

Bridging the Gap between Decision and Logits in Decision-based Knowledge Distillation for Pre-trained Language Models

Qinhong Zhou^{1,3}, Zonghan Yang^{1,3}, Peng Li^{2,4,†}, Yang Liu^{1,2,3,4,†}

¹Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua University, Beijing, China

²Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China

³Beijing National Research Center for Information Science and Technology

⁴Shanghai Artificial Intelligence Laboratory, Shanghai, China

Abstract

Conventional knowledge distillation (KD) methods require access to the internal information of teachers, e.g., logits. However, such information may not always be accessible for large pre-trained language models (PLMs). In this work, we focus on decision-based KD for PLMs, where only teacher decisions (i.e., top-1 labels) are accessible. Considering the information gap between logits and decisions, we propose a novel method to estimate logits from the decision distributions. Specifically, decision distributions can be both derived as a function of logits theoretically and estimated with test-time data augmentation empirically. By combining the theoretical and empirical estimations of the decision distributions together, the estimation of logits can be successfully reduced to a simple root-finding problem. Extensive experiments show that our method significantly outperforms strong baselines on both natural language understanding and machine reading comprehension datasets.¹

1 Introduction

Various natural language processing (NLP) tasks have witnessed promising performance from large pre-trained language models (PLMs) (Devlin et al., 2019; Liu et al., 2019; Raffel et al., 2020; Brown et al., 2020). However, PLMs are usually computationally expensive and memory intensive, hindering their deployment on resource-limited devices. Knowledge distillation (KD) (Hinton et al., 2015) is a popular technique to transfer knowledge from large PLMs to lightweight models. Previous KD works utilize various types of internal information from the teacher model, such as output logits (Sanh et al., 2019; Tang et al., 2019; Liu et al., 2020), hidden states (Sun et al., 2019b; Jiao et al., 2020),

and attention maps (Li et al., 2020). In real-world applications, however, these types of information are sometimes not accessible due to commercial and privacy issues (Brown et al., 2020; Ouyang et al., 2022). Specifically, large-scale PLMs usually only provide *decisions* (i.e., top-1 labels) to users. Motivated by this scenario, we investigate the task of *decision-based* KD (Wang, 2021) for PLMs, in which only decisions of teacher predictions are available.

The information gap between teacher decisions and its internal states is the major challenge for the task. A straightforward approach for decision-based KD is to treat teacher decisions as ground truth labels and use these labels to train a student model (Zhang et al., 2022; Sanyal et al., 2022). However, previous work reveals that logits contain rich knowledge (Hinton et al., 2015), relying only on decisions obviously suffers from information loss. To alleviate the problem, Wang (2021) proposes the DB3KD method to generate pseudo soft labels according to the sample’s robustness. However, DB3KD requires that the input of a model can be modified continuously (e.g., image), which hinders its application on PLMs as their inputs are discrete tokens. Therefore, how to fill the information gap under the discrete input setting remains a challenging problem.

Fortunately, the development of test-time data augmentation for discrete input (Liu, 2019; Shleifer, 2019; Xu et al., 2022) brings hope for resolving the challenge. The basic idea is to modify selected tokens in a piece of text under certain constraints to generate augmented samples and estimate or improve the desired properties of a model based on its behaviors on these samples. Test-time data augmentation has been shown to be effective for uncertainty estimation (Ayhan and Berens, 2018; Smith and Gal, 2018; Wang et al., 2019), adversary robustness (Xu et al., 2022), and so on. Is it possible to narrow down the information gap

† Peng Li (lipeng@air.tsinghua.edu.cn) and Yang Liu (liuyang2011@tsinghua.edu.cn) are corresponding authors.

¹Our code is available at <https://github.com/THUNLP-MT/DBKD-PLM>

with test-time data argumentation in decision-based KD for PLMs?

In this work, we propose a novel decision-based KD method for PLMs. As illustrated in Figure 1, our method is capable of estimating the teacher logits for classes even without observed decisions, narrowing down the information gap between decision and logits. Specially, we estimate the logits by combining test-time data argumentation and non-centred orthant probability estimation. On the one hand, we can obtain an empirical estimation of the decision distribution around a sample by test-time data argumentation. On the other hand, we can also derive a theoretical formula for the decision distribution as a non-centred orthant probability, which is a function of logits. As a result, the problem of logits estimation can be reduced to finding the root of the equation that the function takes the value of the empirical estimation. Extensive experiments on various natural language understanding and machine reading comprehension datasets demonstrate the effectiveness of our proposed method, which outperforms strong baselines significantly. Moreover, quantitative analysis reveals that our method obtains better estimation of logits, narrowing down the information gap.

2 Related Work

Decision-based Knowledge Distillation. To advance conventional knowledge distillation (KD) to more challenging black-box model scenarios, Wang (2021) first propose the problem of decision-based KD, where only teacher decisions (i.e., top-1 labels) are accessible to students. They address the problem by estimating the soft label (analogy to output probabilities) of a sample based on its distance to the decision boundary, which involves continuous modification to the original input fed to the black-box model. Instead, Zhang et al. (2022) and Sanyal et al. (2022) synthesize pseudo data in continuous space and leverage decisions of the teacher model on these data directly. In this work, we focus on decision-based KD for PLMs. Unfortunately, the original inputs of the PLMs are discrete tokens which can not be continuously modified. Therefore, these methods are not applicable to our scenario.

Decision-based KD is also related to black-box KD (Orekondy et al., 2019; Wang et al., 2020) and distillation-based black-box attacks (Zhou et al., 2020; Wang et al., 2021; Truong et al., 2021;

Kariyappa et al., 2021; Yu and Sun, 2022). Both of them involve distilling a student model from black-box models. However, these works generally assume that the score-based outputs are accessible. Decision-based KD focuses on a more challenging scenario where only top-1 labels are accessible.

Test-Time Data Augmentation. Test-time data augmentation is a common technique in computer vision (Krizhevsky et al., 2009; Simonyan and Zisserman, 2015; He et al., 2016; Wang et al., 2019; Lyzhov et al., 2020; Shanmugam et al., 2021) and is also feasible for natural language processing (NLP) (Liu, 2019; Shleifer, 2019; Xu et al., 2022). Although differing in final purpose, test-time and training-time data augmentation share a large portion of common techniques in NLP. Due to the discrete nature of language, one line of work conducts augmentation by modifying tokens based on rules (Şahin and Steedman, 2018; Wei and Zou, 2019; Chen et al., 2020a) or models (Sennrich et al., 2016; Yang et al., 2020; Quteineh et al., 2020; Anaby-Tavor et al., 2020), and another line of work operates in embedding or representation space (Chen et al., 2020b; Cheng et al., 2020; Chen et al., 2021; Wei et al., 2022). In this work, as we do not have access to the teacher model, we follow the first line of work to conduct test-time data augmentation.

3 Background

Knowledge Distillation (KD) is a technique that aims to transfer knowledge from the teacher model to the student model by aligning certain statistics, usually the logits, of the student to those of the teacher. Given input x , we denote the pre-softmax logits vector of the teacher and student as \mathbf{z} and \mathbf{v} , respectively. The process of KD involves minimizing the Kullback-Leibler (KL) divergence between the probabilities induced from \mathbf{z} and \mathbf{v} as follows:

$$\mathcal{L}_{\text{KD}} = \text{KL}(\text{softmax}(\mathbf{v}/\tau) || \text{softmax}(\mathbf{z}/\tau)), \quad (1)$$

where τ is the temperature hyper-parameter. The student model is trained by minimizing the loss function

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda \mathcal{L}_{\text{KD}}, \quad (2)$$

where \mathcal{L}_{CE} is the cross entropy loss over the ground-truth label, and λ is the scaling factor used for balancing the importance of the two losses.

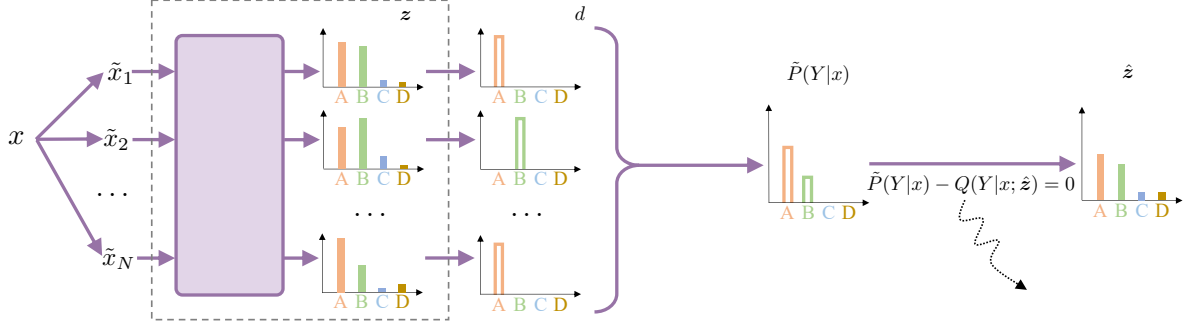


Figure 1: Overview of the proposed framework. For each training sample, the inaccessible teacher logits z contain rich information. However, the teacher model only provides decisions, which lost the information contained in most categories. We bridge this gap in two steps. First, we conduct test-time data augmentation on the original input x and estimate the decision distribution $\tilde{P}(Y|x)$ with teacher decisions d on the augmented data. Then, given the decision distribution $\tilde{P}(Y|x)$, we calculate the logits estimation \hat{z} by solving the equation $\tilde{P}(Y|x) - Q(Y|x; \hat{z}) = 0$, where $Q(Y|x; \hat{z})$ is the theoretical estimation of the decision distribution. As a result, the information contained in original teacher logits z is partly recovered in the estimated logits \hat{z} .

4 Methodology

4.1 Overview

In the decision-based scenario, the z item in Eq. 1 is not accessible. Instead, the PLM API returns the model decision $d = \arg \max_{1 \leq j \leq L} z_j$, where z_j is the j -th logit in z , and L denotes the dimension of output label space. Obviously, d only carries the information of “the j -th logit is the largest”. In contrast, z contains richer information. For example, comparing $z_1 = [0.6, 0.3]$ and $z_2 = [0.9, 0.1]$, although they correspond to the same decision, i.e., $d = 1$, they also imply that the second sample seems more likely to be of the first class. Therefore, there is a big information gap between logits and decisions, and our proposed method aims at narrowing down the gap by finding a better estimation of the logits.

Figure 1 shows the framework of our proposed method. The key idea is combining the empirical and theoretical estimations of the conditional decision distribution $P(Y|x)$, where Y is a random variable denoting decision, to form an equation whose solution is the logits. Specially, first we leverage test-time data augmentation to generate N augmented samples for the given sample x and collect the teacher decisions for them. Then, we build an empirical estimation $\tilde{P}(Y|x)$ for $P(Y|x)$ based on the decisions. Next, we derive a theoretical estimation $Q(Y|x; \hat{z})$ parameterized by the true logits \hat{z} for $P(Y|x)$ and form the following equation

$$\tilde{P}(Y|x) - Q(Y|x; \hat{z}) = 0. \quad (3)$$

Finally, by solving for \hat{z} , we get the estimated

logits and the student model can be trained by conventional KD (Eq. 2).

Compared with the existing decision-based KD method (Wang, 2021), our method leverages data augmentation instead of binary search and optimization on the input samples. Therefore, it is applicable to discrete inputs which can hardly be searched or optimized.

4.2 Empirical Estimation of the Conditional Decision Distribution

In this section, we will introduce how to get $\tilde{P}(Y|x)$ in Eq. 3. Given a sample x and a teacher model \mathcal{M}_θ parameterized by θ , we first generate N augmented samples $\mathcal{X} = \{\tilde{x}_i = F(x, i)\}_{i=1}^N$ with a test-time data augmentation function $F(\cdot, \cdot)$. Then the teacher decisions $\mathcal{D} = \{d_i = \mathcal{M}_\theta(\tilde{x}_i)\}$ are collected. Finally, $P(Y|x)$ is approximated as

$$\tilde{P}(Y|x) \approx \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{d_i}, \quad (4)$$

where $\mathbb{1}_{d_i} \in \{0, 1\}^L$ is a L -dimensional one-hot vector whose d_i -th element is 1, and L is the number of categories.

$F(\cdot, \cdot)$ plays a crucial role in the process and an ideal $F(\cdot, \cdot)$ should satisfy three requirements. First, it should conserve true labels (Wang et al., 2019). Second, it should have a low computational cost, since it will be computed repetitively. Third, the degree of noise introduced by $F(\cdot, \cdot)$ should be quantizable and controllable, crucial for the following steps. Thus, following Wei and Zou (2019), we define $F(\cdot, \cdot)$ as an operation randomly sam-

Algorithm 1: Teacher Logits Estimation

Data: Input text x .

Result: Teacher logits estimation \hat{z}

Require: Teacher model \mathcal{M}_θ , data augmentation transformation function $F(\cdot, \cdot)$, augmented data number N , maximal iteration number m , error bound ϵ , hyper-parameter σ , label number L .

// Empirical estimation of the conditional decision distribution

```
1:  $\{\tilde{x}_n\}_{n=1}^N \leftarrow \{F(x, n)\}_{n=1}^N$ 
2:  $\{d_n\}_{n=1}^N \leftarrow \{\mathcal{M}_\theta(\tilde{x}_n)\}_{n=1}^N$ 
3:  $\tilde{P}(Y|x) \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{d_i}$ 
   // Solving equation to obtain  $\hat{z}$ 
4:  $k \leftarrow 0$ 
5:  $\hat{z} \leftarrow \mathbf{0}$  //  $\hat{z} = [\hat{z}_i, 1 \leq i \leq L]$ 
6: Initialize  $\mathbf{R} \in \mathbb{R}^{(L-1) \times (L-1)}$  whose diagonal elements are  $\sigma^2$  and the others are  $2\sigma^2$ 
7: repeat
8:    $\mathbf{p} \leftarrow \mathbf{0}$  //  $\mathbf{p} = [p_i, 1 \leq i \leq L]$ 
9:   for  $i = 1, i \leq L$  do
10:     $\boldsymbol{\mu} \leftarrow [\hat{z}_i - \hat{z}_1, \dots, \hat{z}_i - \hat{z}_{i-1}, \hat{z}_i - \hat{z}_{i+1}, \dots, \hat{z}_i - \hat{z}_L]$ 
11:     $\mathbf{B} \leftarrow \text{CholeskyDecompose}(\mathbf{R})$ 
12:     $p_i \leftarrow \text{RecursiveIntegration}(\boldsymbol{\mu}, \mathbf{B})$ 
13:     $i \leftarrow i + 1$ 
14:   end for
15:    $k \leftarrow k + 1$ 
16:    $\hat{z} \leftarrow \tilde{P}(Y|x) - \mathbf{p} + \hat{z}$ 
17: until  $|\mathbf{p} - \tilde{P}(Y|x)| \leq \epsilon$  or  $k = m$ 
18: return  $\hat{z}$ 
```

pled from synonym replacement, random insertion, random swap, and random deletion operations.

4.3 Theoretical Estimation of the Conditional Decision Distribution

In this section, we will introduce how to get the theoretical estimation $Q(Y|x; \hat{z})$ parameterized by \hat{z} in Eq. 3. The outline of the derivation is that we assume the logits are sampled from an L -dimensional distribution $P(\mathbf{Z}|x)$, where $\mathbf{Z} = [Z_j]$ is an L -dimensional random variable denoting logits. Then $Q(Y = i|x; \hat{z})$ is equal to the probability that the i -th dimension of \mathbf{Z} takes the largest value, which can be calculated mathematically from $P(\mathbf{Z}|x)$.

Following the above outline, we have

$$Q(Y = i|x; \hat{z}) = P\left(Z_i = \max_{1 \leq j \leq L} Z_j | x\right). \quad (5)$$

To derive the above probability, we reformulate it

in terms of orthant probability. First, we introduce an $L - 1$ dimensional auxiliary random variable $\mathbf{U} = [U_j]$, which is defined as

$$U_j = \begin{cases} Z_i - Z_j & (j < i) \\ Z_i - Z_{j+1} & (j \geq i) \end{cases}, 1 \leq j \leq L - 1. \quad (6)$$

Note that the i -th dimension of \mathbf{Z} is eliminated due to $Z_i - Z_i = 0$. Then Eq. 5 can be rewritten as a non-centred orthant probability distribution

$$Q(Y = i|x; \hat{z}) = P(U_j \geq 0, 1 \leq j \leq L - 1). \quad (7)$$

To simplify the calculation of the probability in Eq. 7, we assume \mathbf{Z} follows a multivariate Gaussian distribution with mean \hat{z} and covariance matrix $\boldsymbol{\Sigma}$, i.e., $\mathbf{Z} \sim \mathcal{N}(\hat{z}, \boldsymbol{\Sigma})$. Then we have $\mathbf{U} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{R})$, where

$$\boldsymbol{\mu}_j = \begin{cases} \hat{z}_i - \hat{z}_j & (j < i) \\ \hat{z}_i - \hat{z}_{j+1} & (j \geq i) \end{cases}, 1 \leq j \leq L - 1. \quad (8)$$

And Eq. 7 can be calculated by the following multiple integrations

$$\int_0^{+\infty} \dots \int_0^{+\infty} \phi_{L-1}(\mathbf{U}; \boldsymbol{\mu}, \mathbf{R}) dU_1 \dots dU_{L-1}, \quad (9)$$

where $\phi_{L-1}(\mathbf{U}; \boldsymbol{\mu}, \mathbf{R})$ is the probability density function.

We leverage the recursive algorithm proposed by Miwa et al. (2003) to solve the above integrations. Taking $L = 4$ as an example, the major steps of the algorithm are as follows. First, we decompose the covariance matrix \mathbf{R} as $\mathbf{R} = \mathbf{B}\mathbf{B}^T$ via Cholesky decomposition, where \mathbf{B} is a lower triangular matrix. Then we have $\mathbf{U} = \mathbf{B}\mathbf{M} + \boldsymbol{\mu}$, where $\mathbf{M} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{L-1})$ and \mathbf{I}_{L-1} is an identity matrix of dimension $L - 1$. Next, $Q(Y = i|x; \hat{z})$ can be further decomposed as

$$\begin{aligned} Q(Y = i|x; \hat{z}) &= P(U_j \geq 0, 1 \leq j \leq 3) \\ &= P(b_{11}M_1 + \mu_1 \geq 0, \\ &\quad b_{21}M_1 + b_{22}M_2 + \mu_2 \geq 0, \\ &\quad b_{31}M_1 + b_{32}M_2 + b_{33}M_3 + \mu_3 \geq 0), \end{aligned} \quad (10)$$

where b_{ij} denotes the elements in the \mathbf{B} matrix, M_j denotes the j -th elements of the random variable \mathbf{M} , and μ_j denotes j -th the elements of $\boldsymbol{\mu}$. Finally, the required probability is given when $b_{ij} > 0$

$$Q(Y = i|x; \hat{z}) = \int_{-\frac{\mu_1}{b_{11}}}^{+\infty} f_1(t)\phi(t)dt, \quad (11)$$

where $\phi(t)$ is the standard normal probability density function, and f_1 is defined as

$$f_1(s) = \int_{\frac{-\mu_2 - b_{21}s}{b_{22}}}^{+\infty} f_2(s, t) \phi(t) dt, \quad (12)$$

$$f_2(s_1, s_2) = \int_{\frac{-\mu_3 - b_{31}s_1 - b_{32}s_2}{b_{33}}}^{+\infty} \phi(t) dt. \quad (13)$$

Algorithm 1 summarizes the entire procedure of our proposed framework, where line 12 refers to the integration steps in Eq. 11 to 13. We provide the proofs of the integration steps in Appendix A.4. In practice, we assume Σ is a diagonal matrix and $\Sigma_{ii} = \sigma^2$ to simplify the calculation, where σ is a hyper-parameter of our algorithm.

5 Experiments

5.1 Experimental Settings

Datasets and Evaluation Metrics. We evaluate our method on machine reading comprehension (MRC) and natural language understanding (NLU) datasets. For MRC, two widely used multiple-choice datasets RACE (Lai et al., 2017) and DREAM (Sun et al., 2019a) are used. For NLU, we select sentiment analysis dataset SST-2 (Socher et al., 2013), linguistic acceptability dataset CoLA (Warstadt et al., 2019), paraphrasing dataset MRPC (Dolan and Brockett, 2005) and QQP (Chen et al., 2017), and natural language inference (NLI) datasets RTE (Bentivogli et al., 2009), MNLI (Williams et al., 2018), and QNLI (Rajpurkar et al., 2016) as representative datasets. Following previous works (Lai et al., 2017; Sun et al., 2019a; Wang et al., 2018), we report Matthews correlation coefficient for CoLA, F1 and accuracy for MRPC and QQP, and accuracy for all the other datasets. For each experiment, the model is evaluated on the validation set once an epoch, and the checkpoint achieving the best validation results is evaluated on the test set. The results averaged over five random seeds are reported for the MRC datasets. Due to the submission quota, we only report results for one trial for the NLU tasks.

Baselines. We compare our method with the following four baselines:

- *Hard*: We regard the teacher decisions as the ground truth labels and train the student model solely with the cross entropy loss.
- *Noisy Logits* (Wang, 2021): The student model is trained via the KD objective (Eq. 2)

with the teacher logits replaced with randomly sampled soft labels.

- *Smooth*: We apply label smoothing (Szegedy et al., 2016) with a smoothing factor 0.1 on the teacher decision of the original sample, and use the smoothed decision as teacher prediction probability. This a straightforward approach to generate soft labels from teacher decisions.

To better investigate the upper bound of our method, we also leverage the following three baselines from Wang (2021):

- *Student CE* (Wang, 2021): The student model is trained using only the cross entropy loss calculated from the ground-truth labels.
- *Standard KD*: The student model is optimized with the standard KD objective (Eq. 2). Note that the teacher model is used as a white-box model in this baseline.
- *Surrogate*: Following Wang (2021), we train the student model via KD with a surrogate teacher, simulating training a lightweight, white-box teacher model for knowledge distillation.

Implementation Details. We implement the teacher model as the *finetuned* 12-layer BERT model (BERT_{BASE}) or 24-layer BERT model (BERT_{LARGE}) for each task. The student model is a 4-layer or 6-layer BERT-style model. Following previous KD works (Sun et al., 2019b; Li et al., 2021b), we initialize the student model from the *raw* 12-layer BERT model. We adopt EDA (Wei and Zou, 2019) as the tool for test-time data augmentation.

5.2 Results on MRC Datasets

Experimental results on the MRC datasets RACE and DREAM are shown in Table 1 and Table 2, respectively. In each table, we report three sets of results with different teacher and student architectures. For example, the string “12L→4L” in the tables means we leverage the finetuned 12-layer BERT model (BERT_{BASE}) as the teacher, and the 4-layer BERT-style model as the student. Note that *Teacher* and *Standard KD* serve as the upper bounds of our method. Therefore, we do not directly compare our method with them. From these results, we can observe that:

Methods	DB	12L→4L			24L→4L			12L→6L		
		Middle	High	All	Middle	High	All	Middle	High	All
Teacher		68.25	61.21	66.74	71.73	64.29	69.94	68.25	61.21	66.74
Standard KD		54.12	48.95	53.08	53.12	50.24	53.45	61.69	53.74	59.28
Student CE		51.03	47.25	49.94	51.03	47.25	49.94	59.11	51.35	56.80
Surrogate		51.28	47.90	<u>51.03</u>	51.28	<u>47.90</u>	<u>51.03</u>	<u>59.89</u>	50.81	56.49
Hard	✓	51.60	<u>47.93</u>	50.29	<u>51.59</u>	46.93	50.78	59.64	51.63	<u>56.98</u>
Noisy Logits	✓	50.57	46.30	49.86	<u>50.57</u>	46.30	49.86	58.97	51.12	<u>55.60</u>
Smooth	✓	<u>51.71</u>	46.93	49.56	50.68	46.80	50.31	59.68	<u>51.71</u>	56.30
Ours	✓	52.81	48.89	52.17	52.98	49.06	52.08	61.10	52.81	58.01

Table 1: Results on the test sets of RACE. Decision-based (DB) methods are checked in the second column. BERT-style architectures are used for the teacher and student models. And the string “12L→4L” denotes the teacher has 12 layers and the student has 4 layers. The best results besides those of *Teacher* and *Standard KD* are in **bold**, and the second best ones are underlined. “Teacher” denotes the performance of the teacher model, respectively. The reported results are the average of five different random seeds.

Methods	DB	12L→4L	24L→4L	12L→6L
Teacher		60.44	61.69	60.44
Standard KD		51.89	53.37	54.97
Student CE		51.00	51.00	53.62
Surrogate		51.11	51.11	53.57
Hard	✓	50.09	49.98	51.31
Noisy Logits	✓	<u>51.13</u>	51.13	<u>53.78</u>
Smooth	✓	51.04	<u>52.03</u>	53.43
Ours	✓	51.64	52.74	54.50

Table 2: Results on the test set of DREAM. The reported results are the average of five different random seeds.

(1) Our proposed method outperforms baselines consistently and significantly. The performance gap between our method and the second best baselines except *Teacher* and *Standard KD* are from 0.96 to 1.39 on the RACE datasets and from 0.51 to 0.72 on the DREAM dataset, indicating that narrowing down the information gap between logits and decisions are effective for decision-based KD.

(2) Surprisingly, our proposed method achieves comparable results with *Standard KD* under a few settings. *Standard KD* treats the teacher as a white-box model and is an intuitive upper bound of our method. However, the smallest performance gap between our method and *Standard KD* is 0.06 (12L→4L on RACE-High) on the RACE datasets and is 0.25 (12L→4L) on the DREAM dataset. Moreover, among all the twelve pairs of results, there is a third of them with a gap of less than 0.50. These results further justify the effectiveness of our proposed method. And we argue that this is mainly due to our better estimation of the logits.

(3) None of the baselines besides *Standard KD* can consistently achieve better results than training

the student model without KD (*Student CE*), indicating that decision-based KD is a challenging task and the lost information from decisions compared with logits is essential. *Hard* performs slightly better than *Student CE* on the RACE datasets but significantly worse than *Student CE* on the DREAM dataset. We conjecture that this is because the teacher models have significantly better results on the RACE datasets than on the DREAM dataset, i.e., *Hard* can only work well with strong teacher models, whose decisions may be less noisy and the information gap between decisions and logits is smaller. Our method can be viewed as a special form of logits smoothing. However, both *Noisy Logits* and *Smooth* only achieves comparable or worse results than *Student CE*, indicating straightforward logits smoothing is not effective.

(4) Our method benefits from both better teacher models and larger student models. When the teacher grows larger (12L→4L v.s. 24L→4L), our method achieves a 1.10 performance gain on the DREAM dataset. Meanwhile, when the student models grow from 4 to 6 layers (12L→4L v.s. 12L→6L), the performance gains on all datasets are remarkably larger. The same trend is also observed for other baselines, suggesting that improving the capacity of the student model is a simple yet effective way to improve the performance of decision-based KD.

5.3 Results on NLU Datasets

Table 3 shows the results on the NLU datasets. First, our proposed method achieves the best results among all decision-based baselines, justifying that our method is generalizable to a large range of NLU

Methods	DB	RTE (Acc.)	MRPC (F1 / Acc.)	CoLA (Matt.)	QNLI (Acc.)	SST-2 (Acc.)	MNLI-m / mm (Acc.)	QQP (F1 / Acc.)	Average
Teacher		66.2	87.3 / 82.3	53.7	90.9	93.6	84.4 / 83.5	71.4 / 89.1	79.1
Standard KD		63.3	82.9 / 75.0	22.7	85.6	89.6	78.8 / 77.6	69.1 / 88.0	71.0
Student CE		63.2	81.2 / 69.8	21.5	85.2	89.2	78.4 / 76.7	67.5 / 87.3	69.9
Surrogate		63.6	<u>82.5</u> / 74.8	18.9	85.2	<u>89.4</u>	<u>78.4</u> / 77.3	67.8 / 87.4	70.2
Hard	✓	63.2	<u>82.5</u> / 74.8	20.7	<u>85.6</u>	89.2	<u>78.1</u> / <u>77.2</u>	<u>68.2</u> / <u>87.6</u>	70.4
Noisy Logits	✓	63.3	81.6 / 74.0	21.8	85.3	88.5	78.1 / 76.7	67.5 / 87.4	70.2
Smooth	✓	<u>63.4</u>	<u>82.4</u> / <u>75.1</u>	<u>22.2</u>	85.1	89.2	78.0 / 77.0	67.9 / 87.6	<u>70.6</u>
Ours	✓	<u>63.4</u>	82.9 / 75.2	23.7	85.7	89.5	78.6 / 77.1	68.5 / 88.0	71.1

Table 3: Results from the GLUE test server. We show the evaluation metrics under the task names. We use BERT_{BASE} model as the teacher and a 4-layer BERT-style model as the student.

tasks. Although our method does not outperform *Surrogate* on the RTE and MNLI-mm datasets, the gap is only 0.2. Second, the performance gap between ours and *Standard KD* is also small, providing extra evidence that our method estimates the teacher logits well. Third, all the baselines excluding *Teacher* and *Standard KD* have comparable performance, suggesting that decision-based KD is also challenging for NLU tasks. Above all, in conjunction with the results on the MRC datasets, we can conclude that our method is effective for diverse tasks and model architectures.

5.4 Analysis on Logits Estimation

We have conjectured that the good performance of our method comes from better logits estimation. To justify this assumption, we conduct a quantitative analysis in this section. We compute the mean squared errors (MSEs) between the soft labels generated from each method after softmax and teacher predictions on the training set of RACE-High². As shown in Figure 2, the soft labels generated by our method are the closest to the teacher predictions among all methods. However, it should be noted that the probabilities (or logits) of the teacher are not perfect, as it does not achieve perfect final performance on the dataset. Therefore, the MSEs have a positive correlation with the final performance but are not oracle indicators.

5.5 Ablation Study

This section consists of a series of experiments aimed at validating the contributions of different components in our method. First, we compare our method with its two variants in Figure 3: (1) *w/o*

²As *Hard* produces zero probabilities, Kullback–Leibler divergence is not applicable. A temperature (τ in Eq. 1) of 10 is leveraged when needed.

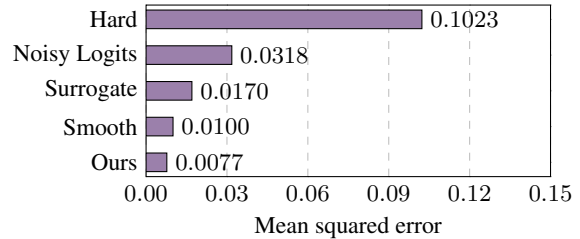


Figure 2: Mean squared errors (MSEs) between the estimated probabilities and the ground-truth probability produced by the teacher model on the training set of RACE-High.

Empirical Estimations. In this variant, we replace the $\tilde{P}(Y|x)$ in Eq. 4 with the teacher decision on original data to skip the empirical estimation step. (2) *w/o Theoretical Estimation.* We replace the $Q(Y|x; \hat{z})$ term in Eq. 3 with $\text{softmax}(z)$ to skip the theoretical estimation step. For each dataset, we also count the percentage of empirical estimations $\tilde{P}(Y|x)$ being one-hot vectors, which means that teacher decisions are consistent on augmented inputs. According to the results, we find that the performance drops for both variants on all datasets, indicating the necessity of empirical estimation and theoretical estimation. Interestingly, we also find a positive correlation between the percentage of one-hot $\tilde{P}(Y|x)$ and the performance degradation from *w/o Theoretical Estimation* variant. This phenomenon highlights the capability of the theoretical estimation step to estimate teacher logits and narrow the information gap between decisions and logits even without observed decisions.

Second, we further analyze the effect of empirical estimation by changing the sampling times N in Eq. 4. As shown in Figure 4, when N increases, the performance of our method first increases and then stabilizes. Considering a larger N leads to

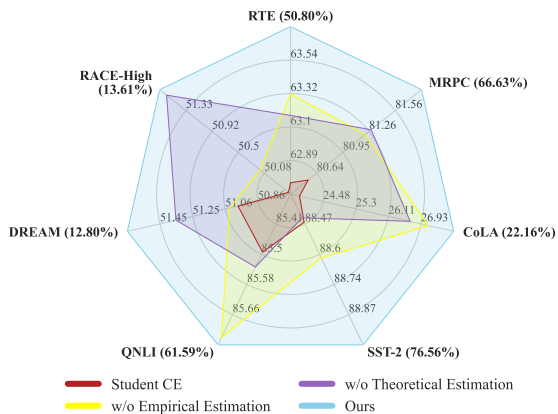


Figure 3: Ablation studies for the empirical estimation step (Section 4.2) and the theoretical estimation step (Section 4.3) on the dev set of each dataset. We also include the results from *Student CE* method for comparison. The percentages after dataset names denote the ratio of empirical estimations $\hat{P}(Y|x)$ being one-hot vectors.

more queries to the teacher model, N should be as low as possible without compromising the model performance. Therefore, $N = 10$ is the optimal choice according to the results.

Finally, we investigate the contribution of Eq. 3, which combines the empirical estimation and the theoretical estimation together. In our framework, the root \hat{z} of the equation is found by fixed-point iteration, and its precision is controlled by the error bound ϵ . Results in Figure 5 show a negative correlation between ϵ and KD performance. As ϵ increases from 10^{-4} to 10^{-1} , the performance of our method slightly drops. When ϵ increases to 1, which means the logits estimation becomes extremely inaccurate, the performance drops dramatically and is close to the performance of *Student CE* method.

5.6 Computational Cost Analysis

The additional computational cost of our method compared to *Standard KD* consists of two parts. The first part is test-time data augmentation, which necessitates multiple queries to the teacher model for each training sample. In this paper, we set the default number of augmented samples N per training case to 10. The second part is solving Eq. 3 using the empirical estimation of decision distribution $\tilde{P}(Y|x)$, which is made negligible by pre-building a lookup table from $\tilde{P}(Y|x)$ to logits estimation \hat{z} before KD. In total, the additional cost of our method mainly comes from 10 queries made to the teacher model per training sample. By

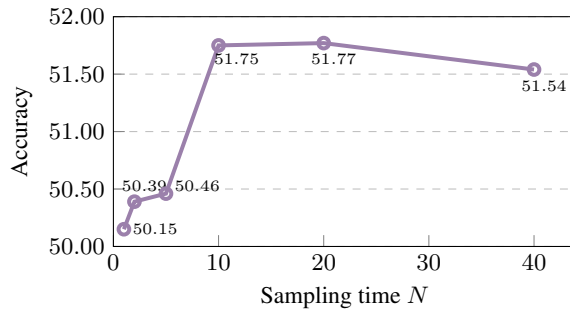


Figure 4: Comparison of different sampling times N of test-time data augmentation in the empirical estimation step. Averaged results on the dev set of RACE-High over 5 different random seeds are reported.

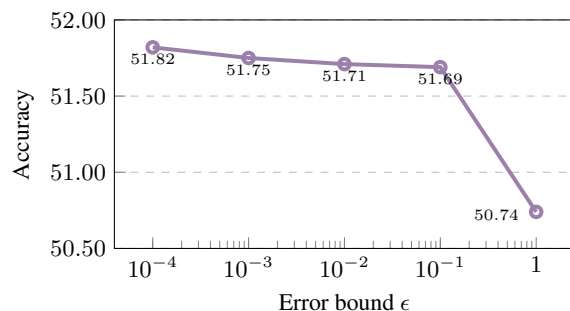


Figure 5: Comparison of different error bounds ϵ of teacher logits estimation in the theoretical derivation step. Averaged results on the dev set of RACE-High over 5 different random seeds are reported.

contrast, the existing soft label generation method DB3KD (Wang, 2021) requires 1,000 to 20,000 queries to the teacher model per training sample.

6 Conclusion

We introduce a novel decision-based KD method, which bridges the information gap between teacher decisions and logits by estimating teacher logits. In contrast to existing solutions for decision-based KD, our method is applicable to NLP tasks with discrete inputs. Extensive experiments over various tasks and model architectures demonstrate the effectiveness of our proposed method.

One future direction for the decision-based KD is the exploration of other NLP tasks, such as neural machine translation, text generation, and question answering. The other direction is KD from non-NN models to NN models, which benefits the training of NN models with additional information from a wider range of models. Unlike conventional KD, decision-based KD does not require internal information from NN models and is promising for solving this problem.

Limitations

This study has two main limitations. The first limitation is its reliance on the assumption that teacher logits on augmented data follow a Gaussian distribution. This assumption is used in the derivation of teacher logits in Section 4.3. However, in practice, teacher logits may not strictly follow a Gaussian distribution. It is challenging to estimate teacher logits under more realistic assumptions, which requires thorough investigations on the distribution of teacher logits and more complex computations for logits estimation.

The second limitation is that our method still requires access to the training dataset of the downstream tasks. In this paper, we focus on KD when teacher PLMs only return decisions. However, our method is not capable of KD without publicly available training data, which is a more challenging scenario for decision-based KD. We believe training a data generation model (Wang, 2021; Zhang et al., 2022; Sanyal et al., 2022) might be useful for such cases.

Ethics Statement

In ethical considerations, our method risks being used as a means of model stealing. Therefore, defensive techniques against the proposed method are required. However, it also has significant positive implications. On the one hand, it can serve as a powerful tool for research on model extraction attacks, thereby promoting the advancement of related studies. On the other hand, it has practical applications in real-world scenarios. For instance, a company may prefer to use a smaller model due to cost considerations, and our method allows for the easy distillation of smaller models without requiring white-box access to larger models. Additionally, our method can be used to distill non-NN models into NN models, reducing the number of model types that need to be maintained and simplifying operation and maintenance.

Acknowledgement

This work is supported by the National Key R&D Program of China (2022ZD0160502) and the National Natural Science Foundation of China (No. 61925601, 62276152, 62236011). We thank all anonymous reviewers for their valuable comments and suggestions on this work. We also thank Shuo Wang and Xiaoyue Mi for their suggestions on the writing.

References

- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. Do not have enough data? deep learning to the rescue! In *AAAI 2020*.
- Murat Seckin Ayhan and Philipp Berens. 2018. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. In *MIDL 2018*.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2009. The fifth PASCAL recognizing textual entailment challenge. In *TAC 2009*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Matiusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *NeurIPS 2020*.
- Guandan Chen, Kai Fan, Kaibo Zhang, Boxing Chen, and Zhongqiang Huang. 2021. Manifold adversarial augmentation for neural machine translation. In *Findings of the ACL 2021*.
- Hannah Chen, Yangfeng Ji, and David Evans. 2020a. Finding Friends and flipping frenemies: Automatic paraphrase dataset augmentation using graph theory. In *Findings of the EMNLP 2020*.
- Jiaao Chen, Zichao Yang, and Diyi Yang. 2020b. Mix-Text: Linguistically-informed interpolation of hidden space for semi-supervised text classification. In *ACL 2020*.
- Zihang Chen, Hongbo Zhang, Xiaoji Zhang, and Leqi Zhao. 2017. Quora question pairs.
- Yong Cheng, Lu Jiang, Wolfgang Macherey, and Jacob Eisenstein. 2020. AdvAug: Robust adversarial augmentation for neural machine translation. In *ACL 2020*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL 2019*.
- Bill Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *IWP 2005*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR 2016*.

- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. TinyBERT: Distilling BERT for natural language understanding. In *Findings of EMNLP 2020*.
- Sanjay Kariyappa, Atul Prakash, and Moinuddin K Qureshi. 2021. MAZE: Data-free model stealing attack using zeroth-order gradient estimation. In *CVPR 2021*.
- Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. RACE: Large-scale reading comprehension dataset from examinations. In *EMNLP 2017*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. 2021. Datasets: A community library for natural language processing. In *EMNLP 2021: System Demonstrations*.
- Jianquan Li, Xiaokang Liu, Honghong Zhao, Ruifeng Xu, Min Yang, and Yaohong Jin. 2020. BERT-EMD: Many-to-many layer mapping for BERT compression with earth mover’s distance. In *EMNLP 2020*.
- Lei Li, Yankai Lin, Deli Chen, Shuhuai Ren, Peng Li, Jie Zhou, and Xu Sun. 2021a. CascadeBERT: Accelerating inference of pre-trained language models via calibrated complete models cascade. In *Findings of EMNLP 2021*.
- Lei Li, Yankai Lin, Shuhuai Ren, Peng Li, Jie Zhou, and Xu Sun. 2021b. Dynamic knowledge distillation for pre-trained language models. In *EMNLP 2021*.
- Bo Liu. 2019. Anonymized BERT: An augmentation approach to the gendered pronoun resolution challenge. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*.
- Weijie Liu, Peng Zhou, Zhiruo Wang, Zhe Zhao, Haotang Deng, and Qi Ju. 2020. FastBERT: a self-distilling BERT with adaptive inference time. In *ACL 2020*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach.
- Alexander Lyzhov, Yuliya Molchanova, Arsenii Ashukha, Dmitry Molchanov, and Dmitry Vetrov. 2020. Greedy policy search: A simple baseline for learnable test-time augmentation. In *Conference on Uncertainty in Artificial Intelligence*. PMLR.
- George A. Miller. 1995. Wordnet: A lexical database for english. *Commun. ACM*.
- Tetsuhisa Miwa, AJ Hayter, and Satoshi Kuriki. 2003. The evaluation of general non-centred orthant probabilities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. 2019. Knockoff nets: Stealing functionality of black-box models. In *CVPR 2019*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback.
- Husam Quteineh, Spyridon Samothrakis, and Richard Sutcliffe. 2020. Textual data augmentation for efficient active learning on tiny datasets. In *EMNLP 2020*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100, 000+ questions for machine comprehension of text. In *EMNLP 2016*.
- Gözde Gül Şahin and Mark Steedman. 2018. Data augmentation via dependency tree morphing for low-resource languages. In *EMNLP 2018*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. In *NeurIPS 2019 Workshop on Energy Efficient Machine Learning and Cognitive Computing*.
- Sunandini Sanyal, Sravanti Addepalli, and R Venkatesh Babu. 2022. Towards data-free model stealing in a hard label setting. In *CVPR 2022*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *ACL 2016*.
- Divya Shanmugam, Davis Blalock, Guha Balakrishnan, and John Gutttag. 2021. Better aggregation in test-time augmentation. In *ICCV 2021*.
- Sam Shleifer. 2019. Low resource text classification with ulmfit and backtranslation.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *ICLR 2015*.

- Lewis Smith and Yarin Gal. 2018. Understanding measures of uncertainty for adversarial example detection.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP 2013*.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019a. Dream: A challenge data set and models for dialogue-based reading comprehension. *TACL 2019*.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019b. Patient knowledge distillation for bert model compression.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *CVPR 2016*.
- Raphael Tang, Yao Lu, Linqing Liu, Lili Mou, Olga Vechtomova, and Jimmy Lin. 2019. Distilling task-specific knowledge from bert into simple neural networks.
- Jean-Baptiste Truong, Pratyush Maini, Robert J. Walls, and Nicolas Papernot. 2021. Data-free model extraction. In *CVPR 2021*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *EMNLP 2018 Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- Dongdong Wang, Yandong Li, Liqiang Wang, and Boqing Gong. 2020. Neural networks are more productive teachers than human raters: Active mixup for data-efficient knowledge distillation from a blackbox model. In *CVPR 2020*.
- Guotai Wang, Wenqi Li, Michael Aertsen, Jan Deprest, Sébastien Ourselin, and Tom Vercauteren. 2019. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing*.
- Wenxuan Wang, Bangjie Yin, Taiping Yao, Li Zhang, Yanwei Fu, Shouhong Ding, Jilin Li, Feiyue Huang, and Xiangyang Xue. 2021. Delving into data: Effectively substitute training for black-box attack. In *CVPR 2021*.
- Zi Wang. 2021. Zero-shot knowledge distillation from a decision-based black-box model. In *ICML 2021*.
- Alex Warstadt, Amanpreet Singh, and Samuel Bowman. 2019. Neural network acceptability judgments. *TACL 2019*.
- Jason Wei and Kai Zou. 2019. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. In *EMNLP 2019*.
- Xiangpeng Wei, Heng Yu, Yue Hu, Rongxiang Weng, Weihua Luo, and Rong Jin. 2022. Learning to generalize to more: Continuous semantic augmentation for neural machine translation. In *ACL 2022*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL 2018*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *EMNLP 2020: System Demonstrations*.
- Lei Xu, Laure Berti-Equille, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2022. In situ augmentation for defending against adversarial attacks on text classifiers. In *KDD 2022 Workshop on Adversarial Learning Methods for Machine Learning and Data Mining*.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. Generative data augmentation for common-sense reasoning. In *Findings of the EMNLP 2020*.
- Mengran Yu and Shiliang Sun. 2022. FE-DaST: Fast and effective data-free substitute training for black-box adversarial attacks. *Computers & Security*.
- Jie Zhang, Chen Chen, Jiahua Dong, Ruoxi Jia, and Lingjuan Lyu. 2022. QEKD: Query-efficient and data-free knowledge distillation from black-box models.
- Mingyi Zhou, Jing Wu, Yipeng Liu, Shuaicheng Liu, and Ce Zhu. 2020. DaST: Data-free substitute training for adversarial attacks. In *CVPR 2020*.

A Appendix

A.1 Experiment Details

Dataset Details In this paper, we use seven different datasets, and all of them are in the English language. We downloaded these datasets from the Datasets (Lhoest et al., 2021) library of version 2.4.0, and our use is consistent with their intended use. The other details of the datasets we used are summarized in Table 4.

Name	Number of train / dev / test	License	Domain
RACE-All (Lai et al., 2017)	87,866 / 4,887 / 4,934	unknown	examinations
RACE-Middle (Lai et al., 2017)	25,421 / 1,436 / 1,436		
RACE-High (Lai et al., 2017)	62,445 / 3,451 / 3,498		
DREAM (Sun et al., 2019a)	6,116 / 2,040 / 2,041	unknown	dialogue
RTE (Bentivogli et al., 2009)	2,490 / 277 / 3,000	CC-BY-4.0	news, Wikipedia
MRPC (Dolan and Brockett, 2005)	3,668 / 408 / 1,725		news
CoLA (Warstadt et al., 2019)	8,551 / 1,043 / 1,063		misc.
SST-2 (Socher et al., 2013)	67,349 / 872 / 1,821		movie reviews
QNLI (Rajpurkar et al., 2016)	104,743 / 5,463 / 5,463		Wikipedia

Table 4: The detailed information of datasets we used in this paper.

Model Details We used BERT-like models (Devlin et al., 2019) in our experiments, including BERT_{BASE} (110M parameters), BERT_{LARGE} (340M parameters), 4-layer BERT-like models (53M parameters), and 6-layer BERT-like models (67M parameters). For BERT_{BASE} and BERT_{LARGE}, the raw model checkpoints are obtained from Huggingface Transformers (Wolf et al., 2020) platform. Following Li et al. (2021b), we initialize the 4-layer and 6-layer BERT-like models from the first 4 and 6 layers of the raw BERT_{BASE} model, respectively.

Other Details We finetune the BERT_{BASE} and BERT_{LARGE} models for 4 epochs. Following Li et al. (2021a), we train small 4-layer or 6-layer models for 10 epochs. We use a learning rate of 5×10^{-5} for MRC tasks and 2×10^{-5} for NLU tasks. σ in Algorithm 1 is tuned from $\{0.5, 1, 2, 4\}$, λ in Eq. 2 is tuned from $\{0.2, 0.5, 0.7\}$, and τ in Eq. 1 is tuned from $\{5, 10, 20\}$ expect for our method, which performs better in the range of $\{1, 2, 4\}$. The α parameters of all operations in EDA are sampled from a half-normal distribution, and we adjust the scale of the distribution to align its expectation with the default $\alpha = 0.1$ in EDA. For MRC tasks, following Sun et al. (2019b), we concatenate the input passage and the question with a [SEP] token and append each answer at the end of the question. The random seeds we used for experiments on MRC datasets and ablation studies on NLU datasets are from 1 to 5. The random seed for teacher training and other NLU experiments is 1. The training of a 4-layer student model on one RTX 3090 Ti GPU costs approximately 6.5 hours for our method.

Methods	RACE-High
Student CE	39.79
Noisy Logits	30.19
Surrogate	37.86
Smooth	41.59
Hard	42.92
Ours	44.13

Table 5: Averaged results of generative models over 5 different random seeds. All methods are evaluated on the RACE-High dev set.

A.2 Experimental Results on Generative Language Models

Theoretically, our method can be applied to generative LMs. In this paper, we evaluate the effectiveness of our method on the RACE dataset. We finetune a 12-layer GPT-2 (Radford et al., 2019) teacher to predict the class label (A/B/C/D) given the context, question, and options as prompt. Then we distill the teacher model to a 4-layer GPT-2 student. For the student model, the output vocabulary at the answer position is restricted to class label tokens. Table 5 shows the performance of our method and baseline methods on the dev set of RACE-High. Our method significantly outperforms decision-based baseline methods and the student CE method.

Different from classification tasks, our method will require much more queries to the teacher model on generation tasks because of the following reasons:

1. The label space dimension L in generative tasks is equal to the vocabulary size, which is quite large. The computational cost of building the look-up table will increase.

- For each position i in a sequence, our method estimates the logits in the i -th position given all the tokens before the i -th position. Therefore, to estimate the logits in the entire sequence, we need to sample teacher decisions in each position. As a result, the computational cost will multiply by the sequence length.

Therefore, how to improve the efficiency of applying our method to generation tasks is an interesting future research direction.

A.3 Detailed Analysis of the Data Augmentation Techniques

In this paper, we use the EDA augmentation tool (Wei and Zou, 2019) for each sample, including its α parameter and the following four default augmentation techniques. Given an input sentence, a α parameter, and the sentence length l_{sent} , the four techniques can be describe as following:

- Synonym Replacement:** First, select αl_{sent} words that are not stop words randomly. Second, replace each word with a random WordNet (Miller, 1995) synonym of itself.
- Random Insertion:** First, select a word in the sentence that is not a stop word randomly. Second, find a random synonym of the selected word. Third, insert the synonym into a random position in the sentence. Finally, do the above steps αl_{sent} times.
- Random Swap:** First, choose two words in the sentence randomly. Second, swap the positions of the chosen words. Finally, do the above steps αl_{sent} times.
- Random Deletion:** Remove each word in the sentence with probability α randomly.

In Table 6, we provide further ablation studies on each technique of EDA. We also include the performance of *Surrogate* method which is the best decision-based baseline. According to the results, our method is robust to different data augmentation techniques.

A.4 Proofs

In this section, we provide the detailed proofs of Eq. 11 to 13.

Methods	RACE-High
Ours	51.75
w/o synonym replacement	51.69
w/o random insertion	51.44
w/o random swap	51.75
w/o random deletion	51.60
Surrogate (best baseline)	50.33

Table 6: Averaged results of ablation methods over 5 different random seeds. For each ablation, we remove one of the four data augmentation techniques from EDA and evaluate it on the RACE-High dev set.

In Section 4.3, we assume Σ is a diagonal matrix and $\Sigma_{ii} = \sigma^2$. Therefore, the diagonal elements of B are positive, and Eq. 10 can be rewritten as:

$$\begin{aligned}
 Q(Y = i|x; \hat{z}) &= P(M_1 \geq \frac{-\mu_1}{b_{11}}, \\
 &M_2 \geq \frac{-\mu_2 - b_{21}M_1}{b_{22}}, \\
 &M_3 \geq \frac{-\mu_3 - b_{31}M_1 - b_{32}M_2}{b_{33}}).
 \end{aligned} \tag{14}$$

Then $Q(Y = i|x; \hat{z})$ can be calculated recursively. Eq. 11 is an integration on $M_1 \geq \frac{-\mu_1}{b_{11}}$, while Eq. 12 and Eq. 13 integrate on $M_2 \geq \frac{-\mu_2 - b_{21}M_1}{b_{22}}$ and $M_3 \geq \frac{-\mu_3 - b_{31}M_1 - b_{32}M_2}{b_{33}}$, respectively.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitation Section
- A2. Did you discuss any potential risks of your work?
Ethics Section
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract Section and Introduction Section
- A4. Have you used AI writing assistants when working on this paper?
Grammarly, to check grammar of the whole paper

B Did you use or create scientific artifacts?

Section 5

- B1. Did you cite the creators of artifacts you used?
Section 5
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
In Appendix
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
In Appendix
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
In Appendix
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
In Appendix

C Did you run computational experiments?

In Section 5

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

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

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

The experimental setup and hyperparameter search methods are in Section 5 and Appendix. We do not include the best-found hyperparameter values

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

In section 5 and appendix

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

In section 5 and appendix

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

Left blank.

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

No response.

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

No response.

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

No response.

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

No response.

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

No response.