# MixSKD: Self-Knowledge Distillation from Mixup for Image Recognition

Chuanguang Yang<sup>1,2</sup>, Zhulin An<sup>1(⊠)</sup>, Helong Zhou<sup>3</sup>, Linhang Cai<sup>1,2</sup>, Xiang Zhi<sup>1,2</sup>, Jiwen Wu<sup>1</sup>, Yongjun Xu<sup>1</sup>, and Qian Zhang<sup>3</sup>

<sup>1</sup> Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
<sup>2</sup> University of Chinese Academy of Sciences, Beijing, China

<sup>3</sup> Horizon Robotics

Abstract. Unlike the conventional Knowledge Distillation (KD), Self-KD allows a network to learn knowledge from itself without any guidance from extra networks. This paper proposes to perform Self-KD from image Mixture (MixSKD), which integrates these two techniques into a unified framework. MixSKD mutually distills feature maps and probability distributions between the random pair of original images and their mixup images in a meaningful way. Therefore, it guides the network to learn cross-image knowledge by modelling supervisory signals from mixup images. Moreover, we construct a self-teacher network by aggregating multi-stage feature maps for providing soft labels to supervise the backbone classifier, further improving the efficacy of self-boosting. Experiments on image classification and transfer learning to object detection and semantic segmentation demonstrate that MixSKD outperforms other state-of-the-art Self-KD and data augmentation methods. The code is available at https://github.com/winycg/Self-KD-Lib.

Keywords: Self-Knowledge Distillation, Mixup, Image Recognition

# 1 Introduction

Knowledge Distillation (KD) [15] is an effective paradigm to enable a given student network to generalize better under the guidance of a pre-trained high-performance teacher. The seminal KD guides the student to mimic the predictive class probability distributions (also namely *soft labels*) generated from the teacher. Interestingly, although the soft label assigns probabilities to incorrect classes, the relative probability distribution also encodes meaningful information on similarity among various categories [46]. The soft label has been widely demonstrated as dark knowledge to enhance the student's performance [15,46].

The conventional KD [15,30,38,43,39] relies on a pretrained teacher to provide soft labels. Thus the two-stage process increases the training pipeline and is time-consuming. Recent studies of *Self-KD* have illuminated that a network



Fig. 1. Overview of the basic idea about our proposed MixSKD.

can distill its own knowledge to teach itself without a pretrained teacher. However, exploring knowledge from the network itself is a non-trivial problem. Existing Self-KD works often employ *auxiliary architecture* [31,52,53,18,51] or *data augmentation* [37,48] to capture additional knowledge to enhance the network. The auxiliary-architecture-based method often appends extra branches to the backbone network. Self-KD regularizes these branches and the backbone network to generate similar predictions via knowledge transfer [31,52]. Another vein is to mine meaningful knowledge from external data augmentations. Dataaugmentation-based methods aim to regularize the consistent predictions between distorted versions of the same instance [37] or two instances from the same class [48]. A common point of previous Self-KD methods is that the mined supervisory signal, *e.g.* soft label, is generated from an individual input sample.

In contrast, we propose incorporating Self-KD with image Mixture (namely MixSKD) into a unified framework. Image mixture has been developed as an advanced data augmentation called Mixup [50]. Mixup performs pixel-wise image mixture between two randomly sampled images  $x_i$  and  $x_j$  to construct a virtual sample  $\tilde{x}_{ij}$ . The virtual sample's label  $\tilde{y}_{ij}$  is a linear interpolation between the one-hot ground-truth labels of  $y_i$  and  $y_j$  with the same mixture proportion. Additional Mixup images introduce more samples in the input space by randomly fusing different images. This allows Self-KD to distill richer knowledge towards image mixture over the independent samples. The success behind Mixup is to encourage the network to favour simple linear behaviours in-between training samples [50]. Our MixSKD further takes advantage of this property by modelling supervisory signals from Mixup images in the feature and probability space.

As shown in Fig. 1, we propose to linearly interpolate the probability distributions  $p_i$  and  $p_j$  inferred from the original images  $x_i$  and  $x_j$  to model the soft ensemble label  $p_{ij}$  to supervise the Mixup distribution  $\tilde{p}_{ij}$ . Intuitively,  $p_{ij}$  en-

codes the crude predictions based on the full information in two original images. It can be seen as a pseudo teacher distribution to provide comprehensive knowledge for Mixup-based  $\tilde{p}_{ij}$ . Symmetrically, we also regard  $\tilde{p}_{ij}$  as the soft label to supervise  $p_{ij}$ .  $\tilde{p}_{ij}$  could be regarded as a data-augmented distribution to refine  $p_{ij}$  for learning robust mixed predictions under Mixup and avoid overfitting. The mutual distillation process encourages the network to produce consistent predictions between the pair of original images  $(x_i, x_j)$  and their Mixup image  $\tilde{x}_{ij}$ . Another efficacy compared with the conventional Mixup training is that Self-KD would force the network to generate similar wrong predictions over incorrect classes between  $p_{ij}$  and  $\tilde{p}_{ij}$ . This property of dark knowledge may also be a critical factor demonstrated by previous KD research [46].

Inspired by [52,31,38,39], we append auxiliary feature alignment modules for transforming the feature maps from shallow layers to match the hierarchy with the final feature map. We match these feature maps between the pair of original images  $(x_i, x_j)$  and their Mixup image  $\tilde{x}_{ij}$  via Self-KD, as shown in Fig. 1. Motivated by [36], we propose an adversarial feature Self-KD method that utilizes  $l_2$  loss to force the feature maps to be close and a discriminator loss to simultaneously increase the difficulty of mimicry. This process encourages the network to learn common semantics between matched feature maps. We further attach auxiliary classifiers to these hidden feature maps to output probability distributions and perform probability Self-KD. This further regularizes the consistency of intermediate information.

Over the training graph, we also construct a self-teacher to supervise the final classifier of the backbone network. The self-teacher aggregates all intermediate feature maps and uses a linear classifier to provide meaningful soft labels. The soft label is informative since it assembles all feature information across the network. During the training phase, we introduce auxiliary branches and a self-teacher to assist Self-KD. During the test phase, we discard all auxiliary architectures, resulting in no extra costs compared with the baseline.

Following the consistent benchmark [48,18], we demonstrate that MixSKD is superior to State-Of-The-Art (SOTA) Self-KD methods and data augmentation approaches on image classification tasks. Extensive experiments on downstream object detection and semantic segmentation further show the superiority of MixSKD in generating better feature representations.

The contributions are three-fold: (1) We propose incorporating Self-KD with image mixture into a unified framework in a meaningful manner to improve image recognition. (2) We construct a self-teacher by aggregating multi-stage feature maps to produce high-quality soft labels. (3) MixSKD achieves SOTA performance on image classification and downstream dense prediction tasks against other competitors.

# 2 Related Works

Multi-Network Knowledge Distillation. The conventional multi-network KD often depends on extra networks for auxiliary training in an offline or online

manner. The offline KD [15,11,27,38,44,43,39,23,24] transfers knowledge from a pre-trained teacher to a smaller student. Online KD [54,57,3,41,40,45,39] aims to train multiple student networks and perform knowledge transfer within the cohort in an online manner. A popular mechanism called *Deep Mutual Learning* (DML) [54] or *Mutual Contrastive Learning* (MCL) [40] suggests that an ensemble of students learn significantly better through teaching each other. Our MixSKD also employs mutual mimicry of intermediate feature maps and probability distributions. An important distinction with DML is that we aim to distill information from the view of Mixup within a single network. In contrast to mutual learning, another vein is to construct a virtual online teacher via knowledge ensembling from multiple student networks [57]. Instead of multiple networks, we aggregate multi-stage feature maps from attached branches within a single network to construct a powerful self-teacher. This self-teacher generates excellent soft labels for auxiliary Mixup training.

Self-Knowledge Distillation. Unlike offline and online KD, Self-KD aims to distill knowledge from the network itself to improve its own performance. Self-KD is a non-trivial problem since it does not have explicit peer networks. Therefore, a natural idea is to introduce *auxiliary architecture* to capture extra knowledge. DKS [31] proposes pairwise knowledge alignments among auxiliary branches and the primary backbone. BYOT [52] regards the deepest classifier as a teacher and transfers knowledge to shallow ones. ONE [57] aggregates an ensemble logit from multiple auxiliary networks to provide the soft label. Beyond the logit level, SAD [16] performs top-down and layer-wise attention distillation within the network itself. FRSKD [18] generates a self-teacher network for itself to provide refined feature maps for feature distillation. Another orthogonal aspect is to excavate knowledge from *data augmentation*. DDGSD [37] and CS-KD [48] regularize the consistency of predictive distributions between two different views. DDGSD utilizes two different augmented versions of the same image as two views, while CS-KD leverages two different samples from the same class. Moreover, prediction penalization is also a form of Self-KD. Tf-KD [46] connects label smoothing [32] to Self-KD and applies a manually designed soft label to replace the teacher. Orthogonal to the above studies, we propose Self-KD from the image mixture perspective and achieve the best performance.

Image Mixture. Image mixture has been developed to a robust data augmentation strategy. The seminal Mixup [50] performs global pixel-wise interpolations between two images associated with the same linear interpolation of one-hot labels. Wang *et al.* [35] employ Mixup to augment a few unlabeled images for data-efficient KD. Beyond input space, Manifold Mixup [33] conducts interpolations in the hidden feature representations. The success behind image mixture is to provide mixed training signals for regularizing the network to behave linearly. *However, none of the previous work explores Mixup training signals as extra knowledge to model soft labels for Self-KD*. Our MixSKD incorporates Mixup with Self-KD into a unified framework and regularizes the network more meaningfully. Although previous CS-KD and FRSKD also attempted to combine Mixup to further enhance the performance, they often utilized a straightforward



**Fig. 2.** Overview of the proposed MixSKD over a network with K = 3 stages. We employ a shared architecture between the training graph of the original input images (upper) and that of the Mixup input images (lower). The meta-architectures with the same notation denote the identical one. During the training stage, we regularize the *feature maps* and *class probability distributions* between the original images and Mixup images via distillation losses. During the test stage, we drop all auxiliary components.

instead of a thoughtful way. Yun *et al.*[47] observed that Mixup samples are locally unnatural, confusing the network for object localization. CutMix [47] is proposed to execute patch-wise mixture between two images. We remark that our MixSKD can also combine CutMix to regularize the spatial consistency between patch-wise information mixture. This may become a promising future work.

# 3 Methodology

## 3.1 Formulation of Training Graph

The convolutional neural network (CNN) f for image classification can be composed of a feature extractor  $\phi$  and a linear classifier g, *i.e.*  $f = g \circ \phi$ . For simplicity, we omit *Global Average Pool* between feature extractor and linear classifier. The feature extractor  $\phi$  often contains multiple stages for refining feature hierarchies that is formulated as  $\phi = \phi_K \circ \cdots \circ \phi_2 \circ \phi_1$ , where K is the number of stages. After each stage  $\phi_k$ ,  $k = 1, 2, \cdots, K - 1$ , we insert an auxiliary branch  $b_k$  for feature augmentation. Each  $b_k$  includes a feature alignment module  $\zeta_k$  and a linear classifier  $g_k$ . Given an input sample  $\boldsymbol{x}$ , the outputs of K - 1 auxiliary branches and the backbone network f are expressed as:

$$b_{1}(\boldsymbol{x}) = g_{1} \circ \zeta_{1} \circ \phi_{1}(\boldsymbol{x}),$$

$$b_{2}(\boldsymbol{x}) = g_{2} \circ \zeta_{2} \circ \phi_{2} \circ \phi_{1}(\boldsymbol{x}),$$

$$\cdots$$

$$b_{K-1}(\boldsymbol{x}) = g_{K-1} \circ \zeta_{K-1} \circ \phi_{K-1} \circ \cdots \circ \phi_{1}(\boldsymbol{x}),$$

$$f(\boldsymbol{x}) = g \circ \phi(\boldsymbol{x}).$$
(1)

Here,  $b_1(\boldsymbol{x}), \dots, b_{K-1}(\boldsymbol{x}) \in \mathbb{R}^C$  and  $f(\boldsymbol{x}) \in \mathbb{R}^C$  are predicted logit vectors, where C is the number of classes. The feature alignment module  $\zeta_k$  is to transform the feature map output from the shallow stage to match the feature dimension with the last stage. This is implemented by making each auxiliary branch's path from input to output have the same number of down-samplings as the backbone network. We formulate  $F_k \in \mathbb{R}^{H \times W \times C}$  as the feature map output from the k-th branch  $b_k$  for all  $k = 1, \dots, K-1$  and  $F_K \in \mathbb{R}^{H \times W \times C}$  as the feature map output from the backbone network f. Here, H, W and C denote the height, width and channel number of the feature map, respectively.  $\{F_k\}_{k=1}^{K-1}$  and  $F_K$  are formulated as:

$$F_k(\boldsymbol{x}) = \zeta_k \circ \phi_k \circ \cdots \circ \phi_1(\boldsymbol{x}), \ F_K(\boldsymbol{x}) = \phi(\boldsymbol{x}), \tag{2}$$

where  $k = 1, \dots, K - 1$ . The transformed feature map often encodes global structures and coarse semantic information. The property is crucial for recognition performance by reducing redundant details in shallow features. The detailed architectures of auxiliary branches for various networks are shown in Appendix.

#### 3.2 Task-Guided Classification Loss

We describe how to construct classification loss from *architecture source* and *data source*. We illustrate the detailed overview of MixSKD in Fig. 2.

Classification loss from architecture source. Given an input sample x with the ground-truth label y, we guide the K - 1 auxiliary branches and backbone network f to learn cross-entropy based classification loss:

$$L_{cls\_b\_f}(\boldsymbol{x}, y) = L_{ce}(\sigma(f(\boldsymbol{x})), y) + \sum_{k=1}^{K-1} L_{ce}(\sigma(b_k(\boldsymbol{x})), y),$$
(3)

where  $\sigma$  represents a *softmax* function to normalize the logits to model posterior probability distributions, and  $L_{ce}$  denotes the cross-entropy. This conventional loss enables auxiliary branches and the backbone network to learn general classification capability and semantic features.

Classification loss from data source. Given two different images  $x_i$  and  $x_j$  with ground-truth labels  $y_i$  and  $y_j$ , Mixup [50] performs linear mixture with a combination factor  $\lambda$  to construct a virtual image  $\tilde{x}_{ij}$  with an interpolated one-hot ground-truth label  $\tilde{y}_{ij}$ :

$$\tilde{\boldsymbol{x}}_{ij} = \lambda \boldsymbol{x}_i + (1 - \lambda) \boldsymbol{x}_j, \ \tilde{y}_{ij} = \lambda y_i + (1 - \lambda) y_j.$$
(4)

Here,  $\lambda \sim \text{Beta}(\alpha, \alpha)$  for  $\alpha \in (0, \infty)$  and  $\lambda \in [0, 1]$ . We utilize the input images  $\boldsymbol{x}_i, \boldsymbol{x}_j$  and  $\tilde{\boldsymbol{x}}_{ij}$  to compute the classification task loss using cross-entropy:

$$L_{cls\_mixup} = L_{cls\_b\_f}(\boldsymbol{x}_i, y_i) + L_{cls\_b\_f}(\boldsymbol{x}_j, y_j) + L_{cls\_b\_f}(\tilde{\boldsymbol{x}}_{ij}, \tilde{y}_{ij})$$
(5)

#### 3.3 Feature map Self-KD

Feature maps often contain information about image intensity and spatial correlation. The hidden feature maps often encode the intermediate learning process. We expect to encourage the network to behave consistently in *intermediate feature maps* between the pair of original images  $(\boldsymbol{x}_i, \boldsymbol{x}_j)$  and their *Mixup* image  $\tilde{\boldsymbol{x}}_{ij}$ . Thus we consider linearly interpolating the feature maps between  $\boldsymbol{x}_i$  and  $\boldsymbol{x}_j$ . The interpolated feature maps are formulated as  $\{\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)\}_{k=1}^{K-1}$ :

$$\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \lambda F_k(\boldsymbol{x}_i) + (1 - \lambda) F_k(\boldsymbol{x}_j).$$
(6)

Motivated by the hint loss in FitNet [30], we employ the squared  $l_2$ -norm for mutual alignment between interpolated feature maps  $\{\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j) \in \mathbb{R}^{H \times W \times C}\}_{k=1}^K$  and *Mixup* feature maps  $\{F_k(\tilde{\boldsymbol{x}}_{ij}) \in \mathbb{R}^{H \times W \times C}\}_{k=1}^K$ :

$$L_{feature} = \sum_{k=1}^{K} \frac{1}{HWC} \parallel \tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j) - F_k(\tilde{\boldsymbol{x}}_{ij}) \parallel^2.$$
(7)

Inspired by Adversarial Feature Distillation (AFD) [36], we further introduce K discriminators  $\{D_k\}_{k=1}^K$  for the same-staged K feature maps. The discriminator in the original AFD [36] is used to distinguish features extracted from student or teacher. Instead, our discriminator is to classify the feature map generated from a linear interpolation or a Mixup image. The discriminator loss is formulated as a binary cross-entropy:

$$L_{dis} = -\sum_{k=1}^{K} \log[D_k(\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j))] + \log[1 - D_k(F_k(\tilde{\boldsymbol{x}}_{ij}))].$$
(8)

Here, the  $D_k$  is a two-layer MLP followed by a Sigmoid function. Motivated by the idea of GAN [12], jointly optimizing Eq. (7) and Eq. (8) via an adversarial process enables the network to learn common semantic information meaningfully. The feature Self-KD conducts an adversarial process to match linearly interpolated feature maps from the pair of original images  $(\boldsymbol{x}_i, \boldsymbol{x}_j)$  with feature maps from the Mixup image  $\tilde{\boldsymbol{x}}_{ij}$ . This forces the network to behave linearly in feature representations between training samples.

#### 3.4 Logit Distribution Self-KD

Beyond feature maps, logit-based class probability distributions often represent the final predictive information. We also expect to encourage the network to behave consistently in *logit-based class posterior distributions* between the pair

of original images  $(\boldsymbol{x}_i, \boldsymbol{x}_j)$  and their Mixup image  $\tilde{\boldsymbol{x}}_{ij}$ . The linear interpolations of images should lead to linear interpolations of corresponding logit-based probability distributions with the same mixed proportion. The interpolated logits vectors of K-1 auxiliary branches and the backbone network f are formulated as  $\{\tilde{b}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)\}_{k=1}^{K-1}$  and  $\tilde{f}(\boldsymbol{x}_i, \boldsymbol{x}_j)$ :

$$\tilde{b}_k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \lambda b_k(\boldsymbol{x}_i) + (1 - \lambda) b_k(\boldsymbol{x}_j), \ \tilde{f}(\boldsymbol{x}_i, \boldsymbol{x}_j) = \lambda f(\boldsymbol{x}_i) + (1 - \lambda) f(\boldsymbol{x}_j).$$
(9)

Logit Self-KD over auxiliary branches  $\{b_k\}_{k=1}^{K-1}$ . First, we aim to maximize the consistency of logit-based class probability distributions between linearly interpolated logits  $\{\tilde{b}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)\}_{k=1}^{K-1}$  and Mixup logits  $\{b_k(\tilde{\boldsymbol{x}}_{ij})\}_{k=1}^{K-1}$  via KL-divergence:

$$L_{b\_logit}((\boldsymbol{x}_i, \boldsymbol{x}_j), \tilde{\boldsymbol{x}}_{ij}) = \sum_{k=1}^{K-1} [L_{KL}(\sigma(\tilde{b}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)/T), \sigma(\overline{b_k(\tilde{\boldsymbol{x}}_{ij})}/T)) + L_{KL}(\sigma(b_k(\tilde{\boldsymbol{x}}_{ij})/T), \sigma(\overline{b_k(\boldsymbol{x}_i, \boldsymbol{x}_j)}/T)].$$
(10)

Here,  $L_{KL}$  is the KL-divergence, and T is a temperature following the original KD [15].  $\overline{b_k(\tilde{x}_{ij})}$  and  $\overline{\tilde{b}_k(x_i, x_j)}$  are fixed copies of  $b_k(\tilde{x}_{ij})$  and  $\tilde{b}_k(x_i, x_j)$ , respectively. As suggested by [26], the gradients through  $\overline{b_k(\tilde{x}_{ij})}$  and  $\overline{\tilde{b}_k(x_i, x_j)}$  are not propagated to avoid the model collapse issue.

**Training the self-teacher network**  $h(\cdot)$ . The idea of the self-teacher  $h(\cdot)$  is to aggregate multi-stage ensemble feature maps to construct an excellent classifier. The classifier could provide meaningful soft labels to teach the backbone classifier in an online manner. Given the K feature maps  $\{\mathcal{F}_k\}_{k=1}^K$  with the same dimension to  $h(\cdot)$ , we first concatenate them along the channel dimension and then use a  $1 \times 1$  convolution  $\text{Conv}_{1 \times 1}$  to shrink channels. Then the ensemble feature maps are followed by a linear classifier  $g_h$  to output class probability distributions. The inference graph is formulated as  $h(\{\mathcal{F}_k\}_{k=1}^K) = g_h(\text{Conv}_{1 \times 1}([\mathcal{F}_1, \cdots, \mathcal{F}_K]))$ .

Given the linearly interpolated feature maps  $\{\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)\}_{k=1}^K$  from the pair of original images  $(\boldsymbol{x}_i, \boldsymbol{x}_j)$ , we can derive the soft label  $h(\{\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)\}_{k=1}^K)$ . Given Mixup feature maps  $\{F_k(\tilde{\boldsymbol{x}}_{ij})\}_{k=1}^K$  from the Mixup image  $\tilde{\boldsymbol{x}}_{ij}$ , we can derive the soft label  $h(\{F_k(\tilde{\boldsymbol{x}}_{ij})\}_{k=1}^K)$ . For easy notation, we define  $h(\boldsymbol{x}_i, \boldsymbol{x}_j) =$  $h(\{\tilde{F}_k(\boldsymbol{x}_i, \boldsymbol{x}_j)\}_{k=1}^K)$  and  $h(\tilde{\boldsymbol{x}}_{ij}) = h(\{F_k(\tilde{\boldsymbol{x}}_{ij})\}_{k=1}^K)$ . Motivated by the success of Manifold Mixup [33], the class probabilities inferred from linear interpolated feature maps could also be supervised by interpolated labels. Thus we further train both  $h(\boldsymbol{x}_i, \boldsymbol{x}_j)$  and  $h(\tilde{\boldsymbol{x}}_{ij})$  by the interpolated label  $\tilde{y}_{ij}$  using cross-entropy loss:

$$L_{cls\_h} = L_{ce}(\sigma(h(\boldsymbol{x}_i, \boldsymbol{x}_j)), \tilde{y}_{ij}) + L_{ce}(\sigma(h(\tilde{\boldsymbol{x}}_{ij})), \tilde{y}_{ij}).$$
(11)

Benefiting from Equ. (11),  $h(\cdot)$  would learn meaningful mixed distributions from multi-stage feature maps from the view of information mixture.

Logit Self-KD over the backbone f. We adopt the self-teacher network  $h(\cdot)$  to supervise the final backbone classifier, since  $h(\cdot)$  aggregates multi-branch feature information and leads to better performance. We utilize the Mixup logit  $h(\tilde{x}_{ij})$  as the soft label to supervise the linearly interpolated logit  $\tilde{f}(x_i, x_j)$ . We also employ logit  $h(x_i, x_j)$  from the linearly interpolated feature map to supervise Mixup logit  $f(\tilde{x}_{ij})$ . The mutual distillation loss is formulated as:

$$L_{f\_logit}((\boldsymbol{x}_i, \boldsymbol{x}_j), \tilde{\boldsymbol{x}}_{ij}) = L_{KL}(\sigma(f(\boldsymbol{x}_i, \boldsymbol{x}_j)/T), \sigma(h(\tilde{\boldsymbol{x}}_{ij})/T)) + L_{KL}(\sigma(f(\tilde{\boldsymbol{x}}_{ij})/T), \sigma(h(\boldsymbol{x}_i, \boldsymbol{x}_j)/T)).$$
(12)

Here, the gradients through the soft labels of  $h(\tilde{x}_{ij})$  and  $h(x_i, x_j)$  are not propagated. In theory, the loss  $L_{b,logit}$  and  $L_{f,logit}$  can regularize the consistency between linearly interpolated class probability distributions from the pair of original images  $(x_i, x_j)$  and the distribution from Mixup image  $\tilde{x}_{ij}$ . This encourages the network to make linear predictions in-between training samples over auxiliary branches  $\{b_k\}_{k=1}^{K-1}$  and backbone network f.

#### 3.5 Overall Loss of MixSKD

We summarize the loss terms above into a unified framework:

$$L_{MixSKD} = \underbrace{L_{cls\_mixup}}_{task\ loss} + \underbrace{\beta L_{feature} + \gamma L_{dis}}_{feature\ Self-KD} + \mu(\underbrace{L_{b\_logit} + L_{cls\_h} + L_{f\_logit}}_{logit\ Self-KD}).$$
(13)

Here, we use  $\beta = 1$  to control the magnitude of feature  $l_2$  loss, where we find  $\beta \in [1, 10]$  works well. Besides, we choose  $\gamma = 1$  for discriminator loss and  $\mu = 1$  for cross-entropy or KL divergence losses since they are probability-based forms and in the same magnitude. We perform end-to-end optimization of the training graph. During the test stage, we discard attached auxiliary components, leading to no extra inference costs compared with the original network.

In theory, our proposed MixSKD encourages the network to behave linearly in the latent feature and probability distribution spaces. The linear behaviour may reduce undesirable oscillations when predicting outliers, as discussed by Zhang *et al.* [50]. Moreover, linearity is also an excellent inductive bias from the view of Occam's razor because it is one of the most straightforward behaviours.

#### 4 Experiments

#### 4.1 Experimental Setup

**Dataset.** We conduct experiments on CIFAR-100 [21] and ImageNet [8] as the standard image classification tasks and CUB-200-2011 (CUB200) [34], Standford Dogs (Dogs) [19], MIT Indoor Scene Recognition (MIT67) [28], Stanford Cars (Cars) [20] and FGVC-Aircraft (Air) [25] datasets as the fine-grained classification tasks. For downstream dense prediction tasks, we use COCO 2017 [22] for

**Table 1.** Top-1 accuracy(%) of various Self-KD (the second block) and data augmentation (the third block) methods across widely used networks for CIFAR-100 classification. The numbers in **bold** and <u>underline</u> indicate **the best** and <u>the second-best</u> results, respectively. \* denotes the result from the original paper.

Method	ResNet-18[14]	WRN-16-2[49]	DenseNet-40[17]	HCGNet-A1[42]
Parameters	11.2M	$0.7 \mathrm{M}$	$1.1\mathrm{M}$	1.1M
Baseline	$ 76.24_{\pm 0.07} $	$72.24_{\pm 0.29}$	$74.61_{\pm 0.24}$	$75.58_{\pm 0.30}$
DDGSD [37]	$76.61_{\pm 0.47}$	$72.46 \pm 0.05$	$75.87_{\pm 0.30}$	$76.50 \pm 0.18$
DKS [31]	$78.64_{\pm 0.25}$	$73.73_{\pm 0.22}$	$74.94_{\pm 0.60}$	$78.04_{\pm 0.16}$
BYOT [52]	$77.88_{\pm 0.19}$	$72.97_{\pm 0.34}$	$75.49_{\pm 0.16}$	$76.81_{\pm 0.64}$
SAD [16]	$76.40_{\pm 0.17}$	$72.62_{\pm 0.17}$	$74.77_{\pm 0.12}$	$75.86 \pm 0.57$
Tf-KD [46]	$76.61 \pm 0.34$	$72.66 \pm 0.21$	$74.68 \pm 0.17$	$75.84 \pm 0.23$
CS-KD [48]	$78.01^{*}_{\pm 0.13}$	$73.23 \pm 0.33$	$75.02_{\pm 0.37}$	$76.36_{\pm 0.30}$
FRSKD [18]	$77.71^{*}_{\pm 0.14}$	$73.27^*_{\pm 0.45}$	$74.91_{\pm 0.36}$	$77.26_{\pm 0.11}$
Cutout [9]	$76.66 \pm 0.42$	$73.66 \pm 0.18$	$75.45 \pm 0.33$	$76.63 \pm 0.15$
Mixup [50]	$78.68_{\pm 0.12}$	$73.60_{\pm 0.59}$	$75.55 \pm 0.81$	$77.89_{\pm 0.18}$
Manifold Mixup [33]	$79.29_{\pm 0.20}$	$72.53 \pm 0.08$	$75.19_{\pm 0.38}$	$77.72_{\pm 0.26}$
AutoAugment [6]	$77.97_{\pm 0.17}$	$\underline{74.16}_{\pm 0.22}$	$76.21_{\pm 0.15}$	$78.13_{\pm 0.78}$
RandAugment [7]	$76.86 \pm 0.55$	$73.87 \pm 0.08$	$76.23 \pm 0.22$	$77.56 \pm 0.50$
Random Erase [55]	$76.75 \pm 0.33$	$74.04_{\pm 0.11}$	$75.52 \pm 0.46$	$77.38 \pm 0.14$
CS-KD+Mixup [48]	$79.60^{*}_{\pm 0.31}$	$73.82_{\pm 0.18}$	$\underline{76.32}_{\pm 0.37}$	$78.03_{\pm 0.44}$
FRSKD+Mixup [18]	$78.74^{*}_{\pm 0.19}$	$73.67 \pm 0.36$	$75.56 \pm 0.07$	$78.23 \pm 0.22$
MixSKD (Ours)	80.32 $_{\pm 0.13}$	$74.89_{\pm 0.27}$	<b>76.85</b> $_{\pm 0.19}$	$78.57_{\pm 0.19}$

object detection and Pascal VOC [10], ADE20K [56] and COCO-Stuff-164K [1] for semantic segmentation. The detailed dataset descriptions and training details are provided in the Appendix. We report the mean accuracy with a standard deviation over three runs using the form of  $mean_{\pm std}$  for CIFAR-100 and fine-grained classification.

**Hyper-parameter setup.** As suggested Mixup [50], we use  $\alpha = 0.2$  in MixSKD for ImageNet and  $\alpha = 0.4$  for other image classification datasets. Moreover, we set temperature T = 3 for all datasets.

#### 4.2 Comparison with State-of-the-arts

**Performance comparison on CIFAR-100.** Table 1 shows the accuracy comparison towards state-of-the-art data augmentation and Self-KD methods on CIFAR-100. Because MixSKD integrates Mixup [50], we also compare our method with the latest Self-KD methods of CS-KD [48] and FRSKD [18] combined with Mixup. We use ResNet [14], WRN [49], DenseNet [17] and HCGNet [42] as backbone networks to evaluate the performance. Almost previous data augmentation and Self-KD methods can enhance the classification performance upon the baseline across various networks. Compared to the Mixup, MixSKD achieves an average accuracy gain of 1.23% across four networks. The result indicates that the

**Table 2.** Top-1 accuracy(%) of various Self-KD methods on ResNet-18 for finegrained classification. The numbers in **bold** and <u>underline</u> indicate **the best** and <u>the second-best</u> results, respectively. \* denotes the result from the original paper.

Method	CUB200	Dogs	MIT67	Cars	Air
Baseline	$57.48_{\pm 0.45}$	$66.83_{\pm 0.29}$	$57.81_{\pm 1.42}$	$83.50_{\pm 0.24}$	$77.07_{\pm 0.26}$
DDGSD $[37]$	$56.89_{\pm 0.42}$	$69.24_{\pm 0.84}$	$56.46_{\pm 0.59}$	$85.04_{\pm 0.11}$	$74.91_{\pm 0.97}$
DKS [31]	$63.72 \pm 0.21$	$71.07 \pm 0.07$	$61.50 \pm 0.12$	$86.13 _{\pm 0.31}$	$79.69 {\scriptstyle \pm 0.31}$
BYOT [52]	$61.77 \pm 0.43$	$69.58 \pm 0.20$	$59.03 \pm 0.42$	$85.36 _{\pm 0.18}$	$79.32 \pm 0.45$
SAD [16]	$55.51_{\pm 0.67}$	$66.10_{\pm 0.08}$	$57.46 \pm 0.79$	$82.94_{\pm 0.22}$	$73.62_{\pm 0.68}$
Tf-KD [46]	$57.44_{\pm 0.25}$	$66.57_{\pm 0.33}$	$57.51_{\pm 0.86}$	$83.59 \pm 0.49$	$76.76_{\pm 0.34}$
CS-KD [48]	$66.72^*_{\pm 0.99}$	$69.15^*{\scriptstyle \pm 0.28}$	$59.55^*{\scriptstyle \pm 0.45}$	$86.87_{\pm 0.04}$	$80.92_{\pm 0.44}$
FRSKD [18]	$65.39^{*}_{\pm 0.13}$	$70.77^*{\scriptstyle \pm 0.20}$	$61.74^*{\scriptstyle \pm 0.67}$	$84.73 _{\pm 0.03}$	$78.85 {\scriptstyle \pm 0.55}$
Mixup [50]	$65.53 \pm 0.73$	$69.30_{\pm 0.10}$	$58.83_{\pm 0.77}$	$86.10_{\pm 0.28}$	$79.94_{\pm 0.18}$
CS-KD+Mixup	$69.29^{*}_{\pm 0.64}$	$70.07^*{\scriptstyle \pm 0.14}$	$60.35^*{\scriptstyle \pm 0.85}$	$\underline{87.10}_{\pm 0.30}$	$81.13 \pm 0.45$
FRSKD+Mixup	$ 67.98^*_{\pm 0.58} $	$\underline{71.64}^*{\scriptstyle\pm0.29}$	$\underline{62.11}^*{\scriptstyle\pm0.81}$	$86.25 {\scriptstyle \pm 0.33}$	$79.97{\scriptstyle \pm 0.58}$
MixSKD (Ours)	$ 72.15_{\pm 0.53} $	$72.14_{\pm 0.22}$	$64.10_{\pm 0.45}$	$89.17_{\pm 0.08}$	$82.95_{\pm 0.31}$

superiority of MixSKD is not only attributed to the usage of Mixup. We further demonstrate that MixSKD is superior to the state-of-the-art CS-KD+Mixup and FRSKD+Mixup with average margins of 0.71% and 1.11%, respectively. These results imply that our MixSKD explores more meaningful knowledge from image mixture than the conventional training signals for Self-KD. Moreover, we find that it is hard to say which is the second-best approach since different methods are superior for various architectures. MixSKD further outperforms the powerful data augmentation method AutoAugment with an average gain of 1.04%.

**Performance comparison on fine-grained classification.** Compared with standard classification, fine-grained classification often contains fewer training samples per class and more similar inter-class semantics. This challenges a network to learn more discriminative intra-class variations. As shown in Table 2, we train a ResNet-18 on five fine-grained classification tasks. MixSKD can also achieve the best performance on fine-grained classification. MixSKD outperforms the best-second results with 2.86%, 0.50%, 1.99%, 2.07% and 1.82% accuracy gains on five datasets from left to right. The results verify that MixSKD can regularize the network to capture more discriminative features.

**Performance comparison on ImageNet classification.** ImageNet is a large-scale image classification dataset, which is a golden classification benchmark. As shown in Table 3, MixSKD achieves the best 78.76% and 94.40% top-1 and top-5 accuracies on ResNet-50 compared to advanced data augmentation and Self-KD methods. It surpasses the best-competing FRSKD+Mixup by 0.97% and 0.76% top-1 and top-5 accuracy gains. The results show the scalability of our MixSKD to work reasonably well on the large-scale dataset.

Performance comparison on transfer learning to object detection and semantic segmentation. We use the ResNet-50 pre-trained on ImageNet

Table 3. Performance comparison with data augmentation and Self-KD methods on ImageNet and downstream transfer learning for object detection on COCO 2017 and semantic segmentation on Pascal VOC, ADE20K and COCO-Stuff-164K. All results are re-implemented by ourselves. The numbers in **bold** and <u>underline</u> indicate **the best** and <u>the second-best</u> results, respectively.

	Image Cla	ssification	Detection	Semantic Segmentation			
Method	ImageNet		COCO	Pascal VOC ADE20K COCO-Stu			
	Top-1 $Acc(\%)$	Top-5 $Acc(\%)$	mAP(%)		mIoU(%)		
ResNet-50	77.08	93.20	41.0	77.09	<u>39.72</u>	<u>34.08</u>	
+Mixup [50]	77.51	93.72	<u>41.1</u>	76.71	39.69	33.38	
+Manifold Mixup [33]	77.43	93.68	40.8	75.62	38.07	32.82	
+FRSKD [18]	76.92	92.96	40.9	77.10	39.48	33.63	
+FRSKD+Mixup [18]	<u>77.79</u>	93.64	41.0	76.61	39.39	33.54	
+MixSKD (Ours)	78.76	94.40	41.5	78.78	42.37	37.12	

**Table 4.** Ablation study of loss terms in MixSKD. The first column with all '-' denotes the baseline. All results denote top-1 accuracy over ResNet-50 on ImageNet.

Type	Loss	Various loss combinations of MixSKD							
Task	$L_{cls\_mixup}$	-	$\checkmark$						
Feature Self-KD	$L_{feature}$	-	-	$\checkmark$	$\checkmark$	-	-	-	$\checkmark$
	$L_{dis}$	-	-	-	$\checkmark$	-	-	-	$\checkmark$
Logit Self-KD	Lb_logit	-	-	-	-	$\checkmark$	-	$\checkmark$	$\checkmark$
	$L_{cls\_h}$	-	-	-	-	-	$\checkmark$	$\checkmark$	$\checkmark$
	$L_{f\_logit}$	-	-	-	-	-	$\checkmark$	$\checkmark$	$\checkmark$
Top-1 Accuracy (%)		77.08	77.68	78.14	78.47	78.06	78.23	78.54	78.76

as a backbone over Cascade R-CNN [2] to perform object detection on COCO 2017 [22] and over DeepLabV3 [5] to perform semantic segmentation on Pascal VOC [10], ADE20K [56] and COCO-Stuff-164K [1]. We follow the standard data preprocessing [29,43] and implement object detection over MMDetection [4] with 1x training schedule and semantic segmentation over an open codebase released by Yang *et al.* [43]. As shown in Table 3, MixSKD achieves the best downstream performance and outperforms the best-second results with 0.4% mAP on detection and 1.68%, 2.65% and 3.04% mIoU margins on Pascal VOC, ADE20K and COCO-Stuff segmentation, respectively. In contrast, although other state-of-the-art Self-KD and data augmentation methods achieve good performance on ImageNet, they often transfer worse to downstream recognition tasks. The results indicate that our method can guide the network to learn better transferable feature representations for downstream dense prediction tasks.

# 4.3 Ablation Study and Analysis

Ablation study of loss terms. As shown in Table 4, we conduct thorough ablation experiments to evaluate the effectiveness of each component in MixSKD on the convincing ImageNet dataset:

**Table 5.** Left: top-1 accuracy (%) under FGSM white-box attack with various perturbation weights  $\epsilon$  on CIFAR-100. Middle: top-1 accuracy (%) under sensitivity analysis of temperature T over ResNet-50 on ImageNet. Right: top-1 accuracy (%) under ablation study of using different feature maps from various stages to construct the self-teacher h over ResNet-50 on ImageNet.

Method	$\epsilon = 0.0001$	$\epsilon = 0.001$	$\epsilon = 0.01$	$\epsilon = 0.1$	T Acc	Stage	Acc
WRN-16-2	60.5	57.9	37.6	8.6	1 77.94	1	78.01
+Mixup	63.1	61.1	40.7	11.4	$2 \ 78.21$	2	78.18
+CS-KD+Mixup	63.8	62.1	42.8	13.0	3 78.76	3	78.36
+FRSKD+Mixup	62.9	60.7	40.9	12.2	$4 \ 78.65$	4	78.42
+MixSKD (Ours)	67.7	65.5	<b>46.8</b>	15.3	$5\ 78.37$	All	78.76

(1) Task loss. We regard the  $L_{cls\_mixup}$  as the basic task loss that trains the network with auxiliary architectures using Mixup. It improves the baseline by a 0.60% accuracy gain. Moreover,  $L_{cls\_mixup}$  can guide the network to produce meaningful feature maps and logits from the original images and Mixup images. The information further motivates us to perform Self-KD to regularize the network. (2) Feature map Self-KD. Using  $L_{feature}$  loss to distill K feature maps achieves a 0.46% improvement over the task loss. Adding an auxiliary discriminative loss  $L_{dis}$  further enhances the performance of  $L_{feature}$  with an extra 0.33% gain. The result suggests that using an adversarial mechanism for feature distillation may encourage the network to capture more meaningful semantics than a single  $l_2$ -based mimicry loss. (3) Logit Self-KD. The distillation loss  $L_{b\_logit}$  over auxiliary branches  $\{b_k\}_{k=1}^{K-1}$  results in a 0.38% gain than the task loss. Training the self-teacher network h by  $L_{cls}h$  for providing soft labels to supervise the backbone classifier f by  $L_{f_{-logit}}$  leads to a significant 0.55% gain than the task loss. Overall, combining logit-based losses can enable the network to achieve a 0.86% gain. (4) Overall loss. Combining logit and feature Self-KD can maximize the performance gain with a 1.08% margin over the task loss.

Adversarial Robustness. Adversarial learning has shown that adding imperceptibly small but intentionally worst-case perturbations to the input image fools the network easily. Thus adversarial robustness is a crucial issue for the practical application of neural networks. As shown in Table 5 (left), we compare the robustness trained by various Self-KD methods under FGSM whitebox attack [13]. MixSKD significantly improves the robustness and surpasses other methods consistently across different perturbations. The results imply that MixSKD shows better potential in adversarial defense.

Impact of the temperature T. The temperature T is utilized to scale the predictive logits for KL-divergence based distillation losses in Eq. (10) and Eq. (12). A larger T leads to a smoother probability distribution. As shown in Table 5 (middle), we observe that MixSKD benefits most from T = 3 than others for producing soft distributions.

Effectiveness of multi-stage feature aggregation to construct selfteacher h. The self-teacher utilizes feature maps not only from the last stage



(a) Comparison of log probabilities of predicted labels (b) Miss rate of the predic-(*left*) and ground-truth labels (*right*) over common mis- tion of mixup image  $\tilde{x}_{ij}$  not classified samples trained by baseline and MixSKD. belonging to  $\{y_i, y_j\}$ .

Fig. 3. MixSKD produces more robust predictions on CIFAR-100.

but also from multiple stages to aggregate richer semantics. As shown in Table 5 (right), we demonstrate that multi-stage feature aggregation can generate higherquality soft labels to achieve better performance than any single feature map.

Meaningful predictive distribution. By encouraging linear behaviours between samples, we examine whether our MixSKD forces the network to achieve a meaningful predictive distribution. To this end, we investigate predictive softmax scores inferred from ResNet-18 trained by baseline and MixSKD on CIFAR-100. Fig. 3(a) shows the histogram statistics of log probabilities of predicted labels and ground-truth labels over common misclassified samples. Compared to the baseline, our MixSKD achieves better prediction quality. On the one hand, it decreases over-confident log probabilities of predicted labels over misclassified samples effectively. On the other hand, it assigns larger log probabilities for ground-truth labels even if the misclassification occurs.

Miss error rate in-between samples. As shown in Fig. 3(b), we evaluate the network trained by baseline, Mixup, and our MixSKD on the Mixup image  $\tilde{x}_{ij} = \lambda x_i + (1 - \lambda) x_j$ . A prediction is regarded as "miss" if it does not belong to  $\{y_i, y_j\}$ . Our MixSKD shows the minimum miss rates across various Mixup factors of  $\lambda$  within the range of [0, 1]. The results indicate that MixSKD leads to more robust predictive behaviours in-between samples.

## 5 Conclusion

This paper presents MixSKD, a powerful Self-KD method to regularize the network to behave linearly in feature maps and class probabilities between samples using Mixup images. Moreover, we also construct an auxiliary self-teacher to transfer external ensemble knowledge for the backbone network. The overall MixSKD outperforms state-of-the-art data augmentation and Self-KD methods on computer vision benchmarks. We hope this work will inspire further research towards information mixture to improve the performance of visual recognition. **Acknowledgment.** Zhulin An is the corresponding author.

# References

- Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: CVPR. pp. 1209–1218 (2018)
- Cai, Z., Vasconcelos, N.: Cascade r-cnn: high quality object detection and instance segmentation. IEEE transactions on pattern analysis and machine intelligence 43(5), 1483–1498 (2019)
- Chen, D., Mei, J.P., Wang, C., Feng, Y., Chen, C.: Online knowledge distillation with diverse peers. In: AAAI. pp. 3430–3437 (2020)
- Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
- 5. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 (2017)
- Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 113–123 (2019)
- Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPR Workshops. pp. 702–703 (2020)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE conference on computer vision and pattern recognition. pp. 248–255. IEEE (2009)
- DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552 (2017)
- Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International journal of computer vision 88(2), 303–338 (2010)
- Ge, S., Luo, Z., Zhang, C., Hua, Y., Tao, D.: Distilling channels for efficient deep tracking. IEEE Transactions on Image Processing 29, 2610–2621 (2019)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems 27 (2014)
- 13. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. ICLR (2015)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
- Hou, Y., Ma, Z., Liu, C., Loy, C.C.: Learning lightweight lane detection cnns by self attention distillation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1013–1021 (2019)
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR. pp. 4700–4708 (2017)
- Ji, M., Shin, S., Hwang, S., Park, G., Moon, I.C.: Refine myself by teaching myself: Feature refinement via self-knowledge distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10664– 10673 (2021)

- 16 C. Yang et al.
- Khosla, A., Jayadevaprakash, N., Yao, B., Li, F.F.: Novel dataset for fine-grained image categorization: Stanford dogs. In: Proc. CVPR Workshop on Fine-Grained Visual Categorization (FGVC). vol. 2. Citeseer (2011)
- Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for finegrained categorization. In: ICCV Workshops. pp. 554–561 (2013)
- Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Technical Report (2009)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- Liu, J., Qin, H., Wu, Y., Guo, J., Liang, D., Xu, K.: Coupleface: Relation matters for face recognition distillation. In: Proceedings of the European Conference on Computer Vision (2022)
- 24. Liu, J., Yu, T., Peng, H., Sun, M., Li, P.: Cross-lingual cross-modal consolidation for effective multilingual video corpus moment retrieval. In: NAACL-HLT (2022)
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151 (2013)
- Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. IEEE transactions on pattern analysis and machine intelligence 41(8), 1979–1993 (2018)
- Peng, B., Jin, X., Liu, J., Li, D., Wu, Y., Liu, Y., Zhou, S., Zhang, Z.: Correlation congruence for knowledge distillation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5007–5016 (2019)
- Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 413–420. IEEE (2009)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. IEEE transactions on pattern analysis and machine intelligence 39(6), 1137–1149 (2016)
- Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. ICLR (2015)
- Sun, D., Yao, A., Zhou, A., Zhao, H.: Deeply-supervised knowledge synergy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6997–7006 (2019)
- 32. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
- 33. Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., Bengio, Y.: Manifold mixup: Better representations by interpolating hidden states. In: International Conference on Machine Learning. pp. 6438–6447. PMLR (2019)
- Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset. Technical Report (2011)
- Wang, D., Li, Y., Wang, L., Gong, B.: Neural networks are more productive teachers than human raters: Active mixup for data-efficient knowledge distillation from a blackbox model. In: CVPR. pp. 1498–1507 (2020)
- Wang, Y., Xu, C., Xu, C., Tao, D.: Adversarial learning of portable student networks. In: AAAI. vol. 32 (2018)
- Xu, T.B., Liu, C.L.: Data-distortion guided self-distillation for deep neural networks. In: AAAI. vol. 33, pp. 5565–5572 (2019)
- Yang, C., An, Z., Cai, L., Xu, Y.: Hierarchical self-supervised augmented knowledge distillation. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence. pp. 1217–1223 (2021)

- Yang, C., An, Z., Cai, L., Xu, Y.: Knowledge distillation using hierarchical selfsupervision augmented distribution. IEEE Transactions on Neural Networks and Learning Systems (2022)
- Yang, C., An, Z., Cai, L., Xu, Y.: Mutual contrastive learning for visual representation learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 3045–3053 (2022)
- Yang, C., An, Z., Xu, Y.: Multi-view contrastive learning for online knowledge distillation. In: IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 3750–3754 (2021)
- Yang, C., An, Z., Zhu, H., Hu, X., Zhang, K., Xu, K., Li, C., Xu, Y.: Gated convolutional networks with hybrid connectivity for image classification. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12581–12588 (2020)
- Yang, C., Zhou, H., An, Z., Jiang, X., Xu, Y., Zhang, Q.: Cross-image relational knowledge distillation for semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12319– 12328 (2022)
- Yang, D., Zhou, Y., Zhang, A., Sun, X., Wu, D., Wang, W., Ye, Q.: Multi-view correlation distillation for incremental object detection. Pattern Recognition p. 108863 (2022)
- 45. Yang, Z., Li, Z., Jiang, X., Gong, Y., Yuan, Z., Zhao, D., Yuan, C.: Focal and global knowledge distillation for detectors. In: CVPR. pp. 4643–4652 (2022)
- 46. Yuan, L., Tay, F.E., Li, G., Wang, T., Feng, J.: Revisiting knowledge distillation via label smoothing regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3903–3911 (2020)
- 47. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6023–6032 (2019)
- Yun, S., Park, J., Lee, K., Shin, J.: Regularizing class-wise predictions via selfknowledge distillation. In: CVPR. pp. 13876–13885 (2020)
- 49. Zagoruyko S, K.N.: Wide residual networks. In: Proceedings of the British Machine Vision Conference (2016)
- Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. ICLR (2018)
- 51. Zhang, L., Bao, C., Ma, K.: Self-distillation: Towards efficient and compact neural networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
- Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., Ma, K.: Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In: ICCV. pp. 3713–3722 (2019)
- Zhang, L., Yu, M., Chen, T., Shi, Z., Bao, C., Ma, K.: Auxiliary training: Towards accurate and robust models. In: CVPR. pp. 372–381 (2020)
- Zhang, Y., Xiang, T., Hospedales, T.M., Lu, H.: Deep mutual learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4320–4328 (2018)
- Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: AAAI. vol. 34, pp. 13001–13008 (2020)
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: CVPR. pp. 633–641 (2017)
- 57. Zhu, X., Gong, S., et al.: Knowledge distillation by on-the-fly native ensemble. In: Advances in neural information processing systems. pp. 7517–7527 (2018)