Teach me how to Interpolate a Myriad of Embeddings

Shashanka Venkataramanan[1]   Ewa Kijak[1]   Laurent Amsaleg[1]   Yannis Avrithis[2]

[1]Inria, Univ Rennes, CNRS, IRISA
[2]Institute of Advanced Research in Artificial Intelligence (IARAI), Athena RC

Abstract

Mixup refers to interpolation-based data augmentation, originally motivated as a way to go beyond *empirical risk minimization* (ERM). Yet, its extensions focus on the definition of interpolation and the space where it takes place, while the augmentation itself is less studied: For a mini-batch of size m, most methods interpolate between m pairs with a single scalar interpolation factor λ.

In this work, we make progress in this direction by introducing *MultiMix*, which interpolates an arbitrary number n of tuples, each of length m, with one vector λ per tuple. On sequence data, we further extend to *dense* interpolation and loss computation over all spatial positions. Overall, we increase the number of tuples per mini-batch by orders of magnitude at little additional cost. This is possible by interpolating at the very last layer before the classifier. Finally, to address inconsistencies due to linear target interpolation, we introduce a *self-distillation* approach to generate and interpolate synthetic targets.

We empirically show that our contributions result in significant improvement over state-of-the-art mixup methods on four benchmarks. By analyzing the embedding space, we observe that the classes are more tightly clustered and uniformly spread over the embedding space, thereby explaining the improved behavior.

1 Introduction

Mixup [55] is a data augmentation method that interpolates between pairs of training examples, thus regularizing a neural network to favor linear behavior in-between examples. Besides improving generalization, it has important properties such as reducing overconfident predictions and increasing the robustness to adversarial examples. Several follow-up works have studied interpolation in the *latent* or *embedding* space, which is equivalent to interpolating along a manifold in the input space [42], and a number of nonlinear and attention-based interpolation mechanisms [52, 18, 17, 38, 16]. However, little progress has been made in the augmentation process itself, *i.e.*, the number of examples being interpolated and the number of interpolated examples being generated.

Mixup was originally motivated as a way to go beyond *empirical risk minimization* (ERM) [39] through a vicinal distribution expressed as an expectation over an interpolation factor λ, which is equivalent to the set of linear segments between all pairs of training inputs and targets. In practice however, in every training iteration, a single scalar λ is drawn and the number of interpolated pairs is limited to the size of the mini-batch, as illustrated in Figure 1(a). This is because, if interpolation takes place in the input space, it would be expensive to increase the number of examples per iteration. To our knowledge, these limitations exist in all mixup methods.

In this work, we argue that a data augmentation process should augment the data seen by the model, or at least by its last few layers, as much as possible. In this sense, we follow *manifold mixup* [42] and generalize it in a number of ways to introduce *MultiMix*, as illustrated in Figure 1(b). First, rather than pairs, we interpolate *tuples* that are as large as the mini-batch. Effectively, instead of
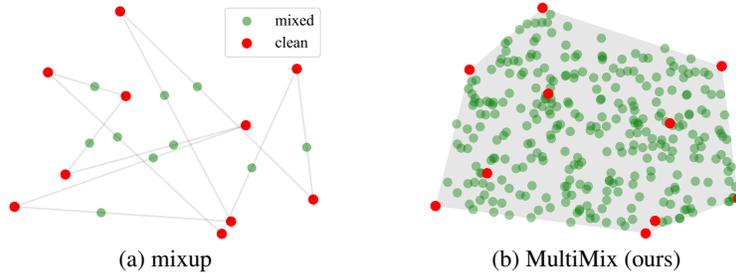
Figure 1: Data augmentation from a mini-batch B consisting of $m = 10$ points in two dimensions. (a) *mixup*: sampling of m points on linear segments between m pairs of points in B, using the same interpolation factor λ. (b) *MultiMix*: sampling of $n = 300$ points in the convex hull of B.

linear segments between pairs of examples in the mini-batch, we sample on their entire *convex hull*. Second, we draw a different vector λ for each tuple. Third, and most important, we increase the number of interpolated tuples per iteration by orders of magnitude by only slightly decreasing the actual training throughput in examples per second. This is possible by interpolating at the deepest layer possible, *i.e.*, just before the classifier, which also happens to be the most effective choice. The interpolated embeddings are thus only processed by a single layer.

Apart from increasing the number of examples seen by the model, another idea is to increase the number of loss terms per example. In many modalities of interest, the input is a *sequence* in one or more dimensions: pixels or patches in images, voxels in video, points or triangles in high-dimensional surfaces, to name a few. The structure of input data is expressed in matrices or tensors, which often preserve a certain spatial resolution until the deepest network layer before they collapse *e.g.* by global average pooling [36, 14] or by taking the output of a classification token [40, 7].

In this sense, we choose to operate at the level of *sequence elements* rather than representing examples by a single vector. We introduce *dense MultiMix*, which is the first approach of this kind in mixup-based data augmentation. In particular, we interpolate densely the embeddings and targets of sequence elements and we also apply the loss densely, as illustrated in Figure 2. This is an extreme form of augmentation where the number of interpolated tuples and loss terms increases further by one or two orders of magnitude, but at little cost.

Finally, linear interpolation of targets, which is the norm in most mixup variants, has a limitation: Given two examples with different class labels, the interpolated example may actually lie in a region associated with a third class in the feature space, which is identified as *manifold intrusion* [12]. In the absence of any data other than the mini-batch, a straightforward way to address this limitation is to devise targets originating in the network itself. This naturally leads to *self-distillation*, whereby a moving average of the network acts as a teacher and provides synthetic soft targets [37], to be interpolated exactly like the original hard targets.

In summary, we make the following contributions:

1. We introduce *MultiMix*, which, given a mini-batch of size m, interpolates an arbitrary number $n \gg m$ of tuples, each of length m, with one interpolation vector λ per tuple— compared with m pairs, all with the same scalar λ for most mixup methods (subsection 3.2).
2. We extend to *dense* interpolation and loss computation over all spatial positions (subsection 3.4).
3. We use *online self-distillation* to generate and interpolate soft targets for mixup—compared with linear target interpolation for most mixup methods (subsection 3.3).
4. We improve over state-of-the-art mixup methods on *image classification*, *robustness to adversarial attacks*, *object detection* and *out-of-distribution detection*. (section 4).

2 Related Work

Mixup   In general, mixup interpolates between pairs of input examples [57] or embeddings [42] and their corresponding target labels. Several follow-up methods mix input images according to

spatial position, either at random rectangles [52] or based on attention [38, 18, 17], in an attempt to focus on a different object in each image. We also use attention in our dense MultiMix variant, but in the embedding space. Other definitions of interpolation include the combination of content and style from two images [16] and the spatial alignment of dense features [41]. Our dense MultiMix variant also uses dense features but without aligning them, hence it can mix a very large number of images and generate even more interpolated data.

Our work is orthogonal to these methods as we focus on the sampling process of augmentation rather than on the definition of interpolation. As far as we are aware, the only methods that mix more than two examples are OptTransMix [60], which involves a complex optimization process in the input space and only applies to images with clean background, and SuperMix [4], which uses a Dirichlet distribution like we do, but interpolates in the input space not more than 3 images, while we interpolate all embeddings of a mini-batch.

Self-distillation Distillation refers to a two-stage knowledge transfer process where a larger teacher model or ensemble is trained before predicting soft targets to train a smaller student model on the same [2, 15, 33, 35, 53] or different [30, 46] training data. The architecture of the two models may be the same with training at multiple stages, for example in continual learning [31, 22]. In self-distillation or co-distillation, not only the models are the same, but the knowledge transfer process is also online, *e.g.* between layers of the same model [56] or between two versions of the model [1], where the teacher parameters may be obtained from the student rather than learned [37]. The latter approach has been successful in self-supervised representation learning [11, 3, 59]. As far as we know, distillation has only been used for mixup as a two stage process between different models [4] and we are the first to use online self-distillation in this context, following [37].

Dense loss functions Although standard in dense tasks like semantic segmentation [28, 13], where dense targets commonly exist, dense loss functions are less common otherwise. Few examples are in few-shot learning [23, 20], where data augmentation is of utter importance, and in unsupervised representation learning, *e.g.* dense contrastive learning [29, 44], learning from spatial correspondences [49, 47] and masked language or image modeling [5, 48, 21, 59]. Some of these methods use dense distillation [49, 59], which is also studied in continual learning [6, 8]. To our knowledge, we are the first to use dense interpolation and a dense loss function for mixup. Our setting is supervised, similar to dense classification [23], but we also use dense distillation [59].

3 Method

3.1 Preliminaries and background

Problem formulation Let $x \in \mathcal{X}$ be an input example and $y \in \mathcal{Y}$ its one-hot encoded target, where $\mathcal{X} = \mathbb{R}^D$ is the input space, $\mathcal{Y} = \{0, 1\}^c$ and c is the total number of classes. Let $f_\theta : \mathcal{X} \to \mathbb{R}^d$ be an encoder that maps the input x to an embedding $z = f_\theta(x)$, where d is the dimension of the embedding. A classifier $g_W : \mathbb{R}^d \to \Delta^{c-1}$ maps z to a vector $p = g_W(z)$ of predicted probabilities over classes, where $\Delta^n \subset \mathbb{R}^{n+1}$ is the unit n-simplex, *i.e.*, $p \geq 0$ and $\mathbf{1}_c^\top p = 1$, and $\mathbf{1}_c \in \mathbb{R}^c$ is an all-ones vector. The overall network mapping is $f := g_W \circ f_\theta$.

Parameters (θ, W) are learned by optimizing over mini-batches. Given a mini-batch of m examples, let $X = (x_1, \ldots, x_m) \in \mathbb{R}^{D \times m}$ be the inputs, $Y = (y_1, \ldots, y_m) \in \mathbb{R}^{c \times m}$ the targets and $P = (p_1, \ldots, p_m) \in \mathbb{R}^{c \times m}$ the predicted probabilities of the mini-batch, where $P = f(X) := (f(x_1), \ldots, f(x_m))$. The objective is to minimize the cross-entropy

$$H(Y, P) := -\mathbf{1}_c^\top (Y \odot \log(P)) \mathbf{1}_m / m \tag{1}$$

of predicted probabilities P relative to targets Y averaged over the mini-batch, where $\odot$ is the Hadamard (element-wise) product. In summary, the mini-batch loss is

$$L(X, Y; \theta, W) := H(Y, g_W(f_\theta(X))). \tag{2}$$

Mixup Mixup methods commonly interpolate pairs of inputs or embeddings and the corresponding targets at the mini-batch level while training. Given a mini-batch of m examples with inputs X and targets Y, let $Z = (z_1, \ldots, z_m) \in \mathbb{R}^{d \times m}$ be the embeddings of the mini-batch, where

$Z = f_\theta(X)$. *Manifold mixup* [42] interpolates the embeddings and targets by forming a convex combination of the pairs with interpolation factor $\lambda \in [0, 1]$:

$$\widetilde{Z} = Z(\lambda I + (1 - \lambda)\Pi) \tag{3}$$

$$\widetilde{Y} = Y(\lambda I + (1 - \lambda)\Pi), \tag{4}$$

where $\lambda \sim \mathrm{Beta}(\alpha, \alpha)$, I is the identity matrix and $\Pi \in \mathbb{R}^{m \times m}$ is a permutation matrix. *Input mixup* [55] interpolates inputs rather than embeddings:

$$\widetilde{X} = X(\lambda I + (1 - \lambda)\Pi). \tag{5}$$

Whatever the interpolation method and the space where it is performed, the interpolated data, *e.g.* $\widetilde{X}$ [55] or $\widetilde{Z}$ [42], replaces the original mini-batch data and gives rise to predicted probabilities $\widetilde{P} = (p_1, \ldots, p_m) \in \mathbb{R}^{c \times m}$ over classes, *e.g.* $\widetilde{P} = f(\widetilde{X})$ [55] or $\widetilde{P} = g_W(\widetilde{Z})$ [42]. Then, the average cross-entropy $H(\widetilde{Y}, \widetilde{P})$ (1) between the predicted probabilities $\widetilde{P}$ and interpolated targets $\widetilde{Y}$ is minimized.

The number of interpolated data is m, same as the original mini-batch data.

3.2 MultiMix

Interpolation Given a mini-batch of m examples with embeddings Z and targets Y, we draw interpolation vectors $\lambda_k \sim \mathrm{Dir}(\alpha)$ for $k = 1, \ldots, n$, where $\mathrm{Dir}(\alpha)$ is the symmetric Dirichlet distribution and $\lambda_k \in \Delta^{m-1}$, that is, $\lambda_k \geq 0$ and $\mathbf{1}_m^\top \lambda_k = 1$. We then interpolate embeddings and targets by taking n convex combinations over all m examples:

$$\widetilde{Z} = Z\Lambda \tag{6}$$

$$\widetilde{Y} = Y\Lambda, \tag{7}$$

where $\Lambda = (\lambda_1, \ldots, \lambda_n) \in \mathbb{R}^{m \times n}$. We thus generalize manifold mixup [42]:

1. from pairs to tuples of length m, as long as the mini-batch: m-term convex combination (6),(7) *vs.* 2-term in (3),(4), Dirichlet *vs.* Beta distribution;

2. from m to an arbitrary number n of tuples: interpolated embeddings $\widetilde{Z} \in \mathbb{R}^{d \times n}$ (6) *vs.* $\mathbb{R}^{d \times m}$ in (3), interpolated targets $\widetilde{Y} \in \mathbb{R}^{c \times n}$ (7) *vs.* $\mathbb{R}^{c \times m}$ in (4);

3. from fixed λ across the mini-batch to a different λ_k for each interpolated item.

Loss Again, we replace the original mini-batch embeddings Z by the interpolated embeddings $\widetilde{Z}$ and minimize the average cross-entropy $H(\widetilde{Y}, \widetilde{P})$ (1) between the predicted probabilities $\widetilde{P} = g_W(\widetilde{Z})$ and the interpolated targets $\widetilde{Y}$ (7). Compared with (2), the mini-batch loss becomes

$$L_M(X, Y; \theta, W) := H(Y\Lambda, g_W(f_\theta(X)\Lambda)). \tag{8}$$

3.3 MultiMix with self-distillation

Networks We use an online self-distillation approach whereby the network $f := g_W \circ f_\theta$ that we learn becomes the *student*, whereas a *teacher* network $f' := g_{W'} \circ f_{\theta'}$ of the same architecture is obtained by exponential moving average of the parameters [37, 11]. The teacher parameters (θ', W') are not learned: We stop the gradient in the computation graph.

Views Given two transformations T and T', we generate two different augmented views $v = t(x)$ and $v' = t'(x)$ for each input x, where $t \sim T$ and $t' \sim T'$. Then, given a mini-batch of m examples with inputs X and targets Y, let $V = t(X), V' = t'(X) \in \mathbb{R}^{D \times m}$ be the mini-batch views corresponding to the two augmentations and $Z = f_\theta(V), Z' = f_{\theta'}(V') \in \mathbb{R}^{d \times m}$ the embeddings obtained by the student and teacher encoders respectively.

Interpolation We obtain the interpolated embeddings $\widetilde{Z}, \widetilde{Z}'$ from Z, Z' by (6) and targets $\widetilde{Y}$ from Y by (7), using the same Λ. The predicted class probabilities are given by $\widetilde{P} = g_W(\widetilde{Z})$ and $\widetilde{P}' = g_{W'}(\widetilde{Z}')$, again obtained by the student and teacher classifiers, respectively.

4
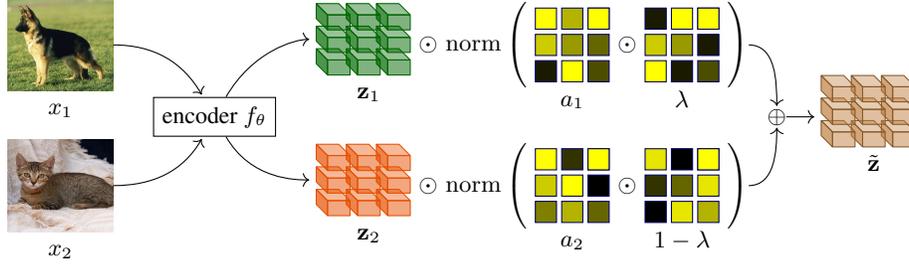
Figure 2: *Dense MultiMix* (subsection 3.4) for the special case $m = 2$ (two examples), $n = 1$ (one interpolated embedding), $r = 9$ (spatial resolution 3×3). The embeddings $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^{d \times 9}$ of input images x_1, x_2 are extracted by encoder f_θ. Attention maps $a_1, a_2 \in \mathbb{R}^9$ are extracted (10), multiplied element-wise with interpolation vectors $\lambda, (1 - \lambda) \in \mathbb{R}^9$ (11) and ℓ_1-normalized per spatial position (12). The resulting weights are used to form the interpolated embedding $\tilde{\mathbf{z}} \in \mathbb{R}^{d \times 9}$ as a convex combination of $\mathbf{z}_1, \mathbf{z}_2$ per spatial position (13). Targets are interpolated similarly (14).

Loss We learn parameters (θ, W) by minimizing a classification and a self-distillation loss:

$$\gamma H(\widetilde{Y}, \widetilde{P}) + (1 - \gamma) H(\widetilde{P}', \widetilde{P}), \tag{9}$$

where $\gamma \in [0, 1]$. The former brings the probabilities $\widetilde{P}$ predicted by the student close to the targets $\widetilde{Y}$, as in (8). The latter brings $\widetilde{P}$ close to the probabilities $\widetilde{P}'$ predicted by the teacher.

3.4 Dense MultiMix

We now extend the previous methodology to the case where the embeddings are structured, *e.g.* in matrices or tensors rather than vectors. This happens *e.g.* with token *vs.* sentence embeddings in NLP and patch *vs.* image embeddings in vision. In practice, this works by removing spatial pooling and rather applying the loss function densely over all tokens/patches. The idea is illustrated in Figure 2. For the sake of exposition, the formulation below uses sets of matrices grouped either by example or by spatial position. In practice, all operations are on tensors.

Preliminaries The encoder is now $f_\theta : \mathcal{X} \to \mathbb{R}^{d \times r}$, mapping the input x to an embedding $\mathbf{z} = f_\theta(x) \in \mathbb{R}^{d \times r}$, where d is the number of channels and r is its spatial resolution—if there are more than one spatial dimensions, these are flattened.

Given a mini-batch of m examples, we have again inputs $X = (x_1, \dots, x_m) \in \mathbb{R}^{D \times m}$ and targets $Y = (y_1, \dots, y_m) \in \mathbb{R}^{c \times m}$. Each embedding $\mathbf{z}_i = f_\theta(x_i) = (z_i^1, \dots, z_i^r) \in \mathbb{R}^{d \times r}$ for $i = 1, \dots, m$ consists of features $z_i^j \in \mathbb{R}^d$ for spatial position $j = 1, \dots, r$. We group features by position in matrices $Z^1, \dots, Z^r$, where $Z^j = (z_1^j, \dots, z_m^j) \in \mathbb{R}^{d \times m}$ for $j = 1, \dots, r$.

Attention Each feature vector will inherit the target of the corresponding input example. However, we also attach a level of confidence according to an attention map. Given an embedding $\mathbf{z} \in \mathbb{R}^{d \times r}$ with target $y \in \mathcal{Y}$ and a vector $u \in \mathbb{R}^d$, the attention map

$$a = h(\mathbf{z}^\top u) \in \mathbb{R}^r \tag{10}$$

measures the similarity of features of $\mathbf{z}$ to u, where h is a non-linearity, *e.g.* softmax or ReLU followed by ℓ_1 normalization. There are different ways to define vector u. For example, $u = \mathbf{z} \mathbf{1}_r / r$ by global average pooling (GAP) of $\mathbf{z}$, or $u = Wy$ assuming a linear classifier with $W \in \mathbb{R}^{d \times c}$, similar to class activation mapping (CAM) [58]. In the case of no attention, $a = \mathbf{1}_r / r$ is uniform.

Given a mini-batch, let $a_i = (a_i^1, \dots, a_i^r) \in \mathbb{R}^r$ be the attention map of embedding $\mathbf{z}_i$ (10). We group attention by position in vectors $a^1, \dots, a^r$, where $a^j = (a_1^j, \dots, a_m^j) \in \mathbb{R}^m$ for $j = 1, \dots, r$.

Interpolation For each spatial position $j = 1, \dots, r$, we draw interpolation vectors $\lambda_k^j \sim \mathrm{Dir}(\alpha)$ for $k = 1, \dots, n$ and define $\Lambda^j = (\lambda_1^j, \dots, \lambda_n^j) \in \mathbb{R}^{m \times n}$. Because input examples are assumed to contribute according to the attention vector $a^j \in \mathbb{R}^m$, we scale the rows of Λ^j accordingly and then

we normalize its columns back to Δ^{m-1} so that they can define convex combinations:

$$M^j = \text{diag}(a^j)\Lambda^j \tag{11}$$

$$\hat{M}^j = M^j \text{diag}(\mathbf{1}_m^\top M^j)^{-1} \tag{12}$$

We then interpolate embeddings and targets by taking n convex combinations over m examples:

$$\widetilde{Z}^j = Z^j \hat{M}^j \tag{13}$$

$$\widetilde{Y}^j = Y \hat{M}^j. \tag{14}$$

This is similar to (6),(7), but there is a different interpolated embedding matrix $\widetilde{Z}^j \in \mathbb{R}^{d \times n}$ as well as target matrix $\widetilde{Y}^j \in \mathbb{R}^{c \times n}$ per position, even though the original target matrix Y is one.

Classifier The classifier is now $g_W : \mathbb{R}^{d \times r} \to \mathbb{R}^{c \times r}$, maintaining the same spatial resolution as the embedding and generating one vector of predicted probabilities per spatial position. This is done by removing average pooling or any down-sampling operation. The interpolated embeddings $\widetilde{Z}^1, \ldots, \widetilde{Z}^r$ (13) are grouped by example into $\widetilde{\mathbf{z}}_1, \ldots, \widetilde{\mathbf{z}}_n \in \mathbb{R}^{d \times r}$, mapped by g_W to predicted probabilities $\widetilde{\mathbf{p}}_1, \ldots, \widetilde{\mathbf{p}}_n \in \mathbb{R}^{c \times r}$ and grouped again by position into $\widetilde{P}^1, \ldots, \widetilde{P}^r \in \mathbb{R}^{c \times n}$.

In the simple case where the original classifier is linear, *i.e.* $W \in \mathbb{R}^{d \times c}$, it is seen as 1×1 convolution and applied densely to each column (feature) of $\widetilde{Z}^j$ for $j = 1, \ldots, r$.

Loss Finally, we learn parameters θ, W by minimizing the weighted cross-entropy $H(\widetilde{Y}^j, \widetilde{P}^j; s)$ of $\widetilde{P}^j$ relative to the interpolated targets $\widetilde{Y}^j$ again densely at each position j, where

$$H(Y, P; s) := -\mathbf{1}_c^\top (Y \odot \log(P))s/(\mathbf{1}_n^\top s) \tag{15}$$

generalizes (1) and the weight vector is defined as $s = \mathbf{1}_m^\top M^j \in \mathbb{R}^n$. This is exactly the vector used to normalize the columns of M^j in (12). The motivation is that the columns of M^j are the original interpolation vectors weighted by attention: A small ℓ_1 norm indicates that for the given position j, we are sampling from examples of low attention, hence the loss is to be discounted.

4 Experiments

4.1 Setup

We use a mini-batch of size $m = 128$ examples in all experiments. For every mini-batch, we apply MultiMix with probability 0.5 or input mixup otherwise. For MultiMix, the default settings are given in subsection 4.5. We follow the experimental settings of AlignMixup [41] and use PreActResnet-18 (R-18) [14] and WRN16-8 [54] as encoder on CIFAR-10 and CIFAR-100 datasets [19]; R-18 on TinyImagenet [50] (TI); and Resnet-50 (R-50) and ViT-S/16 [7] on ImageNet [34]. We use top-1 error (%) as evaluation metric on image classification and robustness to adversarial attacks (subsection 4.2). We also experiment on object detection (subsection 4.3) and out-of-distribution detection, which is in the supplementary material along with more details and results.

4.2 Results: Image classification and robustness

Image classification In Table 1(a) we observe that MultiMix and Dense MultiMix already outperform SoTA on all datasets except CIFAR-10 with R-18, where they are on par with Co-Mixup. The addition of distillation increases the gain and outperforms SoTA on all datasets. Both distillation and dense improve over vanilla MultiMix and their effect is complementary on all datasets. On TI for example, distillation improves by 0.85%, dense by 1.24% and their combination by 1.91%. This combination brings an impressive gain of 2.26% over the previous SoTA – AlignMixup.

In Table 1(b) we observe that on ImageNet with R-50, vanilla MultiMix already outperforms all methods except SoTA AlignMixup. The addition of dense, distillation or both outperforms all SoTA with both R-50 and ViT-S/16. More importantly, it brings an overall gain of 4% over the baseline with R-50 and 3% with ViT-S/16.

DATASET NETWORK	CIFAR-10 R-18	CIFAR-10 W16-8	CIFAR-100 R-18	CIFAR-100 W16-8	TI R-18
Baseline[†]	5.19	5.11	23.24	20.63	43.40
Manifold mixup [42][†]	2.95	3.56	19.80	19.23	40.76
PuzzleMix [18][†]	2.93	2.99	20.01	19.25	36.52
Co-Mixup [17][†]	2.89	3.04	19.81	19.57	35.85
AlignMixup [41][†]	2.95	3.09	18.29	18.77	33.13
MultiMix (ours)	2.97	2.92	18.19	18.57	32.78
+ distil	2.87	2.78	17.72	17.91	31.93
+ dense	2.91	2.89	18.12	18.20	31.54
+ dense + distil	2.81	2.78	17.48	17.66	30.87
Gain	+0.08	+0.21	+0.81	+1.11	+2.26

(a) *Image classification* top-1 error (%) on CIFAR-10/100 and TI (TinyImagenet). R: PreActResnet, W: WRN.

NETWORK METHOD	RESNET-50 SPEED	RESNET-50 ERROR	VIT-S/16 SPEED	VIT-S/16 ERROR
Baseline[†]	1.17	23.68	1.01	26.1
Manifold mixup [42][†]	1.15	22.50	0.97	24.8
PuzzleMix [18][†]	0.84	21.24	0.73	24.3
Co-Mixup [17][†]	0.62	–	0.57	24.1
AlignMixup [41][†]	1.03	20.68	–	–
MultiMix (ours)	1.16	21.19	1.0	24.8
+ distil	1.06	19.88	0.93	23.4
+ dense	0.95	20.63	0.88	23.9
+ dense + distil	0.83	19.79	0.81	23.1
Gain		+0.89		+1.0

(b) *Image classification and training speed* on ImageNet. Top-1 error (%): lower is better. Speed: images/sec ($\times 10^3$): higher is better.

Table 1: *Image classification and training speed.* [†]: reported by AlignMixup. **Bold black**: best; Blue: second best; underline: best baseline. Gain: reduction of error over best baseline. Comparison with additional baselines is given in the supplementary material.

Training speed Table 1(b) shows the training speed of MultiMix and its variants compared with SoTA mixup methods, measured on NVIDIA V-100 GPU, including forward and backward pass. In terms of training speed, the vanilla MultiMix is on par with the baseline, bringing a gain of 2.49%. The addition of distillation is on par with SoTA AlignMixup, bringing a gain of 0.80%. Adding both dense and distillation brings a gain of 0.89% over AlignMixup, while being 19.4% slower. The inference speed is the same for all methods.

Robustness to adversarial attacks We follow the experimental settings of AlignMixup [41] and use $8/255$ l_∞ ϵ-ball for FGSM [10] and $4/255$ l_∞ ϵ-ball with step size $2/255$ for PGD [26] attack. In Table 3 we observe that vanilla MultiMix is already more robust than SoTA on all datasets and settings except FGSM on CIFAR-100 with R-18, where it is on par with AlignMixup. The addition of dense, distillation or both again increases the robustness and shows that their effect is complementary. The overall gain is more impressive than in classification error. For example, against the strong PGD attack on CIFAR-10 with W16-8, the SoTA Co-Mixup improves the baseline by 3.75% and our best result improves the baseline by 9.38%, which is more than double.

4.3 Results: Transfer learning to object detection

We evaluate the effect of mixup on the generalization ability of a pre-trained network to object detection as a downstream task. Following the settings of CutMix [52], we pre-train R-50 on ImageNet with MultiMix and its variants and use it as the backbone for SSD [25] with fine-tuning on Pascal VOC07+12 [9] and Faster-RCNN [32] with fine-tuning on MS-COCO [24].

In Table 2, we observe that, while vanilla Multi-Mix is slightly worse than SoTA AlignMixup, dense and distillation bring improvements over the SoTA on both datasets and are still complementary. This is consistent with classification results. Compared with the baseline, our best setting brings a gain of 2.40% mAP on Pascal VOC07+12 and 3.14% on MS-COCO.

DATASET DETECTOR	VOC07+12 SSD	MS-COCO FASTER R-CNN
Baseline[†]	76.7	33.27
Input mixup[†]	76.6	34.18
CutMix[†]	77.6	35.16
AlignMixup[†]	78.4	35.84
MultiMix (ours)	77.9	35.73
+ distil	78.7	35.97
+ dense	78.5	35.89
+ dense + distil	79.1	36.41
Gain	+0.7	+0.57

Table 2: *Transfer learning* to object detection. Mean average precision (mAP, %): higher is better. [†]: reported by AlignMixup. **Bold black**: best; Blue: second best; underline: best baseline. Gain: increase in mAP.

ATTACK	FGSM					PGD			
DATASET	CIFAR-10		CIFAR-100		TI	CIFAR-10		CIFAR-100	
NETWORK	R-18	W16-8	R-18	W16-8	R-18	R-18	W16-8	R-18	W16-8
Baseline[†]	89.41	88.02	87.12	72.81	91.85	99.99	99.94	99.97	99.99
Manifold mixup [42][†]	77.63	76.11	80.29	56.45	89.25	97.22	98.49	99.66	98.43
PuzzleMix [18][†]	57.11	60.73	78.70	57.77	83.91	97.73	97.00	96.42	95.28
Co-Mixup [17][†]	60.19	58.93	77.61	56.59	–	97.59	96.19	95.35	94.23
AlignMixup [41][†]	54.83	56.20	74.18	55.05	78.83	95.42	96.71	90.40	92.16
MultiMix (ours)	54.19	55.39	75.84	54.58	77.51	94.27	94.83	90.02	91.68
+ distillation	52.55	51.42	73.55	52.77	76.20	92.69	93.90	88.87	90.54
+ dense	54.10	53.33	74.48	53.01	75.57	92.99	92.68	88.60	90.90
+ dense + distillation	52.07	50.17	72.98	52.19	75.18	90.82	90.56	87.58	90.18
Gain	+2.76	+5.95	+1.20	+2.86	+3.65	+4.60	+5.63	+2.82	+1.98

Table 3: *Robustness to FGSM & PGD attacks*. Top-1 error (%): lower is better. [†]: reported by AlignMixup. **Bold black**: best; Blue: second best; underline: best baseline. Gain: reduction of error over best baseline. TI: TinyImagenet. R: PreActResnet, W: WRN. Comparison with additional baselines is given in the supplementary material.

(a) Baseline   (b) Manifold mixup [42]   (c) SaliencyMix [38]   (d) AlignMixup [41]   (e) Dense MultiMix + distillation (ours)
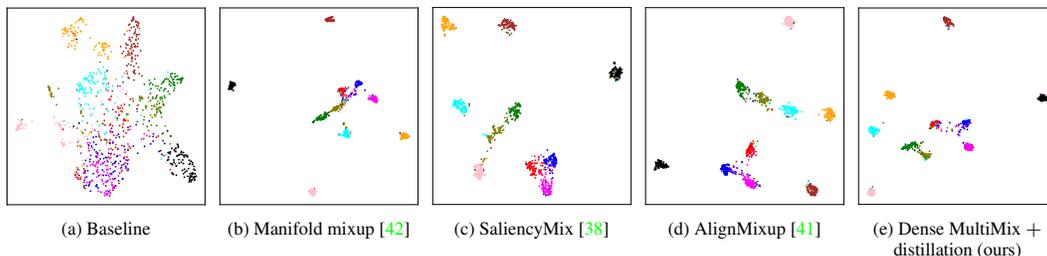
Figure 3: *Embedding space visualization* for 100 test examples per class of 10 randomly chosen classes of CIFAR-100 with PreActResnet-18, using UMAP [27].

4.4 Analysis of the embedding space

Qualitative analysis We qualitatively analyze the embedding space on 10 CIFAR-100 classes in Figure 3. We observe that the quality of embeddings of the baseline is extremely poor with severely overlapping classes, which explains its poor performance on image classification. All mixup methods result in clearly better clustered and more uniformly spread classes. Manifold mixup [42] produces five tightly clustered classes but the other five are still severely overlapping. SaliencyMix [38] and AlignMixup [41] yield four somewhat clustered classes and 6 moderately overlapping ones. Our best setting, *i.e.*, dense MultiMix with distillation, results in five tightly clustered classes and another five somewhat overlapping but less than all competitors. More plots including variants of MultiMix are given in the supplementary material.

Quantitative analysis We also quantitatively assess the embedding space on the CIFAR-100 test set using alignment and uniformity [43]. *Alignment* measures the expected pairwise distance of examples in the same class. Lower alignment indicates that the classes are more tightly clustered. *Uniformity* measures the (log of the) expected pairwise similarity of all examples using a Gaussian kernel as a similarity function. Lower uniformity indicates that classes are more uniformly spread in the embedding space. On CIFAR-100, we obtain alignment 3.02 for baseline, 1.27 for Manifold Mixup [42], 2.44 for SaliencyMix [38], 2.04 for AlignMixup and 0.92 for Dense MultiMix with distillation. We also obtain uniformity -1.94 for the baseline, -2.38 for Manifold Mixup [42], -2.82 for SaliencyMix [38], -4.77 for AlignMixup [41] and -5.68 for dense MultiMix with distillation. These results validate the qualitative analysis of Figure 3.

4.5 Ablations

All ablations are performed using R-18 on CIFAR-100. For MultiMix, we study the effect of the layer where we interpolate, the number of tuples n and a fixed value of Dirichlet parameter α.

More ablations are given in the supplementary material.
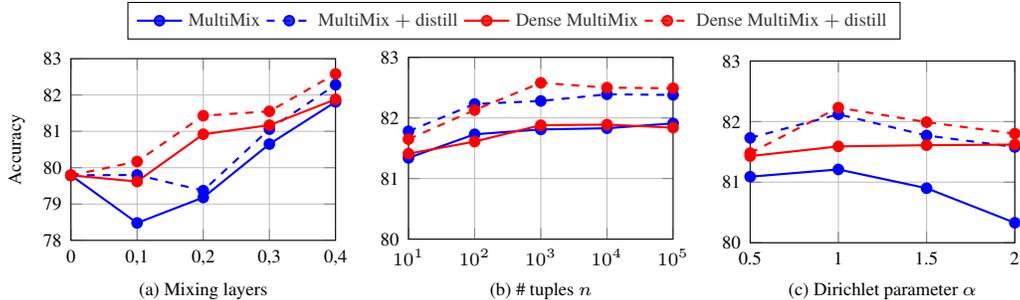
Figure 4: *Ablation study* of MultiMix and its variants on CIFAR-100 using R-18. (a) Interpolation layers (R-18 block; 0: input mixup). (b) Number of tuples n. (c) Dirichlet parameter α.

Interpolation layer For MultiMix, we use the entire network as the encoder f_θ by default, except for the last fully-connected layer, which we use as classifier g_W. Thus, we interpolate embeddings in the deepest layer by default. Here, we study the effect of different decompositions of the network $f = g_W \circ f_\theta$, such that interpolation of embeddings takes place at a different layer. When using distillation, we interpolate at the same layer for both the teacher and the student. In Figure 4(a), we observe that mixing at the deeper layers of the network significantly improves performance. The same behavior is observed when adding dense, distillation, or both. This validates our default choice.

Number of tuples n Since our aim is to increase the amount of data seen by the model, or at least part of the model, it is important to study the number n of interpolated embeddings. We observe from Figure 4(b) that accuracy increases overall with n and saturates for $n \geq 1000$ for all variants of MultiMix. Our best setting, Dense MultiMix with distillation, works best at $n = 1000$. We choose this as default, given also that the training cost increases with n. The training speed as a function of n is given in the supplementary material and is nearly constant for $n \leq 1000$.

Dirichlet parameter α Our default setting for α is to draw it uniformly at random from the interval $[0.5, 2]$ for every interpolation vector (column of Λ) that we draw. Here we study the effect of a fixed value of α. In Figure 4(c), we observe that the best accuracy is achieved for $\alpha = 1$ for most variants of MultiMix, which corresponds to the uniform distribution over the convex hull of the mini-batch embeddings. However, all measurements are lower than the default $\alpha \sim U[0.5, 2]$. For example, from Table 1(a) (CIFAR-100, R-18), dense MultiMix + distillation has accuracy 82.52, compared with 82.23 in Figure 4(c) for $\alpha = 1$.

5 Conclusion

In terms of input interpolation, the take-home message of this work is that, instead of devising smarter and more complex interpolation functions in the input space or the first layers of the representation, it is more beneficial to just perform linear interpolation in the very last layer where the cost is minimal, and then increase as much as possible the number of interpolated embeddings for mixup. This is more in line with the original motivation of mixup as a way to go beyond ERM. In terms of target interpolation, the take-home message is the opposite: instead of linear interpolation of original targets, find new synthetic targets for the interpolated embeddings with the help of the network itself, then interpolate them linearly. This idea fits nicely with self-distillation, which is popular in settings such as self-supervised representation learning and continual learning. Interestingly, self-distillation can be seen as yet another form of augmentation, but in the model space.

A natural extension of this work is the application to settings other than supervised classification, which has been the focus of this work. A limitation is that it is not straightforward to combine the sampling scheme of MultiMix with complex interpolation methods, unless they are fast to compute in the embedding space.

References

[1] Rohan Anil, Gabriel Pereyra, Alexandre Passos, Robert Ormandi, George E Dahl, and Geoffrey E Hinton. Large scale distributed neural network training through online distillation. *arXiv preprint arXiv:1804.03235*, 2018. 3

[2] Cristian Buciluǎ, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *ACM SIGKDD*, 2006. 3

[3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 3

[4] Ali Dabouei, Sobhan Soleymani, Fariborz Taherkhani, and Nasser M. Nasrabadi. Supermix: Supervising the mixing data augmentation. In *CVPR*, 2021. 3, 1, 2, 4, 5

[5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019. 3

[6] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyan Wu, and Rama Chellappa. Learning without memorizing. In *CVPR*, 2019. 3

[7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 6

[8] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. POD-Net: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, 2020. 3

[9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 7

[10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 7

[11] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *NeurIPS*, 2020. 3, 4

[12] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *AAAI*, 2019. 2

[13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 3

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 6

[15] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3

[16] Minui Hong, Jinwoo Choi, and Gunhee Kim. Stylemix: Separating content and style for enhanced data augmentation. In *CVPR*, 2021. 1, 3, 2, 4

[17] Jang-Hyun Kim, Wonho Choo, Hosan Jeong, and Hyun Oh Song. Co-mixup: Saliency guided joint mixup with supermodular diversity. In *ICLR*, 2021. 1, 3, 7, 8, 2, 4

[18] Jang-Hyun Kim, Wonho Choo, and Hyun Oh Song. Puzzle mix: Exploiting saliency and local statistics for optimal mixup. In *ICML*, 2020. 1, 3, 7, 8, 2, 4

[19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 6

[20] Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo. Revisiting local descriptor based image-to-class measure for few-shot learning. In *CVPR*, 2019. 3

[21] Zhaowen Li, Zhiyang Chen, Fan Yang, Wei Li, Yousong Zhu, Chaoyang Zhao, Rui Deng, Liwei Wu, Rui Zhao, Ming Tang, et al. MST: Masked self-supervised transformer for visual representation. In *NeurIPS*, 2021. 3

[22] Z. Li and D. Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, Dec 2018. 3

[23] Yann Lifchitz, Yannis Avrithis, Sylvaine Picard, and Andrei Bursuc. Dense classification and implanting for few-shot learning. In *CVPR*, 2019. 3

[24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 7

[25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 7

[26] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 7

[27] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. Umap: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 2018. 8

[28] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *ICCV*, 2015. 3

[29] Pedro O Pinheiro, Amjad Almahairi, Ryan Benmalek, Florian Golemo, and Aaron Courville. Unsupervised learning of dense visual representations. In *NeurIPS*, 2020. 3

[30] Ilija Radosavovic, Piotr Dollar, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *CVPR*, 2018. 3

[31] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, 2017. 3

[32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 7

[33] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for thin deep nets. In *ICLR*, 2014. 3

[34] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 6

[35] Zhiqiang Shen, Zhankui He, and Xiangyang Xue. MEAL: Multi-model ensemble via adversarial learning. In *AAAI*, 2019. 3

[36] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 2

[37] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2, 3, 4

[38] A F M Uddin, Mst. Monira, Wheemyung Shin, TaeChoong Chung, and Sung-Ho Bae. SaliencyMix: A saliency guided data augmentation strategy for better regularization. In *ICML*, 2021. 1, 3, 8, 2, 4

[39] VN Vapnik. An overview of statistical learning theory. *Neural Networks, IEEE Transactions on*, 10(5):988–999, 1999. 1

[40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2

[41] Shashanka Venkataramanan, Ewa Kijak, Laurent Amsaleg, and Yannis Avrithis. Alignmixup: Improving representation by interpolating aligned features. In *CVPR*, 2022. 3, 6, 7, 8, 1, 2, 4

[42] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *ICML*, 2019. 1, 2, 4, 7, 8, 3

[43] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, 2020. 8

[44] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, 2021. 3

[45] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. 3

[46] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020. 3

[47] Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *CVPR*, 2021. 3

[48] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. *arXiv preprint arXiv:2111.09886*, 2021. 3

[49] Yuwen Xiong, Mengye Ren, Wenyuan Zeng, and Raquel Urtasun. Self-supervised representation learning from flow equivariance. In *ICCV*, 2021. 3

[50] Leon Yao and John Miller. Tiny imagenet classification with convolutional neural networks. Technical report, Standford University, 2015. 6

[51] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 3

[52] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In

ICCV, 2019. 1, 3, 7, 2, 4

[53] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2016. 3

[54] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 6

[55] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 1, 4, 2, 3

[56] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *ICCV*, 2019. 3

[57] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *CVPR*, 2018. 2

[58] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 5, 4

[59] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. iBOT: Image bert pre-training with online tokenizer. In *ICLR*, 2022. 3

[60] Jianchao Zhu, Liangliang Shi, Junchi Yan, and Hongyuan Zha. Automix: Mixup networks for sample interpolation via cooperative barycenter learning. In *ECCV*, 2020. 3

A More experiments

A.1 More on setup

Settings and hyperparameters We train MultiMix and its variants with mixed examples only. We use a mini-batch of size $m = 128$ examples in all experiments. For every mini-batch, we apply MultiMix with probability 0.5 or input mixup otherwise. For input mixup, we interpolate the standard m pairs (5). For MultiMix, we use the entire network as the encoder f_θ by default, except for the last fully-connected layer, which we use as classifier g_W. We use $n = 1000$ tuples and draw a different $\alpha \sim U[0.5, 2.0]$ for each example from the Dirichlet distribution by default. For multi-GPU experiments, all training hyperparameters including m and n are per GPU.

For dense MultiMix, the spatial resolution is 4×4 ($r = 16$) on CIFAR-10/100 and 7×7 ($r = 49$) on Imagenet by default. We obtain the attention map by (10) using GAP for vector u and ReLU followed by ℓ_1 normalization as non-linearity h by default. To predict class probabilities and compute the loss densely, we use the classifier g_W as 1×1 convolution by default; when interpolating at earlier layers, we follow the process described in subsection 3.4. For distillation, both the teacher and student networks have the same architecture. By default, we use $\gamma = \frac{1}{2}$ in (9), that is, equal contribution of original labels and teacher predictions.

CIFAR-10/100 training Following the experimental settings of AlignMixup [41], we train MultiMix and its variants using SGD for 2000 epochs using the same random seed as AlignMixup. We set the initial learning rate to 0.1 and decay it by a factor of 0.1 every 500 epochs. The momentum is set to 0.9 and the weight decay to 0.0001. We use a batch size $m = 128$ and train on a single NVIDIA RTX 2080 TI GPU for 10 hours.

TinyImageNet training Following the experimental settings of PuzzleMix [18], we train MultiMix and its variants using SGD for 1200 epochs, using the same random seed as AlignMixup. We set the initial learning rate to 0.1 and decay it by a factor of 0.1 after 600 and 900 epochs. The momentum is set to 0.9 and the weight decay to 0.0001. We train on two NVIDIA RTX 2080 TI GPUs for 18 hours.

ImageNet training Following the experimental settings of PuzzleMix [18], we train MultiMix and its variants using the same random seed as AlignMixup. We train R-50 using SGD with momentum 0.9 and weight decay 0.0001 and ViT-S/16 using AdamW with default parameters. The initial learning rate is set to 0.1 and 0.01, respectively. We decay the learning rate by 0.1 at 100 and 200 epochs. We train on 32 NVIDIA V100 GPUs for 20 hours.

Tasks and metrics We use top-1 error (%, lower is better) or top-1 accuracy (%, higher is better) as evaluation metric on *image classification* and *robustness to adversarial attacks* (subsection 4.2 and subsection A.2). Additional datasets and metrics are reported separately for *transfer learning to object detection* (subsection 4.3) and *out-of-distribution detection* (subsection A.3).

A.2 More results: Classification and robustness

Using the experimental settings of subsection A.1, we extend Table 1 and Table 3 of subsection 4.2 by comparing MultiMix and its variants with additional mixup methods in Table 4 and Table 5. The additional methods are Input mixup [55], Cutmix [52], SaliencyMix [38], StyleMix [16], StyleCut-Mix [16] and SuperMix [4]. We reproduce SuperMix using the official code[1], which first trains the teacher network using clean examples and then the student using mixed. For fair comparison, we use the same network as the teacher and student models.

In Table 4 and Table 5, we observe that MultiMix and its variants outperform all the additional mixup methods on image classification. Furthermore, they are more robust to FGSM and PGD attacks as compared to these additional methods. The remaining observations in subsection 4.2 are still valid.

DATASET NETWORK	CIFAR-10 R-18	W16-8	CIFAR-100 R-18	W16-8	TI R-18
Baseline[†]	5.19	5.11	23.24	20.63	43.40
Input mixup [55][†]	4.03	3.98	20.21	19.88	43.48
CutMix [52][†]	3.27	3.54	19.37	19.71	43.11
Manifold mixup [42][†]	2.95	3.56	19.80	19.23	40.76
PuzzleMix [18][†]	2.93	2.99	20.01	19.25	36.52
Co-Mixup [17][†]	2.89	3.04	19.81	19.57	35.85
SaliencyMix [38][†]	2.99	3.53	19.69	19.59	33.81
StyleMix [16][†]	3.76	3.89	20.04	20.45	36.13
StyleCutMix [16][†]	3.06	3.12	19.34	19.28	33.49
SuperMix [4]*[†]	4.01	3.87	20.99	20.13	38.55
AlignMixup [41][†]	2.95	3.09	18.29	18.77	33.13
MultiMix (ours)	2.97	2.92	18.19	18.57	32.78
+ distil	2.87	2.78	17.72	17.91	31.93
+ dense	2.91	2.89	18.12	18.20	31.54
+ dense + distil	2.81	2.78	17.48	17.66	30.87
Gain	+0.08	+0.21	+0.81	+1.11	+2.26

(a) *Image classification* top-1 error (%) on CIFAR-10/100 and TI (TinyImagenet). R: PreActResnet, W: WRN.

NETWORK METHOD	RESNET-50 SPEED	ERROR	VIT-S/16 SPEED	ERROR
Baseline[†]	1.17	23.68	1.01	26.1
Input mixup [55][†]	1.14	22.58	0.99	25.3
CutMix [52][†]	1.16	21.40	0.99	25.6
Manifold mixup [42][†]	1.15	22.50	0.97	24.8
PuzzleMix [18][†]	0.84	21.24	0.73	24.3
Co-Mixup [17][†]	0.62	–	0.57	24.1
SaliencyMix [38][†]	1.14	21.26	0.96	24.2
StyleMix [16][†]	0.99	24.06	0.85	25.2
StyleCutMix [16][†]	0.76	22.71	0.71	24.9
SuperMix [4]*[†]	0.92	22.40	–	–
AlignMixup [41][†]	1.03	20.68	–	–
MultiMix (ours)	1.16	21.19	1.0	24.8
+ distil	1.06	19.88	0.93	23.4
+ dense	0.95	20.63	0.88	23.9
+ dense + distil	0.83	19.79	0.81	23.1
Gain		+0.89		+1.0

(b) *Image classification and training speed* on ImageNet. Top-1 error (%): lower is better. Speed: images/sec ($\times 10^3$): higher is better.

Table 4: *Image classification and training speed.* [†]: reported by AlignMixup. *: reproduced, same teacher and student model. **Bold black**: best; Blue: second best; underline: best baseline. Gain: reduction of error over best baseline.

ATTACK	FGSM					PGD			
DATASET NETWORK	CIFAR-10 R-18	W16-8	CIFAR-100 R-18	W16-8	TI R-18	CIFAR-10 R-18	W16-8	CIFAR-100 R-18	W16-8
Baseline[†]	89.41	88.02	87.12	72.81	91.85	99.99	99.94	99.97	99.99
Input mixup [55][†]	78.42	79.21	81.30	67.33	88.68	99.77	99.43	99.96	99.37
CutMix [52][†]	77.72	78.33	86.96	60.16	88.68	99.82	98.10	98.67	97.98
Manifold mixup [42][†]	77.63	76.11	80.29	56.45	89.25	97.22	98.49	99.66	98.43
PuzzleMix [18][†]	57.11	60.73	78.70	57.77	83.91	97.73	97.00	96.42	95.28
Co-Mixup [17][†]	60.19	58.93	77.61	56.59	–	97.59	96.19	95.35	94.23
SaliencyMix [38][†]	57.43	68.10	77.79	58.10	81.16	97.51	97.04	95.68	93.76
StyleMix [16][†]	79.54	71.05	80.54	67.94	84.93	98.23	97.46	98.39	98.24
StyleCutMix [16][†]	58.79	56.12	77.49	56.83	80.59	97.87	96.70	91.88	93.78
SuperMix [4]*[†]	59.98	58.10	78.75	58.19	81.03	97.65	97.20	91.51	92.73
AlignMixup [41][†]	54.83	56.20	74.18	55.05	78.83	95.42	96.71	90.40	92.16
MultiMix (ours)	54.19	55.39	75.84	54.58	77.51	94.27	94.83	90.02	91.68
+ distillation	52.55	51.42	73.55	52.77	76.20	92.69	93.90	88.87	90.54
+ dense	54.10	53.33	74.48	53.01	75.57	92.99	92.68	88.60	90.90
+ dense + distillation	52.07	50.17	72.98	52.19	75.18	90.82	90.56	87.58	90.18
Gain	+2.76	+5.95	+1.20	+2.86	+3.65	+4.60	+5.63	+2.82	+1.98

Table 5: *Robustness to FGSM & PGD attacks.* Top-1 error (%): lower is better. [†]: reported by AlignMixup. *: reproduced, same teacher and student model. **Bold black**: best; Blue: second best; underline: best baseline. Gain: reduction of error over best baseline. TI: TinyImagenet. R: PreActResnet, W: WRN.

TASK	OUT-OF-DISTRIBUTION DETECTION											
DATASET	LSUN (CROP)				iSUN				TI (CROP)			
METRIC	DET ACC	AUROC	AUPR (ID)	AUPR (OOD)	DET ACC	AUROC	AUPR (ID)	AUPR (OOD)	DET ACC	AUROC	AUPR (ID)	AUPR (OOD)
Baseline[†]	54.0	47.1	54.5	45.6	66.5	72.3	74.5	69.2	61.2	64.8	67.8	60.6
Input mixup [55][†]	57.5	59.3	61.4	55.2	59.6	63.0	60.2	63.4	58.7	62.8	63.0	62.1
Cutmix [52][†]	63.8	63.1	61.9	63.4	67.0	76.3	81.0	77.7	70.4	84.3	87.1	80.6
Manifold mixup [42][†]	58.9	60.3	57.8	59.5	64.7	73.1	80.7	76.0	67.4	69.9	69.3	70.5
PuzzleMix [18][†]	64.3	69.1	80.6	73.7	73.9	77.2	79.3	71.1	71.8	76.2	78.2	81.9
Co-Mixup [17][†]	70.4	75.6	82.3	70.3	68.6	80.1	82.5	75.4	71.5	84.8	86.1	80.5
SaliencyMix [38][†]	68.5	79.7	82.2	64.4	65.6	76.9	78.3	79.8	73.3	83.7	87.0	82.0
StyleMix [16][†]	62.3	64.2	70.9	63.9	61.6	68.4	67.6	60.3	67.8	73.9	71.5	78.4
StyleCutMix [16][†]	70.8	78.6	83.7	74.9	70.6	82.4	83.7	76.5	75.3	82.6	82.9	78.4
SuperMix [4][*]	70.9	77.4	80.1	72.3	71.0	76.8	79.6	76.7	75.1	82.8	82.5	78.6
AlignMixup [41][†]	74.2	79.9	84.1	75.1	72.8	83.2	84.1	80.3	77.2	85.0	87.8	85.0
MultiMix (ours)	79.2	82.6	85.2	77.6	75.6	85.1	87.8	83.1	78.3	86.6	89.0	88.2
+ distillation	80.3	84.4	86.3	76.4	79.0	85.6	88.2	84.9	80.7	87.8	89.9	88.2
+ dense	80.8	84.3	85.9	78.0	76.8	85.4	88.0	84.6	81.4	89.0	90.8	88.0
+ dense + distillation	81.0	84.9	86.4	78.2	79.2	86.0	88.5	84.8	81.9	89.3	90.3	88.3
Gain	+6.8	+5.0	+2.3	+3.1	+5.3	+2.8	+4.4	+4.6	+4.7	+4.3	+3.0	+3.3

Table 6: *Out-of-distribution detection* using R-18. Det Acc (detection accuracy), AuROC, AuPR (ID) and AuPR (OOD): higher is better. [†]: reported by AlignMixup. [*]: reproduced, same teacher and student model. **Bold black**: best; Blue: second best; underline: best baseline. Gain: increase in performance. TI: TinyImagenet.

A.3 More results: Out of distribution detection

This is a standard benchmark for evaluating over-confidence. Here, *in-distribution* (ID) are examples on which the network has been trained, and *out-of-distribution* (OOD) are examples drawn from any other distribution. Given a mixture of ID and OOD examples, the network should predict an ID example with high confidence and an OOD example with low confidence, *i.e.*, the confidence of the predicted class should be below a certain threshold.

Following AlignMixup [41], we compare MultiMix and its variants with SoTA methods trained using R-18 on CIFAR-100 as ID examples, while using LSUN [51], iSUN [45] and TI to draw OOD examples. We use detection accuracy, Area under ROC curve (AuROC) and Area under precision-recall curve (AuPR) as evaluation metrics. In Table 6, we observe that MultiMix and its variants outperform SoTA on all datasets and metrics by a large margin. Although the gain of vanilla MultiMix and Dense MultiMix over SoTA mixup methods is small on image classification, these variants significantly reduce over-confident incorrect predictions and achieve superior performance on out-of-distribution detection.

A.4 More ablations

As in subsection 4.5, all ablations here are performed using R-18 on CIFAR-100.

Mixup methods with distillation In subsection 4.2 and Table 4, we observe that distillation significantly improves the performance when used with MultiMix. Here, we also study its effect when applied to SoTA mixup methods.

Given a mini-batch of m examples with inputs X and targets Y, we obtain the augmented views V and V' as discussed in subsection 3.3. We then follow the mixup strategy of each mixup method and obtain the corresponding predicted class probabilities $\widetilde{P}, \widetilde{P}'$ from the student and teacher classifier, respectively. *E.g.*, for manifold mixup [42], we interpolate the embeddings $Z = f_\theta(V), Z' = f_{\theta'}(V')$ using (3) and obtain $\widetilde{P} = g_W(\widetilde{Z})$ and $\widetilde{P}' = g_{W'}(\widetilde{Z}')$. In each case, we obtain the interpolated targets $\widetilde{Y}$ using (4) and train the student network using (9).

---

[1]https://github.com/alldbi/SuperMix

METHOD	VANILLA	+ DISTIL	+ DENSE	+ DENSE + DISTIL
Baseline	76.76	78.28	78.16	79.07
Input mixup [55]	79.79	80.19	80.21	80.54
CutMix [52]	80.63	81.51	81.40	81.61
Manifold mixup [42]	80.20	81.32	80.87	81.47
PuzzleMix [18]	79.99	81.26	80.62	81.44
Co-Mixup [17]	80.19	81.39	80.84	81.69
SaliencyMix [38]	80.31	81.57	81.21	81.73
StyleMix [16]	79.96	81.22	80.76	81.30
StyleCutMix [16]	80.66	81.60	81.41	81.75
SuperMix [4]*	79.01	80.83	80.12	80.83
AlignMixup [41]	81.71	81.80	81.36	81.40
MultiMix (ours)‡	–	–	81.84	82.30
MultiMix (ours)	**81.81**	**82.28**	**81.88**	**82.52**

Table 7: *Image classification* on CIFAR-100 using R-18: The effect of distillation, dense loss and both on SoTA mixup methods. Top-1 accuracy (%): higher is better. *: 'vanilla' refers to teacher pre-training and 'distil' to self-distillation where teacher and student are trained concurrently from scratch. ‡: Instead of dense MultiMix, we only apply the loss densely.

METHOD	u	h	–	+DISTIL
Uniform	–	–	81.33	81.59
Attention (10)	CAM	softmax	81.21	81.45
	CAM	$\ell_1 \circ \mathrm{relu}$	81.63	81.91
	GAP	softmax	81.78	82.01
	GAP	$\ell_1 \circ \mathrm{relu}$	**81.88**	**82.52**

Table 8: *Variants of spatial attention* in dense MultiMix, with and without distillation, on CIFAR-100 using R-18. Top-1 accuracy (%): higher is better. GAP: Global Average Pooling; CAM: Class Activation Maps [58]; $\ell_1 \circ \mathrm{relu}$: ReLU followed by ℓ_1 normalization.

In Table 7, we observe that with distillation, the performance of all SoTA mixup methods improve. For example, the baseline improves by 1.52% accuracy (76.76 → 78.23) and manifold mixup by 1.12% (80.20 → 81.32). On average, we observe a gain of 1% brought by distillation. An exception is AlignMixup [41]: distillation brings a marginal improvement of 0.09% (81.71 → 81.80), making it on-par with vanilla MultiMix.

Mixup methods with dense loss In Table 4 we observe that dense interpolation and dense loss improve vanilla MultiMix. Here, we study the effect of the dense loss when applied to SoTA mixup methods.

Given a mini-batch of m examples, we follow the mixup strategy of the SoTA mixup methods to obtain the mixed embedding $\widetilde{Z}^j \in \mathbb{R}^{d \times m}$ for each spatial position $j = 1, \ldots, r$. Then, as discussed in subsection 3.4, we obtain the predicted class probabilities $\widetilde{P}^j \in \mathbb{R}^{c \times m}$ again for each $j = 1, \ldots, r$. Finally, we compute the cross-entropy loss $H(\widetilde{Y}, \widetilde{P}^j)$ (1) densely at each spatial position j, where the interpolated target label $\widetilde{Y} \in \mathbb{R}^{c \times m}$ is given by (4).

In Table 7, we observe that using a dense loss improves the performance of all SoTA mixup methods. The baseline improves by 1.4% accuracy (76.76 → 78.16) and manifold mixup by 0.67% (80.20 → 80.87). On average, we observe a gain of 0.7% brought by the dense loss. An exception is AlignMixup [41], which drops by 0.35% (81.71 → 81.36). This may be due to the alignment process, whereby the interpolated dense embeddings are not very far from the original.

Finally, we study the effect of using a dense distillation loss on SoTA mixup methods. Here, similarly with (9), the loss has two terms for each spatial position j: the first is the dense cross-entropy loss $H(\widetilde{Y}, \widetilde{P}^j)$ as above and the second is the dense distillation loss $H((\widetilde{P}')^j, \widetilde{P}^j)$, where $\widetilde{P}'$ is obtained by the teacher. In Table 7, we observe that dense distillation further improves the performance of SoTA mixup methods as compared to using the dense loss only.

4

Two-stage distillation Following SuperMix [4], we also study the effect of using a two-stage distillation process with MultiMix, rather than online self-distillation.

In the first stage, we train the teacher using only clean examples for 300 epochs, and we achieve a top-1 accuracy of 75.62%. This is slightly lower than the 76.76% of the baseline from Table 1(a), which is trained for 2000 epochs. In the second stage, we fix the teacher parameters and train the student using the predictions from the teacher network as targets. In particular, we use the second term $H(\widetilde{P}', \widetilde{P})$ of (9), that is, $\gamma = 0$. At inference, the top-1 accuracy drops by 16% ($75.62 \rightarrow 59.77$). This shows that using the setting of SuperMix is not effective, while also being computationally expensive because of the two-stage training.

We also study the effect of training the student with both the interpolated labels $\widetilde{Y}$ (7) and the interpolated predictions $\widetilde{P}'$ of the pretrained teacher as targets. In particular, we use (9) with our default $\gamma = \frac{1}{2}$. At inference, the top-1 accuracy improves by 4.7% compared with the teacher ($75.62 \rightarrow 80.35$). However, the student accuracy of 80.35% is still inferior to our 82.28% by online self-distillation (Table 7). This shows that joint training of teacher and student is beneficial.
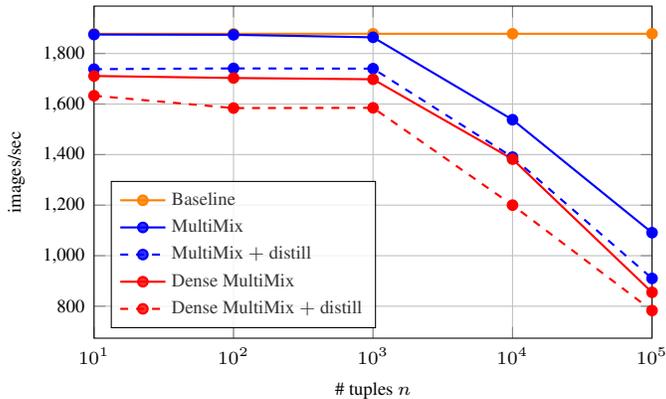
Figure 5: *Training speed* (images/sec) of MultiMix and its variants *vs*. number of tuples n on CIFAR-100 using R-18. Measured on NVIDIA RTX 2080 TI GPU, including forward and backward pass.

Training speed In Figure 5, we analyze the training speed of MultiMix and its variants as a function of number of tuples n. In terms of speed, vanilla MultiMix is on par with the baseline up to $n = 1000$, while bringing an accuracy gain of 5%. The best performing variant—dense Multi-Mix with distillation—is only slower by 15.6% at $n = 1000$ as compared to the baseline, which is arguably worth given the impressive 5.8% accuracy gain. Further increasing beyond $n > 1000$ brings a drop in training speed, due to computing Λ and then using it to interpolate (6),(7). Because $n > 1000$ also brings little performance benefit according to Figure 4(b), we set $n = 1000$ as default for all MultiMix variants.

Dense MultiMix: Spatial attention In subsection 3.4, we discuss different options for attention in dense MultiMix. In particular, no attention amounts to defining a uniform $a = \mathbf{1}_r/r$. Otherwise, a is defined by (10). The vector u can be defined as $u = \mathbf{z}\mathbf{1}_r/r$ by global average pooling (GAP) of $\mathbf{z}$, which is the default, or $u = Wy$ assuming a linear classifier with $W \in \mathbb{R}^{d \times c}$. The latter is similar to class activation mapping (CAM) [58], but here the current value of W is used online while training. The non-linearity h can be softmax or ReLU followed by ℓ_1 normalization ($\ell_1 \circ \text{relu}$), which is the default. Here, we study the affect of these options on the performance of dense Multimix.

In Table 8, we observe that using GAP for u and $\ell_1 \circ \text{relu}$ as h yields the best performance overall. Changing GAP to CAM or $\ell_1 \circ \text{relu}$ to softmax is inferior, more so in the presence of distillation. The combination of CAM with softmax is the weakest, even weaker than uniform attention. CAM may fail because of using the non-optimal value of W while training; softmax may fail because of being too selective. Compared to our best result, uniform attention is clearly inferior, by nearly 1% in the presence of distillation. This validates that the use of spatial attention in dense MultiMix is clearly beneficial. The intuition is the same as in weakly supervised tasks: In the absence of dense targets, assuming the same target of the entire example at every spatial position naively implies that

the object of interest is present everywhere, whereas spatial attention provides a better hint as to where the object may really be.

Dense MultiMix: Spatial resolution   We study the effect of spatial resolution on dense MultiMix. By default, we use a resolution of 4×4 at the last residual block of R-18 on CIFAR-100. Here, we additionally investigate 1×1 (downsampling by average pooling with kernel size 4, same as GAP), 2×2 (downsampling by average pooling with kernel size 2) and 8×8 (upsampling by using stride 1 in the last residual block). We measure accuracy 81.07% for spatial resolution 1×1, 81.43% for for 2×2, 81.88% for 4×4 and 80.83% for 8×8. We thus observe that performance improves with spatial resolution up to 4×4, which the optimal, and then drops at 8×8. This drop may be due to assuming the same target at each spatial position. The resolution 8×8 is also more expensive computationally.