
Self-supervised Label Augmentation via Input Transformations

Hankook Lee¹ Sung Ju Hwang^{2,3,4} Jinwoo Shin^{2,1}

Abstract

Self-supervised learning, which learns by constructing artificial labels given only the input signals, has recently gained considerable attention for learning representations with unlabeled datasets, i.e., learning without any human-annotated supervision. In this paper, we show that such a technique can be used to significantly improve the model accuracy even under fully-labeled datasets. Our scheme trains the model to learn both original and self-supervised tasks, but is different from conventional multi-task learning frameworks that optimize the summation of their corresponding losses. Our main idea is to learn a single unified task with respect to the joint distribution of the original and self-supervised labels, i.e., we augment original labels via self-supervision of input transformation. This simple, yet effective approach allows to train models easier by relaxing a certain invariant constraint during learning the original and self-supervised tasks simultaneously. It also enables an aggregated inference which combines the predictions from different augmentations to improve the prediction accuracy. Furthermore, we propose a novel knowledge transfer technique, which we refer to as self-distillation, that has the effect of the aggregated inference in a single (faster) inference. We demonstrate the large accuracy improvement and wide applicability of our framework on various fully-supervised settings, e.g., the few-shot and imbalanced classification scenarios.

1. Introduction

In recent years, *self-supervised learning* (Doersch et al., 2015) has shown remarkable success in unsupervised repre-

sentation learning for images (Doersch et al., 2015; Noroozi & Favaro, 2016; Larsson et al., 2017; Gidaris et al., 2018; Zhang et al., 2019a), natural language (Devlin et al., 2018), and video games (Anand et al., 2019). When human-annotated labels are scarce, the approach constructs artificial labels, referred to as *self-supervision*, only using the input examples and then learns their representations via predicting the labels. One of the simplest, yet effective self-supervised learning approaches is to predict which transformation t is applied to an input x from observing only the modified input $t(x)$, e.g., t can be a patch permutation (Noroozi & Favaro, 2016) or a rotation (Gidaris et al., 2018). To predict such transformations, a model should distinguish between what is semantically natural or not, and consequently, it learns high-level semantic representations of inputs.

The simplicity of transformation-based self-supervision has encouraged its wide applicability for other purposes beyond unsupervised representation learning, e.g., semi-supervised learning (Zhai et al., 2019; Berthelot et al., 2020), improving robustness (Hendrycks et al., 2019), and training generative adversarial networks (Chen et al., 2019). The prior works commonly maintain two separate classifiers (yet sharing common feature representations) for the original and self-supervised tasks, and optimize their objectives simultaneously. However, this multi-task learning approach typically provides no accuracy gain when working with fully-labeled datasets. This inspires us to explore the following question: *how can we effectively utilize the transformation-based self-supervision for fully-supervised classification tasks?*

Contribution. We first discuss our observation that the multi-task learning approach forces the primary classifier for the original task to be invariant with respect to transformations of a self-supervised task. For example, when using rotations as self-supervision (Zhai et al., 2019), which rotates each image 0, 90, 180, 270 degrees while preserving its original label, the primary classifier is forced to learn representations that are invariant to the rotations. Forcing such invariance could lead to increasing complexity of tasks since the transformations could largely change characteristics of samples and/or meaningful information for recognizing objects, e.g., image classification {6 vs. 9} or {bird vs. bat}.¹ Consequently, this could hurt the overall

¹School of Electrical Engineering, KAIST, Daejeon, Korea
²Graduate School of AI, KAIST, Daejeon, Korea ³School of Computing, KAIST, Daejeon, Korea ⁴AITRICS, Seoul, Korea. Correspondence to: Jinwoo Shin <jinwoos@kaist.ac.kr>.

¹This is because bats hang typically upside down, while birds

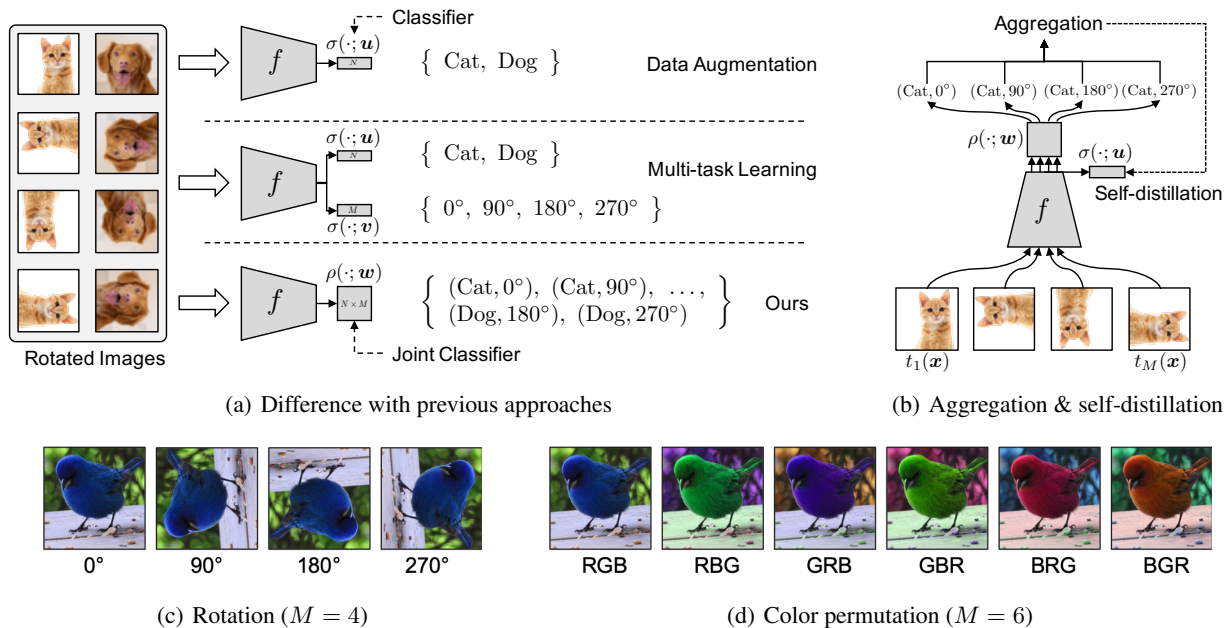


Figure 1. (a) An overview of our self-supervised label augmentation and previous approaches with self-supervision. (b) Illustrations of our aggregation method utilizing all augmented samples and self-distillation method transferring the aggregated knowledge into itself. (c) Rotation-based augmentation. (d) Color-permutation-based augmentation.

representation learning, and degrade the classification accuracy of the primary fully-supervised model (see Table 1 in Section 3.2).

To tackle this challenge, we propose a simple yet effective idea (see Figure 1(a)), which is to learn a single unified task with respect to the *joint* distribution of the original and self-supervised labels, instead of two separate tasks typically used in the prior self-supervision literature. For example, when training on CIFAR10 (Krizhevsky et al., 2009) (10 labels) with the self-supervision on rotation (4 labels), we learn the joint probability distribution on all possible combinations, i.e., 40 labels.

This label augmentation method, which we refer to as *self-supervised label augmentation* (SLA), does not force any invariance to the transformations without assumption for the relationship between the original and self-supervised labels. Furthermore, since we assign different self-supervised labels for each transformation, it is possible to make a prediction by *aggregation* across all transformations at test time, as illustrated in Figure 1(b). This can provide an (implicit) ensemble effect using a single model. Finally, to speed up the inference process without loss of the ensemble effect, we propose a novel *self-distillation* technique that transfers the knowledge of the multiple inferences into a single inference, as illustrated in Figure 1(b).

In our experiments, we consider two types of input transformations for self-supervised label augmentation, *rotation* (4 transformations) and *color permutation* (6 transformations), as illustrated in Figure 1(c) and Figure 1(d), respectively.

To demonstrate the wide applicability and compatibility of our method, we experiment with various benchmark datasets and classification scenarios, including the few-shot and imbalanced classification tasks. In all tested settings, our simple method improves the classification accuracy significantly and consistently. For example, our method achieves 8.60% and 7.05% relative accuracy gains on the standard fully-supervised task on CIFAR-100 (Krizhevsky et al., 2009) and the 5-way 5-shot task on FC100 (Oreshkin et al., 2018), respectively, over relevant baselines.²

2. Self-supervised Label Augmentation

In this section, we provide the details of our self-supervised label augmentation techniques under focusing on the fully-supervised scenarios. We first discuss the conventional multi-task learning approach utilizing self-supervised labels and its limitations in Section 2.1. Then, we introduce our learning framework which can fully utilize the power of self-supervision in Section 2.2. Here, we also propose two additional techniques: *aggregation*, which utilizes all differently augmented samples for providing an ensemble effect using a single model; and *self-distillation*, which transfers the aggregated knowledge into the model itself for accel-

do not.

²Code available at <https://github.com/hankook/SLA>.

erating the inference speed without loss of the ensemble effect.

Notation. Let $\mathbf{x} \in \mathbb{R}^d$ be an input, $y \in \{1, \dots, N\}$ be its label where N is the number of classes, \mathcal{L}_{CE} be the cross-entropy loss function, $\sigma(\cdot; \mathbf{u})$ be the softmax classifier, i.e., $\sigma_i(\mathbf{z}; \mathbf{u}) = \exp(\mathbf{u}_i^\top \mathbf{z}) / \sum_k \exp(\mathbf{u}_k^\top \mathbf{z})$, and $\mathbf{z} = f(\mathbf{x}; \boldsymbol{\theta})$ be an embedding vector of \mathbf{x} where f is a neural network with the parameter $\boldsymbol{\theta}$. We also let $\tilde{\mathbf{x}} = t(\mathbf{x})$ denote an augmented sample using a transformation t , and $\tilde{\mathbf{z}} = f(\tilde{\mathbf{x}}; \boldsymbol{\theta})$ be the embedding of the augmented sample $\tilde{\mathbf{x}}$.

2.1. Multi-task Learning with Self-supervision

In transformation-based self-supervised learning (Doersch et al., 2015; Noroozi & Favaro, 2016; Larsson et al., 2017; Gidaris et al., 2018; Zhang et al., 2019a), models learn to predict which transformation t is applied to an input \mathbf{x} given a modified sample $\tilde{\mathbf{x}} = t(\mathbf{x})$. The common approach to utilize self-supervised labels for other task is to optimize two losses of the primary and self-supervised tasks, while sharing the feature space among them (Chen et al., 2019; Hendrycks et al., 2019; Zhai et al., 2019); that is, the two tasks are trained in a multi-task learning framework. Thus, in the fully-supervised setting, one can formulate the multi-task objective \mathcal{L}_{MT} with self-supervision as follows:

$$\begin{aligned} \mathcal{L}_{\text{MT}}(\mathbf{x}, y; \boldsymbol{\theta}, \mathbf{u}, \mathbf{v}) \\ = \frac{1}{M} \sum_{j=1}^M \mathcal{L}_{\text{CE}}(\sigma(\tilde{\mathbf{z}}_j; \mathbf{u}), y) + \mathcal{L}_{\text{CE}}(\sigma(\tilde{\mathbf{z}}_j; \mathbf{v}), j), \end{aligned} \quad (1)$$

where $\{t_j\}_{j=1}^M$ is pre-defined transformations, $\tilde{\mathbf{x}}_j = t_j(\mathbf{x})$ is a transformed sample by t_j , and $\tilde{\mathbf{z}}_j = f(\tilde{\mathbf{x}}_j; \boldsymbol{\theta})$ is its embedding of the neural network f . Here, $\sigma(\cdot; \mathbf{u})$ and $\sigma(\cdot; \mathbf{v})$ are classifiers for primary and self-supervised tasks, respectively. The above loss forces the primary classifier $\sigma(f(\cdot); \mathbf{u})$ to be invariant to the transformations $\{t_j\}$. Depending on the type of transformations, forcing such invariance may not make sense, as the statistical characteristics of the augmented training samples (e.g., via rotation) could become very different from those of original training samples. In such a case, enforcing invariance to those transformations would make the learning more difficult, and can even degrade the performance (see Table 1 in Section 3.2).

In the multi-task learning objective (1), if we do not learn self-supervision, then it can be considered as a data augmentation objective \mathcal{L}_{DA} as follows:

$$\mathcal{L}_{\text{DA}}(\mathbf{x}, y; \boldsymbol{\theta}, \mathbf{u}) = \frac{1}{M} \sum_{j=1}^M \mathcal{L}_{\text{CE}}(\sigma(\tilde{\mathbf{z}}_j; \mathbf{u}), y). \quad (2)$$

This conventional data augmentation aims to improve upon the generalization ability of the target neural network f by leveraging certain transformations that can preserve their semantics, e.g., cropping, contrast enhancement, and flipping.

On the other hands, if a transformation modifies the semantics, the invariant property with respect to the transformation could interfere with semantic representation learning (see Table 1 in Section 3.2).

2.2. Eliminating Invariance via Joint-label Classifier

Our key idea is to remove the unnecessary invariant property of the classifier $\sigma(f(\cdot); \mathbf{u})$ in (1) and (2) among the transformed samples. To this end, we use a joint softmax classifier $\rho(\cdot; \mathbf{w})$ which represents the joint probability as $P(i, j | \tilde{\mathbf{x}}) = \rho_{ij}(\tilde{\mathbf{z}}; \mathbf{w}) = \exp(\mathbf{w}_{ij}^\top \tilde{\mathbf{z}}) / \sum_{k,l} \exp(\mathbf{w}_{kl}^\top \tilde{\mathbf{z}})$. Then, our training objective can be written as

$$\mathcal{L}_{\text{SLA}}(\mathbf{x}, y; \boldsymbol{\theta}, \mathbf{w}) = \frac{1}{M} \sum_{j=1}^M \mathcal{L}_{\text{CE}}(\rho(\tilde{\mathbf{z}}_j; \mathbf{w}), (y, j)), \quad (3)$$

where $\mathcal{L}_{\text{CE}}(\rho(\tilde{\mathbf{z}}; \mathbf{w}), (i, j)) = -\log \rho_{ij}(\tilde{\mathbf{z}}; \mathbf{w})$. Note that this framework only increases the number of labels, thus the number of additional parameters is negligible compared to that of the whole network, e.g., only 0.4% parameters are newly introduced when using ResNet-32 (He et al., 2016). We also remark that the above objective can be reduced to the multi-task learning objective \mathcal{L}_{MT} (1) when $\mathbf{w}_{ij} = \mathbf{u}_i + \mathbf{v}_j$ for all i, j , and the data augmentation objective \mathcal{L}_{DA} (2) when $\mathbf{w}_{ij} = \mathbf{u}_i$ for all i . From the perspective of optimization, \mathcal{L}_{MT} and \mathcal{L}_{SLA} consider the same set of multi-labels, but the former requires the additional constraint, thus it is harder to optimize than the latter. The difference between the conventional augmentation, multi-task learning and ours is illustrated in Figure 1(a). During training, we feed all M augmented samples simultaneously for each iteration as Gidaris et al. (2018) did, i.e., we minimize $\frac{1}{|B|} \sum_{(\mathbf{x}, y) \in B} \mathcal{L}_{\text{SLA}}(\mathbf{x}, y; \boldsymbol{\theta}, \mathbf{w})$ for each mini-batch B . We also assume that the first transformation is the identity function, i.e., $\tilde{\mathbf{x}}_1 = t_1(\mathbf{x}) = \mathbf{x}$.

Aggregated inference. Given a test sample \mathbf{x} or its augmented sample $\tilde{\mathbf{x}}_j = t_j(\mathbf{x})$ by a transformation t_j , we do not need to consider all $N \times M$ labels for the prediction of its original label, because we already know which transformation is applied. Therefore, we make a prediction using the conditional probability $P(i | \tilde{\mathbf{x}}_j, j) = \exp(\mathbf{w}_{ij}^\top \tilde{\mathbf{z}}_j) / \sum_k \exp(\mathbf{w}_{kj}^\top \tilde{\mathbf{z}}_j)$ where $\tilde{\mathbf{z}}_j = f(\tilde{\mathbf{x}}_j)$. Furthermore, for all possible transformations $\{t_j\}$, we aggregate the corresponding conditional probabilities to improve the classification accuracy, i.e., we train a single model, which can perform inference like an ensemble model. To compute the probability of the *aggregated inference*, we first average pre-softmax activations (i.e., logits), and then compute the softmax probability as follows:

$$P_{\text{aggregated}}(i | \mathbf{x}) = \frac{\exp(s_i)}{\sum_{k=1}^N \exp(s_k)}, \quad (4)$$

where $s_i = \frac{1}{M} \sum_{j=1}^M \mathbf{w}_{ij}^\top \tilde{\mathbf{z}}_j$. Since we assign different labels for each transformation t_j , our aggregation scheme improves accuracy significantly. Somewhat surprisingly, it achieves comparable performance with the ensemble of multiple independent models in our experiments (see Table 2 in Section 3.2). We refer to the counterpart of the aggregation as *single inference*, which uses only the non-augmented or original sample $\tilde{\mathbf{x}}_1 = \mathbf{x}$, i.e., predicts a label using $P(i|\tilde{\mathbf{x}}_1, j=1) = \exp(\mathbf{w}_{i1}^\top f(\mathbf{x}; \boldsymbol{\theta})) / \sum_k \exp(\mathbf{w}_{k1}^\top f(\mathbf{x}; \boldsymbol{\theta}))$.

Self-distillation from aggregation. Although the aforementioned aggregated inference achieves outstanding performance, it requires to compute $\tilde{\mathbf{z}}_j = f(\tilde{\mathbf{x}}_j)$ for all j , i.e., it requires M times higher computation cost than the single inference. To accelerate the inference, we perform self-distillation (Hinton et al., 2015; Lan et al., 2018) from the aggregated knowledge $P_{\text{aggregated}}(\cdot|\mathbf{x})$ to another classifier $\sigma(f(\mathbf{x}; \boldsymbol{\theta}); \mathbf{u})$ parameterized by \mathbf{u} , as illustrated in Figure 1(b). Then, the classifier $\sigma(f(\mathbf{x}; \boldsymbol{\theta}); \mathbf{u})$ can maintain the aggregated knowledge using only one embedding $\mathbf{z} = f(\mathbf{x})$. To this end, we optimize the following objective:

$$\begin{aligned} \mathcal{L}_{\text{SLA+SD}}(\mathbf{x}, y; \boldsymbol{\theta}, \mathbf{w}, \mathbf{u}) &= \mathcal{L}_{\text{SLA}}(\mathbf{x}, y; \boldsymbol{\theta}, \mathbf{w}) \\ &\quad + D_{\text{KL}}(P_{\text{aggregated}}(\cdot|\mathbf{x}) \parallel \sigma(\mathbf{z}; \mathbf{u})) \\ &\quad + \beta \mathcal{L}_{\text{CE}}(\sigma(\mathbf{z}; \mathbf{u}), y), \end{aligned} \quad (5)$$

where β is a hyperparameter and we simply choose $\beta \in \{0, 1\}$. When computing the gradient of $\mathcal{L}_{\text{SLA+SD}}$, we consider $P_{\text{aggregated}}(\cdot|\mathbf{x})$ as a constant. After training, we use $\sigma(f(\mathbf{x}; \boldsymbol{\theta}); \mathbf{u})$ for inference without aggregation.

3. Experiments

We experimentally validate our self-supervised label augmentation techniques described in Section 2. Throughout this section, we refer to data augmentation \mathcal{L}_{DA} (2) as DA, multi-task learning \mathcal{L}_{MT} (1) as MT, and our self-supervised label augmentation \mathcal{L}_{SLA} (3) as SLA for notational simplicity. We also refer baselines which use only random cropping and flipping for data augmentation (without rotation and color permutation) as ‘‘Baseline’’. Note that DA is different from Baseline because DA uses self-supervision as augmentation (e.g., rotation) while Baseline does not. After training with \mathcal{L}_{SLA} , we consider two inference schemes: the single inference $P(i|\mathbf{x}, j=1)$ and the aggregated inference $P_{\text{aggregated}}(i|\mathbf{x})$ denoted by SLA+SI and SLA+AG, respectively. We also denote the self-distillation method $\mathcal{L}_{\text{SLA+SD}}$ (5) as SLA+SD which uses only the single inference $\sigma(f(\mathbf{x}; \boldsymbol{\theta}); \mathbf{u})$.

3.1. Setup

Datasets and models. We evaluate our method on various classification datasets: CIFAR10/100 (Krizhevsky et al., 2009), Caltech-UCSD Birds or CUB200 (Wah et al.,

2011), Indoor Scene Recognition or MIT67 (Quattoni & Torralba, 2009), Stanford Dogs (Khosla et al., 2011), and tiny-ImageNet³ for standard or imbalanced image classification; mini-ImageNet (Vinyals et al., 2016), CIFAR-FS (Bertinetto et al., 2019), and FC100 (Oreshkin et al., 2018) for few-shot classification. Note that CUB200, MIT67, and Stanford Dogs are fine-grained datasets. We use 32-layer residual networks (He et al., 2016) for CIFAR and 18-layer residual networks for three fine-grained datasets and tiny-ImageNet unless otherwise stated.

Implementation details. For the standard image classification datasets, we use SGD with a learning rate of 0.1, momentum of 0.9, and weight decay of 10^{-4} . We train for 80k iterations with a batch size of 128. For the fine-grained datasets, we train for 30k iterations with a batch size of 32 because they have a relatively smaller number of training samples. We decay the learning rate by the constant factor of 0.1 at 50% and 75% iterations. We report the average accuracy of three trials for all experiments unless otherwise noted. When combining with other methods, we use publicly available codes and follow their experimental setups: MetaOptNet (Lee et al., 2019) for few-shot learning, LDAM (Cao et al., 2019) for imbalanced datasets, and FastAutoAugment (Lim et al., 2019) and CutMix (Yun et al., 2019) for advanced augmentation experiments. In the supplementary material, we provide pseudo-codes of our algorithm, which can be easily implemented.

Choices of transformation. Since using the entire input images during training is important for image classification, some self-supervision techniques are not suitable for our purpose. For example, the Jigsaw puzzle approach (Noroozi & Favaro, 2016) divides an input image to 3×3 patches and then computes their embedding separately. Prediction using such embedding performs worse than that using the entire image. To avoid this issue, we choose two transformations that use the entire input image without cropping: *rotation* (Gidaris et al., 2018) and *color permutation*. Rotation constructs $M = 4$ rotated images ($0^\circ, 90^\circ, 180^\circ, 270^\circ$) as illustrated in Figure 1(c). This transformation is widely used for self-supervision due to its simplicity (Chen et al., 2019; Zhai et al., 2019). Color permutation constructs $M = 3! = 6$ different images via swapping RGB channels as illustrated in Figure 1(d). This transformation can be useful when color information is important such as fine-grained classification datasets.

3.2. Ablation Study

Toy example for intuition. To provide intuition on the difficulty of learning an invariant property with respect to certain transformations, we here introduce simple examples: three binary digit-image classification tasks, {1 vs. 9}, {4

³ <https://tiny-imagenet.herokuapp.com/>




Figure 2. Visualization of raw pixels of 1, 4, 6 and 9 in MNIST (LeCun et al., 1998) by t-SNE (Maaten & Hinton, 2008). Colors and shapes indicate digits and rotation, respectively.

Table 1. Classification accuracy (%) of single inference using data augmentation (DA), multi-task learning (MT), and our self-supervised label augmentation (SLA) with rotation. The best accuracy is indicated as bold.

Dataset	Baseline	DA	MT	SLA+SI
CIFAR10	92.39	90.44	90.79	92.50
CIFAR100	68.27	65.73	66.10	68.68
tiny-ImageNet	63.11	60.21	58.04	63.99

Table 2. Classification accuracy (%) of the independent ensemble (IE) and our aggregation using rotation (SLA+AG). Note that a single model requires 0.46M parameters while four independent models do 1.86M parameters. The best accuracy is indicated as bold.

Dataset	Single Model		4 Models	
	Baseline	SLA+AG	IE	IE + SLA+AG
CIFAR10	92.39	94.50	94.36	95.10
CIFAR100	68.27	74.14	74.82	76.40
tiny-ImageNet	63.11	66.95	68.18	69.01

vs. 9}, and {6 vs. 9} in MNIST (LeCun et al., 1998) using linear classifiers based on raw pixel values. As illustrated in Figure 2(a), it is often easier to classify the upright digits using a linear classifier, e.g., 0.2% error when classifying only upright 6s and 9s. Note that 4 and 9 have similar shapes, so their pixel values are closer than other pairs. After rotating digits while preserving labels, the linear classifiers can still distinguish between rotated 1 and 9 as illustrated in Figure 2(b), but cannot between rotated 4, 6 and 9, as illustrated in Figure 2(c) and 2(d), e.g., 13% error when classifying rotated 6s and 9s. These examples show that linear separable data could be no longer linear separable after augmentation by some transformations such as rotation, i.e., explain why forcing an invariant property can increase the difficulty of learning tasks. However, if assigning a different label for each rotation (as we propose in this paper), then the linear classifier can classify the rotated digits, e.g., 1.1% error when classifying rotated 6s and 9s.

Comparison with DA and MT. We empirically verify that our proposed method can utilize self-supervision without loss of accuracy on fully-supervised datasets while data aug-




Figure 3. Training curves of data augmentation (DA), multi-task learning (MT), and our self-supervised label augmentation (SLA) with rotation. The solid and dashed lines indicate training and test accuracy on CIFAR100, respectively.

mentation and multi-task learning approaches cannot. To this end, we train models on generic classification datasets, CIFAR10/100 and tiny-ImageNet, using three different objectives: data augmentation \mathcal{L}_{DA} (2), multi-task learning \mathcal{L}_{MT} (1), and our self-supervised label augmentation \mathcal{L}_{SLA} (3) with rotation. As reported in Table 1, \mathcal{L}_{DA} and \mathcal{L}_{MT} degrade the performance significantly compared to the baseline that does not use the rotation-based augmentation. However, when training with \mathcal{L}_{SLA} , the performance is slightly improved. Figure 3 shows the classification accuracy of training and test samples in CIFAR100 during training. As shown in the figure, \mathcal{L}_{DA} causes a higher generalization error than others because \mathcal{L}_{DA} forces the unnecessary invariant property. Moreover, optimizing \mathcal{L}_{MT} is harder than doing \mathcal{L}_{SLA} as described in Section 2.2, thus the former achieves the lower accuracy on both training and test samples than the latter. These results show that learning invariance to some transformations, e.g., rotation, makes optimization harder and degrades the performance. Namely, such transformations should be carefully handled.

Comparison with independent ensemble. Next, to evaluate the effect of the aggregation in SLA-trained models, we compare the aggregation using rotation with independent ensemble (IE) which aggregates the pre-softmax activations (i.e., logits) over independently trained models.⁴ We here

⁴In the supplementary material, we also compare our method

Table 3. Classification accuracy (%) on various benchmark datasets using self-supervised label augmentation with rotation and color permutation. SLA+SD and SLA+AG indicate the single inference trained by $\mathcal{L}_{\text{SLA+SD}}$, and the aggregated inference trained by \mathcal{L}_{SLA} , respectively. The relative gain is shown in brackets.

Dataset	Baseline	Rotation		Color Permutation	
		SLA+SD	SLA+AG	SLA+SD	SLA+AG
CIFAR10	92.39	93.26 (+0.94%)	94.50 (+2.28%)	91.51 (-0.95%)	92.51 (+0.13%)
CIFAR100	68.27	71.85 (+5.24%)	74.14 (+8.60%)	68.33 (+0.09%)	69.14 (+1.27%)
CUB200	54.24	62.54 (+15.3%)	64.41 (+18.8%)	60.95 (+12.4%)	61.10 (+12.6%)
MIT67	54.75	63.54 (+16.1%)	64.85 (+18.4%)	60.03 (+9.64%)	59.99 (+9.57%)
Stanford Dogs	60.62	66.55 (+9.78%)	68.70 (+13.3%)	65.92 (+8.74%)	67.03 (+10.6%)
tiny-ImageNet	63.11	65.53 (+3.83%)	66.95 (+6.08%)	63.98 (+1.38%)	64.15 (+1.65%)

Table 4. Classification accuracy (%) of SLA+AG based on the set (each row) of composed transformations. We first choose subsets of rotation and color permutation (see first two columns) and compose them where M is the number of composed transformations. ALL indicates that we compose all rotations and/or color permutations. The best accuracy is indicated as bold.

Rotation T_r	Color permutation T_c	M	CUB200
0°	RGB	1	54.24
$0^\circ, 180^\circ$	RGB	2	58.92
ALL	RGB	4	64.41
0°	RGB, GBR, BRG	3	56.47
0°	ALL	6	61.10
$0^\circ, 180^\circ$	RGB, GBR, BRG	6	60.87
ALL	RGB, GBR, BRG	12	65.53
ALL	ALL	24	65.43

use four independent models (i.e., $4\times$ more parameters than ours) since IE with four models and SLA+AG have the same inference cost. Surprisingly, as reported in Table 2, the aggregation using rotation achieves competitive performance compared to the ensemble. When using both IE and SLA+AG with rotation, i.e., the same number of parameters as the ensemble, the accuracy is improved further.

3.3. Evaluation on Standard Setting

We demonstrate the effectiveness of our self-supervised augmentation method on various image classification datasets: CIFAR10/100, CUB200, MIT67, Stanford Dogs, and tiny-ImageNet. We first evaluate the effect of aggregated inference $P_{\text{aggregated}}(\cdot|\mathcal{x})$ in (4) of Section 2.2: see the SLA+AG column in Table 3. Using rotation as augmentation improves the classification accuracy on all datasets, e.g., 8.60% and 18.8% relative gain over baselines on CIFAR100 and CUB200, respectively. With color permutation, the performance improvements are less significant on CIFAR and tiny-ImageNet, but it still provides meaningful gains on fine-grained datasets, e.g., 12.6% and 10.6% relative gain on CUB200 and Stanford Dogs, respectively. In the supple-

with ten-crop (Krizhevsky et al., 2012).

Table 5. Classification error rates (%) of various augmentation methods with SLA+SD on CIFAR 10/100. We train WideResNet-40-2 (Zagoruyko & Komodakis, 2016b) and PyramidNet200 (Han et al., 2017) following the experimental setup of Lim et al. (2019) and Yun et al. (2019), respectively. The best accuracy is indicated as bold.

	CIFAR10	CIFAR100
WRN-40-2	5.24	25.63
+ Cutout	4.33	23.87
+ Cutout + SLA+SD (ours)	3.36	20.42
+ FastAutoAugment	3.78	21.63
+ FastAutoAugment + SLA+SD (ours)	3.06	19.49
+ AutoAugment	3.70	21.44
+ AutoAugment + SLA+SD (ours)	2.95	18.87
PyramidNet200	3.85	16.45
+ Mixup	3.09	15.63
+ CutMix	2.88	14.47
+ CutMix + SLA+SD (ours)	1.80	12.24

mentary material, we also provide additional experiments on large-scale datasets, e.g., iNaturalist (Van Horn et al., 2018) of 8k labels, to demonstrate the scalability of SLA with respect to the number of labels.

Since both transformations are effective on the fine-grained datasets, we also test composed transformations of the two different types of transformations for further improvements. To construct the composed ones, we first choose two subsets T_r and T_c of rotation and color permutation, respectively, e.g., $T_r = \{0^\circ, 180^\circ\}$ or $T_c = \{\text{RGB}, \text{GBR}, \text{BRG}\}$. Then, we compose them, i.e., $T = \{t_c \circ t_r : t_r \in T_r, t_c \in T_c\}$. It means that $t = t_c \circ t_r \in T$ rotates an image by t_r and then swaps color channels by t_c . As reported in Table 4, using a larger set T improves the aggregation inference further. However, under too many transformations, the aggregation performance can be degraded since the optimization becomes too harder. When using $M = 12$ transformations, we achieve the best performance, 20.8% relatively higher than the baseline on CUB200. Similar experiments on Stanford Dogs are reported in the supplementary material.

We further apply SLA+SD (that is faster than SLA+AG in inference) with existing augmentation techniques, Cutout

Table 6. Average classification accuracy (%) with 95% confidence intervals of 1000 5-way few-shot tasks on mini-ImageNet, CIFAR-FS, and FC100. † and ‡ indicates 4-layer convolutional and 28-layer residual networks (Zagoruyko & Komodakis, 2016b), respectively. Others use 12-layer residual networks as Lee et al. (2019). We follow the same experimental settings as Lee et al. (2019) did. The best accuracy is indicated as bold.

Method	mini-ImageNet		CIFAR-FS		FC100	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
MAML [†] (Finn et al., 2017)	48.70±1.84	63.11±0.92	58.9±1.9	71.5±1.0	-	-
R2D2 [†] (Bertinetto et al., 2019)	-	-	65.3±0.2	79.4±0.1	-	-
RelationNet [†] (Sung et al., 2018)	50.44±0.82	65.32±0.70	55.0±1.0	69.3±0.8	-	-
SNAIL (Mishra et al., 2018)	55.71±0.99	68.88±0.92	-	-	-	-
TADAM (Oreshkin et al., 2018)	58.50±0.30	76.70±0.30	-	-	40.1±0.4	56.1±0.4
LEO [‡] (Rusu et al., 2019)	61.76±0.08	77.59±0.12	-	-	-	-
MetaOptNet-SVM (Lee et al., 2019)	62.64±0.61	78.63±0.46	72.0±0.7	84.2±0.5	41.1±0.6	55.5±0.6
ProtoNet (Snell et al., 2017)	59.25±0.64	75.60±0.48	72.2±0.7	83.5±0.5	37.5±0.6	52.5±0.6
ProtoNet + SLA+AG (ours)	62.22±0.69	77.78±0.51	74.6±0.7	86.8±0.5	40.0±0.6	55.7±0.6
MetaOptNet-RR (Lee et al., 2019)	61.41±0.61	77.88±0.46	72.6±0.7	84.3±0.5	40.5±0.6	55.3±0.6
MetaOptNet-RR + SLA+AG (ours)	62.93±0.63	79.63±0.47	73.5±0.7	86.7±0.5	42.2±0.6	59.2±0.5

(DeVries & Taylor, 2017), CutMix (Yun et al., 2019), AutoAugment (Cubuk et al., 2019), and FastAutoAugment (Lim et al., 2019) into recent architectures (Zagoruyko & Komodakis, 2016b; Han et al., 2017). Note that SLA uses semantically-sensitive transformations for assigning different labels, while conventional data augmentation methods use semantically-invariant transformations for preserving labels. Thus, transformations using SLA and conventional data augmentation (DA) techniques do not overlap. For example, the AutoAugment (Cubuk et al., 2019) policy rotates images at most 30 degrees, while SLA does at least 90 degrees. Therefore, SLA can be naturally combined with the existing DA methods. As reported in Table 5, SLA+SD consistently reduces the classification errors. As a result, it achieves 1.80% and 12.24% error rates on CIFAR10/100, respectively. These results demonstrate the compatibility of the proposed method.

3.4. Evaluation on Limited-data Setting

Limited-data regime. Our augmentation techniques are also effective when only few training samples are available. To evaluate the effectiveness, we first construct sub-datasets of CIFAR100 via randomly choosing $n \in \{25, 50, 100, 250\}$ samples for each class, and then train models with and without our rotation-based self-supervised label augmentation. As shown in Figure 4, our scheme improves the accuracy relatively up to 37.5% under aggregation and 21.9% without aggregation.

Few-shot classification. Motivated by the above results in the limited-data regime, we also apply our SLA+AG⁵ method to solve few-shot classification, combined with re-

⁵In few-shot learning, it is hard to define the additional classifier $\sigma(f(\mathbf{x}; \boldsymbol{\theta}); \mathbf{u})$ in (5) for unseen classes when applying SLA+SD.




Figure 4. Relative improvements (%) over baselines under varying the number of training samples per class in CIFAR100.

cent approaches, ProtoNet (Snell et al., 2017) and MetaOptNet (Lee et al., 2019) specialized for this problem. Note that our method augments N -way K -shot tasks to NM -way K -shot when using M -way transformations. As reported in Table 6, ours improves consistently 5-way 1/5-shot classification accuracy on mini-ImageNet, CIFAR-FS, and FC100. For example, we obtain 7.05% relative improvements on 5-shot tasks of FC100. Here, we remark that one may obtain further improvements by applying additional data augmentation techniques to ours (and the baselines), as we shown in Section 3.3. However, we found that training with the state-of-the-art data augmentation technique and/or testing with ten-crop (Krizhevsky et al., 2012) do not always provide meaningful improvements for the few-shot experiments, e.g., the AutoAugment (Cubuk et al., 2019) policy and ten-crop provide marginal (<1%) accuracy gain on FC100 under ProtoNet in our experiments.

Imbalanced classification. Finally, we consider a setting

Table 7. Classification accuracy (%) on imbalanced datasets of CIFAR10/100. Imbalance Ratio is the ratio between the numbers of samples of most and least frequent classes. We follow the experimental settings of Cao et al. (2019). The best accuracy is indicated as bold, and we use brackets to report the relative accuracy gains over each counterpart that does not use SLA.

Imbalance Ratio (N_{\max}/N_{\min})	Imbalanced CIFAR10		Imbalanced CIFAR100	
	100	10	100	10
Baseline	70.36	86.39	38.32	55.70
Baseline + SLA+SD (ours)	74.61 (+6.04%)	89.55 (+3.66%)	43.42 (+13.3%)	60.79 (+9.14%)
CB-RW (Cui et al., 2019)	72.37	86.54	33.99	57.12
CB-RW + SLA+SD (ours)	77.02 (+6.43%)	89.50 (+3.42%)	37.50 (+10.3%)	61.00 (+6.79%)
LDAM-DRW (Cao et al., 2019)	77.03	88.16	42.04	58.71
LDAM-DRW + SLA+SD (ours)	80.24 (+4.17%)	89.58 (+1.61%)	45.53 (+8.30%)	59.89 (+1.67%)

of imbalanced training datasets, where the number of instances per class largely differs and some classes have only a few training instances. For this experiment, we combine our SLA+SD method with two recent approaches, the Class-Balanced (CB) loss (Cui et al., 2019) and LDAM (Cao et al., 2019), specialized for this problem. Under imbalanced datasets of CIFAR10/100, which have long-tailed label distributions, our approach consistently improves the classification accuracy as reported in Table 7 (e.g., up to 13.3% relative gain on an imbalanced CIFAR100 dataset). The results show the wide applicability of our self-supervised label augmentation. Here, we emphasize that all tested methods (including our SLA+SD) have the same inference time.

4. Related Work

Self-supervised learning. For representation learning in unlabeled datasets, self-supervised learning approaches construct artificial labels (referred to as self-supervision) using only input signals, and then learn to predict them. The self-supervision can be constructed in various ways. A simple one of them is transformation-based approaches (Doersch et al., 2015; Noroozi & Favaro, 2016; Larsson et al., 2017; Gidaris et al., 2018; Zhang et al., 2019a). They first modify inputs by a transformation, e.g., rotation (Gidaris et al., 2018) and patch permutation (Noroozi & Favaro, 2016), and then assign the transformation as the input’s label.

Another approach is clustering-based (Bojanowski & Joulin, 2017; Caron et al., 2018; Wu et al., 2018; YM. et al., 2020). They first perform clustering using the current model, and then assign labels using the cluster indices. When performing this procedure iteratively, the quality of representations is gradually improved. Instead of clustering, Wu et al. (2018) assign different labels for each sample, i.e., consider each sample as a cluster.

While the recent clustering-based approaches outperform transformation-based ones for unsupervised learning, the latter is widely used for other purposes due to its simplicity,

e.g., semi-supervised learning (Zhai et al., 2019; Berthelot et al., 2020), improving robustness (Hendrycks et al., 2019), and training generative adversarial networks (Chen et al., 2019). In this paper, we also utilize transformation-based self-supervision, but aim to improve accuracy under full-supervised datasets.

Self-distillation. Hinton et al. (2015) propose a knowledge distillation technique, which improves a network via transferring (or distilling) knowledge of a pre-trained larger network. There are many advanced distillation techniques (Zagoruyko & Komodakis, 2016a; Park et al., 2019; Ahn et al., 2019; Tian et al., 2020), but they should train the larger network first, which leads to high training costs. To overcome this shortcoming, self-distillation approaches, which transfer own knowledge into itself, have been developed. (Lan et al., 2018; Zhang et al., 2019b; Xu & Liu, 2019). They utilize partially-independent architectures (Lan et al., 2018), data distortion (Xu & Liu, 2019), or hidden layers (Zhang et al., 2019b) for distillation. While these approaches perform distillation on the same label space, our framework transfers knowledge between different label spaces augmented by self-supervised transformations. Thus our approach could enjoy an orthogonal usage with the existing ones; for example, one can distill aggregated knowledge $P_{\text{aggregated}}$ (4) into hidden layers as Zhang et al. (2019b) did.

5. Conclusion

We proposed a simple yet effective approach utilizing self-supervision on fully-labeled datasets via learning a single unified task with respect to the joint distribution of the original and self-supervised labels. We think that our work could bring in many interesting directions for future research; for instance, one can revisit prior works on applications of self-supervision, e.g., semi-supervised learning with self-supervision (Zhai et al., 2019; Berthelot et al., 2020). Applying our joint learning framework to fully-supervised tasks other than the few-shot or imbalanced classification task, or

learning to select tasks that are helpful toward improving the main task prediction accuracy, are other interesting future research directions.

Acknowledgements

This work was partly supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)). This work was mainly supported by Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-IT1902-06.

References

- Ahn, S., Hu, S. X., Damianou, A., Lawrence, N. D., and Dai, Z. Variational information distillation for knowledge transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9163–9171, 2019.
- Anand, A., Racah, E., Ozair, S., Bengio, Y., Côté, M.-A., and Hjelm, R. D. Unsupervised state representation learning in atari. In *Advances in Neural Information Processing Systems*, pp. 8766–8779, 2019.
- Berthelot, D., Carlini, N., Cubuk, E. D., Kurakin, A., Sohn, K., Zhang, H., and Raffel, C. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HklkeR4KPB>.
- Bertinetto, L., Henriques, J. F., Torr, P., and Vedaldi, A. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HyxnZh0ct7>.
- Bojanowski, P. and Joulin, A. Unsupervised learning by predicting noise. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 517–526. JMLR. org, 2017.
- Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*, 2019.
- Caron, M., Bojanowski, P., Joulin, A., and Douze, M. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 132–149, 2018.
- Chen, T., Zhai, X., Ritter, M., Lucic, M., and Houlsby, N. Self-supervised gans via auxiliary rotation loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12154–12163, 2019.
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 113–123, 2019.
- Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9268–9277, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- DeVries, T. and Taylor, G. W. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- Doersch, C., Gupta, A., and Efros, A. A. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1422–1430, 2015.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1126–1135, 2017.
- Gidaris, S., Singh, P., and Komodakis, N. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=S1v4N2l0->.
- Han, D., Kim, J., and Kim, J. Deep pyramidal residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5927–5935, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Hendrycks, D., Mazeika, M., Kadavath, S., and Song, D. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in Neural Information Processing Systems*, 2019.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

- Khosla, A., Jayadevaprakash, N., Yao, B., and Fei-Fei, L. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, June 2011.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 1097–1105, 2012.
- Lan, X., Zhu, X., and Gong, S. Knowledge distillation by on-the-fly native ensemble. In *Advances in Neural Information Processing Systems*, pp. 7528–7538. Curran Associates Inc., 2018.
- Larsson, G., Maire, M., and Shakhnarovich, G. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6874–6883, 2017.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lee, K., Maji, S., Ravichandran, A., and Soatto, S. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10657–10665, 2019.
- Lim, S., Kim, I., Kim, T., Kim, C., and Kim, S. Fast autoaugment. In *Advances in Neural Information Processing Systems*, 2019.
- Maaten, L. v. d. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov): 2579–2605, 2008.
- Mishra, N., Rohaninejad, M., Chen, X., and Abbeel, P. A simple neural attentive meta-learner. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B1DmUzWAW>.
- Noroozi, M. and Favaro, P. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pp. 69–84. Springer, 2016.
- Oreshkin, B., López, P. R., and Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 721–731, 2018.
- Park, W., Kim, D., Lu, Y., and Cho, M. Relational knowledge distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3967–3976, 2019.
- Quattoni, A. and Torralba, A. Recognizing indoor scenes. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 413–420. IEEE, 2009.
- Rusu, A. A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., and Hadsell, R. Meta-learning with latent embedding optimization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BJgklhAcK7>.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 4077–4087, 2017.
- Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P. H., and Hospedales, T. M. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1199–1208, 2018.
- Tian, Y., Krishnan, D., and Isola, P. Contrastive representation distillation. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkgpBJrtvS>.
- Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., and Belongie, S. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8769–8778, 2018.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pp. 3630–3638, 2016.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3733–3742, 2018.
- Xu, T.-B. and Liu, C.-L. Data-distortion guided self-distillation for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 5565–5572, 2019.
- YM., A., C., R., and A., V. Self-labelling via simultaneous clustering and representation learning. In *International*

Conference on Learning Representations, 2020. URL <https://openreview.net/forum?id=Hyx-jyBFPr>.

Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. Cutmix: Regularization strategy to train strong classifiers with localizable features. *arXiv preprint arXiv:1905.04899*, 2019.

Zagoruyko, S. and Komodakis, N. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016a.

Zagoruyko, S. and Komodakis, N. Wide residual networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, pp. 87.1–87.12, 2016b. ISBN 1-901725-59-6. doi: 10.5244/C.30.87. URL <https://dx.doi.org/10.5244/C.30.87>.

Zhai, X., Oliver, A., Kolesnikov, A., and Beyer, L. S⁴L: Self-supervised semi-supervised learning. *arXiv preprint arXiv:1905.03670*, 2019.

Zhang, L., Qi, G.-J., Wang, L., and Luo, J. Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2547–2555, 2019a.

Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., and Ma, K. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3713–3722, 2019b.