+$ , let  $p = P_S(y)$ . Applying L'Hôpital's rule to the form  $\frac{\log p}{p^{-\beta}}$ :

$$\lim_{p \rightarrow 0^+} \frac{\log p}{p^{-\beta}} = \lim_{p \rightarrow 0^+} \frac{1/p}{-\beta p^{-\beta-1}} = \lim_{p \rightarrow 0^+} -\frac{p^\beta}{\beta} = 0 \quad (2)$$

Thus, the loss converges to 0, and the term  $P_S(y)^\beta$  effectively ‘vetoes’ the updates for tokens where the student is ignorant.

### A.2 Proof of Theorem 2 (Sharpening Effect)

The forward KL divergence  $D_{\text{KL}}(Q \| P_S)$  is minimized when the student distribution  $P_S$  matches the target  $Q$ . Substituting the definition of  $Q$  from Eq. (4):

$$\begin{aligned} P_S^*(y | x) &= Q(y | x) \\ P_S^* &\propto P_T \cdot (P_S^*)^\beta \\ (P_S^*)^{1-\beta} &\propto P_T \\ P_S^* &\propto P_T^{\frac{1}{1-\beta}} \end{aligned} \quad (3)$$

For  $0 < \beta < 1$ , the exponent  $\gamma = \frac{1}{1-\beta}$  is greater than 1. This confirms that the student converges to a sharpened version of the teacher distribution.

### A.3 Proof of Theorem 3 (Bridge to REINFORCE)

The reverse KL objective is  $J(\theta) = D_{\text{KL}}(P_S \| Q)$ . Treating  $Q$  as a fixed target for the gradient step:

$$\begin{aligned} \nabla J(\theta) &= \nabla \sum_y P_S(y) (\log P_S(y) - \log Q(y)) \\ &= \sum_y \nabla P_S(y) (\log P_S(y) - \log Q(y)) \\ &\quad + \sum_y P_S(y) \nabla \log P_S(y) \\ &= \mathbb{E}_{y \sim P_S} [\nabla \log P_S(y) (\log P_S(y) \\ &\quad - (\log P_T(y) + \beta \log P_S(y)))] \\ &= \mathbb{E}_{y \sim P_S} [\nabla \log P_S(y) \\ &\quad \cdot ((1 - \beta) \log P_S(y) - \log P_T(y))] \end{aligned} \quad (4)$$

The term  $(1 - \beta) \log P_S(y)$  acts as a scaled entropy cost, while  $\log P_T(y)$  serves as the reward signal.