

SAGE: Saliency-Guided Mixup with Optimal Rearrangements

Avery Ma^{1,2*}

ama@cs.toronto.edu

Nikita Dvornik³

n.dvornik@samsung.com

Ran Zhang³

ran.zhang@samsung.com

Leila Pishdad^{4†}

leila.pishdad@mail.mcgill.ca

Konstantinos G. Derpanis^{2,3,5}

kosta@yorku.ca

Afsaneh Fazly³

a.fazly@samsung.com

¹ University of Toronto

² Vector Institute

³ Samsung AI Centre Toronto

⁴ Borealis AI

⁵ York University

Abstract

Data augmentation is a key element for training accurate models by reducing overfitting and improving generalization. For image classification, the most popular data augmentation techniques range from simple photometric and geometrical transformations, to more complex methods that use visual saliency to craft new training examples. As augmentation methods get more complex, their ability to increase the test accuracy improves, yet, such methods become cumbersome, inefficient and lead to poor out-of-domain generalization, as we show in this paper. This motivates a new augmentation technique that allows for high accuracy gains while being simple, efficient (i.e., minimal computation overhead) and generalizable. To this end, we introduce **Saliency-Guided Mixup with Optimal Rearrangements (SAGE)**, which creates new training examples by rearranging and mixing image pairs using visual saliency as guidance. By explicitly leveraging saliency, SAGE promotes discriminative foreground objects and produces informative new images useful for training. We demonstrate on CIFAR-10 and CIFAR-100 that SAGE achieves better or comparable performance to the state of the art while being more efficient. Additionally, evaluations in the out-of-distribution setting, and few-shot learning on mini-ImageNet, show that SAGE achieves improved generalization performance without trading off robustness. Our source code is available at <https://github.com/SamsungLabs/SAGE>.

1 Introduction

Data augmentation (DA) methods synthetically expand a dataset by applying transformations on the available examples, with the goal of reducing overfitting and improving generalization

* Work done during an internship at Samsung AI Centre Toronto

† Work done while at Samsung AI Centre Toronto

© 2022. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.



Figure 1: **Comparison of data augmentation methods.** Thanks to the saliency-guided mixing and image rearrangements, SAGE produces more meaningful and informative scenes, as verified in our experiments.

in models trained on these datasets. In computer vision, conventional DA techniques are typically based on random geometric (translation, rotation and flipping) and photometric (contrast, brightness and sharpness) transformations [1, 3, 14, 18]. While these techniques are already effective, they merely create slightly altered copies of the original images and thus introduce limited diversity in the augmented dataset. A more advanced DA [15, 17] combines multiple training examples into a new image-label pair. By augmenting both the image and the label space simultaneously, such approaches greatly increase the diversity of the augmented set. Consequently, they substantially improve model generalization, without any efficiency overhead, due to their simplicity. Nonetheless, these DA approaches are agnostic to image semantics; they ignore object location cues, and as a result may produce ambiguous scenes with occluded distinctive regions (see Figure 1, Mixup [17] and CutMix [15]).

To account for such shortcomings, a new line of work [8, 11, 12, 21] proposes to explicitly use visual saliency [19] for data augmentation. Typically, a saliency map contains the information about the importance of different image regions for the downstream task. As a result, saliency maps implicitly contain information about objects, their locations and, crucially, about the “informativeness” of image regions. Previous methods [11, 12, 21] take full advantage of the saliency information, and formulate data augmentation as a saliency maximization problem. Given training image patches, their augmentation “assembles” a new image of high visual saliency. This approach greatly improves the test accuracy; however, this comes with a large computation overhead due to the need to maximize saliency at every training step. Moreover, as the augmented images are composed of patches, the resulting scenes are often unrealistic (see Puzzle Mix, Co-Mixup and SaliencyMix in Figure 1), which leads to poor out-of-distribution generalization, as shown later in our experiments. In summary, the existing data augmentation techniques can *either* i) boost the test accuracy, *or* ii) produce a robust model with little computational overhead; there are no methods that can do both.

To address the aforementioned drawbacks, we propose a new augmentation – **Saliency-Guided Mixup with Optimal Rearrangements (SAGE)** – that provides both high accuracy and robustness, and has minimal computation overhead. SAGE is a simple and effective DA technique that uses visual saliency to perform optimal image blending at each spatial location, and optimizes the relative image position such that the resulting visual saliency is maximized. Given two images and their saliency maps, SAGE mixes the images together, such that at each spatial location, the contribution of different images to the mix is proportional to their saliency in that location. The corresponding label is also obtained by interpolating the original labels based on the saliency of the corresponding images. To maximize the resulting saliency of the mix, we find an optimal relative arrangement of the two images prior to the mixing stage. As a

result, SAGE produces smooth and realistic images with clear and distinct foreground objects (see Figure 1), unlike other augmentation techniques. Thanks to our efficient implementation, SAGE has virtually no computation overhead beyond obtaining the saliency information. Furthermore, our computations are partially shared between the saliency masks and the training gradients, which further decreases the amortized training time.

Contributions. We make the following three contributions: (i) We introduce SAGE, a DA method to generate novel training examples by mixing image pairs based on their visual saliency, which promotes discriminative foreground objects in the mix. (ii) SAGE achieves test accuracy better than or comparable to state-of-the-art augmentation techniques, without incurring significant computation overhead. (iii) Through robustness evaluations on perturbed test data, we show that SAGE improves test accuracy without trading off robustness.

2 Related Work

In this section, we review data augmentation techniques that go beyond simple geometrical and color transformations to improve generalization. A popular approach is to synthesize new training input-output pairs by combining information from multiple raw samples. Mixup [27] creates a new image-label pair by linearly interpolating both the input and output space. In contrast, Manifold Mixup [22] and HypMix [16] apply interpolation at the feature level. Others create new training samples by “copy-pasting” patches from one image to another [6, 7, 25]. This class of methods is very efficient and simple to implement. However, a common drawback of these approaches is that they do not take image semantics into account when performing augmentation. This potentially encourages the model to generalize using completely irrelevant information from the new training data, leading to inferior generalization.

To address this problem, recent work explicitly uses visual saliency information in the DA process. KeepAugment [8] leverages input saliency to improve existing DA techniques, e.g., Cutout [9], by always keeping the important regions untouched during augmentation. SaliencyMix [24] improves CutMix [25] by selecting a patch around the peak salient pixel location in the source image and mixing it with the target image. Puzzle Mix formulates DA as an optimization problem, where the objective balances saliency maximization, local smoothness and the optimal transport between data pairs [11]. Co-Mixup [12] extends this idea by encouraging the diversity of the augmentation when mixing a collection of inputs, and thus further complicates the optimization objective. The need to solve the optimization problem at every step significantly slows down the training, which may be prohibitive in some situations. Our saliency-guided method not only reduces this computational overhead, but also generates more plausible augmented images that result in improved test accuracy and out-of-distribution generalization.

3 Technical Approach

The main idea behind SAGE is to synthesize novel images (with their labels) by blending pairs of training samples, using spatial saliency information as guidance for optimal blending. As illustrated in Fig. 2, our method consists of three independent components: i) saliency mask generation (Sec. 3.1), ii) the “Optimal Rearrangement” module (Sec. 3.3), and iii) the “Saliency-guided Mixup” module (Sec 3.2). All chained together, they form our SAGE approach. Below, we elaborate on each of the components and conclude with a discussion on the efficiency of our pipeline in Sec. 3.4.

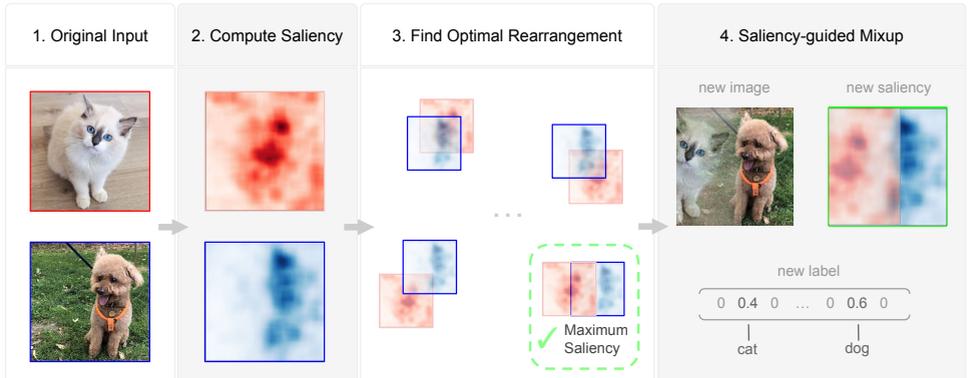


Figure 2: **SAGE overview.** Given the original images, we first compute saliency maps. Next, we find the best rearrangement of the images that maximizes the total saliency (in the green box). Finally, we use our saliency-guided Mixup to fuse the overlapping image parts and derive the new label. As a result, SAGE produces smooth, realistic and informative scenes.

3.1 Computing Saliency Maps

We define the saliency of each image pixel as its importance in making the correct prediction, using a given vision model. More formally, we are given a training sample, (x, y) , where $x \in \mathbb{R}^{d \times d \times 3}$ is an RGB image and $y \in \mathbb{R}^C$ is the corresponding one-hot label vector, a classifier, $f_\theta(\cdot)$, that is the current partially trained model, and our task loss, $\ell(f_\theta(x), y)$, measuring the discrepancy between the classifier’s output and the true label. We define the saliency, $s \in \mathbb{R}^{d \times d}$, as the magnitude of the gradient with respect to the input image,

$$s(x) = |\nabla_x \ell(f_\theta(x), y)|_{l_{2,3}}, \quad (1)$$

where $|\cdot|_{l_{2,3}}$ denotes the l_2 -norm along the third (color) dimension. In practice, the saliency map tends to focus on the foreground objects useful for classification and ignores irrelevant background. Note that our saliency definition differs from others [17, 19] in that we consider the gradient of the full loss, while previous work consider the gradient of the ground-truth class activation with respect to the input image. We find that our definition is advantageous for data augmentation, and additionally allows for more efficient training, as detailed in Sec. 3.4.

3.2 Saliency-guided Mixup

Before describing our Saliency-guided Mixup, we revisit the original Mixup [17]. Mixup creates a new training sample, x' , by linearly mixing pairs of training samples, x_0 and x_1 , i.e., $x' = \lambda \cdot x_0 + (1 - \lambda) \cdot x_1$, and their corresponding labels, i.e., $y' = \lambda \cdot y_0 + (1 - \lambda) \cdot y_1$, where $\lambda \in [0, 1]$. While simple and effective, Mixup has a notable drawback, namely it ignores the image semantics. That is, at every pixel location, the contribution of x_0 and x_1 to the final image is constant. As Fig. 3 (e) shows, this may lead to prominent image regions being suppressed by the background, which is not ideal for data augmentation [17, 18].

To address this shortcoming, we propose Saliency-guided Mixup, where at every image location in x' , the mixing ratio between x_0 and x_1 is different, defined by the saliency of the corresponding image regions. More formally, given two images, x_0 and x_1 , and their saliency

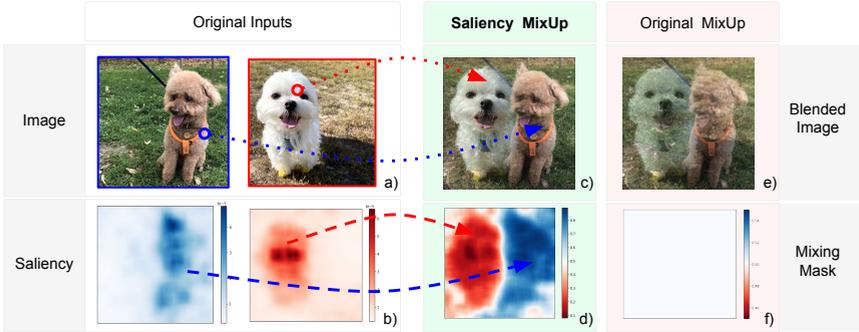


Figure 3: **Comparison between Saliency-guided Mixup and original Mixup.** Given the a) original images with b) saliency maps, our Saliency Mixup computes d) the Mixing Mask M (given by Eq. 2) based on the relative saliency of the inputs. The values of M are represented with a heatmap; blue areas indicate stronger contribution of image 1, red areas correspond to image 2 being more prominent and pale areas indicate more uniform blending. Consequently, salient regions from different images contribute to different locations and result in a realistic, informative output c). In contrast, the original Mixup produces f) a uniform mixing mask (at $\lambda = 0.5$), which results in e) an unrealistic and unclear image.

maps, s_0 and s_1 , we craft a 2D mixing mask, $M \in \mathbb{R}^{d \times d}$, and use it to mix the images:

$$x' = M \odot x_0 + (1 - M) \odot x_1, ; M = \frac{\tilde{s}_0}{\tilde{s}_0 + \tilde{s}_1 + \zeta}, \quad (2)$$

where $x' \in \mathbb{R}^{d \times d \times 3}$, \tilde{s}_0 and \tilde{s}_1 are spatially-normalized and Gaussian-smoothed saliency maps, ζ is a scalar hyperparameter used to avoid division-by-zero and \odot denotes element-wise product. That is, the elements in M are defined as the saliency ratio in different images at the same location. This means that, at any given location, more prominent regions of one image will suppress less salient regions of the other image in the final blend, x' . This strategy largely resolves the issue with the original Mixup and leads to more informative augmentation (see Fig 3 (e)). Lastly, we mix the labels using $y' = \gamma \cdot y_0 + (1 - \gamma) \cdot y_1$, where γ is the mean of the mixing mask, M .

Saliency-guided Mixup, Eq. 2, is most suitable for mixing images that have salient regions in distinct locations. When the maximally salient regions in both images spatially overlap, the mask, M , tends to suppress one or both objects, which leads to uninformative new scenes.

3.3 Optimal Rearrangements via Saliency Maximization

To produce highly-informative augmentations with Eq. 2, even when both images have overlapping salient regions, we propose to shift one image relative to the other prior to mixing. Our objective is to find the shift that maximizes the resulting image saliency. An example of such rearrangements with the resulting augmentations are shown in Fig. 4. In the following, we formalize this shifting process and describe a solution for finding the best rearrangement.

We define the translation operator that shifts a tensor z by $\tau = (\tau_i, \tau_j)$ pixels as

$$\mathcal{T}(z, \tau)[i, j] = \begin{cases} z[i - \tau_i, j - \tau_j], & \text{if } i - \tau_i \in [0, d - 1], j - \tau_j \in [0, d - 1] \\ 0, & \text{otherwise} \end{cases}, \quad (3)$$

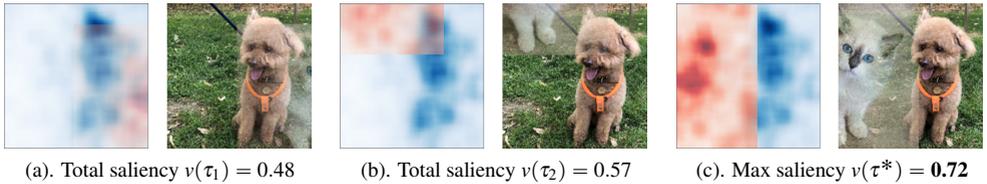


Figure 4: **Possible rearrangements.** In each example, the saliency map corresponding to the rearrangement is shown on the left, the corresponding image (after applying Saliency-guided Mixup) is on the right. The rearrangement maximizing the total saliency is shown in c); clearly, it results in a denser mixed saliency, and produces a more informative image.

where $z[i, j]$ is the value of z at the location (i, j) . Essentially, translation \mathcal{T} shifts all the values in the tensor by the given offset, τ , and zero-pads the empty space.

To quantify how successful a given rearrangement is in resolving the saliency overlap, we measure the total saliency $\mathbb{[\square]}$ after the rearrangement. For a given rearrangement, τ , the total saliency, $v(\tau) \in \mathbb{R}$, is defined as follows:

$$v(\tau) = \sum_{i,j} [M^\tau \odot \tilde{s}_0 + (1 - M^\tau) \odot \mathcal{T}(\tilde{s}_1, \tau)], \quad (4)$$

where $\mathcal{T}(\tilde{s}_1, \tau)$ is the saliency \tilde{s}_1 translated by τ and M^τ is the mixing mask (Eq. 2) computed with \tilde{s}_0 and $\mathcal{T}(\tilde{s}_1, \tau)$. Essentially, the scalar $v(\tau)$ captures the total saliency after the rearrangement (Eq. 3) and fusion (Eq. 4) of the individual saliency tensors. Intuitively, larger total saliency values imply smaller overlap between the salient regions in the shifted images, x_0 and $\mathcal{T}(x_1, \tau)$, and suggests that the resulting mix is more informative. Thus, it is reasonable to look for a rearrangement that maximizes the total saliency. To this end, we propose to find the optimal rearrangement (offset), τ^* , by solving the following: $\tau^* = \operatorname{argmax}_{\tau \in \mathcal{O}} v(\tau)$, where \mathcal{O} is the space of all possible offsets (shown in Fig. 2, step 3).

Finally, we use the obtained optimal rearrangement to generate the augmented sample, x' . This is done by applying our Saliency-guided Mixup to the rearranged image pair (shown in Fig. 2, step 4), i.e., simply plugging the images x_0 and $\mathcal{T}(x_1)$ with the corresponding saliency \tilde{s}_0 and $\mathcal{T}(\tilde{s}_1)$ into Eq. 2. The exact data augmentation algorithm is detailed in the supplement.

3.4 Discussion

One of the advantages of SAGE over other saliency-based augmentations (e.g., $\mathbb{[\square]}$, $\mathbb{[\square]}$) is its efficiency. Here, we elaborate on our pipeline design choices and discuss their complexity.

Saliency-guided Mixup. Compared to the original Mixup blending step, our Saliency-guided Mixup (Sec 3.2) adds a simple element-wise multiplication by the mixing mask. The cost of this operation is negligible to our model’s runtime.

Optimal Rearrangements. As described in Sec. 3.3, to arrive at our final mixture, we consider all possible rearrangements and select the one maximizing the total saliency, Eq. 4. The number of rearrangements grows quadratically with image size and soon becomes the bottleneck. To keep our method efficient, we randomly sample a small portion of all possible arrangements (1% in all experiments), and search among them. In our experiments, this does not affect classification performance, while greatly improving efficiency.

Saliency Computation. Computing saliency requires an extra forward and backward pass of the model. When the existing works $\mathbb{[\square]}$, $\mathbb{[\square]}$ compute saliency masks, they discard all the

Dataset	Model	Vanilla	Mixup	CutMix	Manifold	SaliencyMix	Puzzle Mix	Co-Mixup	SAGE
CIFAR-10	PreActResNet18	95.07	95.97	96.27	96.28	96.15	<u>96.62</u>	96.23	96.95
CIFAR-100	PreActResNet18	76.80	77.40	78.96	78.51	78.85	79.65	<u>79.68</u>	79.91
CIFAR-100	WRN16	78.55	79.83	80.03	79.77	80.16	80.73	80.42	<u>80.45</u>
CIFAR-100	ResNext29	78.77	78.23	77.43	77.97	78.89	79.20	<u>80.27</u>	80.35

Table 1: **Image classification accuracy.** CIFAR-10 and CIFAR-100 results are obtained by averaging over three independent training runs. The best numbers are in bold and the second best numbers are underlined.

intermediate computations and only use the mask itself for DA, which essentially doubles the training time. In contrast, SAGE saves the gradients, g_s , with respect to the model parameters, obtained in the backward pass of saliency computations. These gradients can be combined with the standard gradients, g_a , computed on SAGE-augmented images to perform the final model update with $g = \eta \cdot g_s + (1 - \eta) \cdot g_a$, where $\eta \in [0, 1]$. The hyperparameter, η , effectively controls how much information from the original images is used for updates versus that of the augmented images. This trick allows us to amortize the saliency computations, and reuse the intermediate results for the model updates. Note that this is only possible thanks to our saliency definition (Eq. 1), which differs from the classical one [19].

4 Experiments

We demonstrate the advantage of SAGE in image classification in Sec. 4.1. Sec. 4.2 evaluates SAGE in out-of-distribution generalization, Sec. 4.3 analyzes the efficiency of our pipeline and Sec. 4.4 presents an ablation study of SAGE’s components. Our implementation is largely based on the publicly available repository of Puzzle Mix¹.

4.1 Image Classification

Following previous work [20], we perform evaluations on the CIFAR-10 [23] and CIFAR-100 [23] datasets with the PreActResNet18 [20], ResNext29 [24] and WideResNet16 [26] architectures. For all datasets and models, we follow the optimization schedule described in Puzzle Mix and Co-Mix; training and model details are included in the supplement. For a comprehensive comparison, we use the following DA baselines: (i) Vanilla, i.e., standard data augmentation only, which includes random cropping and horizontal flips, (ii) Mixup [22], (iii) CutMix [25], (iv) Manifold [22], (v) SaliencyMix [22], (vi) Puzzle Mix [22] and (vii) Co-Mixup [22]. Note that all the baseline methods are applied on top of the standard data augmentation. Following previous work [21, 22], we report the results averaged over three independent training runs.

Table 1 summarizes the comparison of SAGE to the baselines, pointing to two key observations. First, the DA techniques utilizing saliency (i.e., SaliencyMix, Puzzle Mix, Co-Mixup and SAGE) substantially outperform other non-saliency-based variants across almost all datasets and architectures. This clear improvement demonstrates that using image semantics for data augmentation leads to better generalization on the test set. Second, among saliency-based methods, SAGE is consistently the best on CIFAR-10; on CIFAR-100,

¹<https://github.com/snu-mlab/PuzzleMix>

Vanilla	Mixup	CutMix	SaliencyMix	Puzzle Mix	Co-Mixup	SAGE
77.9	78.9	78.4	78.6	78.6	79.0	79.8

Table 2: **Few-shot classification accuracy on mini-ImageNet.**

SAGE outperforms Puzzle Mix and Co-Mixup on PreActResNet18 and ResNext29, and has comparable performance on WideResNet. SAGE also outperforms SaliencyMix on all tested architectures on both datasets. We attribute the advantage of SAGE to the fact that our augmented images are smoother and more realistic, combining the advantages of Mixup and the saliency-based methods. This is despite the fact that Puzzle Mix and Co-Mixup are explicitly optimizing for maximum saliency, and have considerably more computational overhead.

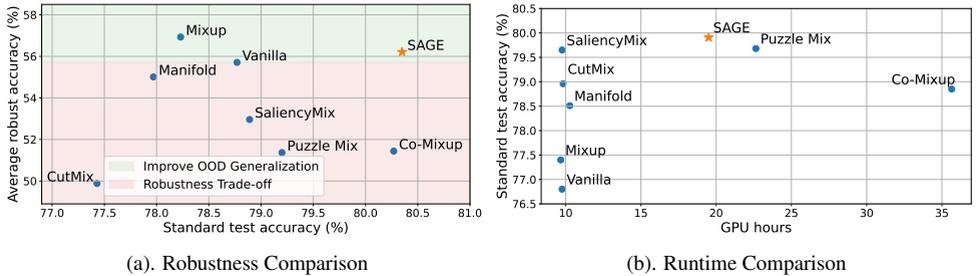
4.2 Out-of-distribution Generalization and Few-shot Adaptation

It is known that different DA techniques may lead to similar test accuracy improvements but have drastically different behavior on out-of-distribution (OOD) data [22]. This phenomenon is attributed to the difference in the quality of the learned representation. Therefore, to further evaluate our approach, we consider generalization in the OOD setting.

In our evaluation, we test the OOD generalization in two scenarios: using corrupted test images (with Gaussian noise or adversarial perturbations [24]) or evaluating generalization to new categories in a few-shot setup [23]. More specifically, we test against three different perturbations: i) Gaussian noise with zero mean and variance of 0.01, ii) ℓ_∞ -norm bounded attack generated using the Fast Gradient Sign Method (FGSM) [9] with $\epsilon = \frac{8}{255}$ and iii) ℓ_2 -norm bounded attack crafted with Fast Gradient Method (FGM) [9] with $\epsilon = 0.5$. Our choice of the attacks and ϵ follows the standard practice used with the robustness benchmarks [1]. To evaluate few-shot adaptation capabilities of our model and test how well the learned representations transfer to novel categories, we perform few-shot classification on the mini-ImageNet dataset [23]. Additional details are provided in the supplement.

To summarize the performance on all three OOD benchmarks, we average the accuracy across the benchmarks, and get a single score quantifying model robustness. Figure 5(a) plots the average OOD accuracy on CIFAR-100, against the standard accuracy on the original test set. We observe a striking difference in the robustness characteristics across different DA methods. Notably, models trained using SAGE are much less sensitive to out-of-distribution shifts compared to the two other saliency-based methods, i.e., Puzzle Mix and Co-Mixup, despite comparable test accuracy improvements. Moreover, the models trained with CutMix, Puzzle Mix and Co-Mixup have worse OOD performance compared to Vanilla training. These methods produce augmentations with unnatural patch-like patterns, which likely leads to unwanted properties of the learned representations. In contrast, Mixup and SAGE fuse images in a homogeneous way, leading to models more robust to various input perturbations. Please refer to the supplement for the full table of results and CIFAR-10 experiments.

To show OOD generalization beyond adversarial attacks, we compare SAGE to other data augmentation techniques for few-shot classification on mini-ImageNet, where the goal is to learn a representation that generalizes to novel categories. We follow the setup from previous work [5], using a single ResNet12 with the prototype classifier. As shown in Table 2, SAGE outperforms other augmentation techniques, including Mixup (the strongest model



(a). Robustness Comparison

(b). Runtime Comparison

Figure 5: **Robustness and efficiency analysis of SAGE.** (a) Robustness versus standard accuracy in OOD generalization. The methods in the green area (i.e., Mixup and SAGE) improve both accuracy and robustness relative to vanilla augmentation, while the others in red (i.e., CutMix, Co-Mixup and Puzzle Mix) improve standard test accuracy at the cost of decreased robustness. (b) Runtime comparison of SAGE and other baselines. We estimate computation cost with a single NVIDIA Tesla T4. For SAGE, there is no noticeable overhead besides the additional forward and backward pass to compute the saliency map which approximately doubles the time of Vanilla training.

on adversarial perturbations). This shows that SAGE is useful in OOD scenarios beyond Gaussian and adversarial perturbations.

4.3 Runtime Analysis

In this section, we compare the training time of different data augmentation methods running on a single NVIDIA Tesla T4. Figure 5(b) plots each method’s average training time (GPU hours) versus accuracy. Notably, the techniques not using saliency (i.e., Mixup, Manifold and CutMix) are as fast as Vanilla, since the data augmentation is performed during data loading, which does not affect the overall training time. SaliencyMix stands apart from the other saliency-based augmentation techniques. This follows because it utilizes an external trained saliency detector based on a shallow pre-deep learning method [18], that is fast but considerably less capable than the deep saliency methods [19] used for the other augmentation techniques. Consequently, SaliencyMix introduces minimal overhead; however, its improvement on classification accuracy is limited. Other saliency-based methods (i.e., PuzzleMix, Co-Mixup and SAGE) are more accurate, yet also significantly slower. Among them, SAGE is the fastest and also the most accurate on CIFAR-100. Based on these observations, we argue that SAGE represents a good trade-off between accuracy and efficiency overall, and is clearly the best choice among the saliency-based methods.

4.4 Ablation Studies

In this section, we further analyze our data augmentation strategy by ablating different design choices in the pipeline. For all the experiments, we use the same setup described in Sec. 4.1 with ResNet18 on CIFAR-10 and CIFAR-100. Please see the supplement for additional ablations.

Saliency-guided Mixup and optimal rearrangements. The two components that make SAGE novel are the Saliency-guided Mixup (Sec. 3.2) and the Optimal Rearrangements

Model	CIFAR-10	CIFAR-100	Search Space	CIFAR-10	CIFAR-100
Vanila	95.07	76.8	1%	96.95	79.91
SAGE w/o SM	96.53	78.89	10%	96.67	79.47
SAGE w/o OR	96.48	78.68	50%	96.58	79.40
SAGE	96.95	79.91	100%	96.69	79.45

Table 3: **Ablation studies of SAGE.** (left) Dissecting the benefit from saliency-guided mixing and optimal rearrangements. Here, SAGE w/o SM (without Saliency-guided Mixup), and SAGE w/o OR (without optimal rearrangements). (right) SAGE’s accuracy depending on the explored rearrangements. The first column indicates the size of the random portion of rearrangements used for data augmentation.

(Sec. 3.3). Here, we evaluate SAGE with some of the components removed or replaced by an existing technique. In particular, we evaluate i) SAGE w/o OR (i.e., without optimal rearrangements) that always performs Saliency-guided Mixup on non-shifted images and ii) SAGE w/o SM (i.e., without Saliency-guided Mixup) for mixing images together that simply replaces one image region with the other image instead of performing smooth saliency-based mixing. Examples of SAGE w/o SM and SAGE w/o OR are included in the supplement. As shown in Table 3, each of the components is important for the final performance and thus justifies their use.

Optimal rearrangements search space. As described in Sec. 3.4, to select a rearrangement, we evaluate a set of locations, and proceed with the one that maximizes saliency. To speed up the search, we only explore a random subset of all rearrangements, 1% in all previous experiments, which suggest that our data augmentation may be sub-optimal. Table 3 shows the model’s performance, depending on the portion of all rearrangements we consider for DA. Surprisingly, using only 1% of the rearrangements works best. While seemingly counterintuitive, we hypothesize the sub-optimal rearrangements act as additional training regularization and introduce more diversity in the augmented data.

5 Conclusion

We proposed SAGE – a new data augmentation approach that integrates visual saliency to produce highly informative training samples. Compared to existing methods, SAGE leads to better test accuracy, and generates more realistic training samples. Moreover, SAGE is the only saliency-based augmentation technique that improves model robustness and OOD performance, while incurring minimal computational overhead. In principle, SAGE is not limited to image classification and can be easily extended to other visual tasks. We believe that SAGE delivers a unique combination of accuracy, robustness and efficiency, and can become the new plug-and-play data augmentation for a wide range of vision tasks.

References

- [1] Francesco Croce, Maksym Andriushchenko, Vikash Sehwal, Edoardo DeBenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. RobustBench: a standardized adversarial robustness benchmark. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track (Round 2)*, 2021.
- [2] Ekin Dogus Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc Le. AutoAugment: Learning augmentation strategies from data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [3] Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. RandAugment: Practical automated data augmentation with a reduced search space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [4] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [5] Nikita Dvornik, Cordelia Schmid, and Julien Mairal. Diversity with cooperation: Ensemble methods for few-shot classification. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
- [6] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
- [7] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin Dogus Cubuk, Quoc Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [8] Chengyue Gong, Dilin Wang, Meng Li, Vikas Chandra, and Qiang Liu. KeepAugment: A simple information-preserving data augmentation approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [11] Jang-Hyun Kim, Wonho Choo, and Hyun Oh Song. Puzzle Mix: Exploiting saliency and local statistics for optimal mixup. In *International Conference on Machine Learning (ICML)*, 2020.
- [12] JangHyun Kim, Wonho Choo, Hosan Jeong, and Hyun Oh Song. Co-Mixup: Saliency guided joint mixup with supermodular diversity. In *International Conference on Learning Representations (ICLR)*, 2021.
- [13] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.

- [14] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.
- [15] Sebastian Montabone and Alvaro Soto. Human detection using a mobile platform and novel features derived from a visual saliency mechanism. In *Image and Vision Computing*, 2010.
- [16] Ramit Sawhney, Megh Thakkar, Shivam Agarwal, Di Jin, Diyi Yang, and Lucie Flek. HypMix: Hyperbolic interpolative data augmentation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [17] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.
- [18] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [19] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [20] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [21] AFM Uddin, Mst Monira, Wheemyung Shin, TaeChoong Chung, Sung-Ho Bae, et al. Saliencymix: A saliency guided data augmentation strategy for better regularization. In *International Conference on Learning Representations (ICLR)*, 2014.
- [22] 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 *International Conference on Learning Representations (ICLR)*, pages 6438–6447, 2019.
- [23] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [24] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [25] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. CutMix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [26] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference (BMVC)*, 2016.

- [27] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, 2018.