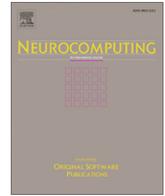




Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Self-paced data augmentation for training neural networks

Tomoumi Takase*, Ryo Karakida, Hideki Asoh

Artificial Intelligence Research Center, National Institute of Advanced Industrial Science and Technology, Tokyo, Japan



ARTICLE INFO

Article history:

Received 30 September 2020

Revised 20 February 2021

Accepted 22 February 2021

Available online 10 March 2021

Communicated by Zidong Wang

Keywords:

Self-paced augmentation

Neural network

Deep learning

Data augmentation

Curriculum learning

Loss function

ABSTRACT

Data augmentation is widely used for machine learning; however, an effective method to apply data augmentation has not been established even though it includes several factors that should be tuned carefully. One such factor is sample suitability, which involves selecting samples that are suitable for data augmentation. A typical method that applies data augmentation to all training samples disregards sample suitability, which may reduce classifier performance. To address this problem, we propose the self-paced augmentation (SPA) to automatically and dynamically select suitable samples for data augmentation when training a neural network. The proposed method mitigates the deterioration of generalization performance caused by ineffective data augmentation. We discuss two reasons the proposed SPA works relative to curriculum learning and desirable changes to loss function instability. Experimental results demonstrate that the proposed SPA can improve the generalization performance, particularly when the number of training samples is small. In addition, the proposed SPA outperforms the state-of-the-art RandAugment method.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Various techniques are used to improve the learning performance of neural networks. For example, data augmentation, which increases the number of samples by adding deformation to the original samples, is used in various applications, e.g., image classification and speech recognition [1,2]. A typical example of data augmentation is affine transformation such as flipping and rotation. Methods based on principal component analysis and discrete cosine transform have been also proposed [43].

However, a significant limitation of data augmentation is that an appropriate method to apply this technique has not yet been established, even though it involves several factors that require careful tuning. Sample suitability is one such factor. In practical application, data augmentation is typically applied to all training samples regardless of sample suitability; however, this can reduce classifier performance. Thus, in this study, we focused on this problem because, to the best of our knowledge, it has not been addressed in previous studies.

In this paper, we propose the simple and practical self-paced augmentation (SPA) method to automatically and dynamically

select suitable samples for data augmentation while training a neural network. The proposed method is easy to use and effective model training. We experimentally confirm the effectiveness of the proposed SPA in experiments using several typical benchmark datasets, including MNIST [3], CIFAR-10 [4], and Tiny ImageNet [5]. In addition, we compare the proposed SPA to the state-of-the-art RandAugment method [6]. We provide a detailed investigation of the effect of SPA on training performance and a simple comparison of prediction accuracy.

The remainder of this paper is organized as follows. In Section 2, we introduce related work, and, in Section 3, we describe our motivations and the procedure of the proposed SPA method. In Section 4, we report experimental results obtained for several numbers of samples, augmentation techniques, and datasets. In Section 5, we discuss two advantages of SPA relative to curriculum learning [7] and desirable changes in loss function instability. Finally, we conclude the paper in Section 6, including a summary of potential future work.

2. Related work

In addition to typical data augmentation operations with affine transformation, several effective techniques have been proposed. For example, mixup [8] generates a new sample from linear

* Corresponding author.

E-mail addresses: takase-tomoumi@aist.go.jp (T. Takase), karakida.ryo@aist.go.jp (R. Karakida), h.asoh@aist.go.jp (H. Asoh).

interpolation of two samples using a mix rate sampled from a beta distribution. In addition, the cutout [9] and random erasing [10] methods apply masks to parts of images, and the random image cropping and patching (RICAP) [11] method randomly crops four images and then patches to generate a new image. These data augmentation techniques are considered to provide a regularization effect [44,45].

Recently, reducing the heuristics of data augmentation has attracted increasing attention. For example, the AutoAugment method [12] searches appropriate data augmentation policies using reinforcement learning. Although this provides high generalization performance, parameter search using reinforcement learning prior to training requires significant time. Lim et al. [13] reported that the parameter search requires 15,000 GPU-hours on the ImageNet [5] dataset. They proposed the Fast AutoAugment method, which reduced the time cost to 450 GPU-hours. However, the computation time remains considerably longer than that of conventional augmentation (CA) methods; thus, the Fast AutoAugment is not considered an easy replacement for CA methods. One of other interesting studies on data augmentation was conducted by Zheng et al. [42], who proposed an efficient data augmentation framework and demonstrated that data augmentation in testing time implicit model ensemble, improving generalization performance.

Dynamically optimizing hyperparameters during training is effective relative to reducing computation time and constructing a practical method. Previous studies have addressed dynamic optimization of hyperparameters for neural networks other than data augmentation. For example, a previous study proposed a method to determine the learning rate dynamically such that the loss function on training data is minimized [14]. Other studies have proposed methods to dynamically tune the optimizer when training using meta-learning, which is a method to learn how to learn base training from higher-order training [15,16]. Note that data augmentation should also be tuned dynamically during training based on some robust strategy.

RandAugment [6] randomly and dynamically selects augmentation techniques during training using a smaller search space than AutoAugment. As a result, this method accelerates training and significantly reduces computation compared to AutoAugment. However, this method only selects the suitable data augmentation techniques and tunes the hyperparameters in those methods, i.e., this method does not address the sample suitability problem. Identifying suitable samples for data augmentation and applying augmentation to those samples is significant relative to improving generalization performance. Therefore, we addressed the sample suitability problem in this study.

3. Proposed self-paced augmentation

3.1. Motivation

As described in the Section 1, data augmentation is conventionally applied to all training samples consistently during the entire process training. This is natural; however, this does not necessarily improve training performance because it does not consider sample suitability. This can be intuitively understood in Fig. 1, where the training loss transition when data augmentation was used is greater than that when data augmentation was not used. Thus, effective samples for data augmentation must be identified according to changing training conditions.

To determine sample suitability, we focused on the following two points. Here, one perspective we considered is curriculum learning [7], which is a strategy that transitions training from easy to difficult samples in a gradual manner. This is discussed and ver-

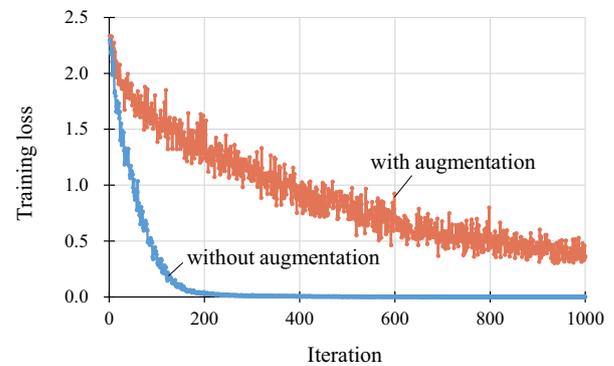


Fig. 1. Training loss transitions when data augmentation was and was not used. Here, a small CNN was trained on the CIFAR-10 dataset. Flip and random erasing data augmentation techniques were applied to all training samples. This result is part of the results shown in Fig. 5.

ified in detail in Sections 5.1 and 5.2. We also considered loss function behavior. As shown in Fig. 1, training loss was high and unstable when data augmentation was applied, i.e., the training loss fluctuated significantly in each iteration. This training loss trend does not work effectively relative to realizing high generalization performance, which is discussed and verified in detail in Sections 5.3 and 5.4. We designed the proposed SPA such that training is performed well in light of these considerations.

3.2. Procedure

Differing from conventional application of data augmentation, which is applied to all training samples, the proposed SPA method dynamically determines and selects samples that are suitable for data augmentation during training. The procedure of the proposed SPA is shown in Fig. 2.

The procedure was divided into two steps. In Step 1, samples are input to a neural network prior to applying data augmentation to samples. Then, training losses are calculated using the outputs and labels, and are used for the augmentation process. Note that the weights of the neural network are not updated at this time. In Step 2, data augmentation is applied to only samples that produce training loss greater than predetermined threshold λ . λ is fixed in this method; however, the set of samples selected for data augmentation dynamically changes as training proceeds because the training loss for each sample changes during training. A detailed analysis of this process is given in Section 5. Then, regardless of whether data augmentation is applied, all samples in a minibatch are input to the neural network and the weights are updated. This overall process is performed for each minibatch.

Although Fig. 2 represents an example in which only rotation was applied to the samples, several data augmentation techniques, e.g., flip, rotation, and mixup [8], can also be applied to samples simultaneously. In addition, samples do not always change when data augmentation is applied because the degree of augmentation is determined randomly each time, e.g., if horizontal flip is applied, a sample is flipped with a probability of 0.5 and keeps unchanged with a probability of 0.5.

The procedure of this algorithm of the proposed method is outlined in Algorithm 1, where f and g are forward and back propagation, respectively, Loss is the loss function, $DA(x_n, y_n)$ is the sample to which data augmentation was applied, and θ is the neural network parameters. Although the same samples are input twice to the neural network in a single minibatch, only forward propagation is performed in the first input, which does not require significant computation time.

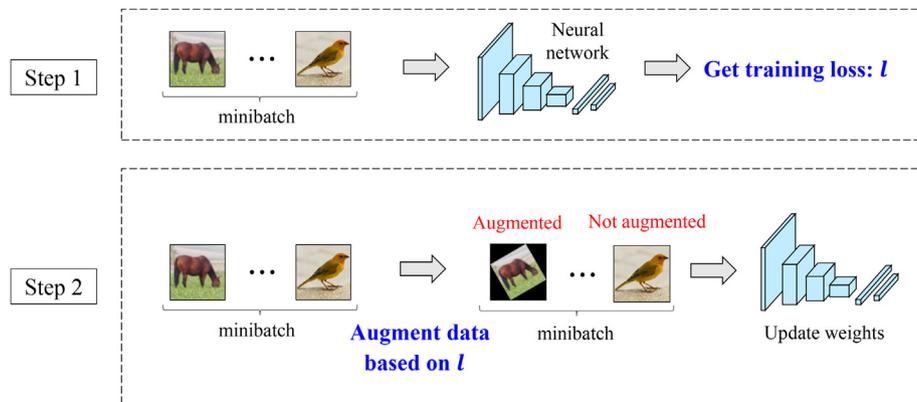


Fig. 2. SPA training procedure. Steps 1 and 2 are performed for each minibatch.

4. Experiments

4.1. Experimental setup

Algorithm 1 SPA Training

```

1: Input:  $(\mathbf{x}, \mathbf{y})$ 
2: Output:  $\theta$ 
3:  $l \leftarrow \text{Loss}(f(\mathbf{x}; \theta), \mathbf{y})$ 
4: for  $n$  from 0 to  $N - 1$  do
5:   if  $l_n \geq \lambda$  then
6:      $(x_n, y_n) \leftarrow \text{DA}(x_n, y_n)$ 
7:   end if
8: end for
9:  $l \leftarrow \text{Loss}(f(\mathbf{x}; \theta), \mathbf{y})$ 
10:  $\theta \leftarrow g(\mathbf{x}; \theta; l)$ 
11: return  $\theta$ 

```

In our experiments to evaluate the proposed SPA method, we focused on investigating the sample suitability problem in data augmentation, which is a new problem not considered by existing methods. Therefore, the main target of comparison was the general use of data augmentation, which augments all training samples. Thus, we compared the performances of the proposed SPA method to that of CA, which augments all training samples, and training with no augmentation (NA), which correspond to $\lambda = 0$ and $\lambda = \infty$ in the proposed SPA, respectively. For SPA, we considered several values for threshold parameter λ (e.g., 0.01, 0.1, and 1).

Acquiring a large number of training samples is often difficult in terms of cost and availability. Data augmentation is used to train a model with a sufficient number of samples in such cases; thus, we focused on verifying data augmentation performance on a small number of training samples. In addition, the sample suitability problem addressed by the proposed SPA method may not be overly problematic in experimental setups with a large number of training samples. However, this is certainly relevant in setups with a small number of samples. Here, in addition to the original datasets with all full training samples, we used datasets prepared by randomly extracting part of the training samples from the original datasets.

To demonstrate the effectiveness of the proposed SPA on various datasets, we used several data augmentation techniques on several image benchmarks for machine learning, i.e., the CIFAR-10 and CIFAR-100 [4], MNIST [3], Fashion-MNIST [17], SVHN [18], STL-10 [19], and Tiny ImageNet [5] datasets. Training was

always performed with a batch size of 100 for reduced datasets and with a batch size of 256 for full datasets.

Here, the initial weights were sampled from uniform distribution, and we used the Adam optimizer [20] with an initial learning rate of 0.001 for reduced datasets and used the momentum stochastic gradient descent (SGD) with an initial learning rate of 0.005 for SVHN (core) and with that of 0.1 for other full datasets. The small CNN was trained for 1000 epochs, and other CNNs were trained for only 200 epochs. The test accuracy on the original test data was calculated for each epoch. The best test accuracies obtained during training were compared, and we evaluated the mean and standard error for five trials with different initial weight values. Note that data augmentation was not applied to the test data.

We also conducted experiments on several data augmentation techniques. With the flip method, horizontal and vertical flips were applied independently at a probability of 0.5. For cropping, 80% of each image was randomly cropped and enlarged to the original size. With translation, each image was translated with an offset of $tx/128$, where t is a value in the range $[-25, 25]$ and the image size is $x \times x$. With rotation, images were rotated using a random value in the range $[0, 360]$. Then, the images were downsized appropriately. In cutout [9], the mask size was set to 0.5, and, in random erasing [10], the mask size was set to a random value from 0.02 to 0.4 of the image size, and the aspect ratio was selected randomly in the range $[0.3, 3]$. In mixup [8] and RICAP [11], parameter α of the beta distribution to determine the mix rate was set to 1.

We compared the generalization performance of SPA to that of the conventional methods in four experiments (Sections 4.2–4.5). In the first experiment, we investigated the effects of the number of training samples. Here, we used a convolutional neural network (CNN) with a [conv-conv-pool-conv-conv-pool-fc] structure, where conv is a 3×3 convolution layer, pool is a 2×2 pooling layer, and fc is a fully connected layer (referred to as “small CNN” in the experimental results). In this CNN, batch normalization [21] and the rectified linear unit [22] were inserted after each layer, with the exception of the final layer, where the softmax function was used. We trained the model with flip or translation techniques on the CIFAR-10 or MNIST datasets.

In the second experiment, we used a small number of training samples. Here, we evaluated several data augmentation techniques on the CIFAR-10, Fashion-MNIST, and SVHN datasets and only used the flip and RICAP methods (which demonstrated high accuracy on the CIFAR-10 dataset) on the STL-10, CIFAR-100, and Tiny-ImageNet datasets. Considering the number of classes, we used 10,000 samples for CIFAR-100, 20,000 samples for Tiny-ImageNet, and 1000 samples for the other datasets. For the models, we used ResNet34 [23] for CIFAR-100 and Tiny ImageNet, and we used the small CNN for the other models.

In the third experiment, we trained WideResNet28-10 [24] using various data augmentation techniques on the full CIFAR-10, Fashion-MNIST, and SVHN (core) datasets. The computation time was also compared between SPA and CA when CIFAR-10 was used.

In the fourth experiment, we compared the proposed SPA method to RandAugment [6], which is a dynamic augmentation method. Although it has demonstrated excellent performance on the full datasets and large CNNs with well-tuned hyperparameters, the effect on a small number of datasets has not been investigated sufficiently. Thus, we compared the performance of SPA to that of RandAugment on 1000 training samples and full datasets. Here, the CIFAR-10, Fashion-MNIST, and SVHN datasets were used in to facilitate the comparison of these methods. With RandAugment, we selected and applied one data augmentation technique dynamically from 14 basic techniques, e.g., equalizing the image histogram and inverting all pixels above a threshold value of magnitude according to the settings used in the original study [6], and we used several values (0, ..., 10) for the magnitude of data augmentation.

4.2. Effects of number of samples

The results shown in Table 1 indicate that the proposed SPA method tends to be effective when the number of samples is relatively small. In addition, the proposed method outperformed CA. In addition, the accuracy values obtained by the CA method were always worse than those of NA in the experiments with the MNIST dataset and the flip method because numbers can be confused with other numbers when flip is applied; thus, flip is unsuitable for images of such digits. This demonstrates that data augmentation does not always improve generalization performance because some augmentation techniques are not suitable for particular datasets. In addition, knowing the suitability of data augmentation for each dataset in advance is often difficult, as shown in the subsequent experiments. However, even in such cases, the accuracies

of the proposed SPA were close to or greater than those of NA. This implies that SPA can mitigate deterioration in generalization performance when unsuitable augmentation is applied because the proposed SPA method does not always apply augmentation to all samples during training.

4.3. Experiments on reduced datasets

The results obtained with a small number of training samples on the various datasets are shown in Table 2. The results verify that the proposed SPA method outperformed CA and NA in most cases. Overall, high accuracy was generally achieved when λ was set to 0.1.

Knowing the suitability of augmentation in advance is often difficult. For example, the rotation technique and its parameter value we used did not demonstrate high performance overall, and CA showed significantly reduced accuracy compared to NA. However, the proposed SPA tended to demonstrate high performances when the rotation technique was used. In addition, although cases that combined two augmentation techniques, e.g., flip and crop, did not demonstrate higher accuracy than when a single method was applied, e.g., only crop, the proposed SPA was effective in such cases.

4.4. Experiments on full datasets

The results obtained using various data augmentation techniques on the full CIFAR-10, Fashion-MNIST, and SVHN datasets are given in Table 3. As shown, the proposed SPA outperformed CA and NA in several cases; however, it did not always outperform these methods. When training with full datasets, a sufficient number of training samples were used; thus, the performance of SPA could be comparable with that of the CA and NA methods. Although the effectiveness of SPA is obviously smaller compared to cases with a small number of training samples, it is possible that training models using the proposed SPA will yield better

Table 1

Best test accuracy (%) when training the small CNN using several number of training samples. Values in parentheses are the standard error. Numbers in bold indicate the highest accuracy, and those in underline indicate considerably greater values than other methods. A data augmentation is regarded as being unsuitable when the accuracy for CA is considerably smaller than that for NA. "Number" denotes the number of training samples, and "trans." denotes the translation method.

Dataset	Number	Aug.	SPA ($\lambda = 0.1$)	CA ($\lambda = 0$)	NA ($\lambda = \infty$)	Unsuitable
CIFAR-10	100	flip	30.87 (± 0.33)	30.84 (± 0.25)	28.52 (± 0.56)	
		translation	32.28 (± 0.48)	33.00 (± 0.48)	28.52 (± 0.56)	
	500	flip	45.09 (± 0.21)	43.84 (± 0.30)	42.23 (± 0.36)	
		translation	48.29 (± 0.59)	49.11 (± 0.75)	42.23 (± 0.36)	
	1000	flip	52.78 (± 0.38)	50.19 (± 0.35)	48.47 (± 0.49)	
		translation	56.96 (± 0.42)	56.16 (± 0.28)	48.47 (± 0.49)	
	5000	flip	66.44 (± 0.41)	63.51 (± 0.39)	64.01 (± 0.19)	
		translation	72.80 (± 0.18)	70.35 (± 0.13)	64.01 (± 0.19)	
	10,000	flip	69.86 (± 0.14)	68.10 (± 0.09)	69.28 (± 0.35)	✓
		translation	75.98 (± 0.11)	75.08 (± 0.15)	69.28 (± 0.35)	
	(all) 50,000	flip	77.71 (± 0.13)	77.70 (± 0.12)	78.32 (± 0.11)	✓
		translation	83.02 (± 0.07)	83.11 (± 0.18)	78.32 (± 0.11)	
MNIST	100	flip	85.71 (± 0.65)	74.25 (± 0.90)	84.12 (± 0.92)	✓
		translation	92.14 (± 0.18)	90.02 (± 0.42)	84.12 (± 0.92)	
	500	flip	96.81 (± 0.11)	89.82 (± 0.14)	96.69 (± 0.02)	✓
		translation	97.72 (± 0.05)	97.20 (± 0.07)	96.69 (± 0.02)	
	1000	flip	97.77 (± 0.04)	93.36 (± 0.02)	97.77 (± 0.10)	✓
		translation	98.52 (± 0.02)	98.23 (± 0.04)	97.77 (± 0.10)	
	5000	flip	98.88 (± 0.05)	96.91 (± 0.05)	98.80 (± 0.04)	✓
		translation	99.15 (± 0.02)	99.00 (± 0.01)	98.80 (± 0.04)	
	10,000	flip	99.27 (± 0.01)	97.86 (± 0.03)	99.18 (± 0.04)	✓
		translation	99.40 (± 0.02)	99.21 (± 0.01)	99.18 (± 0.04)	
	(all) 60,000	flip	99.58 (± 0.02)	98.76 (± 0.06)	99.58 (± 0.04)	✓
		translation	99.60 (± 0.01)	99.62 (± 0.01)	99.58 (± 0.04)	

Table 2

Best test accuracy (%) when training the small CNN using various data augmentation techniques with a small number of samples from various datasets. We evaluated SPA using several λ values, CA, and NA. The results for NA have been added under each dataset in Dataset column. Values in parentheses after accuracies are the standard error. 1 k and 10 k denote the number of training samples. Numbers in bold indicate the highest accuracy, and underlined values indicate considerably greater values than those of CA and NA.

Dataset	Augmentation	SPA ($\lambda = 0.01$)	SPA ($\lambda = 0.1$)	SPA ($\lambda = 1$)	CA ($\lambda = 0$)	
CIFAR-10 (1 k) (NA: 48.47 (± 0.49))	flip	51.62 (± 0.36)	<u>52.78</u> (± 0.38)	52.28 (± 0.20)	50.19 (± 0.35)	
	crop	53.40 (± 0.55)	53.52 (± 0.60)	51.66 (± 0.51)	52.71 (± 0.47)	
	translation	<u>57.63</u> (± 0.29)	56.87 (± 0.39)	54.05 (± 0.26)	56.04 (± 0.36)	
	rotation	40.67 (± 0.37)	47.72 (± 0.48)	<u>50.37</u> (± 0.29)	39.33 (± 0.19)	
	mixup [3]	46.53 (± 0.23)	<u>50.47</u> (± 0.22)	49.94 (± 0.18)	46.45 (± 0.19)	
	cutout [4]	49.15 (± 0.21)	<u>49.55</u> (± 0.30)	49.40 (± 0.20)	48.17 (± 0.28)	
	random erasing [5]	<u>50.81</u> (± 0.50)	50.57 (± 0.38)	49.63 (± 0.58)	49.32 (± 0.39)	
	RICAP [6]	53.80 (± 0.28)	<u>57.73</u> (± 0.10)	54.81 (± 0.41)	54.00 (± 0.27)	
	flip and crop	55.31 (± 0.44)	<u>55.60</u> (± 0.58)	51.89 (± 0.45)	54.31 (± 0.52)	
	translation and rotation	38.04 (± 0.39)	42.20 (± 0.57)	43.50 (± 0.43)	36.36 (± 0.58)	
	Fashion-MNIST (1 k) (NA: 83.79 (± 0.17))	flip	84.82 (± 0.19)	84.83 (± 0.16)	84.32 (± 0.19)	84.82 (± 0.18)
		crop	84.12 (± 0.11)	<u>84.37</u> (± 0.14)	83.64 (± 0.23)	80.00 (± 0.21)
translation		83.65 (± 0.16)	<u>84.31</u> (± 0.15)	84.10 (± 0.16)	83.47 (± 0.21)	
rotation		79.46 (± 0.26)	82.61 (± 0.21)	83.75 (± 0.15)	76.16 (± 0.27)	
mixup [3]		<u>85.30</u> (± 0.15)	84.71 (± 0.18)	84.37 (± 0.19)	84.55 (± 0.16)	
cutout [4]		84.70 (± 0.10)	84.52 (± 0.10)	83.98 (± 0.11)	84.77 (± 0.09)	
random erasing [5]		83.92 (± 0.19)	83.78 (± 0.18)	83.83 (± 0.17)	83.81 (± 0.20)	
RICAP [6]		83.47 (± 0.05)	<u>85.72</u> (± 0.19)	85.04 (± 0.17)	83.47 (± 0.17)	
flip and crop		84.12 (± 0.16)	<u>84.63</u> (± 0.11)	83.83 (± 0.12)	79.82 (± 0.21)	
translation and rotation		75.81 (± 0.12)	80.45 (± 0.19)	82.49 (± 0.19)	73.36 (± 0.23)	
SVHN (1 k) (NA: 78.76 (± 0.38))		flip	75.58 (± 0.41)	77.20 (± 0.33)	78.64 (± 0.23)	71.29 (± 0.50)
		crop	78.47 (± 0.36)	<u>80.13</u> (± 0.23)	79.53 (± 0.21)	73.90 (± 0.26)
	translation	78.01 (± 0.53)	80.09 (± 0.36)	<u>80.17</u> (± 0.35)	73.28 (± 0.53)	
	rotation	69.26 (± 0.48)	76.14 (± 0.25)	<u>78.93</u> (± 0.23)	51.47 (± 0.65)	
	mixup [3]	75.31 (± 0.10)	79.82 (± 0.18)	<u>79.89</u> (± 0.24)	75.87 (± 0.34)	
	cutout [4]	79.22 (± 0.30)	79.16 (± 0.33)	79.04 (± 0.29)	77.28 (± 0.34)	
	random erasing [5]	78.75 (± 0.39)	78.69 (± 0.39)	78.75 (± 0.37)	78.65 (± 0.37)	
	RICAP [6]	72.10 (± 0.32)	78.38 (± 0.44)	<u>80.39</u> (± 0.31)	72.32 (± 0.65)	
	flip and crop	73.67 (± 0.61)	78.10 (± 0.47)	79.12 (± 0.38)	66.18 (± 0.84)	
	translation and rotation	50.15 (± 0.47)	71.66 (± 0.72)	75.67 (± 0.54)	41.02 (± 1.05)	
	STL-10 (1 k) (NA: 49.58 (± 0.42))	flip	<u>54.71</u> (± 0.72)	54.24 (± 0.65)	52.38 (± 0.36)	53.43 (± 0.43)
		RICAP	59.62 (± 0.35)	<u>61.04</u> (± 0.38)	52.88 (± 0.48)	59.99 (± 0.38)
CIFAR-100 (10 k) (NA: 32.62 (± 0.27))	flip	38.19 (± 0.25)	<u>40.08</u> (± 0.21)	38.86 (± 0.11)	38.64 (± 0.29)	
	RICAP	50.03 (± 0.76)	<u>51.81</u> (± 0.39)	47.18 (± 0.36)	49.60 (± 0.46)	
Tiny-ImageNet (20 k) (NA: 25.10 (± 0.57))	flip	30.53 (± 0.20)	30.33 (± 0.21)	<u>31.58</u> (± 0.44)	30.60 (± 0.21)	
	RICAP	36.99 (± 0.54)	36.76 (± 0.55)	35.48 (± 0.44)	35.93 (± 0.62)	

Table 3

Best test accuracy (%) when training WideResNet28-10 using all samples from CIFAR-10, Fashion-MNIST, and SVHN datasets. We evaluated SPA using several λ values, CA, and NA. The results for NA have been added under each dataset in Dataset column. Values in parentheses after accuracies are the standard error. Numbers in bold indicate the highest accuracy, and underlined values indicate considerably greater values than those of CA and NA.

Dataset	Augmentation	SPA ($\lambda = 0.01$)	SPA ($\lambda = 0.1$)	SPA ($\lambda = 1$)	CA ($\lambda = 0$)
CIFAR-10 (NA: 92.43 (± 0.11))	mixup [3]	<u>92.66</u> (± 0.07)	92.48 (± 0.09)	92.65 (± 0.09)	92.24 (± 0.03)
	cutout [4]	94.01 (± 0.10)	93.05 (± 0.06)	92.62 (± 0.07)	93.99 (± 0.06)
	random erasing [5]	92.45 (± 0.08)	92.38 (± 0.06)	92.30 (± 0.06)	92.30 (± 0.03)
	RICAP [6]	95.94 (± 0.04)	86.47 (± 0.54)	68.19 (± 2.13)	<u>96.29</u> (± 0.07)
	flip and crop	<u>95.05</u> (± 0.05)	<u>95.05</u> (± 0.07)	93.03 (± 0.08)	90.91 (± 0.27)
Fashion-MNIST (NA: 93.85 (± 0.07))	translation and rotation	91.37 (± 0.09)	<u>93.79</u> (± 0.04)	92.04 (± 0.13)	23.55 (± 0.87)
	mixup [3]	94.39 (± 0.05)	94.39 (± 0.03)	94.44 (± 0.06)	94.45 (± 0.07)
	cutout [4]	94.49 (± 0.04)	94.19 (± 0.04)	94.23 (± 0.05)	<u>94.64</u> (± 0.04)
	random erasing [5]	93.90 (± 0.04)	93.89 (± 0.03)	93.87 (± 0.07)	93.88 (± 0.05)
	RICAP [6]	94.81 (± 0.08)	94.72 (± 0.03)	91.37 (± 0.37)	<u>95.06</u> (± 0.04)
	flip and crop	94.09 (± 0.03)	<u>94.62</u> (± 0.04)	93.93 (± 0.08)	87.76 (± 0.18)
	translation and rotation	93.57 (± 0.04)	93.02 (± 0.08)	93.88 (± 0.05)	93.88 (± 0.04)
SVHN (NA: 96.45 (± 0.02))	mixup [3]	96.50 (± 0.03)	96.57 (± 0.03)	<u>97.14</u> (± 0.04)	96.43 (± 0.03)
	cutout [4]	<u>97.05</u> (± 0.02)	96.63 (± 0.02)	96.86 (± 0.02)	96.24 (± 0.03)
	random erasing [5]	96.54 (± 0.04)	96.54 (± 0.04)	96.47 (± 0.04)	96.55 (± 0.03)
	RICAP [6]	97.33 (± 0.03)	93.87 (± 0.19)	70.49 (± 2.80)	<u>97.47</u> (± 0.04)
	flip and crop	96.68 (± 0.04)	97.05 (± 0.05)	<u>97.07</u> (± 0.06)	95.27 (± 0.06)
	translation and rotation	96.61 (± 0.04)	97.06 (± 0.04)	<u>97.27</u> (± 0.03)	88.50 (± 0.55)

performance than CA. In particular, as demonstrated in the results for the translation and rotation techniques, SPA still modified performance deterioration when unsuitable data augmentation techniques were applied.

In the experiments on the full CIFAR-10 dataset, the computation time when using SPA was 140.36 (± 0.11) minutes while that when using CA was 108.77 (± 0.18) minutes, and the ratio was 1.29 (± 0.00), which is not significant increase.

4.5. Comparison with RandAugment

We then compared the test accuracies of the proposed SPA to those of the RandAugment method [6]. With 1000 samples from the CIFAR-10 dataset, the test accuracies of RandAugment were always less than 40.00%, and the highest accuracy was 37.74% with an augmentation magnitude of 5. This accuracy is much less than that of the proposed SPA (and even CA or NA), as shown in Table 2. Thus, the proposed SPA method was more effective than RandAugment with a small number of training samples in our experimental setting.

In addition, we evaluated the test accuracies when training using all samples from several datasets. The test accuracies for RandAugment on the CIFAR-10, Fashion-MNIST, and SVHN datasets were 96.39%, 94.06%, and 96.34%, respectively. The best results obtained by the proposed SPA were 94.27%, 94.83%, and 97.34%, respectively (Table 3), and the RandAugment method obtained the best accuracy on the CIFAR-10 dataset; however, the proposed SPA obtain higher accuracy on the Fashion-MNIST and SVHN datasets. Thus, we consider that the proposed SPA is also effective on full datasets, even though it is the most effective with a small number of training samples. Against our expectations, the proposed SPA demonstrated high performance overall even though fixed data augmentation techniques were during training, whereas the RandAugment method randomly selected one of 14 techniques for iteration.

5. Discussion

Here, we discuss why the proposed SPA method performed well in consideration of two well-established training strategies, i.e., curriculum learning [7] and desirable changes to loss function instability [25]. First, we introduce each strategy and describe how SPA is based on these strategies. We then demonstrate the effect of these strategies on SPA through experiments.

5.1. Curriculum learning

Curriculum learning is a training strategy based on the difficulty of samples [7], where training proceeds from easy to difficult samples in a gradual manner to approximate human learning. Baby-step learning [26], which is an improved curriculum learning technique, trains a model using previously used samples, whereas curriculum learning replaces easy samples with difficult samples during training.

These methods require that samples be grouped manually into several sets according to their difficulty, which is achieved using prior knowledge that depends on the data. In contrast, self-paced learning (SPL) [27] automatically creates a boundary surface between difficult and easy samples. Here, samples that produce large losses are considered difficult samples (and vice versa). Thus, sample difficulty is flexible in SPL, although it is fixed during training in the original curriculum learning.

When $\mathbf{v}^{(i)}$ denotes N -dimensional values with 1 or 0 that represent whether N samples are used for training at the i th epoch during training with SPL, $\mathbf{v}^{(i+1)}$ can be expressed as follows:

$$\mathbf{v}^{(i+1)} = \min_{\mathbf{v}^{(i)}} \left(\sum_{n=0}^{N-1} v_n^{(i)} L(y_n, f(\mathbf{x}_n, \theta)) - \eta \sum_{n=0}^{N-1} v_n^{(i)} \right), \text{ s.t. } \mathbf{v}^{(i)} \in [0, 1]^N \quad (1)$$

where \mathbf{x}_n and y_n are the inputs and label of the n th sample, respectively, θ represents the model parameters at that time, $f(\mathbf{x}_n, \theta)$ is the predicted label, L is the loss function, N is the number of samples, $v_n^{(i)}$ is a value (1 or 0) that defines whether the n th sample was used for training at the i th epoch, and η is a parameter to control the training rate. If each $v_n^{(i)}$ becomes 1 to minimize the second term, $\sum_{n=0}^{N-1} v_n^{(i)} L(y_n, f(\mathbf{x}_n, \theta))$ in the first term increases. Thus, v_n for samples with small $L(y_n, f(\mathbf{x}_n, \theta))$, i.e., easier samples, tend to become 1 to minimize $E_A(\theta; \mathbf{v}; \eta)$. Parameter η is increased gradually during training to increase the number of samples used for training and to use more difficult samples.

Several studies have used or improved SPL in the context of neural networks [28,29]. A common strategy in the above methods related to curriculum learning is training on easier samples in the early training stage, which is important to realize improved generalization and faster convergence [30–32]. In addition, Cirik et al. [33] demonstrated experimentally that curriculum learning on long short-term memory networks [34] improves significantly when smaller models are used and helps more when the amount of training data is limited.

5.2. Curriculum learning in SPA

The proposed SPA method differs from the original curriculum learning because SPA uses difficult samples by applying data augmentation and easy samples in the early stage. In addition, SPA does not use difficult samples in the final stage. However, SPA applies data augmentation strategically based on a unique curriculum that partially uses the idea of original curriculum learning relative to the following two points; thus, curriculum learning is one factor that constructs SPA rather than a comparative method.

The first point is to use easier samples for training in the early stage than CA uses. As shown in Fig. 1, applying data augmentation makes samples more difficult; thus, SPA can use easy samples for training by applying data augmentation to only difficult samples. Thus, the samples used for training in the early stage in SPA are easier than those in CA, which only uses difficult samples in the stage by applying data augmentation to all samples.

The other point is that sample selection for data augmentation in SPA is achieved using a procedure that is similar to that used in SPL. In SPA, sample difficulty is assessed using the loss function value when each sample is input to the neural network, and this technique is also employed in SPL [27]. When $\mathbf{v}^{(i)}$ denotes N -dimensional values with 1 or 0 that represent whether data augmentation is applied to each of N samples at the i th epoch during training with SPA, $\mathbf{v}^{(i+1)}$ can be expressed based on Eq. (1) as follows:

$$\mathbf{v}^{(i+1)} = \min_{\mathbf{v}^{(i)}} \left(\sum_{n=0}^{N-1} (1 - v_n^{(i)}) L(y_n, f(\mathbf{x}_n, \theta)) + \eta \sum_{n=0}^{N-1} v_n^{(i)} \right), \text{ s.t. } \mathbf{v} \in [0, 1]^N \quad (2)$$

Here, $v_n^{(i)}$ tends to become 1 for a sample with large $L(y_n, f(\mathbf{x}_n, \theta))$. Data augmentation is applied to samples with $v_n^{(i)} = 1$, and the number of such samples can be reduced by increasing η during training.

Unlike Eq. (1), $(1 - v_n^{(i)})$ weighs the first term, and the sign of the second term is changed. The difference in the first term comes from the fact that “easy” samples are used for training in SPL, whereas “difficult” samples are used for augmentation in SPA.

The difference in the second term comes from the fact that the number of samples with $v_n^{(i)} = 1$ increases during training in SPL, whereas it decreases in SPA. While the proposed SPA method is considered to follow a well-established idea of SPL in terms of sample selection, the proposed method involves much more than simply using SPL, i.e., it is specifically designed for data augmentation.

In implementation, rather than using Eq. (2), we decided on $v^{(i+1)}$ by comparing $L(y_n, f(\mathbf{x}_n, \theta))$ to λ , which is a hyperparameter. As shown in Fig. 3, a sample producing training loss that is greater than threshold λ is considered a difficult sample, and data augmentation is applied to it (and vice versa). Here, appropriate λ depends on the datasets and models; thus, it should be investigated by attempting several λ values. As a whole, many difficult samples exist in the early stage; however, the number of difficult samples decreases as training proceeds because the training loss for each sample tends to decrease. Recall that the procedure of SPA in detailed in Algorithm 1.

We confirmed that the effect of curriculum learning appears in actual training. Here, we counted the numbers of samples according to training losses before and after applying data augmentation at epoch 20. The histograms of the numbers are shown in Fig. 4. When using SPA with $\lambda = 0.1$ (Fig. 4(a)), the number of samples with training loss less than 0.1 did not decrease by applying data augmentation because those samples were not augmented in SPA, whereas, with CA (Fig. 4(b)), the number of samples with training loss less than 0.1 decreased greatly. Thus, as expected, easy examples without data augmentation can be used in the early stage of SPA training.

In addition, we trained the small CNN with 1000 samples from the CIFAR-10 dataset by applying data augmentation to randomly selected samples to demonstrate that applying data augmentation to difficult samples contributes to improved test accuracy. Here, the number of samples to which data augmentation was applied during training was the same as the number of samples with $v_n = 1$ in SPA. The best test accuracy was obtained when data augmentation was applied randomly with $\lambda = 0.1$ was 49.12% (± 0.35), whereas that for SPA with the same λ value was 52.78% (± 0.38). This demonstrates that training in the early stage with easy samples without data augmentation is effective.

5.3. Loss function behavior

Although minimizing training loss is required to obtain high generalization performance, generally, the parameter with a small training loss does not always produce high generalization performance if it is in a poor minimum. Previous studies [35,36] have indicated that generalization performance is degraded when the parameter is trapped in a sharp minimum.

In stochastic gradient descent, a minibatch is used, and the loss function with parameter θ is expressed as $L = f(\mathbf{x}, \theta)$ using variable inputs \mathbf{x} . This can be considered as optimizing different loss functions in each time step; however, it can also be considered as optimizing a fixed average loss function with a parameter that moves stochastically [37]. Takase et al. [25] referred to this fluctuation in loss function from one step to another as loss function instability.

Intuitively, if we can control the instability and change the loss function from unstable to stable during training (Fig. 3), this will

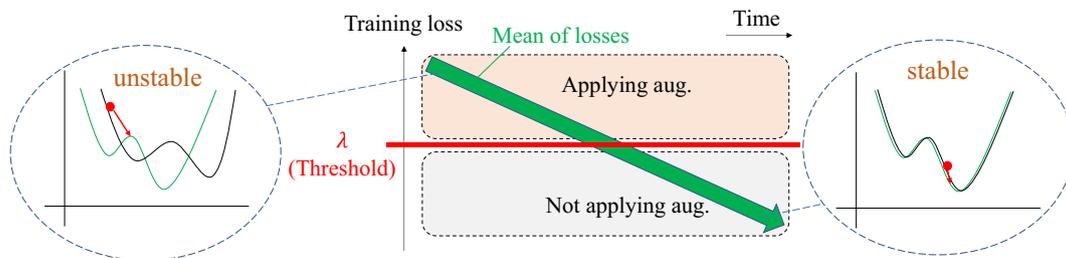


Fig. 3. Criterion by which data augmentation is applied in SPA. The concept of curriculum learning and desirable change of loss function instability is summarized. “Aug.” denotes data augmentation.

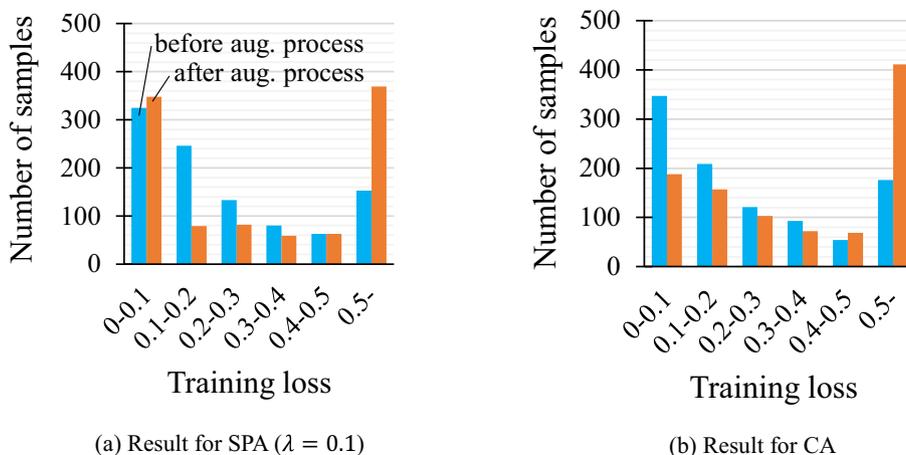


Fig. 4. Histograms of the number of training samples according to training losses. Numbers before and after the augmentation process at epoch 20 are shown for SPA and CA. Random erasing was used as the data augmentation technique.

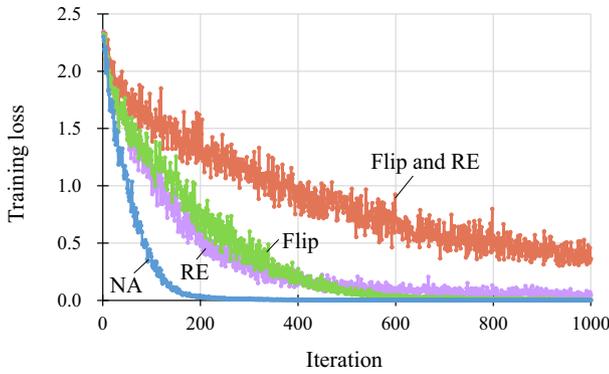


Fig. 5. Effects of data augmentation on training loss instability. RE denotes random erasing. In cases other than NA, data augmentation was applied to all training samples.

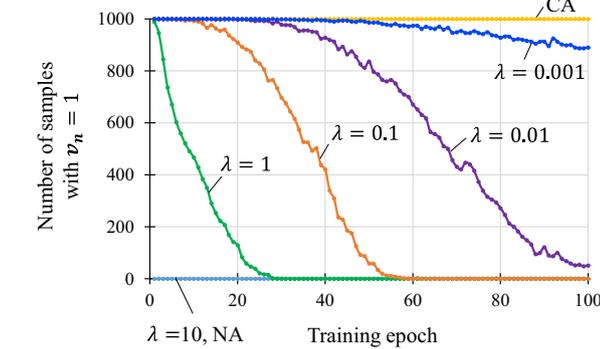
help to avoid becoming trapped in poor local minima, e.g., sharp minima, and facilitate stable convergence. In fact, Takase et al. [25] focused on the fact that a small minibatch size destabilizes the loss function, and they proposed a method that increases the batch size gradually during training. This facilitates searching a wide range of loss functions in the early stage, and then, the search range is reduced gradually, which results in improved generalization performance. This is similar to the idea of simulated annealing [38–40], although it controls the behavior of the parameter rather than the function.

5.4. Loss function behavior in SPA

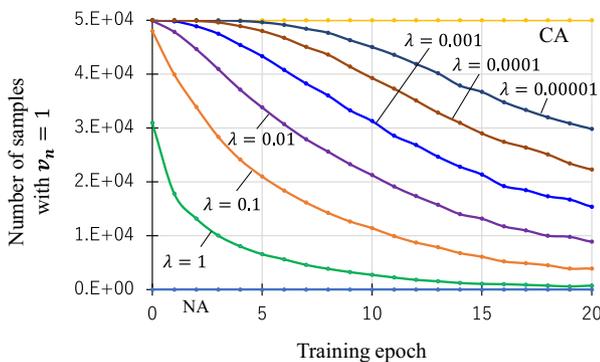
The proposed SPA minimizes the original loss function without data augmentation in the final training stage. However, it generalizes better than NA, which suggests that the transient dynamics of the proposed SPA in the early training stage will provide some advantage relative to finding better global minima.

With this expectation in mind, we confirm that small training loss and changing the loss function from instable to stable are observed in SPA-based training. First, we demonstrate that data augmentation increases training loss and destabilizes the loss function. Here, we trained the small CNN using several data augmentation techniques with the full CIFAR-10 dataset. The transition plotted in Fig. 5 shows that data augmentation increases training loss and destabilizes the loss function because the training loss increases eventually and changes more significantly for each minibatch compared to the no data augmentation case. Using both flip and random erasing provided the greater training loss and the most instable loss function.

This result implies that we can control the value and instability of the loss function by devising how to apply data augmentation. Here, a desirable strategy is to apply data augmentation to many samples in the early stage and then gradually reduce the number of samples for data augmentation. This can greatly reduce training

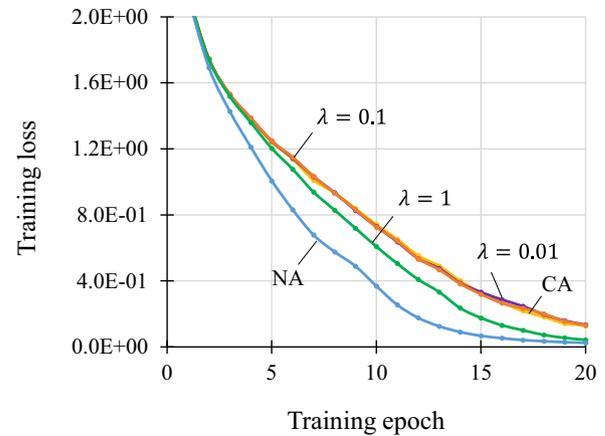


(a) Training for the small CNN on 1,000 samples

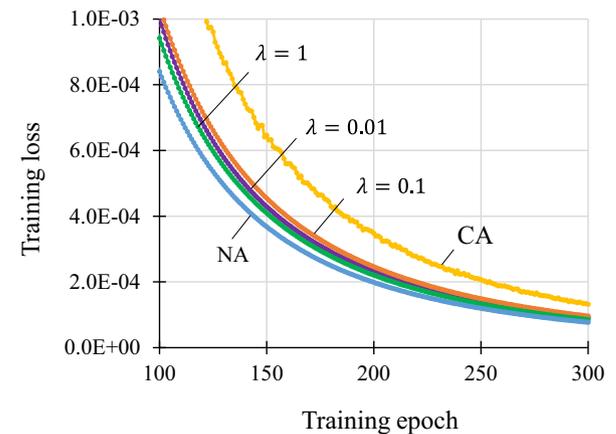


(b) Training for WideResNet28-10 on all samples

Fig. 6. Transitions of the number of samples with $v_n = 1$ during training. The flip augmentation technique was applied to the CIFAR-10 dataset.

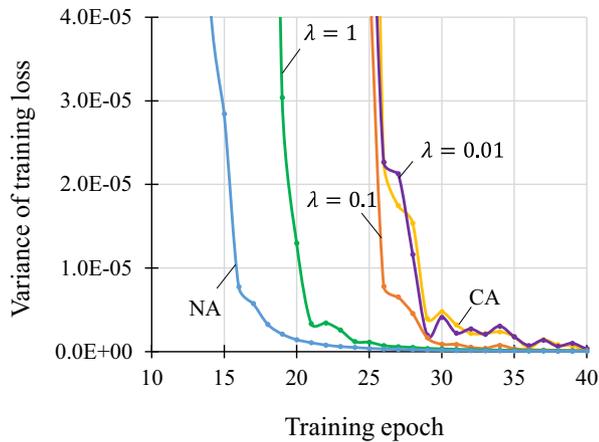


(a) The early stage

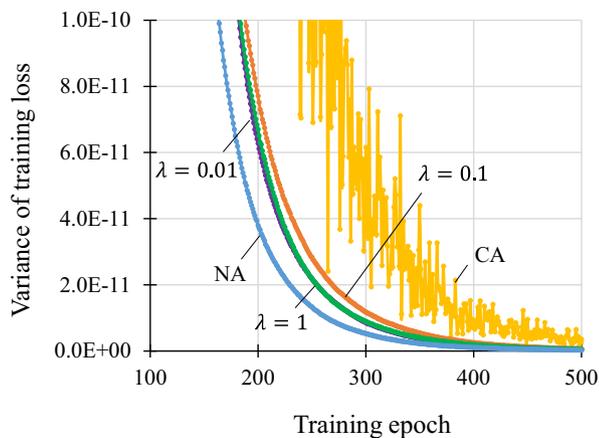


(b) The final stage

Fig. 7. Transitions of training loss for each minibatch. Effect of λ on training loss was investigated. The small CNN was trained with 1000 samples from the CIFAR-10 dataset.



(a) The early stage



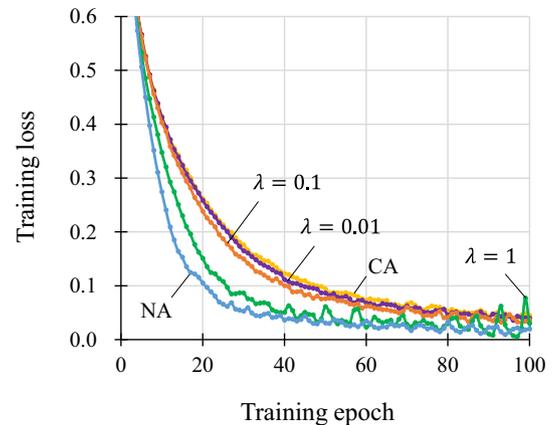
(b) The final stage

Fig. 8. Transitions of variance in training loss for each minibatch. The effect of λ on the loss function instability is investigated. The small CNN was trained with 1000 samples from the CIFAR-10 dataset.

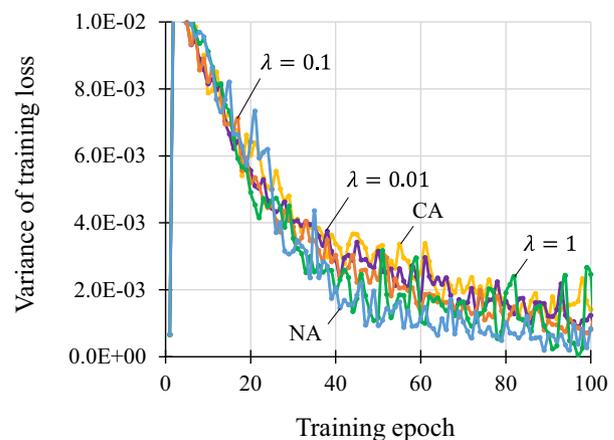
loss and make the loss function unstable in the early stage and stable in the final stage. Eq. (2), i.e., the equation for SPA, includes this desirable change in loss function instability and the effect of curriculum learning. As discussed in Section 5.2, the mean training losses tend to decrease during training because training is performed using gradient descent. Therefore, the number of samples for data augmentation tends to gradually decrease.

We confirmed this was achieved in actual training. Fig. 6 plots the transitions of the number of samples with $v_n = 1$, i.e., samples to which data augmentation was applied, during the first one tenth in the trainings with the flip augmentation technique. These results demonstrate that the number of samples with $v_n = 1$ decreased gradually with appropriate λ values, and smaller λ values decreased the rate of reduction, approaching the transition for CA. Although samples with $v_n = 1$ were very few (or none) depending on the λ value, this does not necessarily lead to reduced accuracy, as discussed in a previous study [41], which insisted that using less augmented data in the final training stage can refine the model and improve accuracy.

Fig. 7 shows the sum of training losses for the minibatches when training with a small number of samples. Although the transition of SPA resembles that of CA in the early stage (Fig. 7(a)), it resembles that of NA in the final stage (Fig. 7(b)), and it became more similar to that of NA as the λ values increased. In reference



(a) Training loss



(b) Variance of training loss

Fig. 9. Training loss and transitions of the variance in training loss for each minibatch. The small CNN was trained with all samples from the CIFAR-10 dataset.

to the literature [25], we also evaluated training loss instability according to the variance in training loss. Fig. 8 shows the transitions of the variance in training loss for each minibatch. Here, to facilitate reliable evaluation, training loss was measured at varying positions according to the neural network parameters rather than at a fixed position. Similar to the results shown in Figs. 6 and 7, SPA demonstrated transitions that are similar to those of CA and NA in the early and final stages, respectively. Thus, loss function instability changed from unstable to stable when training with SPA using a small number of training samples. The transient dynamics of SPA are expected to facilitate avoiding poor or sharp local minima (Section 5.3), although this may be just one of the factors of the proposed SPA method that enhances generalization performance.

We conducted the same experiments using all training samples. Fig. 9 shows the transitions of the training loss and variance in loss for each minibatch when training the small CNN on the CIFAR-10 dataset. Trends that differ from those for a small number of training samples can be observed in these results. In the transition of training loss shown in Fig. 9(a), the loss for SPA is nearly the same as that of CA in the final stage. In the transition of the variance of training loss shown in Fig. 9(b), SPA demonstrates a similar trend with CA and NA. In addition, all methods showed large fluctuation. Therefore, the effect of SPA was small; thus, significant improvements to test accuracy cannot always be expected, as shown in

Table 3. As SPA smoothly connects CA and NA, and exploits the advantages of both methods, it performs the most effectively in the conditions such as Figs. 7 and 8 that the trends in training loss of CA and NA are greatly different.

6. Conclusions

In this paper, we have proposed SPA, which dynamically selects samples for data augmentation. The proposed SPA method only applies data augmentation to samples that produce a training loss greater than fixed threshold λ , and this technique was designed based on curriculum learning and desirable changes in loss functions.

The results of extensive experiments on several datasets with different data augmentation techniques verify the effectiveness of the proposed SPA method, which outperformed both the CA and NA methods in most cases. In particular, SPA demonstrated better performance in two practical situations, i.e., when only a small amount of training data is obtained and when unsuitable data augmentation techniques are used. In addition, in such situations, the proposed SPA significantly outperformed the state-of-the-art RandAugment method. Our experimental results confirm that the effects of curriculum learning and desirable changes in loss functions are evident when training using the proposed SPA method.

Overall, we believe that we have addressed the problem of selecting suitable training samples for data augmentation, which enables more appropriate application of data augmentation techniques. However, several open issues remain relative to the heuristic settings of data augmentation. For example, the parameters of each data augmentation technique, e.g., mix rate in mixup, should be adjusted automatically and dynamically. In addition, although several λ values were used to evaluate the proposed SPA in our experiments, and $v_n = 0.1$ tended to yield high accuracy, incorporating λ optimization into a training system using SPA is expected to be more useful. Thus, in future, we plan to address these issues to improve the generalization performance of neural networks using data augmentation.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work was supported by a project of JSPS KAKENHI Grant Number JP20K19888.

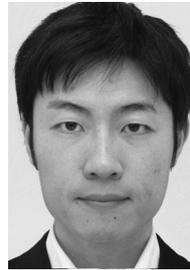
References

- [1] T. Ko, V. Peddinti, D. Povey, S. Khudanpur, Audio augmentation for speech recognition, in: *Proceedings of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2015, pp. 3586–3589.
- [2] L. Perez, J. Wang, The Effectiveness of Data Augmentation in Image Classification using Deep Learning, arXiv preprint arXiv: 1712.04621 (2017).
- [3] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (11) (1998) 2278–2324.
- [4] A. Krizhevsky, G. Hinton, Learning multiple layers of features from tiny images, Technical Report, University of Toronto (2009).
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, ImageNet: a large-scale hierarchical image database, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [6] E. D. Cubuk, B. Zoph, J. Shlens, Q. V. Le, Randaugment: Practical automated data augmentation with a reduced search space, In the *CVPR Workshop* (2020).
- [7] Y. Bengio, J. Louradour, R. Collobert, J. Weston, Curriculum learning, in: *Proceedings of the Annual International Conference on Machine Learning (ICML)*, 2009, pp. 41–48.
- [8] H. Zhang, M. Cisse, Y.N. Dauphin, D. Lopez-Paz, Mixup: beyond empirical risk minimization, *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [9] T. DeVries, G. W. Taylor, Improved regularization of convolutional neural networks with cutout, arXiv preprint arXiv: 1708.04552 (2017).
- [10] Z. Zhong, L. Zheng, G. Kang, S. Li, Y. Yang, Random erasing data augmentation, *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*, 2020.
- [11] R. Takahashi, T. Matsubara, K. Uehara, Data augmentation using random image cropping and patching for deep CNNs, *IEEE Trans. Circ. Syst. Video Technol.* (2020).
- [12] E.D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, Q.V. Le, AutoAugment: learning augmentation strategies from data, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 113–123.
- [13] S. Lim, I. Kim, T. Kim, C. Kim, S. Kim, Fast AutoAugment, In *Advances in Neural Information Processing Systems (NIPS)* (2019).
- [14] T. Takase, S. Oyama, M. Kurihara, Effective neural network training with adaptive learning rate based on training loss, *Neural Networks* 101 (2018) 68–78.
- [15] M. Andrychowicz, M. Denil, S. Gómez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, N. de Freitas, Learning to learn by gradient descent by gradient descent, In *Advances in Neural Information Processing Systems (NIPS)* (2016) 3981–3989.
- [16] S. Hochreiter, A.S. Younger, P.R. Conwell, Learning to learn using gradient descent, in: *Proceedings of the International Conference on Artificial Neural Networks (ICANN)*, 2001, pp. 87–94.
- [17] H. Xiao, K. Rasul, R. Vollgraf, Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms, arXiv preprint arXiv: 1708.07747 (2017).
- [18] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, A.Y. Ng, Reading digits in natural images with unsupervised feature learning, *The NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [19] A. Coates, A. Ng, H. Lee, An analysis of single layer networks in unsupervised feature learning, in: *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011, pp. 215–223.
- [20] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization, *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [21] S. Ioffe, C. Szegedy, Batch normalization: accelerating deep network training by reducing internal covariate shift, *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [22] X. Glorot, A. Bordes, Y. Bengio, Deep sparse rectifier neural networks, in: *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011, pp. 315–323.
- [23] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [24] S. Zagoruyko, N. Komodakis, Wide residual networks 87.1–87.12, *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.
- [25] T. Takase, S. Oyama, M. Kurihara, Why does large batch training result in poor generalization? A comprehensive explanation and a better strategy from the viewpoint of stochastic optimization, *Neural Comput.* 30 (7) (2018) 2005–2023.
- [26] V. I. Spitskovsky, H. Alshawi, D. Jurafsky, Baby Steps: How “Less is More” in Unsupervised Dependency Parsing, In the *NIPS Workshop on Grammar Induction, Representation of Language and Language Learning* (2009) 1–10.
- [27] M.P. Kumar, B. Packer, D. Koller, Self-paced learning for latent variable models, *Adv. Neural Inf. Processing Syst. (NIPS)* (2010) 1189–1197.
- [28] T. H. Kim, J. Choi, ScreenerNet: Learning Self-Paced Curriculum for Deep Neural Networks, arXiv preprint arXiv: 1801.00904 (2018).
- [29] H. Li, M. Gong, Self-paced Convolutional Neural Networks, In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (2017) 2110–2116.
- [30] J.L. Elman, Learning and development in neural networks: the importance of starting small, *Cognition* 48 (1993) 71–99.
- [31] K.A. Krueger, P. Dayan, Flexible shaping: how learning in small steps helps, *Cognition* 110 (2009) 380–394.
- [32] D. Rohde, D. Plaut, Language acquisition in the absence of explicit negative evidence: how important is starting small?, *Cognition* 72 (1999) 67–109.
- [33] V. Cirik, E. Hovy, L-P. Morency, Visualizing and Understanding Curriculum Learning for Long Short-Term Memory Networks, arXiv preprint arXiv: 1611.06204 (2016).
- [34] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780.
- [35] E. Hoffer, I. Hubara, D. Soudry, Train longer, generalize better: closing the generalization gap in large batch training of neural networks, in: *Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 1731–1741.
- [36] Z. Yao, A. Gholami, Q. Lei, K. Keutzer, M.W. Mahoney, Hessian-based Analysis of Large Batch Training and Robustness to Adversaries, in: *Advances in Neural Information Processing Systems (NIPS)*, 2018, pp. 4949–4959.
- [37] L. Bottou, Large-scale machine learning with stochastic gradient descent, in: *Proceedings of the International Conference on Computational Statistics (COMPSTAT)*, 2010, pp. 177–186.
- [38] V. Černý, Thermodynamical approach to the traveling salesman problem: an efficient simulation algorithm, *J. Optim. Theory Appl.* 45 (1985) 41–51.
- [39] S. Kirkpatrick, C.D. Gelatt, M.P. Vecchi, Optimization by simulated annealing, *Science* 220 (1983) 671–680.

- [40] L.R. Rere, M.I. Fanany, A.M. Arymurthy, Simulated annealing algorithm for deep learning, *Procedia Comput. Sci.* 72 (2015) 137–144.
- [41] Z. He, L. Xie, X. Chen, Y. Zhang, Y. Wang, Q. Tian, Data Augmentation Revisited: Rethinking the Distribution Gap between Clean and Augmented Data, arXiv preprint arXiv: 1909.09148 (2019).
- [42] Q. Zheng, M. Yang, X. Tian, N. Jiang, and D., Wang, A Full Stage Data Augmentation Method in Deep Convolutional Neural Network for Natural Image Classification, *Discrete Dynamics in Nature and Society* Volume 2020 (2020).
- [43] L. Nanni, S. Brahmam, S. Ghidoni, and G. Mahuolo, General Purpose (GenP) Bioimage Deep Ensemble Combining New Data Augmentation Techniques and Handcrafted Features, arXiv preprint arXiv: 1904.08084 (2019).
- [44] C. Zhang, S. Bengio, M. Hardt, B. Recht, O. Vinyals, Understanding deep learning requires rethinking generalization, *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- [45] A. Hernández-García and P. König, Data augmentation instead of explicit regularization, arXiv preprint arXiv: 1806.03852 (2018).



Tomoumi Takase received his B.Eng. and M.Eng. degrees from the University of Tokyo in 2012 and 2014, and Ph.D. degree from Hokkaido University in 2018. He is currently a postdoctoral fellow at National Institute of Advanced Industrial Science and Technology (AIST). His research interests include deep learning and data augmentation.



Ryo Karakida received the B.S., M.S., and D.S. degrees from The University of Tokyo, in 2012, 2014, and 2017, respectively. In 2017, he joined the Artificial Intelligence Research Center of the National Institute of Advanced Industrial Science and Technology, Tokyo, Japan. His current research interests include applied mathematics, statistical mechanics, and information geometry for artificial and biological neural networks. He received the Young Researcher Award from IEEE CISJ in 2015 and the Best Paper Award at the 24th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN) in 2016.



Hideki Asoh received his B.Eng. in mathematical engineering and M.Eng. in information engineering from the University of Tokyo, in 1981 and 1983 respectively. In April 1983, he joined in the Electrotechnical Laboratory. From 1993 to 1994 he stayed at German National Research Center for Information Technology as a visiting research scientist. From 2015 to 2019 he was a deputy director of Artificial Intelligence Research Center at National Institute of Advanced Industrial Science and Technology (AIST). His research interests are in constructing intelligent systems which can learn through interactions with the real-world.