

# Extracurricular Learning: Knowledge Transfer Beyond Empirical Distribution

Hadi Pouransari  
Apple

mpouransari@apple.com

Mojan Javaheripi  
UCSD\*

mojavahe@ucsd.edu

Vinay Sharma  
Apple

sharma.vinay@apple.com

Oncel Tuzel  
Apple

ctuzel@apple.com

## Abstract

*Knowledge distillation has been used to transfer knowledge learned by a sophisticated model (teacher) to a simpler model (student). This technique is widely used to compress model complexity. However, in most applications the compressed student model suffers from an accuracy gap with its teacher. We propose extracurricular learning, a novel knowledge distillation method, that bridges this gap by (1) modeling student and teacher output distributions; (2) sampling examples from an approximation to the underlying data distribution; and (3) matching student and teacher output distributions over this extended set including uncertain samples. We conduct rigorous evaluations on regression and classification tasks and show that compared to the standard knowledge distillation, extracurricular learning reduces the gap by 46% to 68%. This leads to major accuracy improvements compared to the empirical risk minimization-based training for various recent neural network architectures: 16% regression error reduction on the MPIIGaze dataset, +3.4% to +9.1% improvement in top-1 classification accuracy on the CIFAR100 dataset, and +2.9% top-1 improvement on the ImageNet dataset.*

## 1. Introduction

Training an accurate model in a supervised learning setup requires a large model capacity and a large labeled dataset. In practice, both requirements cannot be perfectly satisfied: we have limited labeled data, and model size is bounded by the computational budget that is determined by the hardware that runs the model. Knowledge transfer/distillation and data augmentation methods have been developed to address the challenges with computational cost and data scarcity. We briefly discuss these methods, which are also the building blocks of this work.

**Knowledge distillation.** Overparameterized neural networks learn better representations that lead to better

generalization accuracy [1]. For example, both the PyramidNet-110 model [23] and the larger PyramidNet-200 model achieve perfect accuracy on the CIFAR100 [32] training set, while the latter has 3% higher generalization accuracy. This motivated transferring the “knowledge” encoded in the more accurate larger model to the smaller one. Knowledge Distillation [8, 28] (KD) established an important mechanism through which one model (typically of higher capacity, called teacher) can train another model (typically a smaller model that satisfies the computational budget, called student). KD has been implemented in many machine learning tasks, for example image classification [28], object detection [12, 61], video labeling [70], natural language processing [56, 41, 53, 36, 57], and speech recognition [11, 55, 37].

The idea of KD is to encourage the student to imitate teacher’s behavior over a set of data points, called transfer-set. For example, in classification, the teacher’s output includes not only the correct class index (the argmax of softmax generated probabilities), but also additional information regarding similarities to other classes.

The amount of additional information can be quantified by the entropy of the class probabilities produced by the teacher. A teacher with small training loss produces low entropy outputs over the dataset, making KD less effective. Previously proposed remedies for this issue include matching the logits of student and teacher [8], increasing the entropy by smoothing teacher’s output [28], encouraging the student to match its intermediate feature maps to that of the teacher [48], or explicitly training a teacher with high entropy outputs [44]. In this work, we show that using a transfer-set containing uncertain examples along with modeling uncertainty of the teacher addresses this issue, and help bridging the gap between the student and teacher accuracies.

**Data augmentation.** Lack of sufficient labeled data is another challenge in supervised learning. There are several data augmentation approaches to tackle this challenge. These methods exploit domain knowledge to transform training examples to generate more data [51, 33, 17], learn a data generation policy [15, 14, 29, 35, 71, 62], augment the

\*Work is done while doing internship at Apple.

intermediate features of the model [20, 65], or find difficult examples using adversarial training [63]. Some of the recent methods [67, 59, 66, 26, 21] mix two or more data points from the empirical distribution to generate new data points. Alternatively, instead of manually designed transformations, generative models [43, 58, 31, 22] could be utilized to sample new training examples. Here, we use samples from an approximation to the data distribution to construct an improved KD algorithm. Note that unlike classical data augmentation methods that require labels for the augmented data, in our KD framework we only need unlabeled samples.

To bridge the gap between the teacher and the student models, in this work we present a novel KD method, Extracurricular Learning (XCL). Our method is motivated by the following two arguments. First, modeling the output distribution (rather than point estimates) of teacher is important for knowledge transfer as it provides additional information for student. For regression, we explicitly model the output distribution of the teacher as a Gaussian and transfer it to the student model. For classification, the output is already encoded as a categorical distribution. Second, if student exactly matches the teacher’s output on the entire input domain, we are guaranteed to bridge the accuracy gap. This is infeasible in practice due to the student’s limited capacity and optimization imperfections. As such, we propose to match the student and teacher on an extended transfer-set beyond the empirical distribution, particularly where the teacher has high uncertainty. We investigate various approximations of the data distribution to synthesize new examples (the extracurricular material) with high teacher uncertainty for KD. Our main contributions are:

- We empirically show that uncertain samples and uncertainty estimation result in significant improvements in KD generalization accuracy.
- We introduce XCL: a combination of modeling student and teacher output distributions, sampling (uncertain) data points from an approximate data distribution, and KD over this extended transfer-set. XCL does not require additional unlabeled samples and or hyper-parameter tuning.
- XCL reduces the accuracy gap between the student and the teacher by 46% to 68% compared to standard KD. Compared to best practice supervised learning baselines, XCL provides **16%** reduction in regression error on the MPIIGaze dataset and **+3.4%** (PyramidNet), **+4.6%** (ResNet), **+9.1%** (BinaryNet) top-1 classification accuracy improvement on the CIFAR100 dataset, and **+2.9%** (ResNet) on the ImageNet dataset.

## 2. Preliminaries

In supervised learning, we seek parameters  $\theta$  of a parameterized function  $f_\theta$  (e.g., weights of a neural network)

to minimize the expected risk:

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p} [l(f_\theta(\mathbf{x}), \mathbf{y})], \quad (1)$$

where  $p(\mathbf{x}, \mathbf{y})$  is the joint distribution of (example, label) pairs and  $l(\cdot)$  is the loss function determining how close  $f_\theta(\mathbf{x})$  and  $\mathbf{y}$  are. For almost every practical problem,  $p$  is not available, yet a finite set of training data points  $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$  is given. The empirical risk approximation of (1) substitutes  $p$  with empirical distribution  $p_\delta = 1/n \sum_{i=1}^n \delta(\mathbf{x} = \mathbf{x}_i, \mathbf{y} = \mathbf{y}_i)$ , where  $\delta(\mathbf{x} = \mathbf{x}_i, \mathbf{y} = \mathbf{y}_i)$  is a Dirac mass function located at  $(\mathbf{x}_i, \mathbf{y}_i)$ . This leads to the Empirical Risk Minimization (ERM):

$$\min_{\theta} \frac{1}{n} \sum_i l(f_\theta(\mathbf{x}_i), \mathbf{y}_i) \quad (2)$$

In KD [28], a student model  $f_\theta$  is encouraged to match the output of a teacher  $\tau$  on the training set:

$$\min_{\theta} \frac{1}{n} \sum_i l(f_\theta(\mathbf{x}_i), \tau(\mathbf{x}_i)) \quad (3)$$

$\tau$  in (3) can be a single more powerful model or an ensemble of several models. In the original KD [28] an average of losses in (2) and (3) is used.

KD is widely studied for the classification task, where  $\mathbf{y}_i$  is a one-hot vector that indicates the true class of  $\mathbf{x}_i$ . The teacher output  $\tau(\mathbf{x}_i)$ , however, is a soft-label. Components of  $\tau(\mathbf{x}_i)$  encode similarities of  $\mathbf{x}_i$  to other classes [28], which encapsulate additional information compared to  $\mathbf{y}_i$ . Hence, training a model with soft-labels from a stronger teacher instead of one-hot labels leads to accuracy gain.

## 3. Effect of Uncertainty on KD Performance

In this section, we present two sets of experiments that illustrate the importance of (1) uncertain data points in the transfer-set; and (2) modeling teacher and student uncertainties on KD.

**Uncertain samples from data distribution are important for KD.** In the first experiment, we randomly divide the data available in the training set into two disjoint subsets represented by  $\mathcal{A}$  and  $\mathcal{B}$ .  $\mathcal{A}$  corresponds to half of the data that is used to train the teacher while  $\mathcal{B}$  denotes the held-out set. We quantify data uncertainty (aleatoric uncertainty) of examples in  $\mathcal{B}$  using the entropy of predictions generated by the teacher trained over  $\mathcal{A}$ . We then split  $\mathcal{B}$  into two equally sized disjoint subsets  $\mathcal{B} = \mathcal{H} \cup \mathcal{L}$ , corresponding to samples with high and low uncertainties, respectively.

For a soft-label  $\mathbf{y}$  the normalized entropy is defined as

$$\hat{H}(\mathbf{y}) = - \sum_{j=1}^c \mathbf{y}^j \log \mathbf{y}^j / \log c, \quad (4)$$

where superscript  $j$  refers to the  $j$ ’th component of the vector, and  $c$  is the number of classes.  $\hat{H}(\mathbf{y})$  varies between 0 and

dataset	transfer set	entropy (%)	ERM (%)	KD (%)
CIFAR100	$\mathcal{H}$	56.4	52.7	70.7
	$\mathcal{L}$	5.5	61.9	65.3
ImageNet	$\mathcal{H}$	23.1	65.2	73.4
	$\mathcal{L}$	1.8	65.2	70.7

Table 1: Effect of uncertainty of the transfer-set on student’s top-1 validation accuracy. Teacher’s top-1 accuracy is 74.7% and 76.1% for CIFAR100 and ImageNet, respectively.

1, and denotes label uncertainty for a sample. It can also be interpreted as the amount of additional information encoded in soft labels compared to one-hot labels.

In Table 1, we report average normalized entropies of the transfer-sets, and validation accuracies of students models trained using ERM (using one-hot ground truth labels) and KD (using soft-labels from the teacher model). For both datasets, we observe a transfer-set with higher uncertainty results in more effective KD, and therefore higher validation accuracy for student. This trend does not hold for ERM. Note that we can artificially create a transfer-set containing uncertain samples, for example, a transfer-set consisting of Gaussian noise. However, uncertain samples that are not from the data distribution are not suitable for KD. Please see Table 8 for this experiment.

As observed, for an effective KD we need *uncertain samples from the underlying data distribution*. Our intuition is that uncertain samples are located close to the decision boundaries of the teacher model, similar to the argument in [27]. Therefore, they better characterize teacher’s decision boundaries compared to samples with low uncertainty (that are far away from the decision boundaries).

**Uncertainty modeling is important for KD.** In the second experiment, we analyze the effect of additional information contained in teacher’s output distribution on KD using an extreme transfer-set. Let  $\mathcal{Z} \subset \mathcal{B}$  be the set of all samples for which the teacher is incorrect<sup>1</sup>, i.e., the argmax of teacher’s prediction points to a wrong class. In Table 2, we present results of training student models using ground truth labels (ERM) versus soft-labels from teacher (KD). Using ERM over  $\mathcal{Z}$  results in very low accuracy. In contrast, when we use soft-labels from the teacher, the student model achieves surprisingly high accuracy. This illustrates that uncertainty modeling (captured through teacher’s output distribution) is crucial for KD, particularly when using difficult uncertain examples.

## 4. Extracurricular Learning (XCL)

Motivated by the observations discussed in Section 3, we develop an improved KD algorithm, dubbed extracurricular

<sup>1</sup> $|\mathcal{Z}|$  constitutes  $\sim 12\%$  of the training data for both benchmarks.

dataset	entropy (%)	ERM (%)	KD (%)
CIFAR100	58.3	14.2	61.3
ImageNet	26.7	15.6	58.5

Table 2: Student top-1 validation accuracy trained over a transfer-set on which the teacher top-1 accuracy is 0%.

learning (XCL), for regression and classification tasks which utilizes uncertainty estimation and extended transfer-sets.

### 4.1. KD Using Uncertainty Modeling

Uncertainty estimation is important for effective knowledge transfer since (1) it provides student with not only point estimates of teacher’s output, but also the full distribution (thus, more is learned from the teacher); and (2) it prevents over-penalizing student on samples that teacher is not confident about.

**Uncertainty estimation for classification.** In the standard classification task, uncertainty is already modeled, where the teacher’s output is as a categorical distribution capturing the conditional probability of the label  $y$  given  $x$ . The label distribution provides the student with the uncertainty associated to data points. Examples of data uncertainties are when different classes are present in an image, or when there is ambiguity due to occlusion. The student model is trained to minimize the average Kullback-Leibler (KL) divergence from its predicted categorical distribution to that of the teacher.

**Uncertainty estimation for regression.** A trivial extension of KD to regression tasks is replacing the ground truth labels (regression targets) with the teacher predictions. Recently, Saputra *et al.* [50] explored some variations of KD for regression. However, these methods lack uncertainty modeling, which we show is a key property for effective KD. We introduce a KD algorithm for regression incorporating uncertainty estimation.

We model the heteroscedastic uncertainties (uncertainties that depend on each example  $x_i$ ) for regression tasks, similar to [42, 30]. Specifically, for a data point  $x_i$  we assume the model outputs  $f_\theta(x_i) = (\mu_i, \sigma_i^2)$  approximating the conditional probability  $p(y|x_i)$  with a Gaussian  $\mathcal{N}(\mu_i, \sigma_i^2)$ . Hence,  $\mu_i$ ’s regress  $y_i$ ’s and  $\sigma_i$ ’s indicate the uncertainties. To estimate  $\mu_i$ ’s and  $\sigma_i$ ’s without having access to “uncertainty labels”, we learn model parameters by minimizing the Negative Log Likelihood (NLL) loss:

$$\begin{aligned} l(f_\theta(x_i), y_i) &= \frac{1}{2\sigma_i^2} \|\mu_i - y_i\|_2^2 + \frac{1}{2} \log \sigma_i^2 \\ &= \frac{1}{2} \exp(-s_i) \|\mu_i - y_i\|_2^2 + \frac{1}{2} s_i \end{aligned} \quad (5)$$

In practice, the model predicts log variance  $s_i = \log \sigma_i^2$  for numerical stability. This can be simply implemented by adding an additional output to the last layer of a neural

network. The computational overhead for uncertainty estimation is negligible.

In our framework, both student and teacher predict the output conditional distribution as Gaussians:  $\mathcal{N}(\mu_i, \sigma_i^2)$  and  $\mathcal{N}(\mu_i^\tau, \sigma_i^{\tau^2})$ , respectively. We train the teacher using (5). We then distill the teacher’s knowledge to the student by minimizing the KL divergence between two Gaussians

$$\begin{aligned} l(f_\theta(\mathbf{x}_i), \tau(\mathbf{x}_i)) &= D_{\text{KL}}(\mathcal{N}(\mu_i^\tau, \sigma_i^{\tau^2}) \parallel \mathcal{N}(\mu_i, \sigma_i^2)) \\ &= \frac{1}{2} [\exp(s_i^\tau - s_i) + \exp(-s_i) \|\mu_i^\tau - \mu_i\|_2^2 - (s_i^\tau - s_i) - 1] \end{aligned} \quad (6)$$

over the transfer-set.

In Section 5.1, we show KD for regression using the loss in (6) improves student accuracy significantly compared to methods that do not account for uncertainties.

## 4.2. Data Distribution Modeling

XCL extends knowledge transfer to data points beyond the empirical distribution. We approximate the data distribution  $p(\mathbf{x})$  by  $q(\mathbf{x})$ , construct a transfer-set by sampling from  $q$ , and perform KD over this set:

$$\min_{\theta} \mathbb{E}_{\mathbf{x} \sim q} [l(f_\theta(\mathbf{x}), \tau(\mathbf{x}))] \quad (7)$$

In classification,  $l(\cdot)$  is the KL divergence between two categorical distributions, and in regression it is between two Gaussians as in (6).

We can also interpret XCL as expected risk minimization (1) over an approximation of the joint (example, label) distribution  $p(\mathbf{x}, \mathbf{y})$ . XCL approximates the expected risk more accurately compared to the ERM in (2) by deploying a more accurate approximation of the joint distribution  $p(\mathbf{x}, \mathbf{y}) = p(\mathbf{x})p(\mathbf{y}|\mathbf{x})$ . Specifically, (7) can be obtained from (1) if we approximate the data distribution  $p(\mathbf{x})$  by  $q(\mathbf{x})$ , and the label conditional distribution  $p(\mathbf{y}|\mathbf{x})$  by teacher’s output distribution (Gaussian in regression and categorical in classification), and define the loss function  $l(f_\theta(\mathbf{x}), \mathbf{y})$  to be the negative log likelihood.

$q(\mathbf{x})$  can be any function that approximates the data distribution, e.g., unlabeled data, generative models [43, 58, 31, 22], data augmentation [51, 33, 17], data mixing [67, 59, 66, 26, 21], vicinal distribution [10], etc.

Compared to KD on the empirical distribution as in (3), in XCL we match student and teacher on much more data points. Specifically, when  $q(\mathbf{x})$  is a good approximation to data distribution, we encourage the student to imitate teacher’s output on high density regions, which helps transferring knowledge of the teacher to the student.

In Section 3, we empirically showed that uncertain samples in the data distribution are more effective to distill teacher’s knowledge to student. In this section, we propose two approximations to data distribution,  $q(\mathbf{x})$  in (7). Our choices of  $q(\mathbf{x})$  can be efficiently sampled from to construct a transfer-set that extends the empirical data distribution and

includes uncertain data points.

**XCL-Mix: Samples from mixing in pixel space.** In our first data distribution approximation, we model the data manifold as convex combinations of pairs of empirical data samples. To sample from this distribution, we randomly select two data points,  $\mathbf{x}_i$  and  $\mathbf{x}_j$  ( $i, j \in \{1, n\}$ ), and blend them based on a random  $\lambda$  as follows:

$$\begin{aligned} \mathbf{x} &= \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j & (\mathbf{x}_i, \mathbf{y}_i), (\mathbf{x}_j, \mathbf{y}_j) &\sim \hat{p}_\delta \\ \lambda &\sim \text{unif}[0, 1] \end{aligned} \quad (8)$$

where  $\hat{p}_\delta$  refers to the empirical distribution with standard augmentations and normalization. Few samples from this distribution are shown in Figure 2. As shown, for  $\lambda$  values away from 0 and 1 we sample highly uncertain points. This approach is similar to the augmentation method in [67], however, here we only sample data points, and do not use the interpolated labels.

**XCL-GAN: Samples from data manifold using GAN.**

The above method makes a crude approximation to the data distribution by mixing samples in the pixel space. We can provide a better approximation by explicitly modeling the data manifold using a generative model and sampling from it. For this purpose, we utilize a conditional generative adversarial network (GAN) [22] to model the data distribution and sample from it. The generative model  $G$ , given a  $d$ -dimensional latent variable  $\mathbf{z}$ , and a one-hot class vector  $\mathbf{e}_i$  corresponding to class  $i$ , generates a sample  $\mathbf{x} = G(\mathbf{z}; \mathbf{e}_i)$  from class  $i$ . The generative model also provides an explicit way to sample uncertain points when given a mixed class vector:

$$\begin{aligned} i, j &\sim \text{unif}\{1, c\} \\ \mathbf{x} &= G(\mathbf{z}; \lambda \mathbf{e}_i + (1 - \lambda) \mathbf{e}_j) & \lambda &\sim \text{unif}[0, 1] \\ \mathbf{z} &\sim \mathcal{N}_d(\mathbf{0}, \mathbb{I}) \end{aligned} \quad (9)$$

In our experiments, we used BigGAN [7] trained on CIFAR100 and ImageNet datasets. Few samples from this generative model are shown in Figure 3. As shown, by mixing class vectors we sample highly uncertain points. Note that, during distillation we use a union of empirical distribution and samples from the generative model as our transfer-set.

We discuss alternative sampling choices in Section 6.1.

## 5. Experiments

We conducted experiments on regression (2D gaze estimation) and classification (CIFAR100 and ImageNet) tasks and compared XCL with ERM and alternative KD formulations. For ERM we also report results using mixing based data augmentation methods MixUp [67] and CutMix [66]. In MixUp, the dataset is augmented by random convex combination of data points in pixel space.



For a sample  $x = \lambda x_i + (1 - \lambda)x_j$  MixUp uses linear interpolation  $y = \lambda y_i + (1 - \lambda)y_j$  to assign a label to  $x$ , where  $y_i$  and  $y_j$  are one-hot labels corresponding to  $x_i$  and  $x_j$ , respectively. In CutMix, a pair of images are blended by replacing a rectangular block of  $x_i$  with that of  $x_j$ . CutMix also uses linear interpolation to assign a label to a blended image. For XCL, we separately report results using the two data distribution approximations introduced in Section 4.2, denoted by XCL-Mix corresponding to sampling from blended images in pixel space, and XCL-GAN corresponding to sampling from a conditional GAN. When comparing XCL to KD, we also report the gap between the teacher and student accuracies (and % of gap reduction compared to standard KD).

### 5.1. 2D Gaze Estimation

We evaluate XCL on a regression task, human eye-gaze estimation, that is to predict the 2D gaze orientation vector given the image of an eye. We used the MPIIGaze dataset [68, 69] that contains 45,000 annotated eye images of 15 persons. We followed the leave-one-person-out setup similar to the original works [68, 69] by splitting the data to 20% validation and 80% training sets. We used LeNet [34] as the student and PreAct-ResNet [24] as the teacher. Our training setup matches the accuracies reported in the original works [68, 69]. We run each experiment three times with a different random seed.

In Table 3, we report the estimated angle error (in degrees). The baseline method (ERM) has 7.41 error, which can be reduced to 6.88 using MixUP data augmentation. Standard KD obtains 6.65 error, that is more than 10% reduction compared to ERM. We also report results using Attentive Imitation Loss (AIL) [50], a recent method that controls the extent of knowledge transfer at each data point based on teacher’s error. In our experiments AIL did not improve the student accuracy compared to KD. When we incorporate our proposed uncertainty modeling in Section 4.1 to KD (KD+Uncertainty), the student error reduces to 6.36, which corresponds to 36% gap reduction compared to KD. As an ablation, we also show results of KD with data distribution approximation, but without uncertainty modeling (denoted as KD+Mix in Table 3). This setup results in 6.36 angle error. The results are significantly improved using XCL (KD with uncertainty estimation and data distribution approximation), where XCL-Mix and XCL-GAN achieve 42% and 53% teacher-student accuracy gap reductions, respectively. For all cases in Table 3, we used the same teacher (with uncertainty estimation) that has an average angle error of 5.84 degrees. Note that, all methods shown in Table 3 are re-implemented, trained, and tested with identical setups.

In Figure 1, we plot the average predicted uncertainties by the teacher and two student models (trained using

method	angle error	gap
ERM	7.41±0.03	N/A
ERM+MixUp [67]	6.88±0.09	
KD [28]	6.65±0.03	0.81 (-)
KD+AIL [50]	6.74±0.06	0.90 (+11%)
KD+Uncer.	6.36±0.05	0.52 (-36%)
KD+Mix	6.36±0.07	0.52 (-36%)
XCL-Mix	6.31±0.03	0.47 (-42%)
XCL-GAN	<b>6.22±0.01</b>	<b>0.38 (-53%)</b>

Table 3: Gaze angle estimation using a LeNet student model. Teacher is a PreAct-ResNet with 5.84 angle error.

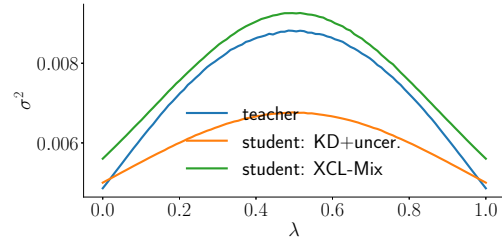


Figure 1: Uncertainty estimation values for the gaze estimation task versus the mixing coefficient.

KD+Uncertainty and XCL-Mix), as a function of mixing coefficient  $\lambda$  defined in (8). As expected, as we mix the images more ( $\lambda$  close to 0.5) all models predict higher uncertainties. Note that, this intuitive prediction is obtained without an explicit supervision for uncertainties. In addition, when we use XCL, the student model imitates teacher on a better data distribution approximation, and therefore has closer uncertainty estimation to the teacher on average.

### 5.2. CIFAR100 Classification

We evaluate the performance of XCL for image classification task on the CIFAR100 dataset [32], containing 100 classes with 50k and 10k images in the training and test sets, respectively. For fair comparison, we reimplemented all benchmarked methods and trained with identical setups. To compute accuracies, we first compute the median over the last 10 epochs, and then average the results over 8 independent runs with different random initializations. The standard-deviation of accuracy for different initializations is denoted by  $\pm\text{std}$ .

**ResNet-18:** We followed the same setup as [17] to train the ResNet-18 model [24]. We trained all methods for  $2 \times$  longer iterations compared to [17], which led to a slightly improved baseline. The reason for longer training is, by using our data distribution approximation we can sample infinite number of examples, therefore the training saturates later. To obtain an accurate teacher  $\tau$ , we use the ensemble method [18] and train a committee consisting of 8 models using CutMix data augmentation [66]. The teacher’s output

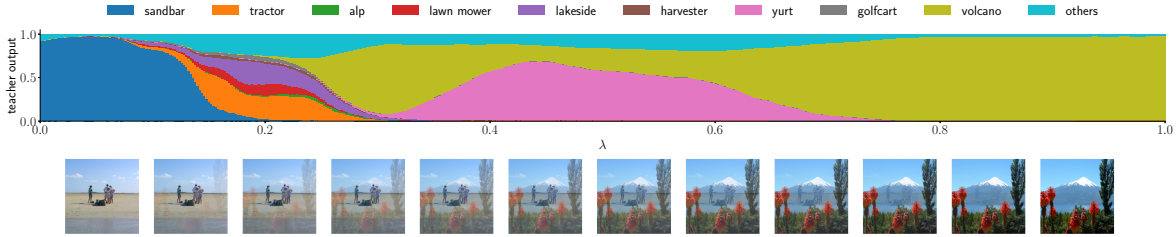


Figure 2: Examples of teacher outputs over a trajectory connecting two images in pixel space.

is the ensemble average of the committee members’ outputs. The ensemble model has top-1 test accuracy of **84.6%**.

The results in Table 4 demonstrate that both XCL-Mix and XCL-GAN significantly reduce the teacher-student accuracy gap compared to the standard KD with the same teacher (by **67 %** and **46 %**, respectively). This leads to major improvements over the ERM training (**+4.6%**) and the data mixing methods (MixUp and CutMix) that use linear interpolation for labels ( **$\sim +4\%$** ).

In Table 4, we also report the average normalized entropy,  $\hat{H}$  defined in (4), over the training/transfer-set of each method. For the standard KD method, soft-labels are obtained from a teacher trained over the empirical distribution. Since the teacher overfits the empirical distribution (becomes overconfident), the uncertainty of the teacher on the empirical distribution is an underestimation<sup>2</sup>. As a result, using empirical distribution to distill the teacher’s knowledge to the student is ineffective when the teacher overfits, which is also observed in [28]. XCL remedies the overfitting (uncertainty underestimation) problem by performing KD over an extended dataset sampled from a data distribution approximation on which teacher’s uncertainties are better quantified.

An alternative is to artificially increase  $\hat{H}(y)$  by applying Label Smoothing (LS) [54]. We applied LS to match the average normalized entropy to that of XCL’s transfer-set. In Table 4, we see LS slightly improves the baseline accuracy. However, LS is worse than XCL by more than 4% while having the same average entropy. We also applied smoothing by using a temperature parameter in KD [28]. KD with temperature and LS required exhaustive hyper-parameter tuning. We found that using temperature can improve performance of KD by 1%, which is still 2% worse than XCL without any parameter tuning.

**PyramidNet-200:** We evaluate the performance of XCL on a higher capacity architecture, PyramidNet-200 [23], which obtains the state-of-the-art results on CIFAR100 dataset. We used the training setup in [66], and obtained close accuracies. The teacher is an ensemble of 8 models

<sup>2</sup>As shown in the Table 4, the average entropy of the teacher over the empirical distribution is 10.5%. To analyze the overfitting, we computed the same measure over the test set, which is 27.8%.

method	entropy (%)	top-1 (%)	top-1 gap (%)	top-5 (%)
ERM	0	78.5±0.3		93.9±0.2
+MixUp [67]	10.8	79.2±0.2	N/A	93.9±0.2
+CutMix [66]	10.8	79.3±0.2		94.7±0.2
+LS [54]	28.2	78.8±0.2		93.9±0.2
KD [28]	10.5	80.0±0.2	4.6 (-)	95.5±0.1
XCL-Mix	28.2	<b>83.1±0.2</b>	<b>1.5 (-67%)</b>	<b>96.7±0.1</b>
XCL-GAN	29.6	82.1±0.2	<b>2.5 (-46%)</b>	96.2±0.1

Table 4: Evaluation on the CIFAR100 dataset using ResNet-18. Teacher is an ensemble of 8 ResNet-18 models with 84.6% top-1 accuracy. All results are reproduced.

method	top-1 (%)	top-1 gap (%)	top-5 (%)
ERM	82.9±0.4		96.3
+MixUp	83.4±0.1	N/A	95.7
+CutMix	84.3±0.3		96.7
KD	83.8±0.2	3.7 (-)	96.1
XCL-Mix	<b>86.3±0.1</b>	<b>1.2 (-68%)</b>	<b>97.6</b>

Table 5: Evaluation on CIFAR100 dataset using PyramidNet-200 model. Teacher top-1 test accuracy is 87.5%.

trained with CutMix, having a top-1 test accuracy of **87.5%**. As shown in Table 5, compared to standard KD with the same teacher, XCL significantly (by **68%**) reduces teacher-student accuracy gap.

**Quantized Networks:** We evaluate the performance of XCL to train an extremely compressed student, a Binary-Weight [13, 46] ResNet-18. This network has  $\sim 20\times$  smaller size compared to the full-precision (32-bit) model. We use the training setup as described in [39]. Teacher is an ensemble of 8 full-precision ResNet-18 models trained with CutMix, with a top-1 accuracy of 84.6%. As shown in Table 6, compared to standard KD with the same teacher, XCL significantly (by **52%**) reduces teacher-student accuracy gap.

### 5.3. ImageNet Classification

The ImageNet 2012 dataset [49] consists of  $\sim 1.3$  million training examples and a validation set with 50,000 images from 1,000 classes. We followed the training setup in [25] and used 300 epochs for all ImageNet experiments [66].

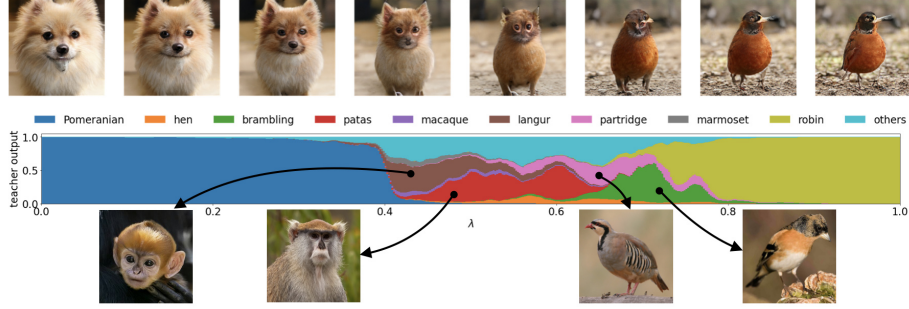


Figure 3: Examples of teacher outputs over a trajectory on BigGAN manifold connecting two pure label (one-hot class vector) images. For the prominent classes we show a sample image for comparison.

method	top-1 (%)	top-1 gap (%)	top-5 (%)
ERM	70.2±0.2		90.5
+MixUp	72.7±0.2	N/A	90.2
+CutMix	75.2±0.2		92.7
KD	74.0±0.2	10.6 (-)	91.9
XCL-Mix	<b>79.5±0.1</b>	<b>5.1 (-52%)</b>	<b>95.2</b>

Table 6: Evaluation on the CIFAR100 dataset using Binary-Weight ResNet-18. Teacher top-1 test accuracy is 84.6%.

method	val top-1	val top-1 gap	val top-5	V2-A top-1	V2-B top-1	V2-C top-1
ERM	79.0		94.5	76.0	67.5	80.6
+MixUp [67]	79.7	N/A	94.8	77.1	68.2	81.5
+CutMix [66]	80.6		95.2	77.1	69.2	81.7
KD [28]	80.7	2.6 (-)	94.3	77.4	68.6	82.1
XCL-Mix	<b>81.9</b>	<b>1.4 (-46%)</b>	<b>95.8</b>	<b>79.0</b>	<b>70.6</b>	<b>83.3</b>
XCL-GAN	81.6	1.7 (-35%)	95.6	78.4	70.5	82.8

Table 7: ResNet-101 evaluation (%) on the ImageNet dataset. Teacher is an ensemble of 4 ResNet-152-D, with top-1 accuracy of 83.3%. The std of XCL val top-1 is  $\simeq 0.1$ .

The model is a regular ResNet-101 architecture [24]. We use an ensemble of 4 ResNet-152-D [25] models trained with CutMix, having a top-1 validation accuracy of 83.3%. In addition to the regular validation set of the ImageNet dataset, we evaluated the performance of the models on three recently introduced test sets for ImageNet, called ImageNetV2 [47] that are collected with different sampling strategies: Threshold-0.7 (V2-A), Matched-Frequency (V2-B), and Top-Images (V2-C). Results are shown in Table 7.

Compared to standard KD with the same teacher, XCL-Mix reduces student-teacher validation accuracy gap by 46%. Similarly, on all other test sets, XCL obtains significant improvements compared to ERM, data mixing methods, and standard KD.

## 6. Analysis of XCL

In this section, we analyze alternative distribution approximations, effect of transfer-set size on distillation,

sampling method	top-1 (%)
Standard Gaussian Image	1.0 ± 0.0
Gaussian Noise Augmentation	79.9 ± 0.2
XCL-Mix	83.1 ± 0.2
XCL-GAN (no mixing)	81.7 ± 0.1
XCL-GAN	82.1 ± 0.2

Table 8: Analysis of different sampling methods for XCL using the CIFAR100 dataset. Teacher has a top-1 test accuracy of 84.6%. For comparison, the student trained with ERM on real data has a top-1 accuracy of  $78.5 \pm 0.3$ .

and present teacher model as a non-linear interpolation method. For all experiments we use the CIFAR100 dataset and ResNet18 architecture as described in Section 5.2.

### 6.1. Alternative distribution approximations.

The analyzed choices of data distribution approximations,  $q$ , are: Standard Gaussian image where each pixel is sampled from  $\mathcal{N}(0, 1)$ ; Pixel-wise Gaussian noise  $\mathcal{N}(0, 0.02)$  added to the empirical distribution; XCL-Mix; XCL-GAN with and without mixing class conditioning vectors (mixed-class vs. one-hot class vectors are input to GAN to generate data). As seen in Table 8, better approximations of the data distribution, e.g., mixing methods, result in better knowledge transfer and higher student accuracy compared to uninformative distributions such as the pixel-wise Gaussian. In these experiments GAN’s approximation of the data distribution is slightly worse than mixing because some modes of the distribution are not recovered by GAN and there are few visual artifacts [7]. Besides, we observe XCL-GAN performs better when mixed class vectors are used. This is consistent with the usefulness of uncertainty in the transfer-set as discussed in Section 3 for real data.

### 6.2. Effect of dataset size

**Class balanced case.** In Figure 4 we plot test error as a function of transfer-set size when the number of samples from different classes are equal. We observe that XCL

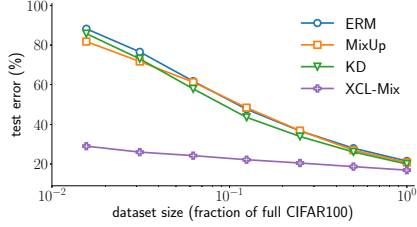


Figure 4: Test error as a function of dataset size. Number of examples from all 100 classes is the same in each case.

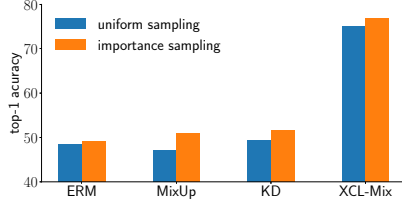


Figure 5: Test accuracy for an imbalanced CIFAR100 dataset training. The number of data samples in 80 random classes is reduced by a factor of 10.

improvement is even more pronounced as the transfer-set size gets smaller. For example, when we use  $1/64$  of the samples in the CIFAR100 as the transfer-set, ERM, MixUp, and KD obtain 11.8%, 14.2%, 18.2% test accuracies, respectively. For the same setup, XCL reaches 71% test accuracy, i.e.,  $5\times$  improvement.

**Class imbalanced case.** We artificially create a class imbalanced dataset: for 80 randomly selected classes out of 100 classes of the CIFAR100 dataset, we use only 10% of samples. In Figure 5 we report test accuracies of models trained with ERM, MixUp, KD, and XCL. XCL compensates for the class imbalance by transferring teacher’s knowledge over additional examples sampled from the approximation of data distribution and outperforms all other training methods by at least a 25% margin. For each training method, we also report results when training is performed using importance sampling: the sampling probability of each instance is inversely proportional to its class population. We observe improvement by up to 4% for all training methods when using importance sampling.

### 6.3. Teacher as a Nonlinear Interpolation Function

Several augmentation schemes such as MixUp [67] and CutMix [66] mix images and use linear interpolation to obtain labels of the mixed images. In this section, we argue the teacher model is a better interpolation function.

Figure 2 demonstrates the transition of teacher’s output probability distribution as the mixing coefficient  $\lambda$  changes. Linear interpolation overconfidently assigns a label to  $x$  with similarities to  $x_i$  and  $x_j$  that add up to 1. In contrast, labels obtained by XCL for an  $x$  between  $x_i$  and  $x_j$  include

similarities to all classes (not just to classes of  $x_i$  and  $x_j$ ). For example, in Figure 2, at  $\lambda = 0.5$  the teacher predicts  $x$  to be classified as a ‘yurt’ which is close to the visual features in the mixed image but is not equal to any of the original labels, i.e., ‘sandbar’ or ‘volcano’. Similarly, in Figure 3 we show transition between two images sampled on the image manifold by interpolating the class conditional probabilities of BigGAN (between two one-hot class vectors). As shown, the teacher output for the samples in the middle of the manifold could be different from the end-point classes (‘pomeranian’ and ‘robin’). However, qualitatively, the generated images are similar to the predicted class distribution.

## 7. Other Related Works

There are extensions of KD that match intermediate feature maps in addition to teacher outputs, e.g., FitNets [48]. In our experiments, intermediate supervision produced marginal improvement compared to the standard KD (79.3%), which is significantly lower than XCL (83.1%). In several works, multi-stage KD was proposed to improve both teacher and student by training a sequence of models [19, 40, 3]. Both intermediate supervision and multi-stage KD techniques are complementary to our framework, and could be incorporated to further reduce the accuracy gap.

There are several recent semi supervised learning methods [45, 64, 6, 52, 5, 2, 9, 60] that produce pseudo labels for unlabeled data using a model trained on a limited labeled set. The extended dataset is then used to train the target model. However, unlike XCL, these methods require additional real samples.

KD has been used to provide a fast approximation to Bayesian Neural Networks [4, 38, 16]. Bayesian Neural Networks implicitly estimate the uncertainty via Monte-Carlo sampling of the network parameters. In our framework, we explicitly model and learn the output distribution of the teacher and utilize it over an extended transfer-set to reduce the KD gap.

## 8. Conclusion

We introduce XCL, a framework for KD that incorporates a combination of (1) uncertainty estimation, (2) data distribution approximation, and (3) imitating the teacher output distribution using an extended transfer-set including highly uncertain points from the approximate data distribution. This results in an easy-to-use algorithm that provides the state-of-the-art accuracies for knowledge distillation (both for classification and regression tasks) without need for additional dataset or hyper-parameter tuning. Experiments on MPIIGaze, CIFAR100, and ImageNet datasets show that XCL achieves state-of-the-art accuracies for KD.



## References

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pages 6155–6166, 2019. 1
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. *arXiv preprint arXiv:1908.02983*, 2019. 8
- [3] Hessam Bagherinezhad, Maxwell Horton, Mohammad Rastegari, and Ali Farhadi. Label refinery: Improving imagenet classification through label progression. *arXiv preprint arXiv:1805.02641*, 2018. 8
- [4] Anoop Korattikara Balan, Vivek Rathod, Kevin P Murphy, and Max Welling. Bayesian dark knowledge. In *Advances in Neural Information Processing Systems*, pages 3438–3446, 2015. 8
- [5] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. 8
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5050–5060, 2019. 8
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018. 4, 7
- [8] Cristian Bucilua, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, 2006. 1
- [9] Paola Cascante-Bonilla, Fuwen Tan, Yanjun Qi, and Vicente Ordonez. Curriculum labeling: Self-paced pseudo-labeling for semi-supervised learning. *arXiv preprint arXiv:2001.06001*, 2020. 8
- [10] Olivier Chapelle, Jason Weston, Léon Bottou, and Vladimir Vapnik. Vicinal risk minimization. In *Advances in neural information processing systems*, pages 416–422, 2001. 4
- [11] Yevgen Chebotar and Austin Waters. Distilling knowledge from ensembles of neural networks for speech recognition. In *Interspeech*, pages 3439–3443, 2016. 1
- [12] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems*, pages 742–751, 2017. 1
- [13] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pages 3123–3131, 2015. 6
- [14] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 113–123, 2019. 1
- [15] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. *arXiv preprint arXiv:1909.13719*, 2019. 1
- [16] Yufei Cui, Wuguannan Yao, Qiao Li, Antoni B Chan, and Chun Jason Xue. Accelerating monte carlo bayesian inference via approximating predictive uncertainty over simplex. *arXiv preprint arXiv:1905.12194*, 2019. 8
- [17] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 1, 4, 5
- [18] Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000. 5
- [19] Tommaso Furlanello, Zachary C Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. *arXiv preprint arXiv:1805.04770*, 2018. 8
- [20] Xavier Gastaldi. Shake-shake regularization. *arXiv preprint arXiv:1705.07485*, 2017. 2
- [21] Chengyue Gong, Tongzheng Ren, Mao Ye, and Qiang Liu. Maxup: A simple way to improve generalization of neural network training. *arXiv preprint arXiv:2002.09024*, 2020. 2, 4
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2, 4
- [23] Dongyoon Han, Jiwhan Kim, and Junmo Kim. Deep pyramidal residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5927–5935, 2017. 1, 6
- [24] 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. 5, 7
- [25] 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 Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2019. 6, 7
- [26] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019. 2, 4
- [27] Byeongho Heo, Minsik Lee, Sangdoo Yun, and Jin Young Choi. Knowledge distillation with adversarial samples supporting decision boundary. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3771–3778, 2019. 3
- [28] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2, 5, 6, 7
- [29] Daniel Ho, Eric Liang, Ion Stoica, Pieter Abbeel, and Xi Chen. Population based augmentation: Efficient

- learning of augmentation policy schedules. *arXiv preprint arXiv:1905.05393*, 2019. 1
- [30] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584, 2017. 3
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2, 4
- [32] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 1, 5
- [33] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1, 4
- [34] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 5
- [35] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast autoaugment. In *Advances in Neural Information Processing Systems*, pages 6662–6672, 2019. 1
- [36] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. Improving multi-task deep neural networks via knowledge distillation for natural language understanding. *arXiv preprint arXiv:1904.09482*, 2019. 1
- [37] Yuchen Liu, Hao Xiong, Zhongjun He, Jiajun Zhang, Hua Wu, Haifeng Wang, and Chengqing Zong. End-to-end speech translation with knowledge distillation. *arXiv preprint arXiv:1904.08075*, 2019. 1
- [38] Andrey Malinin, Bruno Mlodozeniec, and Mark Gales. Ensemble distribution distillation. *arXiv preprint arXiv:1905.00076*, 2019. 8
- [39] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Training binary neural networks with real-to-binary convolutions. *arXiv preprint arXiv:2003.11535*, 2020. 6
- [40] Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant: Bridging the gap between student and teacher. *arXiv preprint arXiv:1902.03393*, 2019. 8
- [41] Subhabrata Mukherjee and Ahmed Hassan Awadallah. Distilling transformers into simple neural networks with unlabeled transfer data. *arXiv preprint arXiv:1910.01769*, 2019. 1
- [42] David A Nix and Andreas S Weigend. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, volume 1, pages 55–60. IEEE, 1994. 3
- [43] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016. 2, 4
- [44] 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. 1
- [45] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4119–4128, 2018. 8
- [46] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016. 6
- [47] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? *arXiv preprint arXiv:1902.10811*, 2019. 7
- [48] 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. 1, 8
- [49] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 6
- [50] Muhamad Risqi U Saputra, Pedro PB de Gusmao, Yasin Almalioglu, Andrew Markham, and Niki Trigoni. Distilling knowledge from a deep pose regressor network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 263–272, 2019. 3, 5
- [51] Patrice Y Simard, David Steinkraus, John C Platt, et al. Best practices for convolutional neural networks applied to visual document analysis. In *Icdar*, volume 3, 2003. 1, 4
- [52] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020. 8
- [53] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019. 1
- [54] 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. 6
- [55] Ryoichi Takashima, Sheng Li, and Hisashi Kawai. An investigation of a knowledge distillation method for ctc acoustic models. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5809–5813. IEEE, 2018. 1
- [56] Raphael Tang, Yao Lu, Linqing Liu, Lili Mou, Olga Vechtomova, and Jimmy Lin. Distilling task-specific knowledge from bert into simple neural networks. *arXiv preprint arXiv:1903.12136*, 2019. 1
- [57] Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Well-read students learn better: The impact of student initialization on knowledge distillation. *arXiv preprint arXiv:1908.08962*, 2019. 1
- [58] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image

- generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016. 2, 4
- [59] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, Aaron Courville, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. *arXiv preprint arXiv:1806.05236*, 2018. 2, 4
- [60] 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. 8
- [61] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019. 1
- [62] Longhui Wei, An Xiao, Lingxi Xie, Xin Chen, Xiaopeng Zhang, and Qi Tian. Circumventing outliers of autoaugment with knowledge distillation. *arXiv preprint arXiv:2003.11342*, 2020. 1
- [63] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan Yuille, and Quoc V Le. Adversarial examples improve image recognition. *arXiv preprint arXiv:1911.09665*, 2019. 2
- [64] Qizhe Xie, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Self-training with noisy student improves imagenet classification. *arXiv preprint arXiv:1911.04252*, 2019. 8
- [65] Yoshihiro Yamada, Masakazu Iwamura, Takuya Akiba, and Koichi Kise. Shakedrop regularization for deep residual learning. *arXiv preprint arXiv:1802.02375*, 2018. 2
- [66] Sangdoo 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 International Conference on Computer Vision*, pages 6023–6032, 2019. 2, 4, 5, 6, 7, 8
- [67] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 2, 4, 5, 6, 7, 8
- [68] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Appearance-based gaze estimation in the wild. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4511–4520, June 2015. 5
- [69] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Mpiigaze: Real-world dataset and deep appearance-based gaze estimation. *IEEE transactions on pattern analysis and machine intelligence*, 41(1):162–175, 2017. 5
- [70] Giulio Zhou, Subramanya Dulloor, David G Andersen, and Michael Kaminsky. Edf: Ensemble, distill, and fuse for easy video labeling. *arXiv preprint arXiv:1812.03626*, 2018. 1
- [71] Barret Zoph, Ekin D Cubuk, Golnaz Ghiasi, Tsung-Yi Lin, Jonathon Shlens, and Quoc V Le. Learning data augmentation strategies for object detection. *arXiv preprint arXiv:1906.11172*, 2019. 1