

# Understanding the Role of Mixup in Knowledge Distillation: An Empirical Study

Hongjun Choi, Eun Som Jeon, Ankita Shukla, Pavan Turaga

Geometric Media Lab

School of Arts, Media and Engineering, Arizona State University

School of Electrical, Computer and Energy Engineering, Arizona State University

hchoi71@asu.edu, ejeon6@asu.edu, ashuk120@asu.edu, pturaga@asu.edu

## Abstract

*Mixup is a popular data augmentation technique based on creating new samples by linear interpolation between two given data samples, to improve both the generalization and robustness of the trained model. Knowledge distillation (KD), on the other hand, is widely used for model compression and transfer learning, which involves using a larger network’s implicit knowledge to guide the learning of a smaller network. At first glance, these two techniques seem very different, however, we found that “smoothness” is the connecting link between the two and is also a crucial attribute in understanding KD’s interplay with mixup. Although many mixup variants and distillation methods have been proposed, much remains to be understood regarding the role of a mixup in knowledge distillation. In this paper, we present a detailed empirical study on various important dimensions of compatibility between mixup and knowledge distillation. We also scrutinize the behavior of the networks trained with a mixup in the light of knowledge distillation through extensive analysis, visualizations, and comprehensive experiments on image classification. Finally, based on our findings, we suggest improved strategies to guide the student network to enhance its effectiveness. Additionally, the findings of this study provide insightful suggestions to researchers and practitioners that commonly use techniques from KD. Our code is available at <https://github.com/hchoi71/MIX-KD>.*

## 1. Introduction

Deep neural networks have achieved impressive performance on a wide range of tasks including language translation [30, 26], image classification [34, 6], and speech recognition [2, 19]. To further improve the model’s efficiency and performance, a large number of training tech-

niques have been proposed such as mixup-augmentation [38, 35] and knowledge distillation [8]. In specific, mixup [38] is a commonly used data augmentation technique based on using convex combinations of samples, and their labels. This technique was introduced to improve generalization, as well as increase the robustness against adversarial examples. Recently, there has been an increasing interest in reducing the model size while preserving comparable performance, which narrows the gap between large networks and small networks. KD is one of the promising methods for this demand [8]. The goal of KD is to exploit the ability to learn concise knowledge representation (logit or feature) from a larger model and then embed such knowledge into a smaller model. For example, in image classification, deep neural networks produce class probability by using softmax function that converts logit  $f_i$  into a probability  $p_i$  by comparing  $f_i$  with other logits as follows:  $p_i = \frac{\exp(f_i/T)}{\sum_j \exp(f_j/T)}$ . In the conventional knowledge distillation [8], the temperature  $T$  is utilized to generate a softer distribution of pseudo-probabilities among the output classes, where a higher temperature increases the entropy of the output, thus, providing more information to learn for the student model.

**Motivation:** On the surface, mixup and KD are very different, however, we found that “smoothness” turns out to be the connecting link between the two and a very important attribute to understanding KD’s interplay with mixup. Intuitively, the KD involves a student mimicking smoothed probability distribution of the teacher, whereas mixup artificially introduces smoothness (the labels are not strictly one-hot) in a model via linear combinations of the labels. Although mixup-augmentation and KD are common techniques in training networks in various applications, the interplay between the two has not been well explored. In this paper, we investigate in detail the impact of mixup-augmentation on knowledge distillation.

To develop insight into several interesting behaviors of

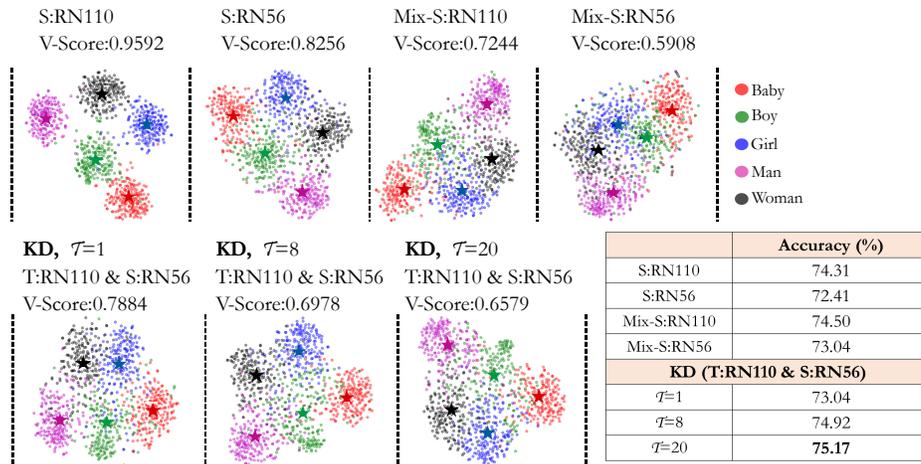


Figure 1. Feature representation of CIFAR100 extracted from the penultimate layer. Here, we only illustrate the train set as the student and teacher networks are trained on the same train set. In the first row, we observe that the higher capacity model (ResNet110 shown as RN110) promotes a tighter clustering in each class (higher V-score) than ResNet56 (RN56). Meanwhile, mixup-trained models (Mix-S:RN110 and Mix-S:RN56) disperse features of the same instances. The second row of the figure shows that increasing temperature  $\mathcal{T}$  in KD has a similar effect on the projections, resulting in dispersed feature representations. Here,  $\star$  indicates the mean point of each cluster. The table shows the corresponding test accuracy. Best viewed in color.

the network, we provide various visualizations in feature and logit levels. For example, to analyze how mixup enforces a feature representation in-between each class, we pick 5 semantically similar classes from the CIFAR100 dataset (baby, boy, girl, man, and woman). Then, we project the features extracted from the penultimate layer into 2-D using tsne [29] as seen in Figure 1. The first row of the figure represents the feature representation for two networks trained from scratch with mixup (Mix-S:ResNet110 or Mix-S:ResNet56) and without mixup (S:ResNet110 or S:ResNet56). Generally, we observe that the higher capacity model (S:ResNet110) encourages the deep networks to learn tight projections in each class while the lower capacity model (S:ResNet56) learns more diffuse projections. This is also verified by clustering metrics, such as the V-score [21], where a higher value implies better clustering. Interestingly, the mixup-trained models disperse features in similar classes, compared to the one trained without mixup, even though a mixup-trained model shows better test accuracy due to an increase in generalization on unseen data. Meanwhile, we found similar observations on the feature representation from distilled models when a higher temperature is used (i.e., distilling more softer logit) as seen in the second row of Figure 1. From the table in Figure 1, high temperature  $\mathcal{T}$  is a default choice to improve the performance. In this way, the teacher transfers more information to the student even though higher temperatures superficially promote feature dispersion in similar classes. However, transferring a high quality of supervision to a student network is also crucial as it can guide the student to learn discriminative representations from the superior teacher network.

Thus, we seek a way to increase supervisory signals from teacher to student, without compromising on performance. This is where mixup presents a way forward.

In short, the crux of our insight is that if a teacher is trained with smoothed data from mixup, then further ‘smoothing’ at high temperatures during distillation can be avoided. This ensures stronger supervisory signals while enjoying the benefits of dataset augmentation via mixup. We summarize our contributions as follows:

1. We provide new insights into devising improved strategies for learning a student model, by a deeper understanding of the behavior of features/logits and networks trained with mixup during KD.
2. To reduce the criticality of choosing a proper ‘temperature’, we develop a simple rescaling scheme that brings different statistical characteristics in logits between teacher and student to a similar range while preserving the relative information between classes, thereby the temperature  $\mathcal{T}$  is no longer used.
3. We identified that strongly interpolated mixup pairs impose extra smoothness on the logits, thus we can generate only a few mixup pairs in a batch, called partial mixup (PMU), and yet achieve comparable and even better performance with this variant in KD.

## 2. Background

We first introduce the background of mixup and KD through a simple mathematical description. Given the training data  $\mathcal{D} = \{(x_1, y_1), \dots, (x_n, y_n)\}$ , the goal of the classification task is to learn classifier  $f : \mathcal{X} \rightarrow \mathbb{R}^k$  by mapping

input  $x \in \mathcal{X} \subseteq \mathbb{R}^d$  to label  $y \in \mathcal{Y} = \{1, 2, \dots, K\}$ . Let  $\mathcal{L}(f(x), y)$  be the loss function that measures how poorly the classifier  $f(x)$  predicts the label  $y$ .

**Mixup Augmentation [38]** In mixup augmentation, two samples are mixed together by linear interpolation as follows:  $\tilde{x}_{ij}(\lambda) = \lambda x_i + (1 - \lambda)x_j$ , and  $\tilde{y}_{ij}(\lambda) = \lambda y_i + (1 - \lambda)y_j$ , where  $\lambda \in [0, 1]$  follows the distribution  $P_\lambda$  where  $\lambda \sim \text{Beta}(\alpha, \alpha)$ . Then, the mixup loss function can be described as

$$\mathcal{L}_{mix}(f) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}_{\lambda \sim P_\lambda} [\mathcal{L}(f(\tilde{x}_{ij}(\lambda)), \tilde{y}_{ij}(\lambda))], \quad (1)$$

where  $\mathcal{L}$  represents the cross-entropy loss function in this study. Specifically, a hyper-parameter  $\lambda$  in Equation 1 is used to specify the extent of mixing. In other words, the control parameter  $\alpha$  in a beta distribution commands the strength of interpolation between feature-target pairs, i.e., the high  $\alpha$  generating strongly interpolated samples.

A fair amount of variants of mixup have been proposed [31, 35, 11]. The general strategy of these mixing-based methods is intrinsically similar in that they extend the training distribution by blending the images and mixing their labels proportionally. Thus, in this study, we only pay attention to the conventional mixup to investigate the interplay between mixup and knowledge distillation [38].

**Knowledge Distillation [8]** In knowledge distillation, given a pre-trained teacher model  $f^T$  on the dataset in advance, the student model  $f^S$  is trained over the same set of data by utilizing concise knowledge generated by  $f^T$ . In specific, once the teacher network is trained, its parameter is frozen during the training in KD, and then, the student network is trained by minimizing the similarity between its output and the soft labels generated by the teacher network. To this end, we minimize the divergence between logits of the student and the teacher as follows:

$$\mathcal{L}_{kd}(f^T, f^S) = \frac{1}{n} \sum_{i=1}^n KL(S(\frac{f^T(x_i)}{\mathcal{T}}), S(\frac{f^S(x_i)}{\mathcal{T}})) \quad (2)$$

where  $S$  indicates the softmax function,  $KL$  measures the Kullback-Leiber divergence, and output logits of the model smoothen by temperature,  $\mathcal{T}$ .

**Data Augmentation in Knowledge Distillation** Recently, several works have utilized data augmentation and achieved promising results in that augmented samples enable the networks to learn relaxed knowledge from different views in distillation frameworks [33, 32, 12]. Across the broadly available methods, they often do not provide insight into the inner working of the models. Meanwhile, recent works [15, 23] have studied KD’s compatible/incompatible views with label smoothing [25] to provide such insight into deep models via empirical analysis. Instead, we are interested in the underlying mechanism of augmentation in the light

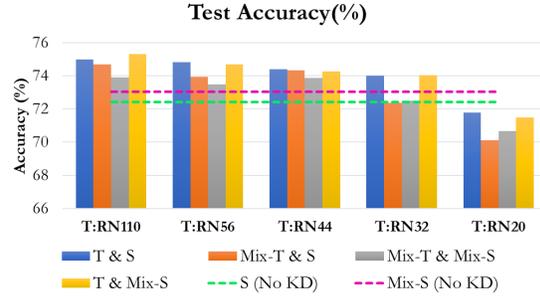


Figure 2. The student (fixed RN56) test accuracy on CIFAR100 under different teacher networks with  $\mathcal{T} = 4$ . Here, T and S denote the standard teacher and the standard student. Mix-T and Mix-S are mixup-trained teachers and mixup-trained students, respectively. Mixup-trained teachers guide lower-performance students than the one with standard teacher models, and lower-capacity teacher models generally distill inferior students.

of the KD process. Our study spans double aspects that attempt to present both compatible and incompatible views through comprehensive empirical analysis. Further, based on the observations, we suggest a better learning strategy to enhance the network’s performance.

### 3. Key findings from Mixup and KD interplay

In this section, we discuss our main findings. We first begin by referring to Figure 2. Here, we describe four possible scenarios where mixup-augmentation could be involved in KD as follows; standard teacher and standard student (T&S), mixup-trained teacher and standard student (Mix-T&S), mixup-trained teacher and mixup-trained student (Mix-T&Mix-S), and standard teacher and mixup-trained student (T&Mix-S) under the same temperature being  $\mathcal{T} = 4$ . We fix the student model as ResNet56 (RN56) and evaluate it by varying the teacher models from ResNet20 (RN20) to ResNet110 (RN110).

As seen in Figure 2, we can make two observations: First, the student performance with the help of a mixup-trained teacher (Mix-T&S, Mix-T&Mix-S) always shows less effectiveness than with the help of a standard teacher (T&S, T&Mix-S) under the same settings of temperature even if the mixup-trained teacher itself shows better test accuracy than the teacher trained without mixup. Second, generally, higher capacity teacher models distill better students, but with a lower capacity teacher (T:RN44 & S:RN56), the student performance still shows improvement in accuracy, compared to the vanilla student model (green dashed line, S (No KD)). Based on these observations, this paper aims to investigate the following questions: 1) *Why does a mixup-trained teacher model impair the student’s effectiveness in KD?* We answer this question in section 3. Then, 2) *what can we do to improve the effectiveness of knowledge distillation when mixup-augmentation is applied?* We address this question in section 4.

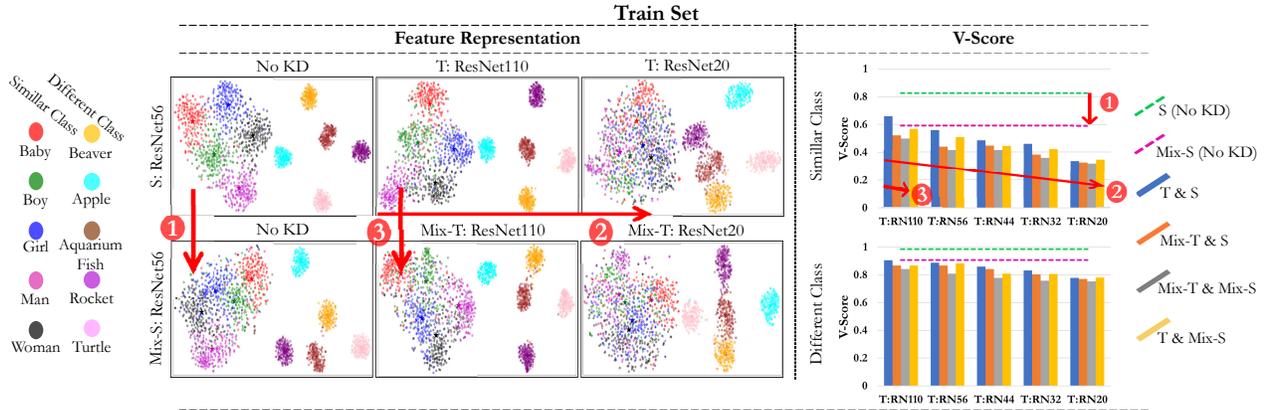


Figure 3. Feature representations of the penultimate layer with various combinations (T:standard teacher, S:standard student, Mix-T:mixup-trained teacher, Mix-S:mixup-trained teacher) and the corresponding V-scores on the CIFAR100 dataset. We selected two groups, semantically similar classes (Baby, Boy, Girl, Man, and Woman), and semantically different classes (Beaver, Apple, Aquarium Fish, Rocket, and Turtle). Here, the features are extracted from the student networks. **Observation 1)** Mixup augmentation encourages the features of samples from similar classes to be dispersed, while it still well-preserved the feature separability in different classes. **Observation 2)** The use of a lower-accuracy teacher network enables the student to learn by less-discriminative features, degrading the performance of KD. **Observation 3)** In presence of a mixup-trained teacher, the dispersion of feature representation undermines the benefit of supervision by the teacher.  $\star$  represents the center point of each cluster. Best viewed in color.

**Observation 1) Mixup vs Non-Mixup.** To investigate the effect of the network trained with mixup-augmentation, we selected a few classes and divided them into two groups: (1) semantically different classes (Beaver, Apple, Aquarium Fish, Rocket, and Turtle) and (2) semantically similar classes (Baby, Boy, Girl, Man, and Woman), all from CIFAR100. For example, in Figure 3, the left figure illustrates the feature representation of the penultimate layer on train sets. If we look at the number 1 in the red circle, the projections of a mixup-trained model from similar classes are more dispersed while the ones of different classes are still well-preserved in their structure. This information loss can also be measured by clustering metric, V-score on the right histogram of the figure, resulting in a drastic drop in V-score for similar classes.

**Observation 2) Distilling from low-accuracy teacher models.** If the student network is settled, retaining the fine quality of supervision (i.e., the performance of a teacher network) is crucial in training better students. As followed in the number 2 in the red circle, the projections of the student with the lower-accuracy teacher are notably dispersed, significant drop in V-score, eventually impairing the student’s performance. This implies that the better student is distilled with the help of discernible features given by the high-capacity teacher.

**Observation 3) Distilling from a mixup-trained teacher network.** Observation 1) showed that the mixup-trained model scatters the features in similar classes. We now look at the case where a mixup-trained teacher conveys the knowledge to the student. In KD, as the student and teacher models are trained on the same train set, we argue that a student trained with supervision by a mixup-trained teacher

cannot take advantage of learning superior knowledge due to feature scattering. As shown by number 3 in this figure, feature representations of similar classes in the student network eventually become more spread out, and V-scores also drop on both train and test sets. We provide the visualizations of train and test sets in the supplementary material.

**Observation 4) and 5) Logit representation.** Unlike the feature representations shown in observations 1)-3), we further visualize the probability distribution of the student network. First, we average probabilities of all classes on the train and test of CIFAR100 and illustrate the mean distribution of the examples that belong to the same category to show the model’s prediction of that category. To compare with quantitative measurement, we also provide the average accuracy and entropy values computed across all examples in Figure 4, where the entropy is popularly used to measure the smoothness of the distribution. The higher entropy value is the smoother distribution.

Here, we observe two intriguing phenomena; **Observation 4)** A mixup-trained model produces softer output logits, illustrated by the short red bars in both train and test set. From this observation, we surmise that a mixup-augmentation involved in training the student in KD contributes extra smoothness to the logits. Further, the student learned from a standard teacher outperforms the mixup-trained teacher in terms of accuracy (74.98% vs 73.92% and 71.78% vs 70.60%) under the same settings of the temperature  $\mathcal{T} = 4$ . Conclusively, when distilling from a mixup-trained teacher, the use of high temperature adversely impacts the accuracy of the student. **Observation 5)** When the inferior knowledge produced by the low-accuracy teacher (T:RN20) transfers to the student, the confidence level on

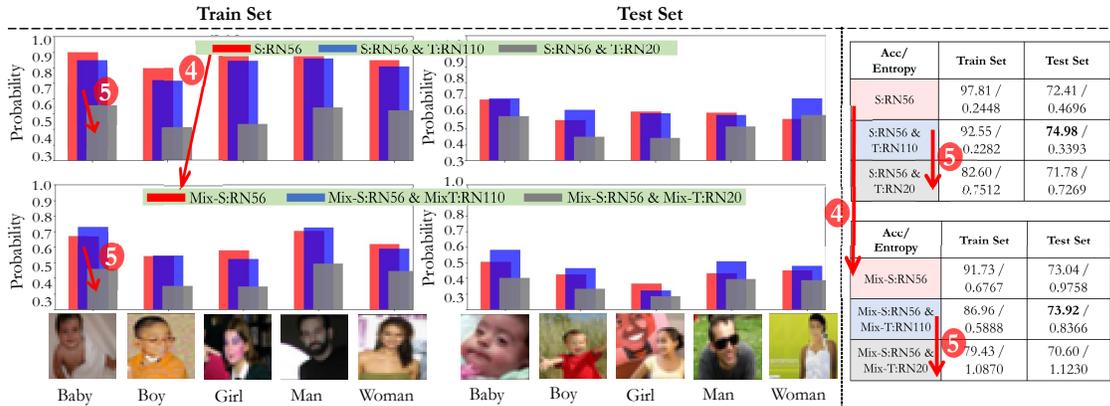


Figure 4. Probability distributions of the following configurations: with/without a mixup-trained model (S:RN56/Mix-S:RN56), the standard student model trained with the help of different capacity teachers (T:RN110 & S:RN56, T:RN20 & S:RN56), and the student model trained with mixup in the presence of different mixup-trained teachers (Mix-T:RN110 & Mix-S:RN56, Mix-T:RN20 & Mix-S:RN56). We show the mean distributions for five similar classes in CIFAR100 on the left and provide the average accuracy and entropy values for all examples. **Observation 4**) When the model is trained with the mixup, the confidence in predictions decreases on both sets (also verified by higher entropy value). **Observation 5**) Also, in the case of distilling from lower-capacity teachers, the confidence in predictions is much lower than the distilled model from high-capacity teachers. Best viewed in color.

the prediction significantly falls as seen in the gray bars on both sets, leading to a remarkable decrease in test accuracy on both cases (74.98%  $\rightarrow$  71.78%, 73.92%  $\rightarrow$  70.60%). It implies that transferring decent quality of supervision to the student is crucial to successful knowledge distillation.

**Observation 6) Mix-T&S at low temperature.** From observation 5), while increased  $\mathcal{T}$  is believed to be helpful to produce better representation for KD, we remark that in the presence of a mixup-trained teacher, an increased  $\mathcal{T}$  can be an adverse effect on the performance of KD because of the feature dispersion/exorbitant smoothness in the logit. At this point, one might give rise to the following question. *What if we lower the temperature to lessen smoothness in the logit?*

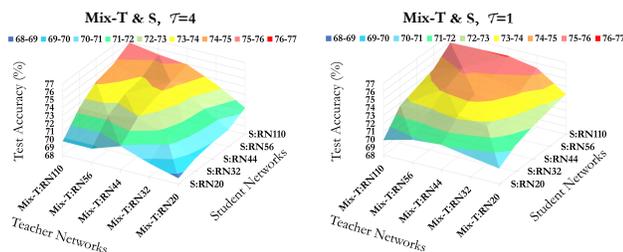


Figure 5. Visualization with the heatmap representation. When distilling from a mixup-trained teacher (Mix-T) at a higher temperature, it transfers inferior knowledge to the student. With lower temperature  $\mathcal{T} = 1$ , it produces better students, covering the relative high-accuracy region in the heatmap.

To investigate how much temperature impacts the test performance in KD, we visualize the test accuracy as a heatmap in Figure 5. In this figure, we explore all combinations of a mixup-trained teacher and student. The left panel of this figure shows test accuracy under the setting  $\mathcal{T} = 4$  and the right panel is for  $\mathcal{T} = 1$ . We note that when a lower

temperature, that makes the logit less smooth, is applied, the better the quality of knowledge that is transferred to the student. This result demonstrates that knowledge from a mixup-trained teacher at high temperatures renders the student model less effective.

#### 4. Effective learning strategies for Mix-KD

Based on our findings, we observe that smoothness in the logit is a critical attribute in distillation, as seen in Figure 5. Since it is difficult to measure suitable smoothness from the perspective of the student, many distillation methods have relied heavily on a naive brute-force search to find the proper temperature. To reduce the criticality of choosing the proper temperature and alleviating excessive smoothness by a strong mixup, we introduce an effective learning strategy to improve the performance of KD.

**Partial mixup.** The control parameter in the beta distribution plays a key role in controlling the strength of interpretation, which also affects the degree of softened output. As  $\alpha \rightarrow 1$ , it provides softer output logits. However, there is a trade-off between avoiding excessive smoothness and improvement in robustness against adversarial attacks in knowledge distillation. To alleviate this issue, we suggest generating only small amounts of mixup pairs used in training, called partial mixup (PMU). For example, PMU=10% refers to only 10% mixup pairs used in a batch and the rests keep untouched. To further understand the behavior of how partial mixup affects output probability, we provide toy examples using 2 classes in the supplementary material.

**Rescaled logits.** Here, we suggest using standard deviation as temperature such that the logit produced by the output layer of the network is rescaled by dividing it by the standard deviation of that logit, thereby the temperature  $\mathcal{T}$  is

no longer hyper-parameter. This way can bring two different statistical characteristics between the teacher and the student logits to similar ranges while it does not hurt the relative structure in-between classes. We underline the importance of rescaling logit since the random mixing portion  $\lambda$  from a beta distribution governs the smoothness degree, yielding irregular smooth output at each iteration. Therefore, we replace the output logit with rescaled one, and the loss function becomes as follows:

$$\mathcal{L}_{kd,r}(\tilde{f}^T, \tilde{f}^S) = \frac{1}{n} \sum_{i=1}^n KL(\mathcal{S}(\frac{f^S(x_i)}{\sigma(f^S(x_i))}), \mathcal{S}(\frac{f^T(x_i)}{\sigma(f^T(x_i))})), \quad (3)$$

where  $\tilde{f} = f/\sigma(f)$ ,  $\mathcal{S}$  indicates the softmax function,  $KL$  is the Kullback-Leibler divergence, and  $\sigma(\cdot)$  is standard deviation of input logit. Then, the final training objective for the student in KD is as follows:

$$\min \mathbb{E}_{(x,y) \sim \mathcal{D}} \mathbb{E}_{\lambda \sim P_\lambda} [\gamma_{kd} \mathcal{L}_{mix}(\tilde{f}^S) + \alpha_{kd} \mathcal{L}_{kd,r}(\tilde{f}^T, \tilde{f}^S)], \quad (4)$$

where  $\gamma_{kd}$  and  $\alpha_{kd}$  are balancing parameters. Note, the hyper-parameter of the partial amount in PMU is not denoted in this equation.

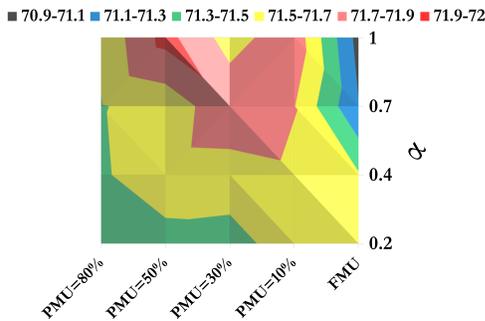


Figure 6. Test accuracy on CIFAR100 under different configurations of  $\alpha$  and the amounts of mixup pairs. The reported accuracy is averaged over 3 runs.

**Choice of  $\alpha$  and PMU.** Further, we investigate how each hyper-parameter (partial amount in PMU and the control parameter  $\alpha$ ) affects the student performance in KD under the following combination, T:RN110 and Mix-S:RN20 on the CIFAR100 dataset. We evaluate the student performance by varying the degree of partial amount from 10% to 80% with different control parameters in the range of  $[0.2, 0.4, 0.7, 1.0]$ . Figure 6 indicates that PMU with high  $\alpha$  generally leads to better students in this setting. The lowest performance is observed in FMU (fully mixup pairs used) with  $\alpha = 1.0$ . We note that PMU does not necessarily aim to outperform FMU, but it plays a regulatory role to adjust smoothness finely, thus helping us to understand the effect of smoothness in KD. We present performance analysis in section 5.1 and also further study how partial mixup responds to adversarial attacks in section 5.3.

## 5. Experiments

In this section, we present the experimental results to validate our conjectures and findings. In the previous section, we noted that (1) a mixup-trained teacher generates an inferior student compared to the teacher trained without mixup under the same temperature (2) but lowering temperature can recover the test accuracy. In KD, mixup augmentation is only useful for the teacher if a mixup-trained teacher provides the student with further benefits such as better knowledge and additional power of robustness. Through our analysis, we see that the expected effect of the use of mixup-trained teachers is not satisfactory. Therefore, we use the teacher trained without mixup for our experiment, i.e., T&Mix-S.

### 5.1. Image classification on CIFAR100 & ImageNet

**Experiments on CIFAR-100:** Table 1 compares top-1 accuracy of various distillation methods and evaluates various network choices for teacher-student. The first two rows of the Table 1 represent many different teacher-student combinations by utilizing the networks as follows: Wide residual networks (Wd-w) [37] where d and w represent depth and width in the networks respectively, MobileNetV2 (MN2) [22], ShuffleNetV1 (SN1) [39]/ShuffleNetV2 (SN2)[14], VGG (VG) [24], and ResNet (RN) [5].

All the models are trained for 240 epochs with a learning rate of 0.05 decayed by 0.1 every 30 epochs after 150 epochs. The balancing parameters  $\gamma_{kd}$  and  $\alpha_{kd}$  are 0.1 and 0.9 for all settings, respectively. In Table 1, we report our results of four different settings as follows, distilled models without mixup (No Mixup), distilled models with PMU (10% and 50% when  $\alpha = 1$ ), and distilled models with full mixup (FMU,  $\alpha = 1$ ). As seen in Table 1, the students trained with PMU consistently outperform the ones trained without mixup. Surprisingly, in some cases (e.g., T:W40-2 & S:W16-2 and T:W40-2 & S:SN1), our student trained with PMU performs better than the teacher.

**Experiments on ImageNet:** Table 2 shows the top-1 accuracy on ImageNet [3]. In this experiment, we choose ResNet34 and ResNet18 [5] as the teacher network and student network, respectively. We train the model for 100 epochs and with initial learning rate is 0.1 decayed by 0.1 at 30, 60, and 80 epochs. The batch size is set as 256. For comparison with other distillation methods, the hyper-parameters of other methods follow their respective papers. The balancing parameters  $\gamma_{kd}$  and  $\alpha_{kd}$  are 0.1 and 0.9 respectively, and we report the partial mixup of 10% and 100% (FMU). We observe that the proposed method with full mixup boosts the top-1 and top-5 accuracy by 2.07% and 1.56% over the baseline and fully mixup shows better performance than 10% partial mixup. We will discuss lower performance cases in the following section.

**Performance analysis:** Based on Table 1, in some of

Table 1. CIFAR-100 test accuracy (%) of student networks trained with several distillation methods. The results of the other distillation methods except \* are quoted from [27]. The best and second-best results are highlighted in **Bold** and **Red**. \* is performed with our implementation based on the author-provided code.

Teacher	W40-2	W40-2	RN56	RN110	RN110	RN32×4	VG13	VG13	RN50	RN50	RN32×4	RN32×4	W40-2
Student	W16-2	W40-1	RN20	RN20	RN32	RN8×4	VG8	MN2	MN2	VG8	SN1	SN2	SN1
Teacher	75.61	75.61	72.34	74.31	74.31	79.42	74.64	74.64	79.34	79.34	79.42	79.42	75.61
Student	73.26	71.98	69.06	69.06	71.14	72.50	70.36	<b>64.6</b>	<b>64.6</b>	70.36	70.50	71.82	70.50
KD[8]	74.92	73.54	70.66	70.67	73.08	73.33	72.98	67.37	67.35	73.81	74.07	74.45	74.83
FitNet[20]	73.58	72.24	69.21	68.99	71.06	73.50	71.02	64.14	63.16	70.69	73.59	73.54	73.73
AT[36]	74.08	72.77	70.55	70.22	72.31	73.44	71.43	59.40	58.58	71.84	71.73	72.73	73.32
SP[28]	73.83	72.43	69.67	70.04	72.69	72.94	72.68	66.30	68.08	73.34	73.48	74.56	74.52
CC[18]	73.56	72.21	69.63	69.48	71.48	72.97	70.71	64.86	65.43	70.25	71.14	71.29	71.38
VID[1]	74.11	73.30	70.38	70.16	72.61	73.09	71.23	65.56	67.57	70.30	73.38	73.40	73.61
RKD[16]	73.35	72.22	69.61	69.25	71.82	71.90	71.48	64.52	64.43	71.50	72.28	73.21	72.21
PKT[17]	74.54	73.45	70.34	70.25	72.61	73.64	72.88	67.13	66.52	73.01	74.10	74.69	73.89
AB[7]	72.50	72.38	69.47	69.53	70.98	73.17	70.94	66.06	67.20	70.65	73.55	74.31	73.34
FT[10]	73.25	71.59	69.84	70.22	72.37	72.86	70.58	61.78	60.99	70.29	71.75	72.50	72.03
NST[9]	73.68	72.24	69.60	69.53	71.96	73.30	71.53	58.16	64.96	71.28	74.12	74.68	74.89
CRD[27]	75.48	74.14	71.16	71.46	73.48	75.51	<b>73.94</b>	<b>69.73</b>	69.11	74.30	75.11	75.65	76.05
ICKD*[13]	75.64	74.18	71.56	71.29	73.49	74.78	73.36	68.61	68.65	73.43	74.96	75.34	76.18
Ours (No mixup)	75.38	73.70	<b>71.85</b>	71.61	73.60	75.46	72.92	67.37	67.72	73.10	73.38	75.06	75.09
Ours (PMU=10%)	<b>76.06</b>	<b>74.42</b>	<b>72.09</b>	<b>71.94</b>	<b>74.07</b>	76.87	73.60	68.52	69.55	74.29	75.89	77.06	76.78
Ours (PMU=50%)	<b>75.87</b>	<b>74.69</b>	71.80	<b>71.78</b>	<b>73.97</b>	<b>77.13</b>	<b>74.00</b>	<b>69.14</b>	<b>69.69</b>	<b>74.61</b>	<b>76.83</b>	<b>77.60</b>	<b>77.18</b>
Ours (FMU)	75.69	73.34	70.98	70.99	73.48	<b>77.25</b>	73.84	68.81	<b>69.80</b>	<b>74.50</b>	<b>77.17</b>	<b>77.92</b>	<b>77.00</b>

Table 2. Top-1 and Top-5 accuracy (%) on ImageNet validation dataset compared with various knowledge distillation methods.

	Teacher	Student	KD[8]	AT[36]	RKD[16]	SP[28]	CC[18]	CRD[27]	Ours(PMU=10%)	Ours(FMU)
Top-1	73.31	69.75	70.66	70.70	70.59	70.79	69.96	71.17	<b>71.38</b>	<b>71.82</b>
Top-5	91.42	89.07	89.88	90.00	89.68	89.80	89.17	90.13	<b>90.40</b>	<b>90.63</b>

the teacher-student combinations, we observe that adding more mixup pairs helps the student achieve higher accuracy than small amounts of pairs. Furthermore, for different architectural styles, the use of full mixup outperforms others in some cases. The presumption from this observation is that the networks from different architectures try to seek their solution paths, which means that the teacher and the student have unlike distributions in the logits, and thus imposing extra smoothness induced by strong augmentation might provide the student with additional information about how the different style teacher represents knowledge. Also, in the case of ImageNet having a large number of classes, the generated knowledge learned by two-mixing images and their labels might produce relatively less informative knowledge compared to CIFAR-100, so a strong smoothness by full mixup is favorable to distilling a better student model in this case. While we heuristically present that the performance is controlled by the amount of smoothness, how much exact smoothness should be imposed across datasets or networks remains an open question that can form the basis for further future work.

## 5.2. Robustness to adversarial examples

One undesirable consequence of models trained is their fragility to adversarial examples [4]. Adversarial examples are made by adding tiny (visually imperceptible) perturba-

Table 3. Classification accuracy against the white-box attack with various perturbations of  $\epsilon$  for every pixel. All methods are trained with full mixup ( $\alpha = 0.2$ ) and our distillation is trained with partial mixup only 10% ( $\alpha = 1.0$ ) on CIFAR100.

eps	FGSM				
	Vanilla(+mixup)	KD(+mixup)	ICKD(+mixup)	CRD(+mixup)	Ours
0.0	69.14(69.42)	70.36(70.65)	71.00(70.36)	71.03(69.25)	<b>72.01</b>
0.001	61.14(63.18)	64.25(65.02)	64.74(64.59)	64.09(63.98)	<b>68.92</b>
0.003	48.80(51.92)	52.98(54.98)	53.35(54.43)	52.91(53.25)	<b>62.21</b>
0.005	39.00(42.59)	44.13(46.28)	44.00(46.40)	43.99(44.20)	<b>55.89</b>
0.01	23.25(27.61)	29.89(32.29)	29.50(31.96)	29.75(29.09)	<b>44.10</b>
I-FGSM					
0.0	69.14(69.42)	70.36(70.65)	71.00(70.36)	71.03(69.25)	<b>72.01</b>
0.001	61.57(63.03)	64.08(64.85)	64.58(64.45)	63.91(63.92)	<b>68.91</b>
0.003	46.79(50.55)	51.43(53.56)	51.72(53.11)	51.39(51.80)	<b>61.53</b>
0.005	34.05(39.19)	39.96(42.99)	39.94(42.97)	40.02(40.78)	<b>53.93</b>
0.01	13.28(18.58)	19.71(23.12)	19.34(23.13)	19.59(19.70)	<b>37.40</b>

tions to legitimate samples to deteriorate the model performance. Unfortunately, many distillation methods have evolved to improve the performance of KD, disregarding attacks. Thus, in this section, we evaluate our model against white-box attacks where we used trained models themselves to generate adversarial examples using two common attacks, the Fast Gradient Sign Method (FGSM) and the Iterative FGSM (I-FGSM) methods [4]. For I-FGSM, we use 10 iterations with an equal step size. The results for both attacks are summarized in Table 3. For distillation methods, the setting of teacher and student is T:RN110 & S:RN20, where they were trained on CIFAR100. We applied mixup augmentation to all methods we explored (see.

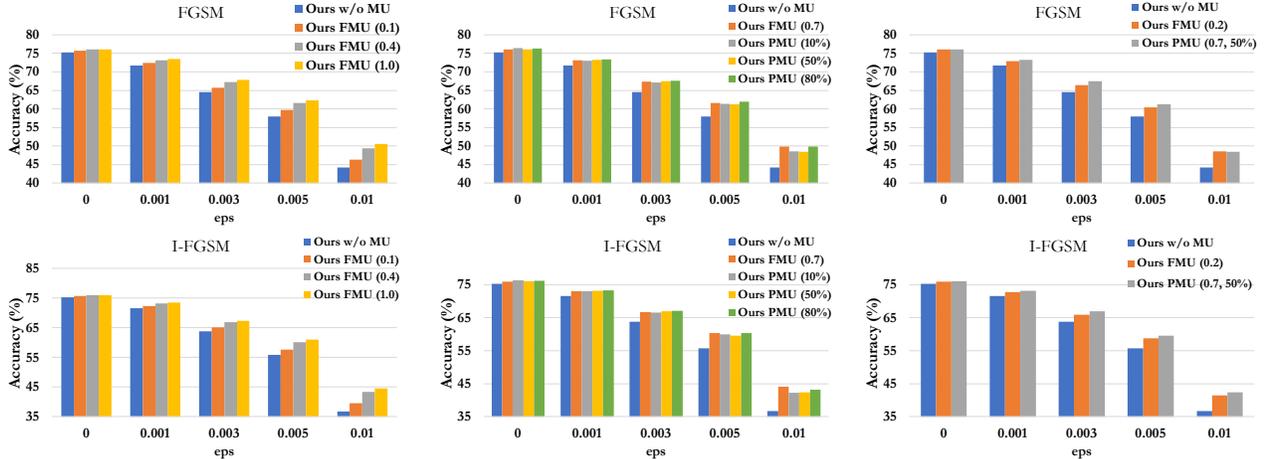


Figure 7. Histogram of test accuracy under the attacks on CIFAR100. The first column indicates comparative results of different  $\alpha$  when full mixup is used. The second column shows the comparison between full mixup and partial mixup with  $\alpha = 0.7$ . And the last figure shows the comparison between FMU (0.2) and PMU (0.7, 50%).

method+mixup in the table). We trained our model using a partial mixup of only 10% with  $\alpha = 1$ . Even with only 10% mixup pairs used, it shows impressively resistant to both attacks. We notice that distillation methods with the mixup, which utilize feature maps such as ICKD and CRD underperform those without mixup. We will show more comparative results of different amounts of mixup pairs in the next section.

### 5.3. Ablation Study

In this section, we conduct an ablation study for hyperparameters with a combination of T:WN40-2 & S:WN16-2 networks. As shown in Table 4, the distilled models with PMU and high  $\alpha$  value generally yields better performance. Further, to investigate how much the degree of partial mixup and  $\alpha$  react against both attacks, we also show test accuracy in Figure 7. As shown in the first and second columns of the histogram, when strong augmentation is involved in training such as full mixup (FMU with a high value of  $\alpha$ ) or the high number of mixup pairs (PMU 80%), it improves robustness. Interestingly, only 10% partial mixup (gray bar in the middle column) defends both attacks well. For the last column of the figure, we selected two distilled models from FMU ( $\alpha = 0.2$ ) and PMU (50%,  $\alpha = 0.7$ ) where they have similar test performance in Table 4. We observe that PMU (50%,  $\alpha = 0.7$ ) shows slightly higher robustness than FMU ( $\alpha = 0.2$ ).

## 6. Conclusions

In this work, we study the role of a mixup in knowledge distillation. We observed that a mixup-trained teacher network produces inferior supervision due to exorbitant smoothness imposed on the features and logits, especially at high temperatures during distillation. Thus, the students experience a reduction in their performance in KD. We sup-

Table 4. Test accuracy under different settings of  $\alpha$  and partial mixup amount on CIFAR100. The accuracy is averaged over 3 runs. M. denotes mixup.

		$\alpha$	0.1	0.2	0.4	0.7	1.0
w/M.	PMU=10%		75.30	75.45	75.89	75.94	75.92
	PMU=30%		75.51	75.40	75.71	<b>76.21</b>	75.95
	PMU=50%		75.50	75.50	75.83	76.02	<b>75.97</b>
	PMU=80%		75.36	75.66	<b>76.05</b>	76.19	75.78
	FMU		<b>75.69</b>	<b>76.01</b>	75.96	75.92	75.75
w/o M.	KD		74.60%				
	Vanilla		73.26%				

port our findings through a series of empirical analyses and large-scale experiments on the image classification task. Our findings provide insight into the inner workings of the distilled model trained with mixup augmentation. These insights allow us to develop an improved learning strategy using rescaled logits and partial mixups.

As we mentioned earlier, various mix-based augmentations have shown their effectiveness for particular tasks. However, these augmentations may tend to create unreasonable training samples as it blends random images [11], which could distort reasonable relative structure among categories. As a result, this potentially yields an unfavorable-smooth effect on logits during distillation. Therefore, developing an augmentation method that automatically selects more reasonable samples ensuring the best-fitting smoothness would further boost the progress in the distillation field. We will further develop this technique in our future work.

## 7. Acknowledgements

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112290073. Approved for public release; distribution is unlimited.

## References

- [1] Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D Lawrence, and Zhenwen Dai. Variational information distillation for knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9163–9171, 2019.
- [2] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjali Kannan, Ron J Weiss, Kanishka Rao, Ekaterina Gonina, et al. State-of-the-art speech recognition with sequence-to-sequence models. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4774–4778, 2018.
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [4] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [6] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2019.
- [7] Byeongho Heo, Minsik Lee, Sangdoon Yun, and Jin Young Choi. Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3779–3787, 2019.
- [8] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [9] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017.
- [10] Jangho Kim, SeongUk Park, and Nojun Kwak. Paraphrasing complex network: Network compression via factor transfer. *arXiv preprint arXiv:1802.04977*, 2018.
- [11] Jang-Hyun Kim, Wonho Choo, Hosan Jeong, and Hyun Oh Song. Co-mixup: Saliency guided joint mixup with super-modular diversity. *arXiv preprint arXiv:2102.03065*, 2021.
- [12] Xingjian Li, Haoyi Xiong, Chengzhong Xu, and Dejing Dou. Smile: Self-distilled mixup for efficient transfer learning. *arXiv preprint arXiv:2103.13941*, 2021.
- [13] Li Liu, Qingle Huang, Sihao Lin, Hongwei Xie, Bing Wang, Xiaojun Chang, and Xiaodan Liang. Exploring inter-channel correlation for diversity-preserved knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8271–8280, 2021.
- [14] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision*, pages 116–131, 2018.
- [15] Rafael Müller, Simon Kornblith, and Geoffrey Hinton. When does label smoothing help? *arXiv preprint arXiv:1906.02629*, 2019.
- [16] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019.
- [17] Nikolaos Passalis and Anastasios Tefas. Learning deep representations with probabilistic knowledge transfer. In *Proceedings of the European Conference on Computer Vision*, pages 268–284, 2018.
- [18] Baoyun Peng, Xiao Jin, Jiaheng Liu, Dongsheng Li, Yichao Wu, Yu Liu, Shunfeng Zhou, and Zhaoning Zhang. Correlation congruence for knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5007–5016, 2019.
- [19] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017.
- [20] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
- [21] Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 410–420, 2007.
- [22] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [23] Zhiqiang Shen, Zechun Liu, Dejia Xu, Zitian Chen, Kwang-Ting Cheng, and Marios Savvides. Is label smoothing truly incompatible with knowledge distillation: An empirical study. *arXiv preprint arXiv:2104.00676*, 2021.
- [24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [25] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [26] Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*, 2019.
- [27] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019.
- [28] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1365–1374, 2019.

- [29] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [31] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *Proceedings of the International Conference on Machine Learning*, pages 6438–6447, 2019.
- [32] Dongdong Wang, Yandong Li, Liqiang Wang, and Boqing Gong. Neural networks are more productive teachers than human raters: Active mixup for data-efficient knowledge distillation from a blackbox model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1498–1507, 2020.
- [33] Huan Wang, Suhas Lohit, Michael Jones, and Yun Fu. Knowledge distillation thrives on data augmentation. *arXiv preprint arXiv:2012.02909*, 2020.
- [34] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020.
- [35] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019.
- [36] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016.
- [37] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [38] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [39] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6848–6856, 2018.