

# Between-class Learning for Image Classification

Yuji Tokozume<sup>1</sup> Yoshitaka Ushiku<sup>1</sup> Tatsuya Harada<sup>1,2</sup>  
<sup>1</sup>The University of Tokyo <sup>2</sup>RIKEN  
 {tokozume,ushiku,harada}@mi.t.u-tokyo.ac.jp

## Abstract

In this paper, we propose a novel learning method for image classification called *Between-Class learning (BC learning)*<sup>1</sup>. We generate between-class images by mixing two images belonging to different classes with a random ratio. We then input the mixed image to the model and train the model to output the mixing ratio. BC learning has the ability to impose constraints on the shape of the feature distributions, and thus the generalization ability is improved. BC learning is originally a method developed for sounds, which can be digitally mixed. Mixing two image data does not appear to make sense; however, we argue that because convolutional neural networks have an aspect of treating input data as waveforms, what works on sounds must also work on images. First, we propose a simple mixing method using internal divisions, which surprisingly proves to significantly improve performance. Second, we propose a mixing method that treats the images as waveforms, which leads to a further improvement in performance. As a result, we achieved 19.4% and 2.26% top-1 errors on ImageNet-1K and CIFAR-10, respectively.<sup>2</sup>

## 1. Introduction

Deep convolutional neural networks (CNNs) [18] have achieved high performance in various tasks, such as image recognition [17, 11], speech recognition [1, 21], and sound recognition [19, 26]. One of the biggest themes of research on image recognition has been network engineering. Many types of image recognition networks have been proposed mainly in ILSVRC competition [17, 23, 25, 11, 28, 29, 12]. Furthermore, training deep neural networks is difficult, and many techniques have been proposed to achieve a high performance: data augmentation techniques [17], new network layers such as dropout [24] and batch normalization [14], optimizers such as Adam [15], and so on. Thanks to these

<sup>1</sup>Our preliminary experimental results on CIFAR-10 and ImageNet-1K were already presented in ILSVRC2017 on July 26, 2017.

<sup>2</sup>The code is publicly available at [https://github.com/mil-tokyo/bc\\_learning\\_image/](https://github.com/mil-tokyo/bc_learning_image/).

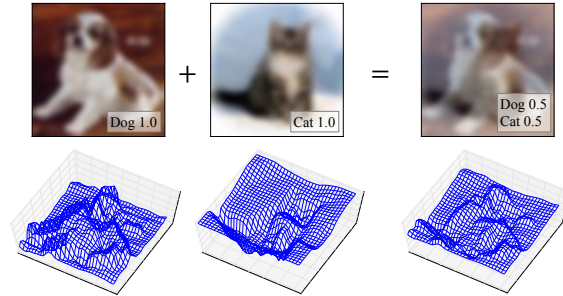


Figure 1. We argue that CNNs have an aspect of treating the input data as waveforms. In that case, a mixture of two images is a mixture of two waveforms. This makes sense for machines, although it does not visually make sense for humans.

research studies, training deep neural networks has become relatively easy, with a stable performance, at least for image classification. At present, a novel approach is needed for further improvement.

In [3], a simple and powerful learning method named *Between-Class learning (BC learning)* was proposed for deep sound recognition. BC learning aims to learn a classification problem by solving the problem of predicting the mixing ratio between two different classes. They generated between-class examples by mixing two sounds belonging to different classes with a random ratio. They then input the mixed sound to the model and trained the model to output the mixing ratio of each class. The advantages of BC learning are not limited only to the increase in variation of the training data. They argued that BC learning has the ability to impose constraints on the feature distributions, which cannot be achieved with standard learning, and thus the generalization ability is improved. They carefully designed the method of mixing two sounds, considering the difference in the sound energies, to achieve a satisfactory performance. As a result, BC learning improved the performance on various sound recognition networks, datasets, and data augmentation schemes, and they achieved a performance surpasses the human level in sound classification tasks.

The question here is whether BC learning also performs well on images. The core idea of BC learning itself, i.e.,

mixing two examples and training the model to output the mixing ratio, can be used irrespective of the modality of input data. BC learning is applicable to sounds, because sound is a kind of wave motion and a mixture of multiple sound data still counts as a sound. However, an image is not a kind of wave motion, and mixing multiple image data does not visually make sense. We show an example of a mixed image in Fig. 1(top). A mixed image loses its objectness and does not count as an image. Therefore, it appears inappropriate to apply BC learning to images.

However, the important thing is not how *humans* perceive the mixed data, but how *machines* perceive them. We argue that CNNs have an aspect of treating input data as waveforms, considering the recent studies on speech and sound recognition and the characteristics of image data as pixel values. We assume that CNNs recognize images by treating them as waveforms in quite a similar manner to how they recognize sounds. Thus, a mixture of two images is a mixture of two waveforms as shown in Fig. 1(bottom), and it would make sense *for machines*. Therefore, what is effective for sounds would also be effective for images.

We thus propose BC learning for images in this paper. First, we propose the simplest mixing method using internal divisions. Surprisingly, this mixing method proves to perform well. Second, we also propose an improved mixing method that treats images as waveform data (BC+). In this method, we subtract the per-image mean value from each image. By doing this, we can treat each image as a zero-mean waveform similar to a sound. We then define the image energy as the standard deviation per image, and mix two images considering the image energy in quite a similar manner to what [3] did for sounds. This mixing method is also simple and easy to implement, and leads to a further improvement in performance.

Experimental results show that BC learning also improves the performance of various types of image recognition networks from a simple network to the state-of-the-art networks. The top-1 error of ResNeXt-101 ( $64 \times 4d$ ) [28] on ImageNet-1K is improved from 20.4% to 19.4% by using the simplest BC learning. Moreover, the error rate of the state-of-the-art Shake-Shake Regularization [9] on CIFAR-10 dataset is improved from 2.86% to 2.26% by using the improved BC learning (BC+). Finally, we visualize the learned features and show that BC learning indeed imposes a constraint on the feature distribution. The contributions of this paper are as follows:

- We applied BC learning [3] to images by mixing two images belonging to different classes and training the model to output the mixing ratio.
- We argued that CNNs have an aspect of treating input data as waveforms, and proposed a mixing method that treats the images as waveforms.

- We conducted experiments extensively and demonstrated the effectiveness of BC learning for images.

The remainder of this paper is organized as follows. In Section 2, we provide a summary of BC learning for sound recognition [3] as a related work. We then propose BC learning for image recognition in Section 3, explaining the relationship with BC learning for sounds. In Section 4, we compare the performance of standard learning and BC learning, and demonstrate the effectiveness of BC learning. Finally, we conclude this paper in Section 5.

## 2. BC learning for sounds

In this section, we describe BC learning for sound recognition [3] as a related work. The contents in this section is a summarization of Section 3 of [3]. Please see [3] for more detailed information.

### 2.1. Overview

In standard learning for classification problems, a single training example is selected from the dataset and input to the model. Then, the the model is trained to output a one-hot label. By contrast, in BC learning, two training examples belonging to different classes are selected from the dataset and mixed with a random ratio. Then, the mixed data is input to the model, and the model is trained to output the mixing ratio of each class. KL-divergence between the outputs of the model and the ratio labels is used as the loss function, instead of the usual cross-entropy loss. Note that mixing is not performed in testing phase.

### 2.2. Mixing method

Let  $\{\mathbf{x}_1, \mathbf{t}_1\}$  and  $\{\mathbf{x}_2, \mathbf{t}_2\}$  be two sets of sounds and one-hot labels belonging to different classes randomly selected from the training dataset. A random ratio  $r$  is generated from the uniform distribution  $U(0, 1)$ , and two sounds and labels are mixed with this ratio. Two labels  $\mathbf{t}_1$  and  $\mathbf{t}_2$  are mixed simply by  $r \mathbf{t}_1 + (1 - r) \mathbf{t}_2$  because BC learning aims to train the model to output the mixing ratio. We then explain how to mix two sounds  $\mathbf{x}_1$  and  $\mathbf{x}_2$ , which should be carefully designed to achieve a satisfactory performance. The simplest method is  $r \mathbf{x}_1 + (1 - r) \mathbf{x}_2$ . Here,  $\frac{r \mathbf{x}_1 + (1 - r) \mathbf{x}_2}{\sqrt{r^2 + (1 - r)^2}}$  is better because sound energy is proportional to the square of the amplitude. However, when the difference in sound pressure level between  $\mathbf{x}_1$  and  $\mathbf{x}_2$  is large, the perception of the sound mixed by this method does not become  $\mathbf{x}_1 : \mathbf{x}_2 = r : (1 - r)$ . In this case, training the model with a label of  $r \mathbf{t}_1 + (1 - r) \mathbf{t}_2$  is inappropriate. To address this problem, they proposed to use a mixing method that considers the sound pressure level of two sounds  $G_1$  and  $G_2$  [dB] so that the auditory perception of the mixed sound becomes  $\mathbf{x}_1 : \mathbf{x}_2 = r : (1 - r)$ :

$$\frac{p \mathbf{x}_1 + (1-p) \mathbf{x}_2}{\sqrt{p^2 + (1-p)^2}}, \quad (1)$$

where  $p = \frac{1}{1 + 10^{\frac{G_1 - G_2}{20}} \cdot \frac{1-r}{r}}.$

### 2.3. How BC learning works

We explain how BC learning improves the classification performance. They argued that BC learning has the ability to impose the following two constraints on the feature distributions learned by the model by training the model to output the mixing ratio between two classes:

- Enlargement of Fisher’s criterion [8] (*i.e.*, the ratio of the between-class distance to the within-class variance).
- Regularization of positional relationship among feature distributions.

They hypothesized that when a mixed sound is input to the model, the feature of the mixed sound is located approximately in near the internally dividing point of the features of original two sounds. This hypothesis came from the fact that linearly-separable features are learned in hidden layers close to the output layer [2] and that humans can perceive which of the two sounds is louder from a digitally mixed sound. They showed that this hypothesis was indeed correct by visualizing the feature distributions of the standard-learned model using PCA. Under this hypothesis, we explain the two constraints that BC learning can impose on the feature distributions.

**Enlargement of Fisher’s criterion.** They argued that BC learning enlarges Fisher’s criterion [8] between any two classes in the feature space. We explain the reason in Fig. 2(top). If Fisher’s criterion between the feature distributions of class A (red) and class B (blue) is small as shown in Fig. 2(upper left), the feature distribution of the sounds obtained by mixing class A and B at a certain ratio (magenta) becomes large, and would have a large overlap with one or both feature distributions of class A and class B. In this case, the model cannot output the mixing ratio for some mixed examples projected onto the overlapping area, and BC learning gives a large loss. To let the model output the mixing ratio and make the loss of BC learning small, Fisher’s criterion should be large as shown in Fig. 2(upper right). In this case, the overlap becomes small, and BC learning gives a small loss. Therefore, BC learning enlarges Fisher’s criterion in the feature space.

**Regularization of positional relationship among feature distributions.** They also argued that BC learning

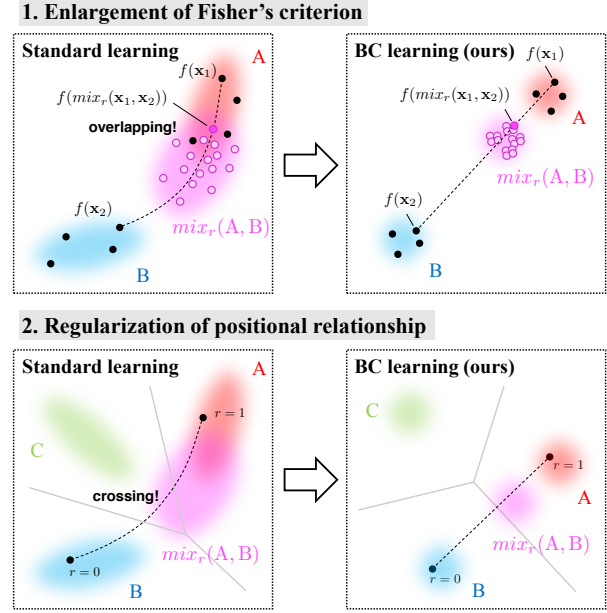


Figure 2. BC learning has the ability to impose constraints on the feature distribution, which cannot be achieved with the standard learning [3]. This figure represents the class distribution in the feature space. The black dashed line represents the trajectory of the feature when we input a mixture of two particular sounds to the model changing the mixing ratio from 0 to 1.

has the effect of regularizing the positional relationship among class feature distributions. We explain the reason in Fig. 2(bottom). If the features of each class are not regularly distributed as shown in Fig. 2(lower left), the decision boundary of class C other than both A and B would appear between class A and class B, and some of the mixed sounds of class A and class B would be misclassified into class C. This is an undesirable situation because there is little possibility that a mixed sound of two classes becomes a sound of other classes. BC learning gives a large loss to this situation because BC learning trains the model to output the mixing ratio between class A and class B. If the features of each class are regularly distributed as shown in Fig. 2(lower right), on the other hand, the decision boundary of class C does not appear between class A and class B, and the model can output the mixing ratio instead of misclassifying the mixed sound as class C. As a result, the loss of BC learning becomes small. Therefore, BC learning has the effect of regularizing the positional relationship of the feature distributions. In this way, they argued that BC learning has the ability to impose constraints on the feature distribution, and thus BC learning improves the generalization ability.

### 3. From sounds to images

In this section, we consider applying BC learning to images. Following BC learning for sounds [3], we select two

training examples from different classes and mix these two examples using a random ratio. We then input the mixed data to the model and train the model to output the mixing ratio. BC learning uses only mixed data and labels and, thus, never uses pure data and labels. How to mix two examples is also important for images. First, we propose the simplest mixing method in Section 3.1. Second, we discuss why BC learning can also be applied to images in Section 3.2. Finally, in Section 3.3, we propose a better version of BC learning, considering the discussion in 3.2.

### 3.1. Simple mixing

Let  $\mathbf{x}_1$  and  $\mathbf{x}_2$  be two images belonging to different classes randomly selected from the training dataset, and  $\mathbf{t}_1$  and  $\mathbf{t}_2$  be their one-hot labels. Note that  $\mathbf{x}_1$  and  $\mathbf{x}_2$  may have already been preprocessed, and they have the same size as the input size of the network. We generate a random ratio  $r$  from  $U(0, 1)$ , and mix two sets of data and labels with this ratio. Because we aim to train the model to output the mixing ratio, we mix two labels simply by:

$$r \mathbf{t}_1 + (1 - r) \mathbf{t}_2. \quad (2)$$

We now explain how to mix  $\mathbf{x}_1$  and  $\mathbf{x}_2$ . In [3], a carefully designed mixing method of two sounds was proposed considering the difference in the sound energies, as we mentioned in Section 2.2. In a sound data, 0 is the absolute center, and the distance from 0 represents the sound energy. However, the pixel values of an image data do not have an absolute center, and there appears to be no concept of energy. We thus first propose the following mixing method as the simplest method using internal divisions:

$$r \mathbf{x}_1 + (1 - r) \mathbf{x}_2. \quad (3)$$

### 3.2. Why BC learning works on images

BC learning is applicable to sounds because a mixed sound still counts as a sound. Sound is a kind of wave motion, and mixing two sounds physically makes sense. Humans can recognize two sounds and perceive which of the two sounds is louder or softer from the digitally mixed sound. However, image data, as pixel values, is not a kind of wave motion *for humans*. Therefore, mixing multiple images does not visually make sense.

However, the important thing is not whether mixing two data physically makes sense, or whether humans can perceive a mixed data, but how a machine perceives a mixed data. We argue that CNNs have an aspect of treating input data as waveforms. In fact, recent studies have demonstrated that CNNs can learn speech and sounds directly from raw waveforms, and each filter learns to respond to a particular frequency area [21, 26, 7]. It is also known that images, as pixel values, can be transformed into components of various frequency areas by using 2-D Fourier transform

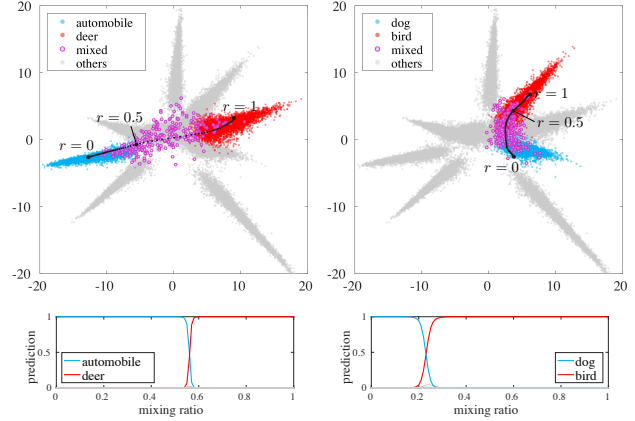


Figure 3. Visualization of the feature space of a standard-learned model using PCA. The features of the mixed images are distributed between two classes.

[22], and some convolutional filters can act as frequency filters (e.g., a Laplacian filter acts as a high-pass filter [5]). Therefore, it is expected that each convolutional filter of a CNN learns to extract the frequency features. In this way, we assume that CNNs recognize images by treating them as waveforms in quite a similar manner to how they recognize sounds. Thus, because a mixture of two images is a mixture of two waveforms *for machines*, what is effective for sounds would also be effective for images.

We visualized the feature distribution of the standard-learned model against mixed data using PCA. We used the output of the 10-th layer of an 11-layer CNN trained on CIFAR-10. We mixed two images using Eqn. 3 and the results are shown in Fig. 3. The magenta dots represent the feature distribution of the mixed images of *automobile* and *deer* (left), *dog* and *bird* (right) with a ratio of 0.5 : 0.5, and the black dotted line represents the trajectory of the feature when we input a mixture of two particular images to the model, changing the mixing ratio from 0 to 1. This figure shows that the mixture of two images is projected onto the point near the internally dividing point of two features, and the features of the mixed images are distributed between two classes, which is the same tendency observed for sounds [3]. Therefore, the same effect as BC learning for sounds, i.e., an enlargement of Fisher’s criterion in the feature space and a regularization of the positional relationship among the feature distributions of the classes, is expected. We compare the feature distributions learned with standard and BC learning and demonstrate that a different shape of feature distribution is learned with BC learning in the visualization experiment.

### 3.3. BC+: Images as waveform data

Here, we consider a new mixing method, which treats images as waveform data. We regard image data as a 2-D

waveform consisting of (R, G, B) vectors. In recent state-of-the-art methods, the input data is normalized for each channel using the mean and standard deviation calculated from the whole training data [28, 13, 9]. In this case, the mean of each image is not equal to 0, and each image data  $\mathbf{x}_i$  is represented as  $\mathbf{x}_i = \mu_i + \mathbf{d}_i$ , where  $\mu_i$  and  $\mathbf{d}_i$  are the static component and wave component, respectively. Here, the simplest mixing method of Eqn. 3 can be rewritten as  $\{r \mu_1 + (1-r) \mu_2\} + \{r \mathbf{d}_1 + (1-r) \mathbf{d}_2\}$ . We assume that the performance improvement with Eqn. 3 is mainly owing to the wave component  $r \mathbf{d}_1 + (1-r) \mathbf{d}_2$  if CNNs treat the input data as waveforms. Moreover, the static component  $r \mu_1 + (1-r) \mu_2$  can have a bad effect because mixing two waveforms generally hypothesizes that the static components of two waveforms are same.

Therefore, we remove the static component by subtracting the per-image mean value (not channel-wise mean). We consider the following mixing method instead of Eqn. 3. By doing this, we can treat each image as a zero-mean waveform similar to a sound.

$$r (\mathbf{x}_1 - \mu_1) + (1-r) (\mathbf{x}_2 - \mu_2) \quad (4)$$

We then apply a scheme similar to that applied to sounds as described in Section 2.2. First, we consider mixing two images with:

$$\frac{r (\mathbf{x}_1 - \mu_1) + (1-r) (\mathbf{x}_2 - \mu_2)}{\sqrt{r^2 + (1-r)^2}}, \quad (5)$$

instead of Eqn. 4, considering that waveform energy is proportional to the square of the amplitude. This process prevents the input variance from decreasing.

Second, we take the difference of image energies into consideration, in order to make the perception of the mixed image  $\mathbf{x}_1 : \mathbf{x}_2 = r : (1-r)$ . We use a new coefficient  $p$  instead of  $r$  and mix two images by  $\frac{p (\mathbf{x}_1 - \mu_1) + (1-p) (\mathbf{x}_2 - \mu_2)}{\sqrt{p^2 + (1-p)^2}}$ . We define  $p$  using the standard deviation per image ( $\sigma_1$  and  $\sigma_2$ ) so that the ratio of amplitude becomes  $\mathbf{x}_1 : \mathbf{x}_2 = r : (1-r)$ . We solve  $p \sigma_1 : (1-p) \sigma_2 = r : (1-r)$  and obtain the proposed mixing method:

$$\frac{p (\mathbf{x}_1 - \mu_1) + (1-p) (\mathbf{x}_2 - \mu_2)}{\sqrt{p^2 + (1-p)^2}}, \quad (6)$$

$$\text{where } p = \frac{1}{1 + \frac{\sigma_1}{\sigma_2} \cdot \frac{1-r}{r}}.$$

The main differences from the mixing method for sounds (Eqn. 1) are subtracting per-image mean values and using standard deviations instead of sound pressure levels. This mixing method is also easy to implement, and experimentally proves to lead to a further improvement in performance, compared to the simplest mixing method of Eqn. 3.

Table 1. Results of ResNeXt-101 ( $64 \times 4d$ ) [28] on ImageNet-1K dataset. BC learning improves the performance when using the default learning schedule. Furthermore, the performance is further improved when using a longer learning schedule, and the single-crop top-1 error is improved by around 1% compared to the default performance reported in [28].

# epochs	Learning	Top-1/top-5 val. error (%)	
		Single-crop	10-crop
100	Standard	20.4 / 5.3 [28]	18.90 / 4.61
	BC (ours)	19.92 / 4.91	18.66 / 4.26
150	Standard	20.44 / 5.25	18.98 / 4.43
	BC (ours)	<b>19.43 / 4.80</b>	<b>18.22 / 4.13</b>




Figure 4. Training curves of ResNeXt-101 ( $64 \times 4d$ ) on ImageNet-1K dataset. The dashed lines represent the training curves when using the default learning schedule, and the solid lines represent the training curves when using a longer learning schedule.

## 4. Experiments

### 4.1. Experiments on ImageNet-1K

First, we compare the performance of standard and BC learning on the 1,000-class ImageNet classification task [20]. In this experiment, we used the simple BC learning. We selected ResNeXt-101 ( $64 \times 4d$ ) [28] as the model for training because it has a state-of-the-art level performance and, moreover, the official Torch [6] training codes are available<sup>3</sup>. To validate the comparison, we incorporated BC learning into these official codes. When using BC learning, we selected two training images and applied the default data augmentation scheme as in described [28] to each image, and obtained two  $224 \times 224$  images. We then mixed these two images with a random ratio selected from  $U(0, 1)$ . In addition to the default learning schedule (# of epochs = 100), we also tried a longer learning schedule (# of epochs = 150). In 100-epochs training, we started training with a learning rate of 0.1 and then divided the learning rate by 10 at the epoch in  $\{30, 60, 90\}$ , as in [28]. In 150-

<sup>3</sup><https://github.com/facebookresearch/ResNeXt>



Table 2. Results on CIFAR-10 and CIFAR-100 datasets. We show the average and the standard error of 5 or 10 trials. BC learning improves the performance of various settings. Note that <sup>†</sup> is trained with a different learning setting from the default.

Model	Learning	Error rate (%) on	
		CIFAR-10	CIFAR-100
11-layer CNN	Standard	6.07 $\pm$ 0.04	26.68 $\pm$ 0.09
	BC (ours)	5.40 $\pm$ 0.07	24.28 $\pm$ 0.11
	BC+ (ours)	<b>5.22 <math>\pm</math> 0.04</b>	<b>23.68 <math>\pm</math> 0.10</b>
ResNet-29 <sup>†</sup> [28]	Standard	4.24 $\pm$ 0.06 / 4.39 [28]	20.18 $\pm$ 0.07
	BC (ours)	3.75 $\pm$ 0.04	19.56 $\pm$ 0.10
	BC+ (ours)	<b>3.55 <math>\pm</math> 0.03</b>	<b>19.41 <math>\pm</math> 0.07</b>
ResNeXt-29 (16 $\times$ 64d) <sup>†</sup> [28]	Standard	3.54 $\pm$ 0.04 / 3.58 [28]	<b>16.99 <math>\pm</math> 0.06</b> / 17.31 [28]
	BC (ours)	<b>2.79 <math>\pm</math> 0.06</b>	18.21 $\pm$ 0.12
	BC+ (ours)	2.81 $\pm$ 0.06	17.93 $\pm$ 0.09
DenseNet-BC ( $k = 40$ ) <sup>†</sup> [13]	Standard	3.61 $\pm$ 0.10 / 3.46 [13]	17.28 $\pm$ 0.12 / 17.18 [13]
	BC (ours)	2.68 $\pm$ 0.03	16.36 $\pm$ 0.10
	BC+ (ours)	<b>2.57 <math>\pm</math> 0.06</b>	<b>16.23 <math>\pm</math> 0.07</b>
Shake-Shake Regularization [9]	Standard	2.86 [9]	<b>15.85 [9]</b>
	BC (ours)	2.38 $\pm$ 0.04	15.90 $\pm$ 0.06
	BC+ (ours)	<b>2.26 <math>\pm</math> 0.01</b>	16.00 $\pm$ 0.10

epochs training, we also started training with a learning rate of 0.1 and then divided the learning rate by 10 at the epoch in {50, 90, 120, 140}. We reported classification errors on the validation set using both single-crop testing [28] and 10-crop testing [17].

The results are shown in Table 1. We also show the training curves in Fig. 4. The performance of BC learning with the default 100-epochs training was significantly improved from that of standard learning. Moreover, the performance of BC learning was further improved with 150-epochs training, while that of standard learning was not improved, and we achieved 19.43%/4.80% single-crop top-1/top-5 validation errors and 18.22%/4.13% 10-crop validation top1/top5 errors. The single-crop top-1 error was improved by around 1% compared to the default performance reported in [28].

**Discussion.** Learning between-class examples among 1,000 classes is difficult, and it tends to require a large number of training epochs. As shown in Fig. 4, the performance on the first 100 epochs of the 150-epochs training of BC learning is worse than the performance of standard learning. Therefore, the learning schedule should be carefully designed. Furthermore, we assume that the usage of curriculum learning [4] would be helpful to speed up the training; namely, at the early stage, we generate a mixing ratio close to 0 or 1 and input relatively pure examples, and we gradually change the distribution of  $r$  to flat.

## 4.2. Experiments on CIFAR

Now, we compare the performance of standard and BC learning on CIFAR-10 and CIFAR-100 datasets [16].

We trained the standard 11-layer CNN, ResNet-29 [28], ResNeXt-29 (16  $\times$  64d) [28], DenseNet (BC,  $k = 40$ ) [13], and Shake-Shake Regularization (S-S-I and S-E-I for CIFAR-10 and CIFAR-100, respectively) [9]. To validate the comparison, we also incorporated BC learning into their original Torch [6] training codes<sup>4</sup>. The 11-layer CNN was also incorporated into one of them. All of these codes use the standard shifting/mirroring data augmentation scheme that is widely used for these two datasets. We added 75 training epochs with a further smaller learning rate (1/10) to the default learning schedule of a total of 300 epochs for ResNet-29, ResNeXt-29, and DenseNet. We show the difference from the default settings in the appendix, as well as the configuration of 11-layer CNN. We trained each model 10 times for the 11-layer CNN and ResNet-29, and 5 times for other networks. We report the average and the standard error of the final top-1 errors.

We summarize the results in Table 2. The performances of all networks on CIFAR-10 were improved with the simple BC learning. Furthermore, with the improved version of BC learning (BC+), which treat the image data as waveforms, the performance was further improved. The best result on CIFAR-10 was 2.26% on Shake-Shake Regularization. The performance was stable, and the error rate of all 5 trials were in the range of 2.25%–2.28%. We do not know whether this result should be counted as the state-of-the-art; however, BC learning proves to be able to improve the performance of various networks, from a simple network to the state-of-the-art networks.

<sup>4</sup><https://github.com/facebookresearch/ResNeXt>, <https://github.com/liuzhuang13/DenseNet>, <https://github.com/xgastaldi/shake-shake>




Figure 5. Training curves on CIFAR-10 (average of all trials).

We also show the training curves in Fig. 5. Note that the training curves represent the average of all trials. Contrary to the training curves on ImageNet-1K, the testing error of BC learning decreases at almost the same speed as the standard learning in the early stage of the training. Furthermore, the last 75 training epochs for 11-layer CNN and DenseNet leads to a lower testing error when using BC learning.

The performances on CIFAR-100 were also improved with BC learning. Although it may be difficult to learn the between-class examples among 100 classes with no improvement to performance on ResNeXt-29 and Shake-Shake Regularization, BC learning shows a significant improvement on 11-layer CNN, ResNet-29, and DenseNet.

**Relationship with data augmentation.** Here, we show the performance when using no data augmentation in Table 3. We show the average of 10 trials. As shown in this table, the degree of improvement in the performance is at the same level as, or even smaller than when using the standard data augmentation, although the variation of training data increases from 50,000 to approximately  $50,000C_2$ . We assume this is because the potential within-class variance is small when using no data augmentation. If the within-class variance of the feature space is small, the variance of the features of the mixed images also becomes small, and the overlap in Fig. 2(top) becomes small. Therefore, the effect of BC learning becomes small as a result. We assume that BC learning is compatible with, or even strengthened by, a strong data augmentation scheme.

Table 3. Comparison when using no data augmentation. The figures between brackets indicate the error rates when using the standard data augmentation.

Model	Learning	Error rate (%) on	
		CIFAR-10	CIFAR-100
11-layer CNN	Standard	9.68 (6.07)	33.04 (26.68)
	BC (ours)	8.38 (5.40)	31.00 (24.28)
ResNet-29 [28]	Standard	8.38 (4.24)	31.36 (20.18)
	BC (ours)	7.69 (3.75)	30.79 (19.56)

### 4.3. Ablation analysis

To understand the part that is important for BC learning, we conducted an ablation analysis following [3]. We trained an 11-layer CNN on CIFAR-10 and CIFAR-100 using various settings. We implemented the codes of ablation analysis using Chainer v1.24 [27]. All results are shown in Table 4, as well as the average of 10 trials.

**Mixing method.** The differences of the improved BC learning (Eqn. 6) from the simplest BC learning (Eqn. 3) are as follows: a) we subtract par-image mean; b) we divide the mixed image by  $\sqrt{r^2 + (1-r)^2}$  considering that waveform energy is proportional to the square of the amplitude; and c) we take the difference of image energies into consideration. We investigated which of them has a great effect. As a result, considering the difference of image energies (c) proved to be of little significance, comparing a+b and a+b+c. This would be because the variance of image energies is smaller than that of sound energies. However, per-image mean subtraction (a) and dividing by the square root (b) are important (a+b+c vs. b+c and a vs. a+b). This result shows that treating the image data as waveforms contributes to the improvement in performance.

**Label.** We compared the different labels that we applied to the mixed image. The performance worsened when we used a single label ( $t = t_1$  if  $r > 0.5$ , otherwise  $t = t_2$ ) and softmax cross entropy loss. It would be inappropriate to train the model to recognize a mixed image as particular class. Using multi label ( $t = t_1 + t_2$ ) and sigmoid cross entropy loss marginally improved the performance, but the proposed ratio label and KL loss performed the best. The model can learn the between-class examples more efficiently when using our ratio label.

**Number of mixed classes.** We investigated the relationship between the performance and the number of classes of images that we mixed. Surprisingly, the performance was improved when we mixed two images belonging to the

Table 4. Ablation analysis. We trained 11-layer CNN on CIFAR-10 and CIFAR-100 using various settings. We report the average error rate of 10 trials.

Comparison of	Setting	Error rate (%) on	
		C-10	C-100
Mixing method	None (Eqn. 3, BC)	5.40	24.28
	a (Eqn. 4)	5.45	24.25
	a+b (Eqn. 5)	<b>5.17</b>	23.72
	a+b+c (Eqn. 6, BC+)	5.22	<b>23.68</b>
	b+c	5.26	23.98
Label	Single	6.35	27.28
	Multi	6.05	26.31
	Ratio (proposed)	<b>5.22</b>	<b>23.68</b>
# mixed classes	$N = 1$	5.98	26.01
	$N = 1$ or 2	5.31	23.79
	$N = 2$ (proposed)	5.22	<b>23.68</b>
	$N = 2$ or 3	<b>5.15</b>	23.78
	$N = 3$	5.32	24.20
Where to mix	Input (proposed)	<b>5.40</b>	24.28
	pool1	5.74	<b>24.09</b>
	pool2	6.52	25.38
	pool3	6.05	27.40
	fc4	6.05	26.70
	fc5	6.12	25.99
Standard learning		6.07	26.68

same class ( $N = 1$ ). Additionally, if we selected two images completely randomly and allowed the two images to be sometimes the same class ( $N = 1$  or 2), the performance was worse than proposed  $N = 2$ , in which two images were always different classes. Because BC learning is a method of providing constraints to the feature distribution of different classes, we should select two images belonging to the different classes. We also tried to use mixtures of three different classes with a probability of 0.5 in addition to the mixtures of two different classes ( $N = 2$  or 3), but the performance was not significantly improved from  $N = 2$ . Moreover, the performance when we used only the mixtures of three different classes ( $N = 3$ ) was worse than that of  $N = 2$  despite the larger variation in training data. We assume that mixing more than two classes cannot efficiently provide constraints to the feature distribution.

**Where to mix.** We investigated what occurs when we mix two examples within the network. Here, we used the simple mixing method of Eqn. 3. The performance was also improved when we mixed two examples at the layer near the input layer. We assume this is because the activations of lower layers can be treated as waveforms because the spatial information is preserved to some extent. Additionally, mixing at the layer close to the output layer had little effect on the performance. It is interesting that the performance was




Figure 6. Visualization of feature distributions using 3-D PCA. BC learning indeed imposes a constraint on the feature distribution.

worsened when we mixed two examples at the middle point of the network (pool2 and pool3 for CIFAR-10 and CIFAR-100, respectively). It is expected that the middle layer of the network extracts features that represent both spatial and semantic information simultaneously, and mixing such features would not make sense for machines.

#### 4.4. Visualization

Finally, we visualize the features learned with standard and BC learning in Fig. 6. We applied PCA to the activations of the 10-th layer of the 11-layer CNN trained on CIFAR-10 against the training data. As shown in this figure, the features obtained with BC learning are spherically distributed, and have small within-class variances, whereas that obtained with standard learning are widely distributed from near to far the decision boundaries. We conducted further analysis on the learned features in the appendix. In this way, BC learning indeed imposes a constraint on the feature distribution, which cannot be achieved with standard learning. We conjecture that is why the classification performance was improved with BC learning.

#### 5. Conclusion

We proposed a novel learning method for image classification called BC learning. We argued that CNNs have an aspect of treating input data as waveforms, and attempted to apply a similar idea to what has been done for sounds. As a result, the performance was significantly improved by simply mixing two images using internal divisions and training the model to output the mixing ratio. Moreover, the performance was further improved with a mixing method that treats the images as waveforms. BC learning is a simple and powerful method that can impose constraints on the feature distribution. We assume that BC learning can be applied not only to images and sounds but also to other modalities.

#### Acknowledgement

This work was supported by JST CREST Grant Number JPMJCR1403, Japan.



## References

- [1] O. Abdel-Hamid, A.-r. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu. Convolutional neural networks for speech recognition. *IEEE/ACM TASLP*, 22(10):1533–1545, 2014. 1
- [2] S. An, F. Boussaid, and M. Bennamoun. How can deep rectifier networks achieve linear separability and preserve distances? In *ICML*, 2015. 3
- [3] Y. Tokozume, Y. Ushiku, and T. Harada. Learning from between-class examples for deep sound recognition. In *ICLR*, 2018. <https://openreview.net/forum?id=BlGi6LeRZ>. 1, 2, 3, 4, 7
- [4] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, 2009. 6
- [5] P. Burt and E. Adelson. The laplacian pyramid as a compact image code. *TCOM*, 31(4):532–540, 1983. 4
- [6] R. Collobert, S. Bengio, and J. Mariéthoz. Torch: a modular machine learning software library. Technical report, Idiap, 2002. 5, 6
- [7] W. Dai, C. Dai, S. Qu, J. Li, and S. Das. Very deep convolutional neural networks for raw waveforms. In *ICASSP*, 2017. 4
- [8] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936. 3
- [9] X. Gastaldi. Shake-shake regularization. In *ICLR Workshop*, 2017. 2, 5, 6
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [12] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. *arXiv preprint arXiv:1709.01507*, 2017. 1
- [13] G. Huang, Z. Liu, K. Q. Weinberger, and L. van der Maaten. Densely connected convolutional networks. In *CVPR*, 2017. 5, 6
- [14] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 1
- [15] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 1
- [16] A. Krizhevsky. Learning multiple layers of features from tiny images. *Tech Report*, 2009. 6
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1, 6
- [18] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. In *IEEE*, 1998. 1
- [19] K. J. Piczak. Environmental sound classification with convolutional neural networks. In *MLSP*, 2015. 1
- [20] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 5
- [21] T. N. Sainath, R. J. Weiss, A. Senior, K. W. Wilson, and O. Vinyals. Learning the speech front-end with raw waveform cldnns. In *Interspeech*, 2015. 1, 4
- [22] K. S. Shanmugam, F. M. Dickey, and J. A. Green. An optimal frequency domain filter for edge detection in digital pictures. *TPAMI*, (1):37–49, 1979. 4
- [23] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 1
- [24] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 15(Jun):1929–1958, 2014. 1
- [25] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 1
- [26] Y. Tokozume and T. Harada. Learning environmental sounds with end-to-end convolutional neural network. In *ICASSP*, 2017. 1, 4
- [27] S. Tokui, K. Oono, and S. Hido. Chainer: a next-generation open source framework for deep learning. In *NIPS Workshop on Machine Learning Systems*, 2015. 7
- [28] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 1, 2, 5, 6, 7
- [29] X. Zhang, Z. Li, C. C. Loy, and D. Lin. Polynet: A pursuit of structural diversity in very deep networks. *arXiv preprint arXiv:1611.05725*, 2016. 1