

Received 7 January 2023, accepted 30 January 2023, date of publication 6 February 2023, date of current version 23 February 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3243108

## RESEARCH ARTICLE

# Observations on K-Image Expansion of Image-Mixing Augmentation

JOONHYUN JEONG<sup>1</sup>, SUNGMIN CHA<sup>2</sup>, JONGWON CHOI<sup>3</sup>, (Member, IEEE), SANGDOO YUN<sup>4</sup>,  
TAESUP MOON<sup>2</sup>, (Member, IEEE), AND YOUNGJOON YOO<sup>1</sup>, (Member, IEEE)

<sup>1</sup>Image Vision Team, NAVER CLOVA, Seongnam 13561, South Korea

<sup>2</sup>Department of Electrical and Computer Engineering, Seoul National University, Seoul 08826, South Korea

<sup>3</sup>Department of Advanced Imaging, Graduate School of AI (GSAIM), Chung-Ang University, Seoul 06973, South Korea

<sup>4</sup>AI Laboratory, NAVER CLOVA, Seongnam 13561, South Korea

Corresponding author: Youngjoon Yoo (youngjoon.yoo@navercorp.com)

This work was supported in part by CLOVA, NAVER Corporation; in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) through the Korean Government (MSIT), through the Artificial Intelligence Graduate School Program, Chung-Ang University, under Grant 2021-0-01341; and in part by the Next Generation AI for Multi-Purpose Video Search under Grant 2021-0-02067.

**ABSTRACT** Image-mixing augmentations (e.g., Mixup and CutMix), which typically involve mixing two images, have become the de-facto training techniques for image classification. Despite their huge success in image classification, the number of images to be mixed has not been elucidated in the literature: only the naive K-image expansion has been shown to lead to performance degradation. This study derives a new K-image mixing augmentation based on the stick-breaking process under Dirichlet prior distribution. We demonstrate the superiority of our K-image expansion augmentation over conventional two-image mixing augmentation methods through extensive experiments and analyses: 1) more robust and generalized classifiers; 2) a more desirable loss landscape shape; 3) better adversarial robustness. Moreover, we show that our probabilistic model can measure the sample-wise uncertainty and boost the efficiency for network architecture search by achieving a 7-fold reduction in the search time. Code will be available at <https://github.com/yjyoo3312/DCutMix-PyTorch.git>.

**INDEX TERMS** Image classification, augmentation, Dirichlet process.

## I. INTRODUCTION

The advent of deep classification networks has emphasized the importance of data augmentation [1], [2], [3]. Proper data augmentation can remedy the performance degradation due to insufficient data and weak robustness to noisy data [4]. Accordingly, many researchers have proposed training strategies to apply the data augmentation methods to the deep classification network.

Among the popular data augmentation methods, image-mixing augmentation methods, especially CutMix [2] exhibited impressive performance in training large-scale deep classification networks. Image-mixing augmentation

methods augment a new image by mixing the two paired images. For example, CutMix mixes the paired images by re-formulating their segments into one image. By applying this simple principle, the image-mixing augmentation successfully improves the performance of deep classification networks in various scenarios. Furthermore, through image-mixing augmentation, the deep learning model becomes robust to corrupted and uncertain data.

However, the mechanism underlying image-mixing augmentation is still not fully understood. Specifically, even the optimal number of images to mix has not been elucidated: in response, the number of  $K$  images was empirically set to 2. Researchers [3], [5] have made naive attempts at  $K$ -image expansions of the augmentation. However, the  $K$ -image expansion attempts have

The associate editor coordinating the review of this manuscript and approving it for publication was Sudipta Roy<sup>1</sup>.

been unsuccessful in terms of classification performance improvements. Here, we aim to answer the following question: Is  $K = 2$ , the number of image, optimal for image-mixing augmentation?

In this study, we derive a novel formulation for generalizing image-mixing augmentation and apply it to obtain improved results for image-mixing augmentation methods. Notably, we find that a mixture of three or more images can further improve the performance of baseline methods using only the paired images. The superiority of the generalized formulation is validated under different classification scenarios. In addition, we test the robustness of our method: the results reveal that our method can drive the model into the widest (flattest) and deepest local minima. In terms of adversarial robustness, we experimentally demonstrate that the proposed image-mixing augmentation methods strengthen adversarial robustness and reveal that the expansion into the K-image case further improves the robustness.

Additionally, we demonstrate that the proposed image-mixing augmentation can be used to characterize and estimate the uncertainty of the data samples. Based on the estimated uncertainty, we acquire the subsampled data pool that can efficiently represent the overall data distribution. We validate the efficiency of our subsampling framework under the proposed scheme on network architecture search (NAS). Notably, our method preserves the performance while achieving 7.7 times higher training speed when using the subsampled data pool as a training set.

Our contribution can be summarized as follows.

- We generalize the image-mixing augmentations for image classification and achieve better generalization ability on unseen data against the baseline methods.
- We experimentally analyze the mechanism behind the better generalization of K-image augmentation by illustrating a loss landscape near the discovered minima. Accordingly, we reveal its ability to achieve convergence to wider and deeper local minima. We also demonstrate that K-image augmentation improves the adversarial robustness of the model.
- We propose a new data subsampling method by measuring sample uncertainty based on the proposed image-mixing augmentation, which is especially beneficial for handling a small number of training samples. We further verify the efficiency of the proposed subsampling method by applying it to NAS.

The rest of the paper is organized as follows. In Section II, we list related studies of image augmentation, data efficiency, and architecture search. Section III describes details of the formulation, implementation, and applications of the proposed augmentation method. Section IV demonstrates the experimental results of classification on CIFAR and ImageNet, adversarial robustness, and NAS from the subsampled dataset by the proposed methods. Finally, we conclude the paper by mentioning the limitations in Section V.

## II. RELATED WORKS

### A. AUGMENTATION

Including augmentation in training classification networks has become standard practice for achieving high performance. Beginning with simple augmentations, such as random crop, flipping, and color jittering, increasingly complex techniques, including including Cutout [1], Mixup [3], CutMix [2], PuzzleMix [6], SaliencyMix [7], and Co-Mixup [5] have been applied. Among the latter, CutMix, Mixup, PuzzleMix, and SaliencyMix typically mix two images, and a recent variant, Co-Mixup, has reportedly achieved an impressive enhancement in classification performance. Co-Mixup also generalized the image-mixing augmentation methods into K-image cases using submodular-supermodular optimization, which involves huge computational cost. Notably, our proposed K-image mixing augmentation methods do not require optimization and thus, have less computational overhead than Co-Mixup while achieving similar performance (Table 2).

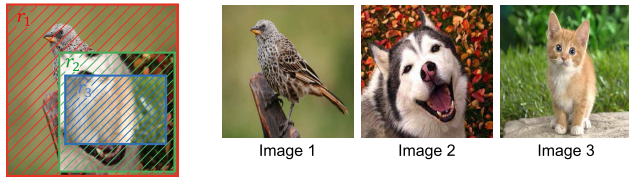
### B. DATA EFFICIENCY

Several approaches, with the aim of efficiently utilizing a training dataset with a semantically important measure, have focused on collecting examples that are considered informative by re-weighting the importance of different training samples: calculating the importance value from additional forward [8] and backward path [9] of training, defining the approximated function [10], or using loss based training scheme [11], [12], [13]. Nevertheless, a criterion based on the hardness of the example cannot be generalized if the samples contain label noise. Reference [14] also showed that hard examples, unlike easy examples, are unsuitable for the initial stages of training. Building upon previous works on measuring the importance of samples, we propose a robust importance subsampling methodology. We apply our subsampling concept to the differentiable search-based NAS and achieve performance improvements in both search time and classification accuracy.

### C. NETWORK ARCHITECTURE SEARCH

Initiative NAS [15], [16], [17], [18], [19] utilizing reinforcement learning (RL) requires significant computational cost so that is difficult to apply them to ImageNet scale dataset. To alleviate the problem, the weight-sharing NAS [20], [21], [22], [23], [24], [25], [26] introduce the *SuperNet* concept, which includes all the operation in the search space and extract the target architecture, *SubNet* from the *SuperNet*. For the extraction of the *SubNet*, [20], [21], [22], [26] propose a gradient-based searching method, which has become dominant in the research field, currently. In this study, we demonstrate the effectiveness of the subsampled data from our proposed DCutMix in NAS by implementing it to PC-DART [21]. Like the other methods, PC-DARTS focuses on designing a cell, and a user can easily adapt the

layer depth during the architecture search phase by appending or removing more of the search cells in the search space.



**FIGURE 1. Example of three-image composition in CutMix case.** We composite red box to green box, green box to blue box with the ratio of  $1 : 1 - v_1$ , and  $1 : 1 - v_2$  as in (4). Consequently, the region proportion of each image fraction  $r_1$  (red diagonal pattern),  $r_2$  (green diagonal pattern), and  $r_3$  (blue diagonal pattern) will correspond to  $\{\phi_1, \phi_2, \phi_3\}$ , which follows Dirichlet distribution. Notably, a low variable anchor image (anchor image with red border) mostly serves as either an easy or hard sample regardless of the occlusion position. In contrast, a highly variable anchor image (anchor image with green border) serves as both an easy and hard sample depending on the random image mixing operation. These highly variable anchor images possibly provide more diverse information during training.

### III. PROPOSED METHOD

In this section, we define a formulation of the proposed K-image mixing augmentation and apply a probabilistic augmentation framework to an image classification task. Accordingly, as a novel method of applying the proposed K-image mixing augmentation, we propose a subsampling method that utilizes the uncertainty measurement in the augmented data samples.

#### A. FORMULATION FOR K-IMAGE MIXING AUGMENTATION

##### 1) K-IMAGE MIXING AUGMENTATION

In this subsection, we formulate the K-image generalization for image mixing augmentation on the image classification task. We consider that augmented sample  $x_c$  is composed of  $x_1, \dots, x_K$ , denoted as:

$$x_c = f_c(x_1, \dots, x_K; \phi_1, \dots, \phi_K), \quad (1)$$

where the function  $f_c(\cdot)$  denotes the composite function, and the term  $\phi = \{\phi_1, \dots, \phi_K\}$  is a mixing parameter denoting the portion of each sample  $x_k$  on the composite sample  $x_c$ . Note that Equation (1) can be considered as the general form of the popular image-mixing augmentations, such as CutMix [2] and Mixup [3] which mix only two images. Specifically, in the case of Mixup (denoted as *DMixup*) the function  $f_c(\cdot)$  is defined by the weighted summation as follows:

$$x_c = \sum_{k=1}^K \phi_k x_k. \quad (2)$$

The mixing parameter  $\phi$  is defined by the Beta distribution in the usual two-image cases (e.g., CutMix [2] and Mixup [3]) Note that Equation (2) can be naturally expanded to the case of the K-image mixing case by applying a Dirichlet distribution  $\phi \sim \text{Dir}(\alpha)$ ,  $\alpha \in \mathbb{R}^K$ . In this case, the composite sample  $x_c$  becomes a random variable for the given hyper-parameter  $\alpha$ .

##### 2) K-IMAGE GENERALIZATION OF CutMix

Based on the above formulation, we define the K-image generalization of Cutmix (denoted as *DCutMix*). Note that the definition of the function  $f_c(\cdot)$  becomes more complicated because the function should contribute all the segments of images  $x_1, \dots, x_K$  to composite image  $x_c$  considering their mixing parameters  $\phi$ . Here, we composite the images proportionally following the stick-breaking process (SBP [27]), with the widely used approach of sampling from Dirichlet distribution.

Assume that  $\phi$  is sampled from the prior distribution  $\text{Dir}(\alpha)$ . The K-image mixing augmentation of CutMix is conducted by compositing the image with respect to the proportion  $\phi_k \in \phi$ , where  $\sum_k \phi_k = 1$ . For sampling  $\phi$  from  $\text{Dir}(\alpha)$ , we use SBP by leveraging an intermediate variable  $v$  as follows. Firstly, let  $v = [v_1, \dots, v_{K-1}] \in \mathbb{R}^{K-1}$  and each  $v_k$  is denoted as follows:

$$v_1 = \phi_1$$

$$v_k = \phi_k / \prod_{j=1}^{k-1} (1 - v_j), \quad k = 2, \dots, K - 1. \quad (3)$$

Note that the variable  $v$  is sampled from the beta distribution ( $\text{Beta}(1, \alpha)$ ) by deriving SBP. Now, we define the image fractions  $r = \{r_1, \dots, r_K\}$  from K different images which constitute to a mixed sample  $x \in \mathbb{R}^{W \times H \times C}$ . Let the function  $\tilde{r} = d(x|v)$  randomly discriminate the image fractions  $\tilde{r} : x \setminus \tilde{r}$  with the area ratio  $v : 1 - v$ , where  $x \setminus \tilde{r}$  denotes the region of  $x$  excluding  $\tilde{r}$ . Consequently, the fractions  $r$  are determined by following equation:

$$r_k = d(x \setminus \sum_{j=0}^{k-1} r_j | v_k), \quad k = 1, \dots, K - 1, \quad (4)$$

where the virtual fraction  $r_0$  and the last fraction  $r_K$  are set to  $\emptyset$  and  $x \setminus \sum_{j=1}^{K-1} r_j$ , respectively. The discrimination function  $d(\cdot)$  determines the exact bounding box coordinates  $r_{k_x}, r_{k_y}, r_{k_w}, r_{k_h}$  of image fraction  $r_k$ , to be located within the bounding box coordinates of former image patch  $r_{k-1}$ . These coordinates are randomly sampled from the uniform distribution with random variable  $\gamma$ , as follows:

$$r_{k_x} \sim \text{Unif}(r_{k-1_x}, r_{k-1_x} + r_{k-1_w} - r_{k_w}),$$

$$r_{k_y} \sim \text{Unif}(r_{k-1_y}, r_{k-1_y} + r_{k-1_h} - r_{k_h}) \quad (5)$$

where its width  $r_{k_w}$  and height  $r_{k_h}$  are determined by  $v_k$  as defined in Equation (4). Note that, in the case of  $k = 1$ ,  $r_{1_x} = 0, r_{1_y} = 0, r_{1_w} = W$  and  $r_{1_h} = H$ . Hence, the composite function  $f_c$  of the *DCutMix* is governed by hyper-parameter  $\alpha$  and random variable  $\gamma$ . An illustration of the proposed K-image mixing augmentation following Equation (4) is presented in Figure 1.

In the subsequent experiment section, we will experimentally demonstrate the advantages of the proposed K-image generalization in terms of loss landscape, adversarial robustness, and classification accuracy.



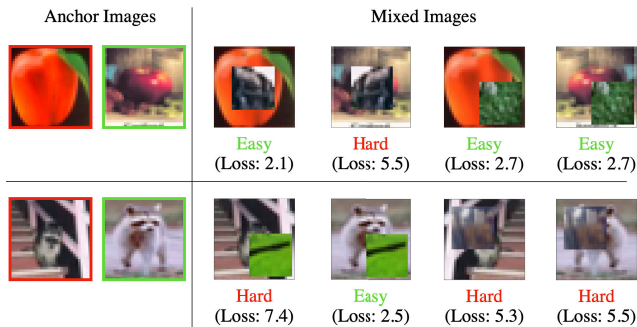
### 3) PROBABILISTIC FRAMEWORK

The overall probabilistic framework of the classification problem, considering the proposed augmentation, can be defined as:

$$\begin{aligned}
 p(l_c|x_c) &= \int p(l_c|x_c, \phi)p(\phi|\alpha)d\phi, \\
 &= \int p(\sum_k\{\phi_k l_k\}|f_c(x; \phi))p(\phi|\alpha)d\phi, \\
 &\cong \frac{1}{N_s} \sum_{\phi^{(j)}} f_W(\{\sum_k \phi_k^{(j)} l_k\}|f_c(x; \phi^{(j)})), \quad (6)
 \end{aligned}$$

where  $(\mathbf{x}, \mathbf{l}) = \{(x_1, l_1), \dots, (x_K, l_K)\}$ ,  $l_c = \sum_k\{\phi_k l_k\}$ , and  $\phi^{(j)}$  is the  $j^{th}$  sample drawn from the Dirichlet prior distribution  $p(\cdot|\alpha)$ . Hereafter, we define the label  $l_i \in \mathbb{R}^L$  as a one-hot indexing variable denoting one of total  $L$  total classes. Based on the derivation from Monte-Carlo dropout [28], we can approximate the distribution  $p(l_c|x_c)$  in the variational function  $f_W(\cdot)$  with regard to several different  $\phi$  and  $\gamma$  samples. The variational function is realized by a classification network, parameterized by  $W$ , with a softmax output. In the case of DCutMix, we additionally consider another variable  $\gamma$  from (6), such as:

$$p(l_c|x_c) \cong \sum_{\gamma^{(i)}} \sum_{\phi^{(j)}} f_W(\{\sum_k \phi_k^{(j)} l_k\}|f_c(x; \phi^{(j)}, \gamma^{(i)})). \quad (7)$$



**FIGURE 2.** Example of composed samples where the loss values are highly variable or invariable depending on the position of occlusion by non-anchor image patches. The anchor images (bordered with green) have highly variable loss values depending on the occlusion. In contrast, the other anchor images (bordered with red) have relatively invariable loss values regardless of the occlusion.

Consequently, from (6) and (7), we can approximate the posterior  $p(l_c|x_c)$  by estimating the predictive mean of network outputs, depending on several differently augmented data sampled for varying  $\phi$  and  $\gamma$  values. Similarly, the uncertainty of a given data sample  $x_c$  for the given classification network can be approximated by calculating the posterior estimated from augmented data samples.

### B. IMPLEMENTATION DETAILS OF DCutMix

In this section, we present the implementation details of DCutMix. We describe the pseudo-code of the mixing process of DCutMix in Algorithm 1. First, we sample variable

$\phi$  from  $Dir(\alpha)$  (see Line 1). For  $K - 1$  iterations, we cut and mix  $K - 1$  image fractions. At each iteration, a mini-batch input and target are shuffled along with the batch dimension. An intermediate variable  $v$  is then selected using SBP sampled from  $\phi$  (see Line 6, 7, 12, and 13). The variable  $v$  determines the width and height of the image patch to be mixed, where the exact position is bounded on the former image patch (see Lines 17 and 18). We then cut an image patch from source images  $x_s$  and mix on  $x_c$  (see Line 19). In Lines 20-27, the soft label is accordingly mixed by  $\lambda$  and  $\lambda_{K-1}$ , which denote the exact area ratio of each mixed image patch.

### C. SUBSAMPLING USING THE MEASURED DATA UNCERTAINTY

As a new method of utilizing the K-image mixing augmentation, we propose a novel subsampling method that considers the data uncertainty obtained from K-image augmentation for the first time. In order to measure the uncertainty of a data sample, we define the loss distribution  $L(l_c|x_c)$  for variously augmented data samples depending on  $\phi$ ,  $\gamma$  and its expectation can be approximated based on (7) as follows:

$$\mathbb{E}[L(l_c|x_c)] \cong \sum_{\gamma^{(i)}} \sum_{\phi^{(j)}} \mathcal{L}(\{\sum_k \phi_k^{(j)} l_k\}|f_c(x; \phi^{(j)}, \gamma^{(i)})), \quad (8)$$

where  $\mathbf{x} = \{x_1, \dots, x_K\}$ , and  $\mathcal{L}$  denotes the cross-entropy loss. The expectation is defined on the space by the random variable  $\phi$  and  $\gamma$ . Similarly, the uncertainty can also be acquired by estimating the variance of the loss distribution  $L$ . Figure 2 shows qualitative examples of uncertainty measurement, given sample data and their mixed images. Noticeably, the diverse tendency of loss values changes for each mixed image, mainly depending on the randomly selected position of occlusion caused by non-anchor image patches.

For measuring the sample-wise uncertainty using the loss distribution, we select an *anchor sample*  $x_i \in \mathbf{x}$  with fixed  $\phi_i$  and then jitter  $\phi \setminus \phi_i$  related to other non-anchor samples  $x \setminus x_i$  to calculate the uncertainty of the anchor sample  $x_i$ . The  $\phi \setminus \phi_i$  are drawn from a conditional Dirichlet distribution  $D(\alpha \setminus \alpha_i)$ , according to its definition. We will term  $L_i = \{\mathcal{L}_{i,m} | m = 1, \dots, M\}$  as the loss distribution for all the mixed images given the anchor  $x_i$ ; the corresponding loss is calculated from (8). The number  $M$  denotes the total number of sampling  $\phi \setminus \phi_i$  from  $D(\alpha \setminus \alpha_i)$ .

Based on the sample-wise uncertainty measurement, we aim to sample the core training data sub-set among the entire training dataset. Presumably, for better generalization of a neural network when training using a small number of data points and image-mixing augmentation, the core training sub-set should consist of the highly uncertain samples which can serve as both easy- and hard-level samples depending on the image-mixing augmentation (e.g., the images bordered with green in Figure 2). Therefore, a new training subset is subsampled by descending order of uncertainty measure. We observed that employing the coefficient of variation (CV) metric, which is defined as  $\frac{\sigma(L_i)}{m(L_i)}$  where  $\sigma(\cdot)$  and  $m(\cdot)$  is the

**Algorithm 1** Pseudo-Code of DCutMix

**Input:** Input images  $x \in \mathbb{R}^{N \times C \times W \times H}$ , Target labels  $l \in \mathbb{R}^{N \times L}$ , Number of mixing images  $K$ , Parameter of Dirichlet distribution  $\alpha \in \mathbb{R}^{N \times K}$

**Output:** Augmented image  $x_c$ , Augmented label  $l_c$

- 1:  $\phi = \text{Dir}(\alpha)$
- 2:  $x_c = x$
- 3: **for**  $k \leftarrow 1$  to  $K - 1$  **do**
- 4:  $x_s, l_s = \text{Shuffle}(x, l)$   
// stick-breaking process from equation (3)
- 5: **if**  $k == 1$  **then**
- 6:  $v = \phi[k]$  {sample intermediate variable  $v$ .}
- 7:  $v_m = 1 - v$
- 8:  $r_{w,p}, r_{h,p} = W, H$  {set the size of bounding box.}
- 9:  $r_{x,p}, r_{y,p} = 0, 0$
- 10:  $l_p = l$  {set a label of the previous image patch.}
- 11: **else if**  $k \leq K - 1$  **then**
- 12:  $v = \phi[k]/v_m$
- 13:  $v_m = v_m \times (1 - v)$
- 14: **end if**  
// bounding box setting, from equation (4) and (5).
- 15:  $r_w = \text{Round}(r_{w,p} * \sqrt{1 - v})$
- 16:  $r_h = \text{Round}(r_{h,p} * \sqrt{1 - v})$
- 17:  $r_x \sim \text{Unif}(r_{x,p}, r_{x,p} + r_{w,p} - r_w)$
- 18:  $r_y \sim \text{Unif}(r_{y,p}, r_{y,p} + r_{h,p} - r_h)$
- 19:  $x_c[:, :, r_x : r_x + r_w, r_y : r_y + r_h] = x_s[:, :, r_x : r_x + r_w, r_y : r_y + r_h]$
- 20:  $\lambda = (r_{w,p} \times r_{h,p} - r_w \times r_h) / W \times H$  {set the mixing ratio.}
- 21: **if**  $k == 1$  **then**
- 22:  $l_c = \lambda \times l_p$  {make a soft label based on the mixing ratio.}
- 23: **else if**  $k < K - 1$  **then**
- 24:  $l_c = l_c + \lambda \times l_p$
- 25: **else**
- 26:  $\lambda_{K-1} = (r_w \times r_h) / (W \times H)$
- 27:  $l_c = l_c + \lambda \times l_p + \lambda_{K-1} \times l_s$
- 28: **end if**
- 29:  $r_{w,p}, r_{h,p} = r_w, r_h$  {record the bounding box and soft label.}
- 30:  $r_{x,p}, r_{y,p} = r_x, r_y$
- 31:  $l_p = l_s$
- 32: **end for**
- 33: **Return**  $x_c, l_c$  {return an augmented image and soft label.}

standard deviation and average of  $L_i$ , is most effective for measuring the uncertainty (See Figure 5 for details).

## 1) SUBSAMPLING DETAILS

We herein describe the implementation details for the proposed subsampling framework. A newly subsampled set  $D$  for each class is defined as follows:

$$D = S(O(L_k) | k = j_1, \dots, j_{N_{intra}}, t), \quad (9)$$

where  $O(\cdot)$  denotes the subsampling measure and  $S(\cdot)$  denotes a sampling function indicating whether data sample  $x_k$  is to be included in  $D$  or not by using the subsampling ratio  $t$ . Here,  $j$  denotes the index of  $N_{intra}$  number of intra-class images where the class labels are equivalent among the others.  $O(\cdot)$  is a proxy for subsampling; data samples are subsampled in order of  $O$ . With regard to the sampling function  $S(\cdot)$ , it samples  $t \times N_{intra}$  data samples based on the sampling measure  $O(\cdot)$ , which falls into two categories: a *deterministic* function sampling top  $t \times N_{intra}$  samples sorted by  $O(\cdot)$ , and an *interval-based* function that collect samples sorted by  $O(\cdot)$  with a fixed interval.

**TABLE 1.** Comparison of DCutMix and DMixup against other augmentations and regularization methods for PyramidNet-110, 200 models on CIFAR-100 dataset (left) and CIFAR-10 dataset (right). The values in the parentheses of the top-1 error denote the reduced amount of the error compared to the vanilla model, where no augmentation or regularization was applied.

Model	# Params	Top-1 Err (%)
PyramidNet-110 ( $\tilde{\alpha} = 64$ ) [29]	1.7 M	19.85
+ Mixup [3]	1.7 M	18.92 (-0.93)
+ CutMix [2]	1.7 M	17.97 (-1.88)
+ DMixup ( $K = 3, \alpha = \frac{1}{3}$ )	1.7 M	<b>18.60 (-1.25)</b>
+ DCutMix ( $K = 5, \alpha = 0.2$ )	1.7 M	<b>16.95 (-2.90)</b>
PyramidNet-200 ( $\tilde{\alpha} = 240$ )	26.8 M	16.45
+ StochDepth [30]	26.8 M	15.86 (-0.59)
+ Label Smoothing [31]	26.8 M	16.73 (+0.28)
+ Cutout [1]	26.8 M	16.53 (+0.08)
+ DropBlock [32]	26.8 M	15.73 (-0.72)
+ Mixup [3]	26.8 M	15.63 (-0.82)
+ Manifold Mixup [33]	26.8 M	15.09 (-1.36)
+ CutMix [2]	26.8 M	14.47 (-1.98)
+ StyleMix [34]	26.8 M	16.37 (-0.08)
+ StyleCutMix [34]	26.8 M	14.17 (-2.28)
+ DMixup ( $K = 3, \alpha = 1$ )	26.8 M	<b>15.07 (-1.38)</b>
+ DCutMix ( $K = 5, \alpha = 1$ )	26.8 M	<b>13.86 (-2.59)</b>
Model		Top-1 Err (%)
PyramidNet-200 ( $\tilde{\alpha} = 240$ )		3.85
PyramidNet-200 + Cutout [1]		3.10 (-0.75)
PyramidNet-200 + Mixup [3]		3.09 (-0.76)
PyramidNet-200 + Manifold Mixup [33]		3.15 (-0.7)
PyramidNet-200 + CutMix [2]		2.88 (-0.97)
PyramidNet-200 + DMixup ( $K = 5, \alpha = 0.2$ )		<b>2.90 (-0.95)</b>
PyramidNet-200 + DCutMix ( $K = 3, \alpha = \frac{1}{3}$ )		<b>2.42 (-1.43)</b>

Regarding the subsampling measure  $O(\cdot)$ , we employed the sample-wise uncertainty measure using Coefficient Variation (CV),  $\frac{\sigma(L_i)}{m(L_i)}$  (where  $L_i$  is derived from (8)). For estimating the sample-wise uncertainty, we set the number of non-anchor images  $x \setminus x_i$  and their Dirichlet sampling parameter  $\alpha \setminus \alpha_i$  as 2 and  $\{\frac{2}{9}, \frac{2}{9}\}$ , respectively. Additionally, we set the total number of sampling  $\phi \setminus \phi_i$  from the Dirichlet distribution, namely  $M$ , as 10.

**TABLE 2.** Comparison of Top-1 error rate and training time for the state-of-the-art image mixing augmentations on CIFAR-100 with various backbone models. (a): PreActResNet18 [35], (b): WRN16-8 [36], (c): ResNeXt29-4-24 [37]. Values in the parentheses denote the ratio of average training time relative to the baseline (vanilla) case.

Model	Vanilla	Mixup	CutMix	PuzzleMix	Co-Mixup	StyleCutMix	DMixup (ours)	DCutMix (ours)
a	23.59 (1x)	22.43 (1.04x)	21.29 (1.10x)	20.62 (2.73x)	19.87 (29.34x)	20.49 (56.52x)	21.68 (1.14x)	20.50 (1.07x)
b	21.70 (1x)	20.08 (1.02x)	20.55 (1.04x)	19.24 (2.25x)	19.15 (13.46x)	18.90 (74.98x)	<b>19.13</b> (1.02x)	<b>18.64</b> (1.01x)
c	21.79 (1x)	21.70 (1.02x)	22.28 (1.05x)	21.12 (2.22x)	19.78 (6.93x)	19.94 (43.48x)	20.09 (1.02x)	20.48 (1.02x)

#### IV. EXPERIMENTS AND DISCUSSION

Here, we experimentally verify the effect of the K-expanded image mixing augmentation. First, we show the improved classification performance after applying our method to CIFAR-10/100 and analyze its advantages in terms of the shape of the loss landscape. Second, on ImageNet, we propose an elaborately designed K-image mixing augmentation that considers the saliency map to overcome label noise. Moreover, we present the experimental result on classification and adversarial robustness for further discussion. Finally, we demonstrate the effectiveness of the proposed data subsampling method and its practical application in NAS.

##### A. EXPERIMENTAL RESULT ON IMAGE CLASSIFICATION

###### 1) CIFAR-10/100

We present classification test results on CIFAR-10 and 100 [38] datasets in Table 1. The results were obtained from the equivalent training and augmentation-specific hyperparameter setup used in [2]. Firstly, Table 1 (left) presents the superiority of DCutMix and DMixup over the other augmentation and regularization methods on the CIFAR-100 dataset. First, for light-weight backbone PyramidNet-110 [29], DCutMix and DMixup improve the performance compared to the baselines (i.e., CutMix and Mixup) by approximately 1% and 0.32%, respectively. For the deeper neural network PyramidNet-200, DCutMix and DMixup achieved the enhancement compared to the baselines, and DCutMix shows the lowest top-1 error compared to other baselines. Secondly, we evaluate our proposed methods on the CIFAR-10 dataset as shown in Table 1 (right). We again observe that DCutMix and DMixup both achieved performance enhancement. Specifically, DCutMix achieved the best performance among the baseline augmentation methods we tested.

###### *a*: COMPARISON TO RECENT AUGMENTATION METHODS

Also, we further compare our DCutMix and DMixup with state-of-the-art image-mixing augmentation methods, including PuzzleMix [6], Co-Mixup [5], and StyleMix [34], in Table 2. As seen in the results, DCutMix achieved better accuracy than PuzzleMix and a competitive accuracy compared to the Co-Mixup. We note that the proposed augmentation provides comparable classification performance with achieving superior calculation time compared with recently published augmentation methods such as StyleMix and Co-Mixup [39]. Co-Mixup and StyleMix each require more training time overhead (over 20 times and 50 times) than

ours due to the high optimization cost. These overall results demonstrate the effectiveness of the proposed K-image generalization for augmentation methods.

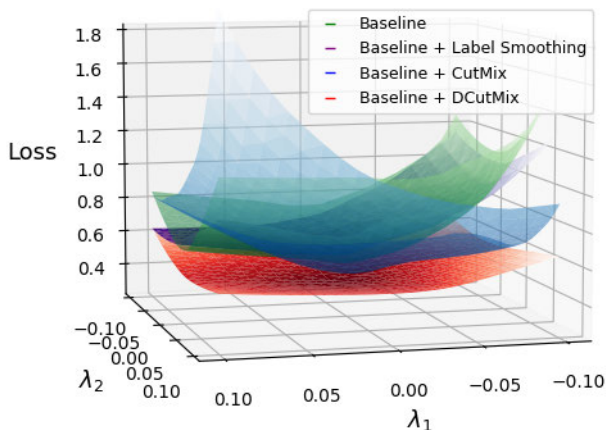
###### *b*: ANALYSIS ON THE SHAPE OF THE LOSS LANDSCAPE

For more explicit investigation, we analyze DCutMix with regard to its loss landscape. Flatness of the loss landscape near local minima has been considered as a key indicator of improved model generalization in various situations in numerous previous studies [40], [41], [42], [43], [44]. Regarding the shape of the loss landscape, convergence to a wide (flat) local minima is generally considered to represent a model with better generalization performance on an unseen test dataset.

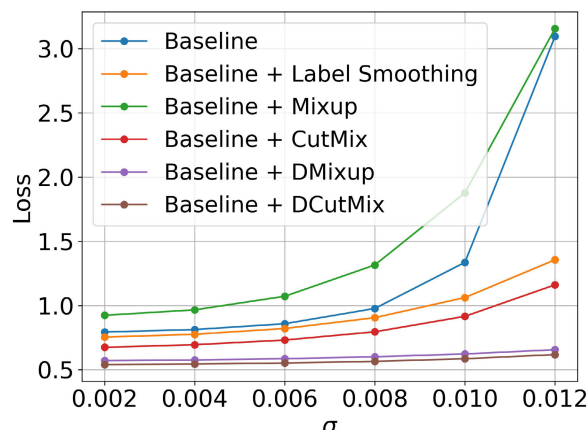
Accordingly, we use the PyHessian [45] framework to obtain the loss landscape patterns of each model, as illustrated in Figure 3. The plotted result shows that DCutMix has the widest loss landscape near local minima among the compared models. Moreover, DCutMix exhibits lower losses overall, denoting good generalization to the unseen test data as well. We further plotted the patterns of loss landscape for each model in Figure 3a by perturbing the model parameters with random Gaussian noise through increasing the degree of variance  $\sigma$  [42]. DCutMix and DMixup clearly exhibited the widest and lowest loss landscape compared to the other methods, including CutMix and Mixup, which are baseline two-image mixing augmentation methods. For DMixup particularly, we observed that the convergence stability was better than that of Mixup. We believe that this reveals the superiority of the proposed K-image mixing augmentation.

From a more analytical point of view, we can hypothesize that the K-image generalization of DCutMix and DMixup, regarding the wide flat local minima, is attributable to their labels being softer than that of CutMix and Mixup. Several researchers have reported that a model trained with an artificially smoothed label can result in the model converging to wide local minima, thus achieving better generalization [41], [42], [44], as in the case of the superior results of Label Smoothing [41] compared to the baseline in Figure 3a and Figure 3b. However, as opposed to the previous regularization methods using *artificially* smoothed label, note that our softened label *directly* reveals the augmented ratio of several images, and hence, we conjecture the tendency can be a key factor why the model trained by our approach converges to lower and wider local minima.





(a) 3-d loss surface plot by perturbing the model across the two directions using PyHessian.



(b) 2-d loss surface plot by perturbing the model across one random direction.

**FIGURE 3.** Comparison of image mixing augmentation and regularization methods in perspective of the loss-surface near local minima. We measured the loss surfaces on the test set of CIFAR-100 with PyramidNet [29].

2) ImageNet

We present ImageNet classification results of DCutMix and DMixup compared to the two-image mixing baselines: CutMix and Mixup. The results are obtained under the equivalent training and augmentation-specific hyper-parameter setup used in [2]. As presented in Table 3 (left), DMixup considerably improved the performance of Mixup and Manifold Mixup by 0.7% and 0.62% respectively, while reducing the top-1 error. However, DCutMix exhibited a higher top-1 error rate compared to CutMix. This result was attributable to DCutMix suffering from the label noise problem, where a background object other than the ground truth class object is contained in the randomly cropped image [46], as shown in Figure 4. Moreover, Table 3 (right) reveals that as the number of mixing images  $K$  increased, the performance of DCutMix deteriorated due to the higher probability of background objects being accumulated.

To address this label noise problem, we devised a more sophisticated mixing method named Saliency-DCutMix, which employs saliency-map information for integration with our DCutMix. First, we obtain a salient image patch by selecting the most salient pixel point of the saliency map as the center point, as suggested in [7]. Here, the width and height of each patch from (3) to ensure the Dirichlet distribution is followed. Consequently, we mixed these salient image patches with SBP, similar to in DCutMix, as given in (4). Figure 4 shows the qualitative examples of DCutMix and Saliency-DCutMix. Samples augmented with DCutMix contain background class objects other than the ground truth object, which could lead to label noise during training. Meanwhile, samples augmented with Saliency-DCutMix reveal that the foreground class objects are mixed without background class objects being included. In Table 3 (right), Saliency-DCutMix indeed exhibits relatively stable and significantly improved performance regardless of  $K$  compared to DCutMix.

**TABLE 3.** Performance of DMixup and DCutMix, and Saliency-DCutMix on ImageNet (left). Impact of  $K$  on DCutMix and Saliency-DCutMix (right).

Model	Top-1 Err (%)	Top-5 Err (%)
ResNet-50 (Baseline)	23.68	7.05
ResNet-50 + Mixup	22.58	6.40
ResNet-50 + Manifold Mixup	22.50	6.21
ResNet-50 + CutMix	21.40	5.92
ResNet-50 + DMixup ( $K = 3, \alpha = \frac{1}{3}$ )	21.88	6.15
ResNet-50 + DCutMix ( $K = 3, \alpha = \frac{1}{3}$ )	21.76	5.91
ResNet-50 + Saliency-DCutMix ( $K = 3, \alpha = \frac{1}{3}$ )	21.38	5.87

Method	$\alpha$	$K$	Top-1 Err (%)	Top-5 Err (%)
ResNet-50 + DCutMix	1	3	21.62	5.92
ResNet-50 + DCutMix	1	4	22.17	6.23
ResNet-50 + DCutMix	1	5	22.02	6.05
ResNet-50 + DCutMix	1	7	22.39	6.17
ResNet-50 + DCutMix	1	9	22.58	6.34
ResNet-50 + Saliency-DCutMix	1	3	21.40	5.87
ResNet-50 + Saliency-DCutMix	1	4	21.55	6.00
ResNet-50 + Saliency-DCutMix	1	5	21.70	5.93
ResNet-50 + Saliency-DCutMix	1	7	21.44	5.93

Furthermore, Saliency-DCutMix achieved higher performance than its baseline two-image mixing augmentation method, CutMix, as demonstrated in Table 3 (left).

a: ImageNet-O

To evaluate the robustness of our proposed model to the out-of-distribution (OOD) data samples, we performed tests on the ImageNet-O dataset [47]. The dataset contains OOD images whose class labels do not belong to 1000 classes of the ImageNet-1K dataset. The most ideal output of a

classification model against the OOD data sample is uniformly predicting all classes with low confidence because a class in the OOD dataset was not considered when training the classification model. These OOD images reliably cause various models to be misclassified with high confidence. To evaluate the robustness of each model against the OOD dataset samples, we measured the area under the precision-recall curve (AUPR) on the ImageNet-O dataset, where a higher AUPR denotes that the model robustly predicted OOD samples with lower confidence. Notably, In Table 4, the model trained without augmentation (Vanilla) exhibited the best AUPR. All the augmentation methods are highly over-confident for the OOD samples, and the results demonstrate the fragility of the augmentation methods when a label distribution shift is present.

**TABLE 4. Performance of various augmentation methods on ImageNet-O. We used ResNet-50 for all the methods. Vanilla denotes ResNet-50 trained without augmentation.**

Model	ImageNet-O AUPR (%)
Vanilla	16.96
Mixup	16.30
DMixup	16.87
CutMix	15.85
DCutMix	16.73
Saliency-DCutMix	16.00

### b: ADVERSARIAL ROBUSTNESS

After the vulnerability of deep neural networks was elucidated by [48], achieving superior classification performance for both non-attacked examples (standard accuracy) and adversarial robustness (robust accuracy) has been considered as the key factor for making the deep neural network truly robust and reliable [49], [50], [51], [52]. To achieve higher adversarial robustness, several competing attack and defense methods have been alternately proposed [49], [50], [51], [52]. In contrast to the above-mentioned studies, [2], [50] have reported that training a model with an input transformation or augmentation enhances the robustness of the model against adversarial examples without adversarial training [50], which suffers from high training cost and a severe trade-off between standard and robust accuracy.

In this section, we demonstrate the additional advantage of K-image mixing augmentation in terms of adversarial robustness. We selected each classification model (ResNet-50) trained by the baseline (w/o augmentation), CutMix, DCutMix, and Saliency-DCutMix using the ImageNet training dataset. To evaluate adversarial robustness against more diverse types of attack, we considered not only white-box attacks as in [2], but also gray- and black-box attacks [51].

Regarding white-box attack where the attacker can freely access to the model's parameter, we use FGSM  $\ell_\infty$  attack [49] with  $\epsilon = 8$ , as in [2]. The black-box attack is more challenging for an attacker because they do not have any information

about the target model to be attacked. For this case, we set a substitute model (*i.e.* ResNet152) and made adversarial examples by attacking it. Gray-box attack is a compromise between white- and black-box attacks: the attacker knows the architecture of the model (*i.e.* ResNet50)) without having access to the weight parameters. Therefore, we generate adversarial examples using a substitute ResNet50 model trained with the ImageNet dataset using a different random seed. For gray- and black-box attacks, we generated adversarial examples of ImageNet validation dataset using a more strong attack method called as PGD  $\ell_\infty$  attack [50] with epsilon  $\epsilon = 8$ .

Table 5 shows top-1 accuracy on given adversarial examples generated by each attack. As proposed in [2], a model trained with CutMix exhibits better adversarial robustness than the baseline case against all types of attacks. DCutMix achieves more improved adversarial robustness against gray- and black-box attacks compared to CutMix. However, this was not the case against the white-box attack. Notably, the white-box attack is the most powerful attack among the attacks. We hypothesized that the label noise problem associated with DCutMix degrades adversarial robustness against a strong attack (white-box), and this hypothesis was indirectly confirmed through experiments on Saliency-DCutMix. We observed that saliency-map-guided DCutmix (Saliency-DCutmix) has stronger adversarial robustness than CutMix and DCutMix for all types of attacks.

For a more sophisticated investigation of robustness against shifts in input data distribution, we evaluated the accuracy of augmentation methods on the ImageNet-A dataset [47]. this dataset contains natural adversarial examples, which cause the wrong classification in an ImageNet-pre-trained model without any adversarial attacks. In Table 5, we found that the performance tendency on ImageNet-A is similar to that of adversarially attacked ImageNet. CutMix exhibited better accuracy than the baseline case and DCutMix. Meanwhile, Saliency-DCutMix improved the CutMix accuracy, exhibiting better generalization on natural adversarial examples and, hence, better robustness on the input data distribution shift.

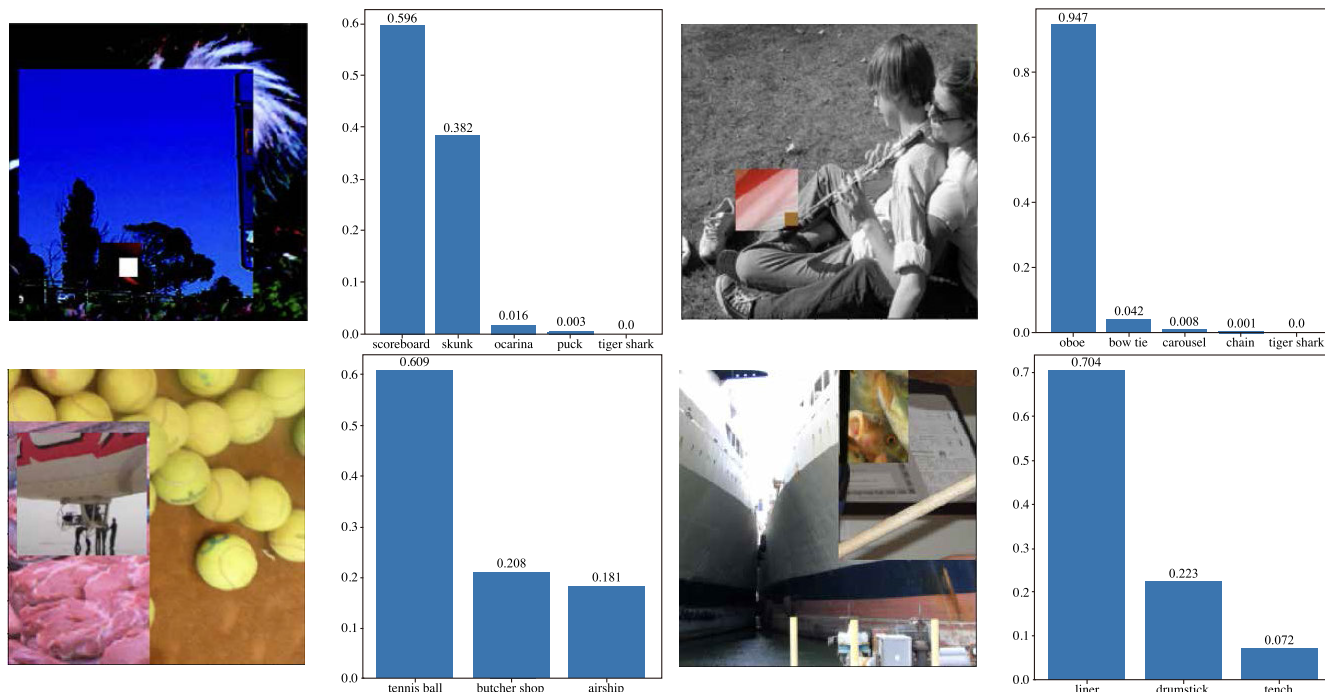
## B. DATA SUBSAMPLING AND APPLICATION

### 1) DATA SUBSAMPLING

We investigate the effect of the proposed subsampling method when trained with DCutMix as an augmentation in Figure 5. In the figure, we compare our data subsampling method with others using different subsampling measures. For all subsequent experiments involving data subsampling, a full 10K CIFAR-100 validation set was used for evaluation, and we reported the averaged results for three independent random seeds using PyramidNet [29].

As shown in Figure 5, sampling the easy-only or hard-only examples based on  $m(L_i)$  shows deteriorated performance compared to the random subsampling. The hard-only subsampling severely suffered from poor generalization. This





**FIGURE 4.** Qualitative example of an augmented image and its soft label distribution for top K classes (top: DCutMix, bottom: Saliency-DCutMix). Note that for the DCutMix case, visually salient information of top K classes is not fully contained (e.g., “Scoreboard” or “bow-tie”) in the augmented image. Meanwhile, visually salient information of all the top K classes is fully contained for the Saliency-DCutMix case.

**TABLE 5.** Top-1 robust accuracy on ImageNet against various adversarial attacks (2nd, 3rd, 4th column) and accuracy on ImageNet-A (last column). The values in parentheses denote accuracy increment compared to the Baseline case.

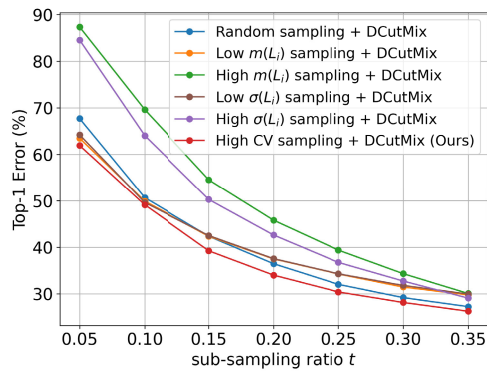
Model	White-box	Gray-box	Black-box	ImageNet-A
	FGSM ( $L_\infty, \epsilon = 8$ )	PGD ( $L_\infty, \epsilon = 8$ )	PGD ( $L_\infty, \epsilon = 8$ )	
ResNet-50 (Baseline)	14.34	40.74	46.01	3.38
+ CutMix	40.89 (+26.55)	43.38 (+2.64)	51.06 (+5.05)	7.46 (+4.08)
+ DCutMix ( $K = 3, \alpha = \frac{1}{3}$ )	32.90 (+18.56)	43.63 (+2.89)	51.73 (+5.72)	5.84 (+2.46)
+ Saliency-DCutMix ( $K = 3, \alpha = \frac{1}{3}$ )	<b>41.37 (+27.03)</b>	<b>44.27 (+3.53)</b>	<b>51.91 (+5.90)</b>	<b>7.69 (+4.31)</b>

**TABLE 6.** Comparison of the state-of-the-art NAS methods on ImageNet under comparably small resource constraints. (.) denotes the proxy dataset where the architecture was searched on.

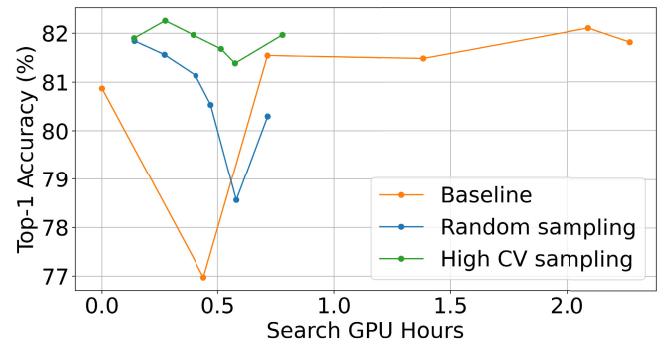
Architecture	Top-1 Err (%)	Top-5 Err (%)	# Params (M)	# FLOPs (M)	Search Cost (GPU days)	Search method
AmoebaNet-C [53]	24.3	7.6	6.4	570	3150	evolution
MnasNet-92 [54]	25.2	8.0	4.4	388	-	RL
ProxylessNAS [26]	24.9	7.5	7.1	465	8.3	gradient-based
SNAS [23]	27.3	9.2	4.3	522	1.5	gradient-based
BayesNAS [55]	26.5	8.9	3.9	-	0.2	gradient-based
PC-DARTS (CIFAR10) [21]	25.1	7.8	5.3	586	0.1	gradient-based
PC-DARTS (ImageNet) [21]	24.2	7.3	5.3	597	3.8	gradient-based
PC-DARTS (CIFAR100) [21]	23.8	7.09	6.3	730	0.1	gradient-based
PC-DARTS (High CV Sub-sampled CIFAR100)	24.3	7.2	5.8	671	<b>0.01</b>	gradient-based

result indicates that subsampling only hard samples where the salient regions were occluded by image-mixing augmentation being applied (see Figure 2) is not desirable

under the constraint of a small number of training samples. In a similar manner, subsampling only easy samples extract the biased data samples that cannot be helpful for better



(a) Top-1 error plotting for data subsampling methods.  $t$  denotes the ratio of images in the subsampled set to those in the whole training set.



(b) Top-1 accuracy plotting of NAS given GPU searching cost on CIFAR-100.

**FIGURE 5.** Experiments on data subsampling: performance of various data subsampling methods (a). Graph (b) shows the speedup of NAS from the proposed subsampling method: searching on the entire CIFAR-100 dataset (Baseline), on randomly subsampled CIFAR-100, and on subsampled CIFAR-100 by high CV measure.

generalization. Moreover, simply employing standard deviation  $\sigma(L_i)$  as subsampling measure induced a similar test error plot as that from the above mean-based sampling methods. On the other hand, our high-CV-based subsampling significantly outperformed the random sampling. Specifically, the test error was 5.79% lower when the number of subsampled training samples was extremely small (i.e.,  $t = 0.05$ ). High CV subsampling enables us to acquire various levels of data samples, from easy to hard. High CV subsampling basically selects the easy data samples, which can frequently become hard samples depending on the image-mixing augmentation. Therefore, High-CV subsampling leads to better performance when training with image-mixing augmentation. We demonstrated the superiority of our subsampling method over other subsampling methods employing uncertainty derived by weight dropout [56] and *K-Center Coreset* sampling [57].

## 2) APPLICATION ON NAS

We further demonstrate the practicality and effectiveness of our proposed data subsampling method on another domain, namely, NAS. Our goal is to reduce the time spent searching the architectures by searching on the subsampled dataset drawn from our framework rather than on the full training dataset. We demonstrated that the architecture search time is greatly reduced without accuracy degeneration. Notably, the data subset subsampled by our algorithm can be applied to any neural architecture search framework, including gradient-based and non-gradient-based search methods. We adopt one of the most computationally efficient and stabilized NAS methods, PC-DARTS [21], as our baseline.

For the searching process, we divided the subsampled (or entire) training dataset into two equal parts, with one for optimizing the network parameters and the other one for optimizing the architecture hyperparameters (i.e.,  $\alpha, \beta$  in [21]). Additionally, we adopted the warm-up strategy

during the search process, where only network parameters are optimized. We freeze the hyperparameters  $\alpha, \beta$  for the first 15 epochs as in [21]. We applied the warm-up strategy for the first five epochs for the baseline method (i.e., searching on the entire dataset) where the number of total searching epochs was 10 (i.e., the left-most point for the Baseline in Figure 5b of the manuscript). We used Tesla V100 GPU to perform the search.

For the evaluation involving training the searched network from scratch, we used the equivalent training hyperparameters as in [21]. The performance of the neural networks searched on the entire CIFAR-100 dataset (baseline) is plotted in Figure 5b. We plotted the performance of neural networks searched on the entire CIFAR-100 dataset (baseline) by adjusting the searching epochs while adjusting the subsampling ratio  $t$  for searching on the randomly subsampled dataset and our subsampled dataset drawn by high CV measure. The searching epochs were adjusted while adjusting the subsampling ratio  $t$  for searching on the randomly subsampled dataset and our subsampled dataset drawn using a high CV measure. The results demonstrate the outstanding efficiency of our subsampling framework in terms of search time and accuracy. Specifically, it (searching on the high-CV subsampled dataset) achieved comparable accuracy with a 7.7-fold reduction in search time compared to other baselines. Furthermore, it consistently outperformed random subsampling given an equivalent number of data samples for searching.

As listed in Table 6, we observed that our framework serves as an effective proxy dataset, and the neural network searched using it is well-generalized on ImageNet. Notably, it reduced the GPU search time to as much as 0.01 d (i.e., 16 min) while achieving comparable or even higher accuracy than that of the models searched with PC-DARTS on the entire CIFAR-10, CIFAR-100, and randomly subsampled ImageNet datasets. Moreover, compared to the other NAS methods,

ours achieved the best accuracy and significantly lower search computational cost.

## V. CONCLUDING REMARKS AND LIMITATION

In this study, we present the advantages of expanding the number of images for image-mixing augmentation based on various experimental results and analyses. First, we propose the generalized method for K-image mixing augmentation motivated by SBP. Second, we demonstrate that the proposed K-image mixing augmentation improves classification performance. Moreover, from a novel perspective, we demonstrated that the key factor behind this improvement is the convergence to wide local minima. Moreover, we empirically found that increasing the number of images for the image mixing augmentation enhances the adversarial robustness of a classification model against various types of adversarial examples. Additionally, we derived a new subsampling method that utilizes the proposed K-image mixing augmentation in a novel way. We experimentally demonstrate that the proposed subsampling method can effectively reduce search time without performance degradation. We believe our observations can inspire new research directions for image mixing augmentation and data subsampling.

## A. LIMITATION

Because our method focuses on setting a probabilistic framework explaining the CutMix augmentation and its potential effectiveness, we did not employ other semantic knowledge, such as spatial attention or saliency map. However, if strictly targeting the SOTA classification performance, it would be a promising future direction to employ the additional information in the augmentation process. Furthermore, employing the idea in other computer vision tasks, such as object detection and segmentation, will enhance the applicability of the method.

## ACKNOWLEDGMENT

(Joonhyun Jeong and Sungmin Cha contributed equally to this work.)

## REFERENCES

- [1] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," 2017, *arXiv:1708.04552*.
- [2] S. Yun, D. Han, S. Chun, S. J. Oh, Y. Yoo, and J. Choe, "CutMix: Regularization strategy to train strong classifiers with localizable features," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6023–6032.
- [3] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "Mixup: Beyond empirical risk minimization," 2017, *arXiv:1710.09412*.
- [4] S. Chun, S. J. Oh, S. Yun, D. Han, J. Choe, and Y. Yoo, "An empirical evaluation on robustness and uncertainty of regularization methods," 2020, *arXiv:2003.03879*.
- [5] J. Kim, W. Choo, H. Jeong, and H. O. Song, "Co-Mixup: Saliency guided joint mixup with supermodular diversity," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–21. [Online]. Available: <https://openreview.net/forum?id=gvxJzw8kW4b>
- [6] J.-H. Kim, W. Choo, and H. O. Song, "Puzzle mix: Exploiting saliency and local statistics for optimal mixup," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 5275–5285.
- [7] A. F. M. S. Uddin, M. S. Monira, W. Shin, T. Chung, and S.-H. Bae, "SaliencyMix: A saliency guided data augmentation strategy for better regularization," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–12. [Online]. Available: <https://openreview.net/forum?id=M0QkvBGTTq>
- [8] A. Katharopoulos and F. Fleuret, "Not all samples are created equal: Deep learning with importance sampling," 2018, *arXiv:1803.00942*.
- [9] M. Ren, W. Zeng, B. Yang, and R. Urtasun, "Learning to reweight examples for robust deep learning," 2018, *arXiv:1803.09050*.
- [10] A. Katharopoulos and F. Fleuret, "Biased importance sampling for deep neural network training," 2017, *arXiv:1706.00043*.
- [11] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *J. Comput. Syst. Sci.*, vol. 55, no. 1, pp. 119–139, Aug. 1997.
- [12] T. Malisiewicz, A. Gupta, and A. A. Efros, "Ensemble of exemplar-SVMs for object detection and beyond," in *Proc. Int. Conf. Comput. Vis.*, Nov. 2011, pp. 89–96.
- [13] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2980–2988.
- [14] M. P. Kumar, B. Packer, and D. Koller, "Self-paced learning for latent variable models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2010, pp. 1189–1197.
- [15] B. Baker, O. Gupta, N. Naik, and R. Raskar, "Designing neural network architectures using reinforcement learning," 2016, *arXiv:1611.02167*.
- [16] Z. Zhong, J. Yan, W. Wu, J. Shao, and C.-L. Liu, "Practical block-wise neural network architecture generation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2423–2432.
- [17] Z. Zhong, Z. Yang, B. Deng, J. Yan, W. Wu, J. Shao, and C.-L. Liu, "Block-QNN: Efficient block-wise neural network architecture generation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 7, pp. 2314–2328, Jul. 2021.
- [18] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," 2016, *arXiv:1611.01578*.
- [19] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8697–8710.
- [20] H. Liu, K. Simonyan, and Y. Yang, "DARTS: Differentiable architecture search," 2018, *arXiv:1806.09055*.
- [21] Y. Xu, L. Xie, X. Zhang, X. Chen, G.-J. Qi, Q. Tian, and H. Xiong, "PC-DARTS: Partial channel connections for memory-efficient architecture search," 2019, *arXiv:1907.05737*.
- [22] B. Wu, K. Keutzer, X. Dai, P. Zhang, Y. Wang, F. Sun, Y. Wu, Y. Tian, P. Vajda, and Y. Jia, "FBNet: Hardware-aware efficient ConvNet design via differentiable neural architecture search," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10734–10742.
- [23] S. Xie, H. Zheng, C. Liu, and L. Lin, "SNAS: Stochastic neural architecture search," 2018, *arXiv:1812.09926*.
- [24] Z. Guo, X. Zhang, H. Mu, W. Heng, Z. Liu, Y. Wei, and J. Sun, "Single path one-shot neural architecture search with uniform sampling," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2020, pp. 544–560.
- [25] X. Zhang, P. Hou, X. Zhang, and J. Sun, "Neural architecture search with random labels," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 10907–10916.
- [26] H. Cai, L. Zhu, and S. Han, "ProxylessNAS: Direct neural architecture search on target task and hardware," 2018, *arXiv:1812.00332*.
- [27] J. Sethuraman, "A constructive definition of Dirichlet priors," *Statistica sinica*, vol. 4, no. 2, pp. 639–650, 1994.
- [28] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [29] D. Han, J. Kim, and J. Kim, "Deep pyramidal residual networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5927–5935.
- [30] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Weinberger, "Deep networks with stochastic depth," in *Proc. ECCV*, 2016, pp. 646–661.
- [31] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2818–2826.
- [32] G. Ghiasi, T.-Y. Lin, and Q. V. Le, "DropBlock: A regularization method for convolutional networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 10750–10760.

- [33] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, D. Lopez-Paz, and Y. Bengio, "Manifold mixup: Better representations by interpolating hidden states," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6438–6447.
- [34] M. Hong, J. Choi, and G. Kim, "StyleMix: Separating content and style for enhanced data augmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14862–14870.
- [35] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2016, pp. 630–645.
- [36] S. Zagoruyko and N. Komodakis, "Wide residual networks," 2016, *arXiv:1605.07146*.
- [37] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1492–1500.
- [38] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Tech. Rep., 2009.
- [39] J. Kim, W. Choo, H. Jeong, and H. O. Song, "Co-mixup: Saliency guided joint mixup with supermodular diversity," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–21. [Online]. Available: <https://openreview.net/forum?id=gvxJzw8kW4b>
- [40] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, "On large-batch training for deep learning: Generalization gap and sharp minima," 2016, *arXiv:1609.04836*.
- [41] G. Pereyra, G. Tucker, J. Chorowski, A. Kaiser, and G. Hinton, "Regularizing neural networks by penalizing confident output distributions," 2017, *arXiv:1701.06548*.
- [42] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4320–4328.
- [43] P. Chaudhari, A. Choromanska, S. Soatto, Y. LeCun, C. Baldassi, C. Borgs, J. Chayes, L. Sagun, and R. Zecchina, "Entropy-SGD: Biasing gradient descent into wide valleys," *J. Stat. Mech., Theory Exp.*, vol. 2019, no. 12, Dec. 2019, Art. no. 124018.
- [44] S. Cha, H. Hsu, T. Hwang, F. P. Calmon, and T. Moon, "CPR: Classifier-projection regularization for continual learning," 2020, *arXiv:2006.07326*.
- [45] Z. Yao, A. Gholami, K. Keutzer, and M. Mahoney, "PyHessian: Neural networks through the lens of the Hessian," 2019, *arXiv:1912.07145*.
- [46] S. Yun, S. J. Oh, B. Heo, D. Han, J. Choe, and S. Chun, "Re-labeling ImageNet: From single to multi-labels, from global to localized labels," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2340–2350.
- [47] D. Hendrycks, K. Zhao, S. Basart, J. Steinhardt, and D. Song, "Natural adversarial examples," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 15262–15271.
- [48] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," 2013, *arXiv:1312.6199*.
- [49] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," 2014, *arXiv:1412.6572*.
- [50] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," 2017, *arXiv:1706.06083*.
- [51] C. Guo, M. Rana, M. Cisse, and L. Van Der Maaten, "Countering adversarial images using input transformations," 2017, *arXiv:1711.00117*.
- [52] Y. Dong, Q.-A. Fu, X. Yang, T. Pang, H. Su, Z. Xiao, and J. Zhu, "Benchmarking adversarial robustness on image classification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 321–331.
- [53] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *Proc. AAAI*, 2019, pp. 4780–4789.
- [54] M. Tan, B. Chen, R. Pang, V. Vasudevan, M. Sandler, A. Howard, and Q. V. Le, "MnasNet: Platform-aware neural architecture search for mobile," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2820–2828.
- [55] H. Zhou, M. Yang, J. Wang, and W. Pan, "BayesNAS: A Bayesian approach for neural architecture search," in *Proc. ICML*, 2019, pp. 7603–7613.
- [56] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian active learning with image data," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1183–1192.
- [57] O. Sener and S. Savarese, "Active learning for convolutional neural networks: A core-set approach," 2017, *arXiv:1708.00489*.



**JOONHYUN JEONG** received the bachelor's degree from the Department of Computer Science and Engineering, Kyung Hee University, South Korea, in 2019. He is currently working with NAVER CLOVA, South Korea. His research interests include model compression and data augmentation in computer vision tasks.



**SUNGMIN CHA** received the B.S. degree in computer engineering from Pukyong National University, Busan, South Korea, in 2016, and the M.S. degree in information and communication engineering from the Daegu Gyeongbuk Institute of Science and Technology (DGIST), Daegu, South Korea, in 2018. He is currently pursuing the Ph.D. degree in electrical and computer engineering with Seoul National University (SNU), Seoul, South Korea. His current research interests include deep neural network-based unsupervised image denoising, continual learning, and adversarial robustness for the deep neural networks.



**JONGWON CHOI** (Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea, in 2012 and 2014, respectively, and the Ph.D. degree in electrical engineering from Seoul National University, Seoul, South Korea, in 2018. From 2018 to 2020, he was with the Research Intelligence Research Center, Samsung SDS, Seoul, as a Research Engineer. In 2020, he joined the Department of Advanced Imaging, Chung-Ang University, Seoul, where he is currently working as an Assistant Professor. His research interests include surveillance system with deep learning, architecture of deep learning, and low-level computer vision algorithms.



**SANGDOON YUN** received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from Seoul National University, Seoul, South Korea, in 2010, 2013, and 2017, respectively. He is currently a Research Scientist with the AI Laboratory, NAVER. His current research interests include computer vision, deep learning, and image classification.





**TAESUP MOON** (Member, IEEE) received the B.S. degree in electrical engineering from Seoul National University (SNU), Seoul, South Korea, in 2002, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 2004 and 2008, respectively. From 2008 to 2012, he was a Scientist with Yahoo! Labs, Sunnyvale, CA, USA. He was a Postdoctoral Researcher with the Department of Statistics, UC Berkeley, from 2012 to 2013.

From 2013 to 2015, he was a Research Staff Member with the Samsung Advanced Institute of Technology (SAIT). From 2015 to 2017, he was an Assistant Professor with the Department of Information and Communication Engineering, Daegu Gyeongbuk Institute of Science and Technology (DGIST). From 2017 to 2021, he was an Associate Professor with the Department of Electrical and Computer Engineering, Sungkyunkwan University (SKKU), Suwon, South Korea. He is currently an Associate Professor with the Department of Electrical and Computer Engineering, SNU. His current research interests include machine/deep learning, signal processing, information theory, and various (big) data science applications.



**YOUNGJOON YOO** (Member, IEEE) received the B.S. degree in electrical and computer engineering and the Ph. D. degree in electrical engineering from Seoul National University, Seoul, South Korea, in 2011 and 2017, respectively. He is currently a Research Scientist with NAVER AI Research, and also leads the Image Vision Team, NAVER CLOVA. His research interests include deep learning for computer vision and probabilistic theory.

...