

SuperMix: Supervising the Mixing Data Augmentation

Ali Dabouei, Sobhan Soleymani, Fariborz Taherkhani, Nasser M. Nasrabadi
 West Virginia University
 {ad0046, ssoleyma, ft0009}@mix.wvu.edu, nasser.nasrabadi@mail.wvu.edu

Abstract

This paper presents a supervised mixing augmentation method termed SuperMix, which exploits the salient regions within input images to construct mixed training samples. SuperMix is designed to obtain mixed images rich in visual features and complying with realistic image priors. To enhance the efficiency of the algorithm, we develop a variant of the Newton iterative method, 65× faster than gradient descent on this problem. We validate the effectiveness of SuperMix through extensive evaluations and ablation studies on two tasks of object classification and knowledge distillation. On the classification task, SuperMix provides comparable performance to the advanced augmentation methods, such as AutoAugment and RandAugment. In particular, combining SuperMix with RandAugment achieves 78.2% top-1 accuracy on ImageNet with ResNet50. On the distillation task, solely classifying images mixed using the teacher’s knowledge achieves comparable performance to the state-of-the-art distillation methods. Furthermore, on average, incorporating mixed images into the distillation objective improves the performance by 3.4% and 3.1% on CIFAR-100 and ImageNet, respectively. The code is available at <https://github.com/alldbi/SuperMix>.

1. Introduction

Despite the revolutionary performance of deep neural networks (DNNs), they easily overfit when the training set is qualitatively or quantitatively deficient [27, 33]. Quality of the data can be interpreted as how well the data is expressive of the true distribution of inputs in the underlying task. This helps the model to learn discriminative patterns likely to occur at inference time. Quantity of the data, on the other hand, allows the model to observe discriminative patterns from different views and generalize the task-specific notions according to the major factors of variation in the input domain. Although analytical analysis of such important properties of the data has remained arduous [16], empirical evaluations on training deep models often highlight a common observation: incorporating more data leads to a better

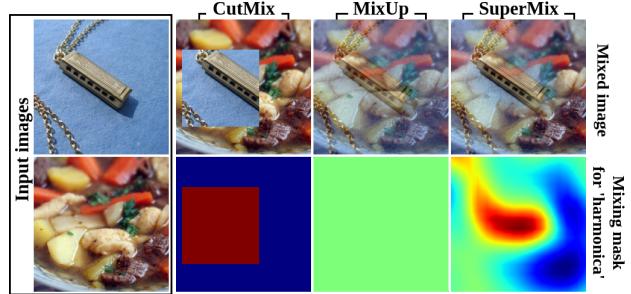


Figure 1: SuperMix combines salient regions in input images to construct unseen data for training.

generalization [25, 13]. Hence, data augmentation has become a fundamental component of the training paradigms, aiming to enlarge the training set by transforming images in the given dataset.

Conventional image data augmentation involves combinations of context-preserving transformations, such as horizontal flip, crop, scale, color manipulation, and cut out [17, 12, 9]. Recently, notable efforts have been devoted to improving the augmentation, *e.g.*, by automating the search for the optimal augmentation policies [4, 20, 5]. The majority of these methods have focused on transforming single images, while ignoring the potentially very useful combination of multiple images for augmentation. To address this shortcoming, several studies have considered combining multiple images to construct novel images [18, 22, 34, 32, 29]. However, these methods either mix images blindly and disregard the salient regions [34, 11, 32, 29] or do not scale to large-scale problems [18]. Furthermore, the current mixing functions are not expressive enough and often suppresses visual patterns by averaging or covering features in one image with the trivial features in another image. The corresponding pseudo labels are also not accurate and constrain the training performance [11].

This paper presents a mixing augmentation approach termed SuperMix, which exploits the salient regions of input images to construct more advantageous mixed data. The supervision for this purpose can be obtained from the target model itself, *i.e.*, self-training [26, 30, 23, 19, 2, 31], or

a more sophisticated model aiming to guide a student network via knowledge transfer [1, 14]. Figure 1 provides a visual comparison of mixed images produced by different methods. In a nutshell, the contributions of the paper are as follows:

- We formalize the problem of supervised mixing augmentation using a set of mixing masks associating the pixel value at each spatial location in the mixed image to the spatial locations in the input images.
- The optimization problem is carefully constrained to assure that the solutions are rich in salient features and comply with the realistic image priors.
- We develop a modified Newton iterative algorithm for SuperMix, suitable for large-scale applications. This approach provides $65\times$ speed-up as compared to SGD on ImageNet.
- We demonstrate that mixed images intrinsically induce smooth predictions, and thus, help reveal knowledge of the teacher model in knowledge distillation.

2. Related work

Data augmentation: Data augmentation aims to improve the generalization of the model by enlarging the train set using transformations preserving the context of inputs in the learning problem. Conventional image transformations for this purpose are horizontal flip, crop, scale, color manipulation, and cut out [17, 12, 9]. A contemporary trend of research on the topic has focused on selecting the best sequence of transformations according to the task, dataset, and learning model. AutoAugment (AA) [4] automated the search for augmentation policies given a predefined set of transformations. Despite the significant performance of AA, it suffers from prohibitive training complexity imposed by Reinforcement Learning. Multiple approaches have attempted to reduce the training complexity by employing more efficient search methods, *e.g.*, density matching in fast AutoAugment (FAA) [20], or population based augmentation (PBA) [15]. RandAugment (RA) [5] have shown that the search space and selection criteria can be significantly simplified by carefully combining random transformations. However, these methods ignore the potentially useful combination of multiple images for augmentation.

Mixing augmentation: Several recent studies have considered employing multiple images for data augmentation [18, 34, 11, 32, 29]. Smart Augmentation [18] proposed merging multiple images from the same class using a DNN trained concurrently with the target model. However, training an additional deep model alongside every target model is resource exhaustive and severely limits the scalability of the approach for large-scale problems. Moreover, the

method is restricted to merge images from the same class which limits the diversity and novelty of visual patterns in the merged images. MixUp [34, 29] combined a pair of images for the augmentation by convex linear interpolation. CutMix [32] proposed overlaying a cropped area of an input image on another image to augment the data. Although MixUp and CutMix have demonstrated notable improvements to the training of object recognition models, they suffer from major shortcomings. First, they often average or replace salient regions in one image with insignificant regions, *e.g.*, background, in another image. Second, due to the lack of supervision the labels computed for the mixed images are not accurate and limits the usefulness of the mixed images. However, SuperMix addresses these issues by extracting the salient regions of inputs and carefully combining them according to the realistic image priors and saliency-preserving constraints.

3. Supervised Mixing Augmentation

Given a training set $\mathcal{D} = \{(x_i, y_i)\}_{i=0}^{N-1}$, mixing methods take a subset $X \subset \mathcal{D}$ to produce the mixed image \hat{x} and the corresponding label \hat{y} . A crucial property of mixed images is that they must reside close to the manifold of the training data since the goal of the mixing is to enlarge the support of the training distribution. Previous mixing methods [34, 29] have considered this requirement by employing operations that preserve local smoothness of images. MixUp [34, 29] combines a pair of images (x_i, x_j) using convex linear interpolation as: $\hat{x} = rx_i + (1 - r)x_j$, where $r \sim \text{Beta}(\alpha, \alpha)$ is a random mixing weight from the symmetric Beta distribution with $\alpha \in (0, \infty)$. Due to the lack of supervision, the soft label for \hat{x} is computed using the same linear interpolation as: $\hat{y} = r\delta(y_i) + (1 - r)\delta(y_j)$, where $\delta(\cdot)$ is the one-hot encoding function. This blind mixing suffers from two shortcomings. First, coefficient r assigns an equal importance to the whole image which can suppress important features by averaging with the background or less important features from the other image. Second, the computed soft label, \hat{y} , does not accurately describe the probability of classes represented by the mixed image and, thus, limits the effectiveness of the augmentation.

3.1. Mixing function

We formalize a general formulation for the augmentation function that allows multiple images to be combined locally. We use a set of mixing masks $M = \{m_i\}_{i=0}^{k-1}$, where $m_i : \Lambda \rightarrow [0, 1]$ associates each spatial location $u \in \Lambda$ in x_i with a scalar value $m_i(u)$. Using the mixing masks, we define the mixing function as:

$$\hat{x} := \sum_{i=0}^{k-1} x_i \odot m_i, \quad (1)$$

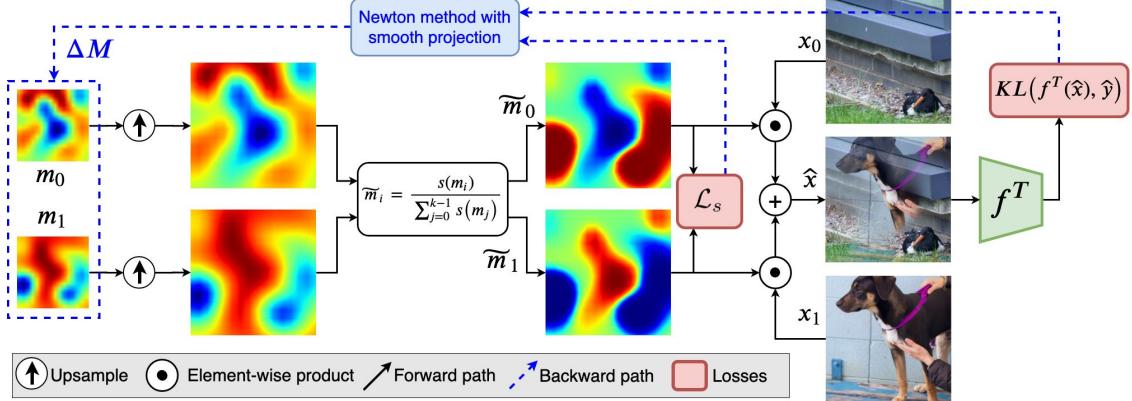


Figure 2: Schematic diagram of the proposed method for mixing $k = 2$ input images using the supervision from f^T .

where x_i is the i^{th} sample in X , the operator \odot denotes the element-wise product, and $\sum_i m_i(u) = 1$ to hold the convexity of the combination. The mixing function recovers MixUp [34] when $k = 2$ and all values in each mask are equal. It also recovers CutMix [32] when $k = 2$ and all values except the cropped area in one of the masks are equal to one. Figure 1 provides a visual comparison of the role of the masks in the mixing augmentation. In the next section, we describe how knowledge of a teacher model can be used to compute M such that the mixed image, \hat{x} , encompasses the rich visual information of images in X .

3.2. Supervised mixing

Let $f^T : \mathbb{R}^{W \times H \times C} \rightarrow [0, 1]^n$ denote the probability vector predicted by the teacher (T) for n classes and f_i^T be the probability for the i^{th} class. We optimize the set of masks M in Equation 1 such that all salient regions in X , according to the knowledge of the teacher, be present in the mixed image, \hat{x} . This can be interpreted as: $f^T(\hat{x}) \approx \hat{y}$, where \hat{y} is high for classes associated with images in X . We formulate the target soft label, \hat{y} , computed in previous approaches [34, 32] for $k = 2$ using the Beta distribution. We generalize for $k \geq 2$ by sampling the mixing coefficients from the Dirichlet distribution. Let $(r_0, \dots, r_{k-1}) \sim \text{Dir}(\alpha)$ be a random sample from the symmetric multivariate Dirichlet distribution with parameter α and size k , we define the target soft label as:

$$\hat{y} := \sum_{i=0}^{k-1} r_i \delta(y^T(x_i)), \quad (2)$$

where $y^T(x_i) = \arg \max_j f_j^T(x_i)$ is the predicted class for $x_i \in X$, and $\delta(\cdot)$ is the one-hot encoding function.

The set of mixing masks can be optimized to minimize the divergence between the output of the teacher model on the mixed image and the target soft label computed in Equation 2. The masks must also hold two additional properties

to comply with the realistic image priors. First, generated images must reside close the manifold of the training data. In practice, this interprets that each mask must be spatially smooth so that the generated images resemble the spatial structure of the inputs. Second, masks must be sparse across the input samples to ensure each spatial location in the output image is assigned merely to a single image which prevents averaging multiple images at each spatial location and suppressing important features. Considering these, the optimization problem for finding the mixing masks can be written as:

$$\begin{aligned} \arg \min_{m_0, \dots, m_{k-1}} & KL(f^T(\hat{x}) || \hat{y}) + \lambda_\sigma \mathcal{L}_\sigma(M) + \lambda_s \mathcal{L}_s(M) \quad \text{s.t.:} \\ & a. 0 \leq m_i(u) \leq 1, \quad b. \sum_i m_i(u) = 1, \end{aligned} \quad (3)$$

where \mathcal{L}_σ is a penalty term for the roughness of masks, e.g., total variation (TV) norm, \mathcal{L}_s is a loss function to encourage sparsity of masks across input samples, and $KL(\cdot || \cdot)$ is the Kullback-Leibler divergence.

Here, we provide an iterative algorithm to solve the optimization problem efficiently. At each iteration t , the convexity conditions can be satisfied by the following normalization:

$$\tilde{m}_i^t = \frac{s(m_i^t)}{\sum_{j=0}^{k-1} s(m_j^t)}, \quad (4)$$

where $s(\cdot)$ is the sigmoid function. Hence, the generalized mixing function in Equation 1 takes the normalized masks to construct \hat{x} . Using the normalized masks, we define the sparsity promoting loss as:

$$\mathcal{L}_s := \frac{1}{kW^H} \sum_{u,i} \tilde{m}_i^t(u) (\tilde{m}_i^t(u) - 1). \quad (5)$$

This loss function encourages the mask values to approach 0 or 1. Since the values of masks at each spatial location sum to 1, due to the normalization in Equation 4, only one of the masks takes the high value to minimize the loss.

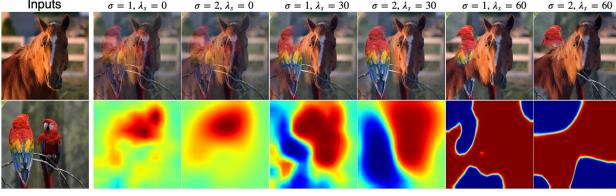


Figure 3: Visualizing the effect of smoothing factor, σ , and the sparsity promoting weight, λ_s , on the mixed images. Masks are estimated using ResNet34 and are associated with the ‘horse’ class.

3.3. Optimization Method

A proper set of mixing masks can be estimated by minimizing the objective of SuperMix as $\mathcal{L}_{SM} = KL + \lambda_\sigma \mathcal{L}_\sigma + \lambda_s \mathcal{L}_s$. A reduced form of this problem has been studied in saliency detection and explanation of DNN predictions by employing SGD [10] or deep generators [6]. However, the current problem is more complex since multiple images are involved in the optimization and the roughness penalty and sparsity promoting loss should be minimized on all the corresponding masks. As we discussed and evaluated in Section 4.4, SGD is very slow and not feasible for solving the problem in case of large-scale image recognition tasks. Furthermore, employing a dedicated deep model to mix data by extending [6] makes the algorithm model-dependent and is not computationally efficient.

We develop a fast and efficient algorithm to optimize the mixing masks based on Newton’s iterative method for finding roots of a nonlinear system of equations in the underdetermined case [21, 24]. Specifically, instead of optimizing \mathcal{L}_{SM} , we optimize $\mathcal{L}'_{SM} = KL + \lambda_s \mathcal{L}_s$ using a smooth projection (SP) [7] that directly satisfies the smoothness of masks. As we analyze later in Section 4.4, this significantly improves the execution time of the mixing. Considering the first-order approximation of \mathcal{L}'_{SM} at M , each mask is updated at iteration t to find the roots as: $m_i^{t+1} \leftarrow m_i^t + \Delta m_i^t$. Here, the update is computed using the Newton’s method as:

$$\Delta M^t = \frac{-|\mathcal{L}'_{SM}|}{\|\nabla \mathcal{L}'_{SM}\|_2^2} \nabla \mathcal{L}'_{SM}, \quad (6)$$

where the gradient is with respect to M^t , the concatenation of $\{m_0^t, \dots, m_{k-1}^t\}$. Since both the divergence and \mathcal{L}_s are nonnegative, $|\mathcal{L}'_{SM}| = \mathcal{L}'_{SM}$. This formulation uses the ℓ_2 -norm projection to compute ΔM^t . We modify it using SP to preserve the smoothness of masks and compute the smooth update as:

$$\widetilde{\Delta M}^t = \frac{-\mathcal{L}'_{SM}}{(g_\sigma * \nabla \mathcal{L}'_{SM})^T \nabla \mathcal{L}'_{SM}} (g_\sigma * \nabla \mathcal{L}'_{SM}), \quad (7)$$

where $g_\sigma * \nabla \mathcal{L}'_{SM}$ is a smoothed version of the gradients using the 2D Gaussian smoothing filter g with the standard

Algorithm 1 SuperMix

```

1: inputs: Classifier  $f^T$ , set of  $k$  images  $X$ ,  

   low-pass filter  $g_\sigma$ .  

2: output: Mixed sample  $\hat{x}$ .  

3:  $Y = \{\text{argmax}_j f_j^T(x_i) : x_i \in X\}$ .  

4: Sample  $(r_0, \dots, r_{k-1})$  from  $\text{Dir}(\alpha)$ .  

5:  $\hat{y} = \sum_{i=0}^{k-1} r_i \delta(y^T(x_i))$ .  

6: Initialize  $(m_0, \dots, m_{k-1}) \leftarrow 0$ ,  

    $\hat{x}^0 \leftarrow \frac{1}{k} \sum_{x_i \in X} x_i$ ,  $t \leftarrow 0$ .  

7: condition = Top- $k$  predicted classes by  $f(\hat{x}^t)$  are not in  $Y$ .  

8: while condition do  

9:    $\mathcal{L}'_{SM} = KL(f^T(\hat{x}^t) || \hat{y}) + \lambda_s \mathcal{L}_s$ .  

10:   $\widetilde{\Delta M}^t = \frac{-\mathcal{L}'_{SM}}{(g_\sigma * \nabla \mathcal{L}'_{SM})^T \nabla \mathcal{L}'_{SM}} g_\sigma * \nabla \mathcal{L}'_{SM}$ .  

11:   $m_i^{t+1} \leftarrow m_i^t + \Delta m_i$  for  $i \in \{0, \dots, k-1\}$ .  

12:   $\tilde{m}_i^{t+1} = s(m_i^{t+1}) / \sum_{j=0}^{k-1} s(m_j^{t+1})$ .  

13:   $\hat{x}^{t+1} \leftarrow \sum_{i=0}^{k-1} x_i \odot \tilde{m}_i^{t+1}$ .  

14:   $t \leftarrow t + 1$   

15: end while  

16: return  $\hat{x}^t$ .

```

deviation σ . It must be noted that all matrices in Equations 6 and 7 are vectorized before the matrix operations, and are reshaped back at the end of the iteration. In addition, due to the smoothness of masks, we optimize a down-sampled set of masks and up-sample them before performing the mixing. Algorithm 1 and Figure 2 demonstrate the detailed algorithm and schematic diagram for SuperMix, respectively.

Termination Criteria: The algorithm terminates when the Top- k predicted classes of $f^T(\hat{x})$ are the same as the predicted class for samples in X . For instance, when X consists of two images recognized as ‘cat’ and ‘dog’, the Top-2 classes in $f^T(\hat{x})$ should be classes of ‘cat’ and ‘dog’. This criterion assures that important features in the input set are visible in the mixed image. Figure 4 provides a visual comparison of the mixed images produced by different methods.

4. Experiments

We evaluate the performance of SuperMix on two tasks of object classification and knowledge distillation [1, 14] using two benchmark datasets of CIFAR-100 [17] and ImageNet [8]. For knowledge distillation, we evaluate SuperMix on two major previous SOTA methods [14, 28] and two mixing augmentation techniques including MixUp and CutMix. For the sake of fair comparison, *pseudo labels for these blind mixing methods are computed using the same teacher employed in SuperMix*. All training experiments use random horizontal flip and random crop as the default augmentations. We perform the algorithm on random sets of input samples drawn from \mathcal{D} to generate \mathcal{D}' . For the sake of brevity, we define the augmentation factor $\kappa = \frac{|\mathcal{D}'|}{|\mathcal{D}|}$ to show the ratio of the size of the mixed dataset over the size

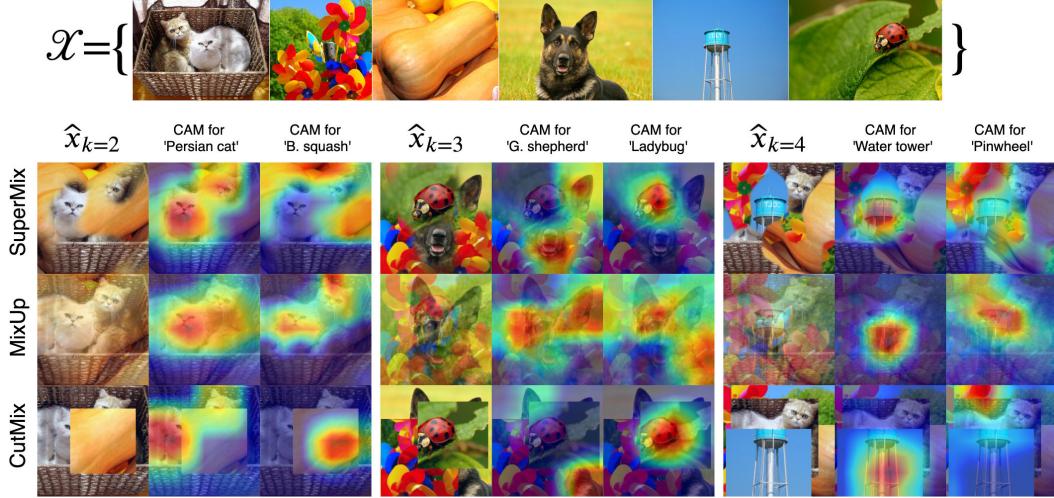


Figure 4: Visual comparison of the mixed images generated by SuperMix, MixUp, and CutMix, with $k \in \{2, 3, 4\}$ on ResNet34. Class activation maps [35] are computed for two classes in mixed images.

Dataset	Model	Base.	Automated aug.			Mixing aug.			SuperMix + RA[5]
			AA[4]	FAA[20]	RA[5]	MixUp	CutMix	SuperMix	
CIFAR-100	WRN-40-2 _a	74.0	79.3	79.4	79.2	77.2	77.9	79.7	79.9
	WRN-28-10	81.2	82.9	82.7	83.3	82.1	82.9	83.6	83.9
	S-S(26 2x96d)	82.9	85.7	85.4	85.6	84.8	85.0	85.5	85.8
ImageNet	ResNet-50	76.3/93.1	77.6/93.8	77.6/93.7	77.6/93.8	77.0/93.4	77.2/93.5	77.6/93.7	78.2/94.0
	ResNet-200	78.5/94.2	80.0/95.0	80.6/95.3	80.4/95.3	79.6/94.8	79.9/94.9	80.8/95.4	81.3/95.6

Table 1: Performance of augmentation methods on CIFAR-100 (Top-1 accuracy) and ImageNet (Top-1/Top-5 accuracy).

of the original dataset.

For knowledge distillation on CIFAR-100, we also consider an additional baseline by using unlabeled data from the training set of ImageNet32x32 [3] (ImgNet32) to construct unlabeled sets. This helps to better evaluate the role of the data provided by the mixing augmentation methods. We use SGD optimizer with an initial learning rate of 0.1 and momentum of 0.9. Weight decay is set to $5e-4$. The learning rate is decayed by 0.1 at epochs 200, 300, 400, and 500, and the maximum number of epochs is set to 600. Since in our experiments $\kappa \geq 1$, the number of epochs according to the mixed dataset will scale with $\frac{1}{\kappa}$ to keep the number of training iterations fixed for all experiments. For instance, when $\kappa = 5$, the maximum number of epochs for the mixed dataset is 120. The batch size is set to 128 and 256 for CIFAR-100 and ImageNet, respectively. For the CIFAR-100 dataset, we set σ of the Gaussian smoothing in SuperMix to 1 and the spatial size of the masks to 8×8 . For ImageNet, σ is set to 2 and the size of masks is set to 16×16 . For all benchmark comparisons, we set $\alpha = 3$ and $\lambda_s = 25$. Moreover, in all experiments, the performance of SuperMix is evaluated by generating 5×10^5 and 10^6 images on CIFAR-100 and ImageNet, respectively, unless otherwise noted. All the hyper-parameters for the distillation experiments are selected according to the exper-

imental setup of [28] and the ablation studies in Section 4.3. Network architectures and settings for baseline methods are provided in the supplemental material.

4.1. Object classification

We follow the standard setup of evaluation for automated augmentation [4, 20, 15] and compare them with SuperMix on the task of object classification. For SuperMix, we first train the target model on the original dataset and then use it to generate mixed data with k equal to 2 and 3 for CIFAR-100 and ImageNet, respectively. Afterward, we train the target model from scratch on the mixture of the augmented data and the original data. Rest of the result are reported from the original papers. As an additional evaluation, we combine SuperMix with RangAugment (RA) [5]. For this purpose, we first mix images using SuperMix and then apply RA with the default parameters [5] for CIFAR-100 and ImageNet. Table 1 presents the results for these experiments. On four out of five experiments, SuperMix provide performance competitive to SOTA approaches of automated augmentation. Furthermore, combining RA with SuperMix further improves the performance of classification across all the experiments. These evaluations highlight the effectiveness of mixing multiple images for data augmentation.

Teacher	WRN-40-2 _b		ResNet56		ResNet110		ResNet32x4	VGG13
Student	WRN-16-2	WRN-40-1	ResNet20	ResNet20	ResNet32	ResNet8x4	VGG8	
Teacher acc.	75.61		72.34		74.31		79.42	74.64
Student acc.	73.26	71.98	69.06	69.06	71.14	72.50	70.36	
Distillation method	KD [14]	74.92	73.54	70.66	70.67	73.08	73.33	72.98
	CRD [28]	75.48	74.14	71.16	71.46	73.48	75.51	73.94
	ImgNet32	74.91	74.80	71.38	71.48	73.17	75.57	73.95
	CE+	76.20*	75.53	72.00	72.27	74.60*	76.73	74.56
	MixUp	76.40*	75.85*	72.33	72.68	74.24	76.81	74.87*
	CutMix	76.93*	76.11*	72.64*	72.75	74.80*	77.16	75.38*
	SuperMix	76.93*	76.11*	72.64*	72.75	74.80*	77.16	75.38*
	KD+	76.52*	75.70*	72.22	72.23	74.24	76.46	75.02*
	MixUp	76.58*	76.10*	72.89*	72.82	74.94*	77.07	75.58*
	CutMix	76.81*	76.45*	72.67*	72.83	74.87*	76.90	75.50*
	SuperMix	77.45*	76.53*	73.19*	72.96	75.21*	77.59	76.03*

Table 2: Classification performance (%) of student models on CIFAR-100. Teacher and student are from the same architecture family but different depth/wideness and capacity. We denote by * results where the student surpasses the teacher performance. Only ImgNet32 uses unlabeled data from an external source. Average over 4 independent runs.

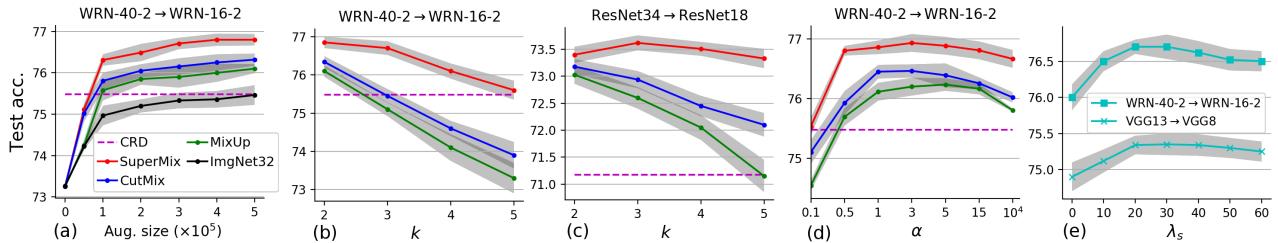


Figure 5: Evaluating the role of augmentation size and hyper-parameters.

4.2. Knowledge Distillation

In addition to KD [14] and CRD [28], we consider a simple method for distillation to highlight the effectiveness of mixing augmentation. In this method, we train the student models to classify mixed images labeled by the teacher model. The labels only show the winner class and does not contain any information regarding the rest of the classes. We refer to this method as Cross-Entropy (CE) distillation.

Results on CIFAR-100: Tables 2 and 3 presents the results for two challenging scenarios of distillation. In the first scenario, teacher and student are from the same family of architectures but have different depth/wideness and capacity. In the second scenario, teacher and student are from completely different network architectures. Employing the simple CE method using the mixed data consistently outperforms previous methods in both distillation scenarios. The data generated by SuperMix demonstrates the best performance across all evaluations, and, on five out of seven teacher-student setups from the same architecture family, students trained on the SuperMix data outperform their teachers. Last four rows in Tables 2 and 3 present the results for knowledge distillation using the original KD [14]. More importantly, results on MixUp, CutMix, and SuperMix demonstrate that they can notably enhance the

performance of the distillation techniques.

These observations highlight three crucial points. First, the limited size of the training set is a major factor constraining the performance of knowledge distillation. According to Table 2, almost all of the students achieve comparable results to CRD when external data of ImgNet32 is provided. Second, mixing augmentation provides more informative data for distillation compared to unlabeled data from an external source. Third, the supervised mixing results in rich images that are highly favorable for knowledge distillation and outperforms blind mixing methods.

Results on ImageNet: We showcase the effectiveness of the mixed data on ImageNet by distilling the knowledge of ResNet-34 into ResNet-18. Table 5 presents the results for the distillation on the ImageNet dataset. Using the simple CE method consistently outperforms the previous SOTA approaches. In five out of eight experiments of distillation using mixed images, the student outperforms the teacher. This demonstrates the scalability and effectiveness of the mixing augmentation for the task of knowledge distillation. Moreover, combining mixed data with the original distillation objective further enhances the distillation performance validating the effectiveness of the mixing augmentation for knowledge transfer in large-scale datasets.

Teacher	VGG13		ResNet50		ResNet32x4		WRN-40-2
Student	MobileNetV2	MobileNetV2	VGG8	ShuffleNetV1	ShuffleNetV2	ShuffleNetV1	
Teacher acc.	74.64		79.34		79.42		75.61
Student acc.	64.60		64.60	70.36	70.50	71.82	70.50
Distillation method	KD [14]	67.37	67.35	73.81	74.07	74.45	74.83
	CRD [28]	69.73	69.11	74.30	75.11	75.65	76.05*
	ImgNet32	68.85	68.01	73.96	76.80	77.56	75.87*
	CE+	MixUp	71.13	71.71	75.41	78.16	77.29*
	CutMix	70.93	70.64	75.84	77.89	79.32	77.50*
	SuperMix	71.65	72.13	76.07	78.47	79.53*	77.92*
	ImgNet32	69.14	68.44	74.32	76.87	77.90	76.23*
	KD+	MixUp	71.29	71.99	75.59	78.22	79.14
	CutMix	71.10	70.93	76.01	77.92	79.53*	77.65*
	SuperMix	71.81	72.40	76.28	78.51	79.80*	78.07*

Table 3: Classification performance (%) of student models on CIFAR-100. Teacher and student models are from different architectures. We denote by * results where the student surpasses the teacher performance. Average over 4 independent runs.

Net	SGD	Newton	
		w/o SP	w/ SP
ImgNet	<i>ET(sec.)</i>	15.41	6.59
	<i>iters</i>	34.5	15.1
CIFAR	<i>ET(sec.)</i>	4.25	1.98
	<i>iters</i>	23.6	11.7
	<i>ET(ms.)</i>	482	97
	<i>iters</i>	19.5	3.7
	<i>ET(ms.)</i>	509	122
	<i>iters</i>	21.8	4.6
			0.2

Table 4: Comparison of execution time.

4.3. Ablation studies

Impact of the size of the training set: In this part, we investigate how the size of the dataset affects the distillation performance by measuring the Top-1 test accuracy of WRN-16-2 versus the augmentation size on CIFAR-100. For all the mixing methods, we set $k = 2$ and $\alpha = 1$, *i.e.*, sampling mixing coefficients from the uniform distribution. Figures 5a presents the results for these evaluations. The distillation performance improves by increasing the augmentation size and plateaus at 5×10^5 . All the datasets generated using mixing augmentations outperform the unlabeled dataset of ImgNet32. This highlights the superiority of mixed images for knowledge transfer compared to unlabeled data from an external source. Based on these observations, we set the size of the mixed dataset to 5×10^5 for all experiments on CIFAR-100.

Impact of k : We evaluate the role of k by conducting experiments on CIFAR-100 and ImageNet datasets. Figures 5b and 5c present the results for this evaluation. A major shortcoming of MixUp and CutMix is that they mix images without any supervision. Including more input images to produce a mixed image increases the chance of incorrect cropping in CutMix, and averaging overlapping features in Mixup. This explains the notable deterioration of the distillation performance in all experiments with $k > 2$ using

these augmentation methods. Both of these incidents degrade the quality and effectiveness of features in the mixed image, which can also be observed from the visual comparisons provided in Figure 4. We observe that the spatial size of the image can limit k . Performance of distillation using SuperMix degrades for $k > 2$ on CIFAR-100. However on ImageNet, $k = 3$ yields the best distillation performance.

Impact of α : Parameter α determines the probability distribution for the presence of each input class in the mixed image. We measure the performance of distillation versus several values of α to identify its optimal value. Figure 5d presents results for these experiments. For $\alpha \rightarrow 0$, the mixing augmentation becomes inactive since only one input category will appear in the augmented images, *i.e.*, $r_0 = 1$ or $r_1 = 1$. For $\alpha \rightarrow +\infty$, the contribution of images become equal, *i.e.*, $r_0 = r_1 = 0.5$. This is more favorable for distillation since both input images contribute equally to the mixed image. For $\alpha = 1$, contribution of each input in the mixed image is selected from the uniform distribution $\text{Unif}(0, 1)$. According to the figures, we select $\alpha = 3$ for all other experiments unless otherwise noted.

Sparsity among masks: The sparsity promoting loss forces each spatial location in the output image to be assigned to only one image in the input set. This improves the mixing performance by preserving the most important features in each spatial location. We evaluate the performance of distillation versus λ_s in Figure 5e. By increasing the weight of sparsity the performance of distillation improves until $\lambda_s \approx 30$. After that the accuracy of masks degrades since the sparsity promoting loss dominates the KL loss. Figure 3 evaluates this phenomenon by visualising the mixing mask versus λ_s .

4.4. Execution time

Here, we compute the execution time of SuperMix. To this aim, we define two baselines for the sake of compar-

	Teacher	Student	KD	CRD	CE	KD	CE	KD	CE	KD
			+MixUp _{k=2}		+CutMix _{k=2}		+SuperMix _{k=2}		+SuperMix _{k=3}	
Top-1	73.31	69.75	70.66	71.17	73.03	73.29	73.18	73.33*	73.42*	73.62*
Top-5	91.42	89.07	89.88	90.13	91.27	91.44	91.36	91.44*	91.51*	91.66*

Table 5: Top-1 and Top-5 classification accuracy of ResNet18 on ImageNet dataset. Results where the student surpasses the teacher performance are marked by *. Average over 4 independent runs.

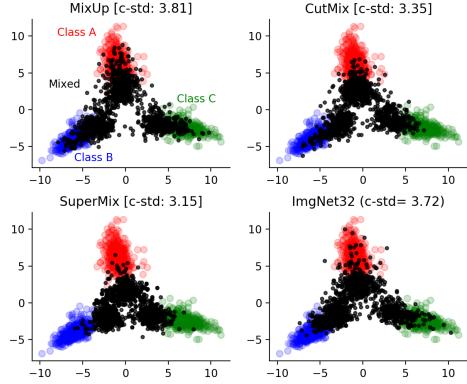


Figure 6: Visualizing representations for the mixed images.

ison. For the first baseline, we use SGD instead of the Newton method to optimize the set of masks. The second baseline is the Newton method without SP. Hence, the optimization in both baselines is performed on $\mathcal{L}_{SM} = KL + \lambda_\sigma \mathcal{L}_\sigma + \lambda_s \mathcal{L}_s$. Inspired by the previous work on saliency detection [10], we use the TV norm for the spatial smoothness loss as: $\mathcal{L}_s = \frac{1}{kWH} \sum_i \sum_{u \in \Lambda} \|\nabla m_i(u)\|_3^3$. Based on experimental observations, we set $\lambda_s = 250$, learning rate of SGD to 0.1. All other parameters are set to the values identified in previous sections. All algorithms are implemented with parallel processing on two NVIDIA Titan RTX with batch size of 128. For further implementation details, please refer to the released code.

Figure 4 presents the results for these comparisons. Newton method with SP, *i.e.*, SuperMix, is at least **65**× faster than SGD on both datasets. Moreover, due to SP which directly satisfied the spatial smoothness condition, SuperMix is at least **19**× faster than the same algorithm when it has to include \mathcal{L}_s .

4.5. Embedding space evaluations

We perform two sets of evaluations on CIFAR-100 to further analyze characteristics of the mixed images. In the first set of experiments, we feed the original data and the mixed images to VGG13 and visualize the output of the logits layer, in 2D for three random classes using PCA. The SuperMix images are generated with $k = 2$. Figure 6 demonstrates these evaluations. Representations for the SuperMix data has less overlap with the distribution of the representations for the original data. This suggests that the Super-

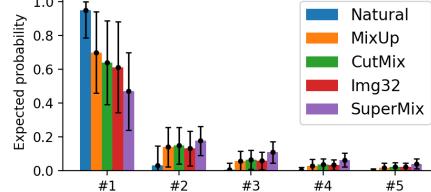


Figure 7: Distribution of top 5 predictions.

Mix data encompass more novel structure compared to the original data, unlabeled data from other mixing methods or an external source. The SuperMix data are harder to classify for the model since the representations are concentrated close to the center of the embedding. To better evaluate this, we compute the class standard deviation (c-std) of representations for each class. The computed values are reported on the top of the corresponding images in Figure 6.

Hinton *et al.* [14] pointed that smoothing out the predictions of a model can better reveal its knowledge of the task. Since SuperMix generates images by combining multiple inputs, the outputs of the model on SuperMix data are intrinsically more smooth compared to that of the other augmentation types. We validate this by computing the average of the sorted Top-5 probability predictions of VGG13 on the original and augmented images of CIFAR-100. As demonstrated in Figure 7, predictions of the target model is significantly smoother on mixed images. Moreover, SuperMix produces the data with the most smooth labels.

5. Conclusion

In this paper, we studied the potential of mixing multiple images using supervision of a teacher for the data augmentation. We proposed SuperMix, a supervised mixing augmentation method that combines salient regions in multiple images to produce unseen training samples. The effectiveness and efficiency of SuperMix is validated through extensive experiments, evaluations, and ablation studies. Specifically, incorporating SuperMix data for distillation enhances the state of the art of knowledge distillation. SuperMix provides comparable performance to the automated augmentation methods, and when combined, notably improves the generalization of the model.

References

- [1] Cristian Bucilu, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541. ACM, 2006.
- [2] Xinlei Chen, Abhinav Shrivastava, and Abhinav Gupta. Neil: Extracting visual knowledge from web data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1409–1416, 2013.
- [3] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*, 2017.
- [4] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 113–123, 2019.
- [5] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical data augmentation with no separate search. *arXiv preprint arXiv:1909.13719*, 2019.
- [6] Piotr Dabkowski and Yarin Gal. Real time image saliency for black box classifiers. In *Advances in Neural Information Processing Systems*, pages 6967–6976, 2017.
- [7] Ali Dabouei, Sobhan Soleymani, Fariborz Taherkhani, Jeremy Dawson, and Nasser M Nasrabadi. Smoothfool: An efficient framework for computing smooth adversarial perturbations. *arXiv preprint arXiv:1910.03624*, 2019.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [9] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [10] Ruth C Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3429–3437, 2017.
- [11] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3714–3722, 2019.
- [12] Dongyoon Han, Jiwhan Kim, and Junmo Kim. Deep pyramidal residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5927–5935, 2017.
- [13] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Patwary, Mostofa Ali, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.
- [14] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [15] Daniel Ho, Eric Liang, Ion Stoica, Pieter Abbeel, and Xi Chen. Population based augmentation: Efficient learning of augmentation policy schedules. *arXiv preprint arXiv:1905.05393*, 2019.
- [16] Le Kang, Peng Ye, Yi Li, and David Doermann. Convolutional neural networks for no-reference image quality assessment. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1733–1740, 2014.
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [18] Joseph Lemley, Shabab Bazrafkan, and Peter Corcoran. Smart augmentation learning an optimal data augmentation strategy. *Ieee Access*, 5:5858–5869, 2017.
- [19] Li-Jia Li and Li Fei-Fei. Optimol: automatic online picture collection via incremental model learning. *International journal of computer vision*, 88(2):147–168, 2010.
- [20] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast autoaugment. In *Advances in Neural Information Processing Systems*, pages 6662–6672, 2019.
- [21] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2574–2582, 2016.
- [22] Luis Perez and Jason Wang. The effectiveness of data augmentation in image classification using deep learning. *arXiv preprint arXiv:1712.04621*, 2017.
- [23] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. Semi-supervised self-training of object detection models. *WACV/MOTION*, 2, 2005.
- [24] Andrzej P Ruszczyński and Andrzej Ruszczyński. *Nonlinear optimization*, volume 13. Princeton university press, 2006.
- [25] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In

Advances in Neural Information Processing Systems, pages 5014–5026, 2018.

- [26] H Scudder. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965.
- [27] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [28] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019.
- [29] Yuji Tokozume, Yoshitaka Ushiku, and Tatsuya Harada. Between-class learning for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5486–5494, 2018.
- [30] Vladimir Vapnik and Vlaminir Vapnik. Statistical learning theory, 1998.
- [31] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, pages 189–196, 1995.
- [32] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. *arXiv preprint arXiv:1905.04899*, 2019.
- [33] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- [34] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [35] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.