
You Only Cut Once: Boosting Data Augmentation with a Single Cut

Junlin Han^{1,2,3} Pengfei Fang^{1,2} Weihao Li¹ Jie Hong^{1,2} Mohammad Ali Armin¹
Ian Reid³ Lars Petersson¹ Hongdong Li²

Abstract

We present **You Only Cut Once (YOCO)** for performing data augmentations. YOCO cuts one image into two pieces and performs data augmentations individually within each piece. Applying YOCO improves the diversity of the augmentation per sample and encourages neural networks to recognize objects from partial information. YOCO enjoys the properties of parameter-free, easy usage, and **boosting almost all augmentations for free**. Thorough experiments are conducted to evaluate its effectiveness. We first demonstrate that YOCO can be seamlessly applied to varying data augmentations, neural network architectures, and brings performance gains on CIFAR and ImageNet classification tasks, sometimes surpassing conventional image-level augmentation by large margins. Moreover, we show YOCO benefits contrastive pre-training toward a more powerful representation that can be better transferred to multiple downstream tasks. Finally, we study a number of variants of YOCO and empirically analyze the performance for respective settings. Code is available at [GitHub](#).

1. Introduction

Deep neural networks have been widely applied to various computer vision tasks such as image classification (He et al., 2016b; Krizhevsky et al., 2012; Simonyan & Zisserman, 2014), semantic segmentation (Long et al., 2015; Chen et al., 2014), and object detection (He et al., 2017; Girshick et al., 2014). However, with the rapid gains in computational resources, neural networks can easily overfit a training set with numerous images (Russakovsky et al., 2015), leading to a large generalization gap on test data.

¹Data61-CSIRO, Canberra, Australia ²Australian National University, Canberra, Australia ³University of Adelaide, Adelaide, Australia. Correspondence to: Junlin Han <junlin.han@data61.csiro.au>.




Figure 1. YOCO cuts one image into two equal pieces, either in the height or the width dimension. The same data augmentations are performed independently within each piece. Augmented pieces are then concatenated together to form one single augmented image. The upper row shows the results of YOCO applied to `Color jitter`, where the two pieces are both augmented, producing a diversified fully-augmented image. The lower row presents the result of employing YOCO to `Horizontal flip`, only the left piece has been augmented, yielding a partially-augmented image.

In addition, neural networks can be effortlessly fooled by adversarial examples or attacks (Hendrycks et al., 2021a; Madry et al., 2018). Thus, to improve the generalization performance and robustness of neural networks, multiple training strategies have been proposed mainly from two views: regularization techniques (Srivastava et al., 2014; Huang et al., 2016; Szegedy et al., 2016) and data augmentations (Krizhevsky et al., 2012). In this work, we focus on improving data augmentations in particular.

Standard data augmentations are mostly performed at the image-level, which yields all-round performance improvements in both generality and robustness. Usually, image-level augmentation preserves semantics globally, following humans' cognitive intuition. Yet, human beings are also able to recognize objects from partial information alone. Patches, *i.e.*, the internal information of images, are strong natural signals. Before the deep learning era, the power of patches has been exploited in many low-level vision (Efros & Leung, 1999; Kervrann & Boulanger, 2006) and high-level vision (Sivic & Zisserman, 2003; Csurka et al., 2004; Lazebnik et al., 2006) works. Recently, splitting one image into multiple non-overlapping patches as the input for neural networks is a key success of Vision Transformer (ViT) (Dosovitskiy et al., 2020). However, how to perform

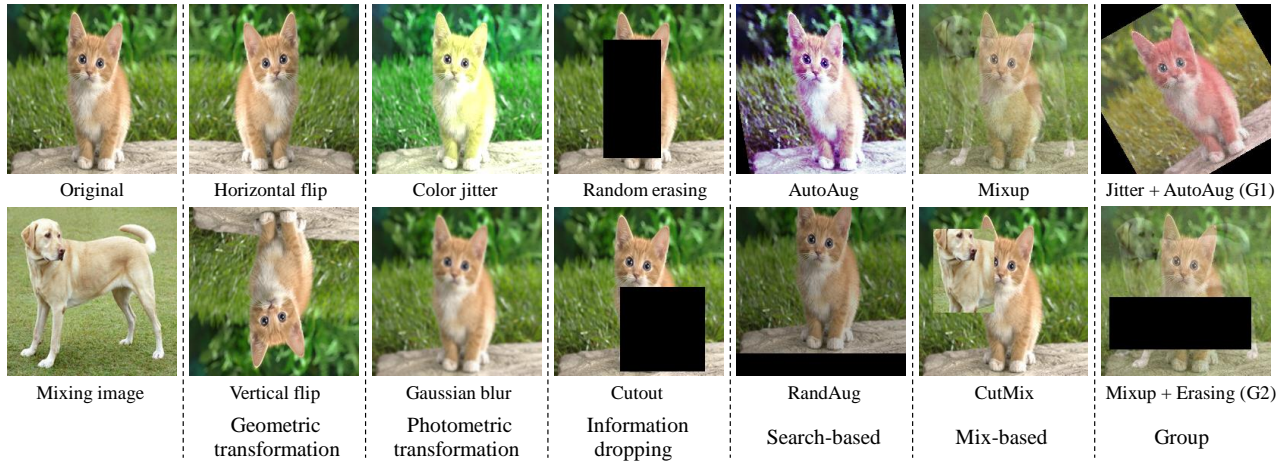


Figure 2. Illustrations of the studied augmentations in classification tasks. We roughly classify data augmentations into five categories. For each category, we study the two most representative augmentations. We also combine two augmentations as a group, forming a new category named group of augmentations.

data augmentations at a non-image level (in other words, patch-level or piece-level) is rarely studied.

From the above, we hypothesize that it is desirable to propose a strategy of performing augmentations beyond the image level. An image can be seen as a combination of multiple patches, or even two pieces. For a specific augmentation, we may perform this augmentation on these pieces individually and combine transformed pieces back to a single image. Such a strategy should increase the diversity in both local regions level as well as at the holistic image level and may also encourage neural networks to share the same cognitive ability of recognizing objects from partial information like humans can.

Following our hypothesis, we propose You Only Cut Once (YOCO), a simple method of performing data augmentations. Specifically, YOCO cuts one image into two equal pieces, in either the height dimension or the width dimension, performs data augmentations individually within each piece, and concatenates the two augmented pieces back together, as shown in Figure 1.

We conduct extensive evaluations of YOCO on various augmentations, neural network architectures, datasets, and challenging tasks. For image classification tasks on CIFAR-10 and CIFAR-100, YOCO can be seamlessly applied to almost all (11 out of 12) data augmentations, where we evaluate YOCO with 7 neural network architectures and 12 augmentations. YOCO boosts 11 augmentations to train classification models with better generalization abilities. For ImageNet, we evaluate generality, partial image recognition, calibration, robustness against adversarial attacks, corruption robustness, and robustness under distribution shifts. The six most representative augmentations are evaluated

(the first row of Figure 2). For every evaluation metric, YOCO outperforms image-level augmentation overall.

Next, we study the effect of YOCO on contrastive pre-training and conduct experiments based on MoCo v2 (Chen et al., 2020b) and SimSiam (Chen & He, 2021). YOCO benefits the contrastive self-supervised representation learning, where a more powerful representation is produced, which can be better transferred to classification, object detection, and instance segmentation tasks.

YOCO is parameter-free, easy to employ, and **boosts almost all augmentations for free**. As the presented YOCO is with the most simplified setting, we study some more complex design choices with more hyper-parameters embedded. Lastly, we analyze and study four questions (1): how does YOCO work? (2): when to employ YOCO? (3): why should we employ YOCO? and (4): does YOCO help other vision tasks?

2. Related work

2.1. Data Augmentation

Data augmentation is an effective strategy for training neural networks. We roughly divide common data augmentations into five categories, and for each category, we choose the two most representative augmentations: (1) Geometric transformation (Horizontal flip and Vertical flip), (2) Photometric transformation (Color jitter and Gaussian blur), (3) Information dropping (Random erasing (Zhong et al., 2020) and Cutout (DeVries & Taylor, 2017)), (4) Search-based (AutoAug (Cubuk et al., 2018) and RandAug (Cubuk et al., 2020)), and (5) Mix-based (Mixup (Zhang et al., 2017) and CutMix (Yun

et al., 2019)). To study the effect of YOCO applied to multiple augmentations, we further combine two augmentations as a group, forming category (6) group of augmentations (Color jitter + AutoAug (G1) and Mixup + Random erasing (G2)). Figure 2 presents all augmentations we studied. We leave feature-level augmentation (Upchurch et al., 2017; Wang et al., 2021) and GAN-based augmentation (Bowles et al., 2018) for future study.

Search-based augmentation utilizes reinforcement learning to search from a pool of augmentation policies for an optimal combination. Mix-based augmentation employs multi-image information by creating mixed input images with soft labels for training. Augmentations are important for multiple vision tasks. For instance, recent contrastive learning methods (He et al., 2020; Chen et al., 2020a;b; Chen & He, 2021) relies heavily on multiple data augmentations to construct contrastive pairs. All these augmentations are previously performed at the image-level, we study how to perform augmentations at a non-image level in multiple vision tasks.

2.2. Patches in Vision

Patches have been widely used as strong signals for various vision tasks, in both conventional methods, and learning-based methods. Applications include texture synthesis (Efros & Leung, 1999), image denoising (Kervrann & Boulanger, 2006), image-to-image translation (Park et al., 2020; Han et al., 2021b), super-resolution (Shocher et al., 2018), bag-of-features based classification (Sivic & Zisserman, 2003; Csurka et al., 2004; Lazebnik et al., 2006). Recently, both CNNs and ViTs leverage patches as the input for classification networks (Brendel & Bethge, 2019; Dosovitskiy et al., 2020).

Patches are also employed in data augmentation. For instance, Patch Gaussian (Lopes et al., 2019) applies Gaussian Blur to only a portion of the images. Cut, Paste and Learn (Dwivedi et al., 2017) generates large annotated instance datasets by cutting object instances and pasting them on backgrounds. Georgakis *et al.* (Georgakis et al., 2017) superimposes 2D object images into real environments images. CutMix (Yun et al., 2019) replaces one patch of an image with a patch from another image. PAA (Lin et al., 2021) extends the setting of AutoAug (Cubuk et al., 2018) by searching augmentation policies in pre-defined patches. Qin *et al.* (Qin et al., 2021) improves the robustness of ViTs through patch-based negative augmentation. However, no existing work has studied how to generally perform the same augmentation at a non-image level.

3. Method

This section presents the proposed YOCO method. Let $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$ denote an image and $a(\cdot)$ denote data augmentations, where $a(\cdot) : \mathbb{R}^{C \times H \times W} \rightarrow \mathbb{R}^{C \times H \times W}$, $\mathbf{X}' = a(\mathbf{X})$. Here \mathbf{X}' is the augmented image and $a(\cdot)$ can be one of any data augmentations or multiple data augmentations. In contrast to standard image-level augmentation, which applies the augmentation to the image directly as $\mathbf{X}' = a(\mathbf{X})$, YOCO first cuts the image into two equally sized pieces, in either height or width dimension with equal probability, as

$$[\mathbf{X}_1, \mathbf{X}_2] = \text{cut}_H(\mathbf{X}), \text{ if } 0 < p \leq 0.5, \text{ s.t. } \mathbf{X}_i \in \mathbb{R}^{C \times \frac{H}{2} \times W}$$

or

$$[\mathbf{X}_1, \mathbf{X}_2] = \text{cut}_W(\mathbf{X}), \text{ if } 0.5 < p < 1, \text{ s.t. } \mathbf{X}_i \in \mathbb{R}^{C \times H \times \frac{W}{2}},$$

where p is sampled from $(0, 1)$ uniformly (*i.e.*, $p \in U(0, 1)$), \mathbf{X}_i are cut pieces, $\text{cut}_H(\cdot)$ and $\text{cut}_W(\cdot)$ represent the cut operation in the height dimension and the width dimension, respectively.

Then $a(\cdot)$ is applied separately within each piece, and augmented pieces are concatenated back together as

$$\mathbf{X}' = \text{concat}[a_1(\mathbf{X}_1), a_2(\mathbf{X}_2)].$$

Augmentations are determined by randomness, which include random probabilities of being applied, random operations, and random magnitudes. Both $a_1(\cdot)$ and $a_2(\cdot)$ are instances of $a(\cdot)$, though applying the same data augmentation $a(\cdot)$, they may behave differently, thereby increasing the diversity in both the local region level and the holistic image level.

YOCO presented so far is in the most simple format, where the number of cuts and the position of the cut are both fixed. We also formulate a more general format of YOCO. In the general setting, You Only Cut Once (YOCO) becomes You Cut Many Times, where M cuts in the height dimension and N cuts in the width dimension are applied, such that $(M + 1) * (N + 1)$ pieces of internal regions can be obtained as $[\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_{(N+1)*M}, \mathbf{X}_{(N+1)*(M+1)}] = \text{cut}^{M*N}(\mathbf{X})$, where the position of cuts can be chosen uniformly or randomly. The identical procedure is also applied to those pieces, by $\mathbf{X}' = \text{concat}_{i=1, j=1}^{M+1, N+1} [a_{i*j}(\mathbf{X}_{i*j})]$.

4. Classification

In this section, we evaluate YOCO for its capability to improve image classifiers in multiple classification settings. Firstly, we study the effect of YOCO on CIFAR-10 and CIFAR-100 (Krizhevsky et al., 2009) datasets among 7 different models (5 CNN and 2 ViT) and 12 augmentations.

You Only Cut Once: Boosting Data Augmentation with a Single Cut

Models	Geometric trans		Photometric trans		Information dropping		Search-based		Mix-based		Group	
	H flip	V flip	Jitter	Blur	Erasing	Cutout	AutoAug	RandAug	Mixup	CutMix	G1	G2
PreResNet18	4.64	5.90	4.79	5.29	4.32	4.66	3.71	4.18	3.23	3.45	3.69	3.63
+ YOCO	5.05	5.52	4.70	5.00	4.28	4.36	3.38	4.08	3.11	3.38	3.56	3.25
Δ	+0.41	-0.38	-0.09	-0.29	-0.04	-0.30	-0.33	-0.10	-0.12	-0.07	-0.13	-0.38
Xception	4.62	5.20	4.52	5.31	4.07	4.15	3.56	4.01	3.70	3.66	4.04	3.41
+ YOCO	4.91	5.11	4.51	4.65	4.07	3.94	3.31	3.95	3.17	3.22	3.32	3.20
Δ	+0.29	-0.09	-0.01	-0.66	0.00	-0.21	-0.25	-0.06	-0.53	-0.44	-0.72	-0.21
DenseNet121	4.58	5.21	4.72	5.09	4.17	4.55	3.83	4.22	3.71	3.53	3.98	3.57
+ YOCO	5.09	5.18	4.70	4.76	4.13	4.25	3.67	3.91	3.28	3.50	3.50	3.36
Δ	+0.51	-0.03	-0.02	-0.33	-0.04	-0.30	-0.16	-0.31	-0.43	-0.03	-0.48	-0.21
ResNeXt50	4.69	5.12	4.74	5.66	4.18	4.93	3.65	4.09	3.69	3.52	3.79	3.66
+ YOCO	5.10	5.03	4.60	5.27	3.96	4.64	3.26	3.91	3.45	3.44	3.26	3.41
Δ	+0.41	-0.09	-0.14	-0.39	-0.22	-0.29	-0.39	-0.18	-0.24	-0.08	-0.53	-0.25
WRN-28-10	3.51	4.10	3.63	3.99	2.92	3.26	2.72	3.23	2.60	2.86	2.75	2.79
+ YOCO	3.53	4.09	3.44	3.97	2.89	2.96	2.62	3.05	2.46	2.82	2.65	2.35
Δ	+0.02	-0.01	-0.19	-0.02	-0.03	-0.30	-0.10	-0.18	-0.14	-0.04	-0.10	-0.44
ViT	4.69	5.51	4.74	5.49	4.32	4.44	3.76	4.13	3.67	3.49	3.82	3.27
+ YOCO	4.92	5.34	4.70	4.95	4.18	4.42	3.54	3.96	3.17	3.21	3.51	3.14
Δ	+0.23	-0.17	-0.04	-0.54	-0.14	-0.02	-0.22	-0.17	-0.50	-0.28	-0.31	-0.13
Swin	4.83	5.27	4.90	5.15	4.26	4.55	3.63	4.15	3.56	3.64	3.74	3.58
+ YOCO	4.85	5.25	4.40	5.12	4.23	4.03	3.54	4.13	3.16	3.22	3.44	3.46
Δ	+0.02	-0.02	-0.50	-0.03	-0.03	-0.52	-0.09	-0.02	-0.40	-0.42	-0.30	-0.12
Average Δ	+0.27	-0.11	-0.14	-0.32	-0.07	-0.28	-0.22	-0.15	-0.34	-0.19	-0.37	-0.25

Table 1. **Test top-1 error rate on CIFAR-10.** We evaluate 12 augmentations on 7 network architectures. Among 84 results, 76 of them are improved by applying YOCO. Improved results are highlighted in blue. Best viewed in color. Standard deviations lie in (0.04, 0.15). Average Δ is the difference between YOCO and image-level augmentation, which is based on the mean of 21 independent runs in total.

Models	Geometric trans		Photometric trans		Information dropping		Search-based		Mix-based		Group	
	H flip	V flip	Jitter	Blur	Erasing	Cutout	AutoAug	RandAug	Mixup	CutMix	G1	G2
PreResNet18	23.36	25.16	23.95	25.96	24.03	25.01	21.94	22.70	21.32	20.05	22.97	21.23
+ YOCO	24.66	24.73	23.62	25.11	23.94	23.41	20.55	22.70	19.62	19.30	21.06	20.92
Δ	+1.30	-0.43	-0.33	-0.85	-0.09	-1.60	-1.39	0.00	-1.70	-0.75	-1.91	-0.31
Xception	23.34	25.16	23.39	26.17	25.01	25.44	22.50	22.87	21.41	20.17	23.02	21.79
+ YOCO	24.46	24.65	23.42	25.51	23.74	23.56	20.29	22.45	19.96	19.25	20.75	20.68
Δ	+1.12	-0.51	+0.03	-0.66	-1.27	-1.88	-2.21	-0.42	-1.45	-0.92	-2.27	-1.11
DenseNet121	23.14	23.71	24.25	26.21	23.97	24.71	22.43	22.66	21.17	19.50	22.80	21.23
+ YOCO	24.21	24.62	23.14	24.68	23.34	23.83	20.63	23.11	19.87	19.26	21.13	20.95
Δ	+1.07	+0.91	-1.11	-1.53	-0.63	-0.88	-1.80	+0.45	-1.30	-0.24	-1.67	-0.28
ResNeXt50	24.60	24.61	23.58	25.86	24.50	24.84	22.32	22.77	20.78	19.72	22.69	21.69
+ YOCO	25.11	24.22	23.76	25.66	23.51	23.77	21.01	22.47	19.74	19.68	20.58	20.55
Δ	+0.51	-0.39	+0.18	-0.20	-0.99	-1.07	-1.31	-0.30	-1.04	-0.04	-2.11	-1.14
WRN-28-10	19.90	21.70	19.54	21.95	19.79	19.73	18.07	19.25	18.93	17.03	18.39	17.74
+ YOCO	20.44	21.34	19.59	21.17	19.65	19.68	17.17	19.29	17.64	17.29	18.89	17.72
Δ	+0.54	-0.36	+0.05	-0.78	-0.14	-0.05	-0.90	+0.04	-1.29	+0.26	-0.50	-0.02
ViT	23.67	24.01	24.53	26.06	24.09	25.10	22.14	22.79	21.15	20.27	22.75	21.64
+ YOCO	24.52	24.74	23.49	25.04	23.20	23.38	20.61	22.34	19.26	19.17	20.85	21.20
Δ	+0.85	+0.73	-1.04	-1.02	-0.89	-1.72	-1.53	-0.45	-1.89	-1.10	-1.90	-0.44
Swin	23.57	24.51	24.04	25.40	24.81	25.11	22.30	23.14	21.42	19.78	22.92	21.26
+ YOCO	24.66	24.37	23.96	25.39	23.22	23.86	20.70	22.62	19.43	19.74	20.75	20.88
Δ	+1.09	-0.14	-0.08	-0.10	-1.59	-1.25	-1.60	-0.52	-1.99	-0.04	-2.17	-0.38
Average Δ	+0.93	-0.03	-0.33	-0.73	-0.80	-1.21	-1.53	-0.17	-1.52	-0.40	-1.79	-0.53

Table 2. **Test top-1 error rate on CIFAR-100,** where 12 augmentations and 7 network architectures are evaluated. YOCO outperforms image-level augmentation in 68 results out of 84 results. blue indicates improved results. Best viewed in color. Standard deviations are between (0.09, 0.41). Average Δ is the difference between YOCO and image-level augmentation over 21 independent trials.

Next, we validate YOCO on ImageNet-1K (henceforth referred to as ImageNet) (Russakovsky et al., 2015). Due to that ImageNet is a large-scale dataset containing more than 1.2M training images from 1K categories, we generally and thoroughly evaluate YOCO on ImageNet to explore its potential.

4.1. CIFAR Experiments

Experimental setup. We train 5 CNN architectures (PreResNet18 (He et al., 2016a), Xception (Chollet, 2017),

DenseNet121 (Huang et al., 2017), ResNeXt50 (Xie et al., 2017), WRN-28-10 (Zagoruyko & Komodakis, 2016)) and 2 ViT models (ViT (Dosovitskiy et al., 2020) and Swin (Liu et al., 2021)) under the same training recipe, which mostly follows the training setting used in Co-Mixup (Kim et al., 2021). That is, models are trained for 300 epochs, with an initial learning rate of 0.1 decayed by a factor 0.1 at epochs 100 and 200. We employ the SGD optimizer with a momentum of 0.9 and a weight decay of 0.0001. We use a batch size of 100, and models are trained on 2 GPUs. Training images are randomly cropped to 32×32 resolution with

zero padding followed by random `Horizontal flip`. Other augmentations are added based on this setup.

Results. We report the *best* results based the mean of three independent trials. Table 1 and Table 2 present the results on CIFAR-10 and CIFAR-100, respectively. On CIFAR-10, among **12** augmentations, YOCO boosts **11** of them, toward a more generalized image classifier. For all **84** results, YOCO improves **76** of them, showing the superiority generalization of YOCO over image-level augmentation. The improvements are non-trivial (average improvement $\geq 0.2\%$) for `Blur`, `Cutout`, `AutoAug`, `Mixup`, `G1`, and `G2`. On CIFAR-100, YOCO also boosts **11** augmentations, with **68** results improved. The improvements of applying YOCO are notable (average improvement $\geq 0.7\%$) for `Blur`, `Erasing`, `Cutout`, `AutoAug`, `Mixup`, and `G1`.

For both CIFAR-10 and CIFAR-100, we also observe that applying YOCO to `Horizontal flip` degrades the performance. We analyse the reasons empirically in Section 7. The computational cost and memory usage of YOCO are both **negligible**. Compared to image-level augmentation, YOCO runs at an identical speed. Thereby, YOCO boosts nearly all augmentations for free.

4.2. ImageNet Experiments

Experimental setup. We follow a simple training recipe used in `CutMix` (Yun et al., 2019). In doing so, we train a standard ResNet50 (He et al., 2016b) for 300 epochs, with an initial learning rate of 0.1, decayed by a factor of 0.1 at epochs 75, 150, and 225. We employ the SGD optimizer with a momentum of 0.9 and models are trained with a weight decay of 0.0001. We employ a batch size of 256, and 4 GPUs are used for training. Training images are randomly cropped to 224×224 resolution followed by random `Horizontal flip` prior to any other augmentations. We evaluate and test 6 representative augmentations as shown in the first row of Figure 2.

4.2.1. EVALUATION PROTOCOL

For models trained on ImageNet, other than generalization, we also evaluate partial image recognition, calibration, robustness against adversarial attacks, corruption robustness, and robustness under distribution shifts.

Partial image recognition. One design goal of YOCO is to inspire neural networks to share the same ability of recognizing objects from partial information like human can. We thus design a non-overlapping 4 crop evaluation. Specifically, we load test images in 512×512 resolution, and images are center cropped to 448×448 resolution, following the most common crop ratio of 0.875. We employ two cuts, one in the height dimension and one in the width dimension, to cut 448×448 images into 4 non-overlapping

224×224 resolution pieces. Four cropped pieces are evaluated and we report the mean of the results.

Calibration. We follow the evaluation protocol used in `PixMix` (Hendrycks et al., 2021b). The calibration task aims to match the empirical frequency of correctness. The posteriors from a model should satisfy

$$\mathbb{P}\left(Y = \arg \max_i f(X)_i \mid \max_i f(X)_i = C\right) = C,$$

where f is an image classifier, X, Y are random variables representing the data distribution. We use RMS calibration error (Hendrycks et al., 2018) and adaptive binning (Nguyen & O’Connor, 2015).

Adversarial attack. We employ two adversarial attacks, FGSM (Goodfellow et al., 2014) and PGD (Madry et al., 2018), to verify the robustness of image classifiers. We borrow the implementation from `Torchattacks` (Kim, 2020). For all attacks, we apply an ℓ_∞ budget of $8/255$. PGD is with 4 steps of optimization. We do not use larger steps of optimization as the accuracy can decline to zero for non-adversarially trained image classifiers.

Corruption robustness. Following `Co-Mixup` (Kim et al., 2021), we evaluate image classifiers on two challenging corrupted test sets, to verify the generalization ability and robustness on unseen environments. The corrupted test sets are created with the following operations: (1) replace the background with a random image and (2) adding `Gaussian noise` to the background.

Distribution shift. We employ `ImageNet-A` (Hendrycks et al., 2021a), one of the most challenging ImageNet test sets containing natural adversarial examples, to verify the performance of image classifiers against input data distribution shifts.

4.2.2. RESULTS

We report both *best* and *last* results. All results are summarized in Table 3. Visualizations on classification results are available in the Appendix D.

Generality. Among 5 tested augmentations, YOCO outperforms image-level augmentation in terms of generalization. Considering its negligible cost, such improvement brought from YOCO is huge. For example, applying YOCO to `Horizontal flip` and `AutoAug` improves the *last* top-1 accuracy by **0.42%** and **0.72%** respectively.

Partial image recognition. YOCO improves the ability of image classifiers in recognizing objects from partial information. YOCO applied to 5 augmentations can reach improved performance. For `Random erasing`, dissimilar to other augmentations, applying YOCO does not involve changing the structure/consistency of images, thereby the partial image recognition ability is not improved.

You Only Cut Once: Boosting Data Augmentation with a Single Cut

Augs	Generalization Clean	Partial Clean	Calibration Clean, RMS↓	Adversarial attacks FGSM attack	PGD attack	Corruptions Random replace	Gaussian noise	Distribution shift ImageNet-A
H flip	77.10/76.59	55.21/56.14	6.42/8.81	17.28/21.01	9.37/14.98	58.90/58.77	73.58/72.92	3.39/4.29
+ YOCO	77.28/77.01	56.27/56.44	7.87/8.19	20.59/21.88	14.21/15.21	59.11/58.91	73.63/72.96	3.80/4.32
△	+0.18/+0.42	+1.06/+0.30	+1.45/-0.62	+3.31/+0.87	+4.84/+0.23	+0.21/+0.14	+0.05/+0.04	+0.41/+0.03
Jitter	77.15/76.87	56.04/56.34	8.06/8.19	18.31/18.80	11.07/11.89	59.11/58.82	73.74/73.38	3.98/4.13
+ YOCO	77.35/77.12	56.41/56.60	7.50/8.07	16.32/17.77	5.96/6.90	59.40/59.32	74.46/74.14	4.29/4.56
△	+0.20/+0.25	+0.37/+0.26	-0.56/-0.12	-1.99/-1.03	-5.11/-4.99	+0.29/+0.50	+0.72/+0.76	+0.31/+0.43
Erasing	77.40/77.09	56.75/56.82	7.67/7.88	22.01/22.67	12.74/13.31	58.75/58.44	73.81/73.48	4.11/4.32
+ YOCO	77.29/77.20	56.66/56.76	7.57/8.00	22.34/22.82	12.93/14.00	58.80/58.85	73.95/73.59	4.12/4.28
△	-0.11/+0.11	-0.09/-0.06	-0.10/+0.12	+0.33/+0.15	+0.19/+0.69	+0.05/+0.41	+0.14/+0.11	+0.01/-0.04
AutoAug	77.55/76.93	55.61/55.05	6.17/6.73	11.51/11.62	2.05/2.18	58.36/58.12	74.45/73.85	4.54/4.69
+ YOCO	77.88/77.65	56.48/56.52	6.04/6.25	12.09/11.73	1.11/0.86	58.90/58.29	76.03/75.70	5.01/4.80
△	+0.33/+0.72	+0.87/+1.47	-0.13/-0.48	+0.58/+0.11	-0.94/-1.32	+0.54/+0.17	+1.58/+1.85	+0.47/+0.11
Mixup	77.72/77.67	55.27/55.40	8.23/9.03	35.92/35.41	7.90/7.44	61.46/61.07	75.47/75.12	8.37/8.59
+ YOCO	77.81/77.74	56.00/55.97	4.04/4.42	39.63/39.67	13.02/13.63	61.14/60.85	75.10/74.88	8.27/8.55
△	+0.09/+0.07	+0.73/+0.57	-4.19/-4.61	+3.71/+4.26	+5.12/+6.19	-0.32/-0.22	-0.37/-0.24	-0.10/-0.04
G1	77.61/77.19	55.57/55.24	6.29/6.19	13.11/13.22	2.34/2.87	58.94/58.99	75.22/74.94	5.01/5.22
+ YOCO	77.76/77.64	56.00/55.97	6.17/6.02	13.62/13.61	1.82/2.61	59.15/59.03	75.73/75.59	5.41/5.47
△	+0.15/+0.45	+0.43/+0.73	-0.12/-0.17	+0.51/+0.39	-0.52/-0.26	+0.21/+0.04	+0.51/+0.65	+0.40/+0.25

Table 3. Results of applying YOCO to train a ResNet50 on ImageNet. We evaluate generalization (top-1 accuracy), partial image recognition (top-1 accuracy), calibration (RMS), robustness against adversarial attacks (top-1 accuracy), corruption robustness (top-1 accuracy), and robustness under distribution shifts (top-1 accuracy). We report both *best* result (left) and *last* result (right). Improved results are highlighted in blue color. Best viewed in color. △ is the difference between YOCO and image-level augmentation.

Calibration. In terms of calibration, YOCO improves 4 augmentations. For Horizontal flip, the RMS of *best* result is increased from **6.42** to **7.87**, while the *last* result conversely benefits from YOCO. Epoch **171** produces the *best* result for image-level augmentation while the *best* result after applying YOCO is at epoch **235**. The inconsistent results are might due to classification models does not suffer from overfitting and overconfidence in the middle stage of training. To conclude, YOCO helps image classifiers to have better-calibrated predictions most of the time.

Adversarial attack. For robustness against adversarial attacks, augmentations that do not involve color transformations usually benefit from YOCO, where YOCO hurts the robustness when applied to Color Jitter, AutoAug, and G1, but is helpful for all other augmentations. YOCO applied to color transformations breaks the color consistency of images. We suppose this property is not good for defending adversarial attacks. YOCO still performs well overall, as applying YOCO can be effectual for other augmentations, where YOCO approximately **doubles** the results of Mixup under PGD attack.

Corruption robustness. YOCO consistently shows better performance when evaluated with challenging corrupted images. 5 augmentations are boosted by YOCO. The improvements are particularly considerable for AutoAug, where the *best* top-1 accuracy is improved with the performance margins of **0.54%** and **1.58%** for Random replacement and Gaussian noise correspondingly. This suggests that when only partial information is available, YOCO with better ability in recognizing partial objects behaves more robustly.

Distribution shift. Testing with the challenging ImageNet-A test set, YOCO again outperforms image-level augmentations overall. Improving the results on ImageNet-A is

difficult, but applying YOCO to augmentations including Horizontal flip, Color jitter, AutoAug, and G1 realizes improved results, showing the superiority of YOCO.

Reported results. We find that the ResNet50 implemented in timm (Wightman, 2019) yields a better performance than Pytorch/Torchvision inbuilt ResNet50 and original ResNet50 implemented in (He et al., 2016b). For a fair comparison, all ImageNet experiments are based on the timm implementation. For image-level augmentation, we report our reproduced results. Our reproduced results are usually higher than previously reported results. For instance, the *best* Top-1 accuracy for Horizontal flip is reported as **77.1%** by us, while being previously (same training recipe) reported as **76.3%** (Yun et al., 2019).

5. Contrastive Learning

One core step of contrastive learning is to create two views from one image by data augmentations. YOCO increases the diversity of data, applying YOCO to contrastive learning yields more challenging views. However, strong augmentations are usually not beneficial to contrastive learning (Chen et al., 2020a). In this section, we verify whether YOCO benefits MoCo v2 (Chen et al., 2020b) and SimSiam (Chen & He, 2021) in the pre-training stage. The learned embedding is evaluated via multiple downstream tasks. Specifically, we evaluate the classification performance on ImageNet, under both the linear protocol and semi-supervised fine-tuning settings. The downstream tasks include VOC (Everingham et al., 2010) detection, COCO (Lin et al., 2014) detection and COCO instance segmentation.

You Only Cut Once: Boosting Data Augmentation with a Single Cut

Method	ImageNet classification			VOC detection			COCO detection			COCO instance seg		
	linear protocol	1% label	10% label	AP ₅₀	AP	AP ₇₅	AP ₅₀	AP	AP ₇₅	AP ₅₀ ^{mask}	AP ^{mask}	AP ₇₅ ^{mask}
MoCo v2	67.5	34.5	61.1	81.9	56.8	62.9	57.3	38.0	40.8	54.1	33.3	35.5
+ YOCO	67.6	34.9	61.5	82.4	56.6	63.6	57.4	38.1	41.6	54.1	33.4	35.5
SimSiam	68.1	17.1	57.3	80.2	54.5	60.0	51.4	33.5	36.1	48.7	29.9	32.0
+ YOCO	68.3	17.3	57.7	80.2	54.5	60.1	53.1	34.6	37.3	50.1	30.8	32.9

Table 4. **Results of contrastive learning.** MoCo v2 and SimSiam are pre-trained for 200 and 100 epochs in ImageNet, respectively. All are based on ResNet50 pre-trained with two 224x224 views. The evaluation metric of ImageNet classification is Top-1 accuracy.

5.1. Experimental Setup

Pre-training. We first apply YOCO to Horizontal flip only. Following the training setting introduced in MoCo v2 and SimSiam, we pre-train them with YOCO for 200 and 100 epochs respectively. The models are pre-trained following the same settings as their own base methods. The backbone is ResNet50. More details of training settings are provided in the Appendix C.

Linear protocol. We freeze the learned embedding and train a linear classifier on top of it. The experimental setting are identical to the same settings as their own base methods.

Fine-tuning. We also fine-tune the learned feature under the semi-supervised setting on ImageNet with 1% and 10% label fractions, where labels are provided by SimCLR (Chen et al., 2020a). The fine-tuning setting follows Jigsaw Clustering (Chen et al., 2021).

VOC detection. We mostly follow the setting used in MoCo. Faster R-CNN (Ren et al., 2015) with the C4-backbone is fine-tuned on VOC 2007 trainval + 2012 train and evaluated on the VOC 2007 test. Due to a GPU limit, all experiments are conducted using 4 GPUs, whereas 8 GPUs are used in MoCo. We halve the batch size from 16 to 8, double the max iterations from 24K to 48K, and scale the base learning rate from 0.02 to 0.01. All reported results are the average over 3 trials under this setting. The results of MoCo v2 and SimSiam baseline are reproduced by us using the corresponding official pre-trained ResNet50 models.

COCO detection and instance segmentation. Following MoCo, the model is Mask R-CNN (He et al., 2017) with a C4 backbone. Similar to VOC, batch size and base learning rate are halved while the maximum iteration is doubled, from 90K to 180K. The results of baselines are also reproduced by us under this setting.

5.2. Results

Table 4 depicts the results. Among multiple tasks, both MoCo v2 and SimSiam benefit from YOCO.

The effect of YOCO applied to other augmentations is also studied, where we gradually increase the number of augmentations that are being applied with YOCO, specifically, we study YOCO applied to (1): Horizontal

Method	ImageNet classification			VOC detection		
	linear protocol	1% label	10% label	AP ₅₀	AP	AP ₇₅
MoCo v2	67.5	34.5	61.1	81.9	56.8	62.9
(1)	67.6	34.9	61.5	82.4	56.6	63.6
(2)	67.3	34.8	61.3	82.5	56.9	63.4
(3)	67.2	34.5	61.2	82.4	56.6	62.7
(4)	67.0	34.4	61.0	82.1	56.0	61.9

Table 5. **The effect of YOCO applied to multiple augmentations in contrastive learning.** We study YOCO applied to (1): Horizontal flip, (2): Horizontal flip + Gaussian blur, (3): Horizontal flip + Gaussian blur + Grayscale, and (4): Horizontal flip + Gaussian blur + Grayscale + Color jitter.

flip, (2): Horizontal flip + Gaussian blur, (3): Horizontal flip + Gaussian blur + Grayscale, and (4): Horizontal flip + Gaussian blur + Grayscale + Color jitter. The results are presented in Table 5, where YOCO applied to more augmentations hurts the learned representation, suggesting that crafting more challenging views in moderation is helpful for contrastive representation learning.

Detailed training settings are provided in the Appendix C.

6. Ablation study

The concept of YOCO is presented in a simple way, where YOCO only involves one cut and the position of that cut is fixed. The general format of YOCO may involve multiple cuts and random positions of cuts. It naturally gives rise to questions like: Will it be beneficial if the position of the cut varies? How many cuts are needed to boost data augmentations? In this section, we aim to explore the optimal setting of YOCO.

Experiments are performed on the CIFAR-100 dataset. All reported results are the mean of PreResNet18 and DenseNet121. We study 5 augmentations (the first row of Figure 2, except for G1).

Position of the cut. We employ the beta distribution $Beta(\alpha, \alpha)$ to sample the position of the cut, where we set α as 0.2, 0.4, 0.6, 0.8, and 1.0. The results are presented in Table 6. Compared to the fixed position, randomly sampled positions can be better applied to Horizontal flip and AutoAug but hurt the performance of being applied to Color jitter, Erasing, and Mixup. As

the position of the cut behaves differently with augmentations, it is hard to select the optimal position. Designing an augmentation-specific strategy to select the optimal position of the cut might be a potential way to further improve the performance of YOCO.

α	Augmentations				
	H flip	Jitter	Erasing	AutoAug	Mixup
0.2	76.2	76.3	75.9	79.8	79.7
0.4	75.9	76.5	75.8	79.8	79.7
0.6	76.5	76.4	76.3	79.9	79.8
0.8	76.4	76.3	75.9	79.8	79.5
1.0	76.4	76.2	75.7	80.2	79.4
YOCO	75.6	76.6	76.4	79.4	80.3

Table 6. **Position of the cut.** Test top-1 accuracy on CIFAR-100 are reported. We randomly sample the position of the cut from a beta distribution $\text{Beta}(\alpha, \alpha)$.

Number of cuts. For both cuts in the height dimension and the width dimension, we set the number of cuts to 1, 2, and 3. We report the results of all combinations in Table 7. For all augmentations, increasing the number of cuts hurts the performance. More cuts further increase the diversity of data, yielding the data to be too complex to learn from. This is particularly obvious when augmentations involve geometric transforms, as the semantic consistency is often corrupted when a large number of cuts are applied. For large-scale datasets such as ImageNet, we also observe a similar trend. Thus, You Only Cut Once is sufficient.

Number H,W	Augmentations				
	H flip	Jitter	Erasing	AutoAug	Mixup
1,1	70.8	76.1	76.0	77.8	79.4
1,2	69.1	75.9	76.0	77.2	79.2
1,3	68.0	75.5	75.9	76.7	79.2
2,1	69.7	76.5	75.9	76.9	80.1
2,2	68.7	75.9	75.6	76.5	79.6
2,3	67.9	75.7	75.6	76.2	79.5
3,1	69.9	75.9	75.8	77.0	80.2
3,2	67.6	75.8	75.5	76.3	79.8
3,3	67.4	75.2	75.5	76.0	79.6
YOCO	75.6	76.6	76.4	79.4	80.3

Table 7. **Number of cuts.** We report test top-1 accuracy on CIFAR-100. The number of cuts are set to 1, 2, or 3 for both height dimension and width dimension.

7. Analysis and Discussion

7.1. How Does YOCO Work?

We classify training data into two main categories, namely original data and augmented data, where the augmented data involves fully-augmented data, partially-augmented data, and diversified fully-augmented data. In most cases, YOCO is capable of creating two additional kinds of augmented data than image-level augmentation, partially-augmented data, and diversified fully-augmented data, as shown in Figure 1. We study the effect of YOCO by exploring the performance of classifiers trained with each kind of data solely. Results are summarized in Table 8.

Combination	H flip	Augmentations		
		Jitter	Erasing	Mixup
Original	72.7	76.8	76.8	76.8
Fully-aug	71.7	76.2	76.5	78.8
Partially-aug	74.2	76.5	76.5	79.6
Diversified fully-aug	72.5	76.3	76.9	79.7

Table 8. **How does YOCO work?** Test top-1 accuracy on CIFAR-100. We train image classifiers with one kind of data solely.

Trained with one kind of data solely, partially-augmented data and diversified fully-augmented data both outperform fully-augmented data. This might be the reason for the superior performance of YOCO, where these two kinds of data further increase the diversity of augmented data, such that neural networks can learn more meaningful knowledge.

7.2. When to Employ YOCO?

We focus on studying when to employ YOCO in classification tasks. From both CIFAR and ImageNet results, YOCO outperforms image-level augmentation among most augmentations. However, YOCO applied to `Horizontal flip` hurts the performance of image classifiers on the CIFAR dataset. CIFAR mostly contains iconic object images, that is, the object, as a foreground part, occupies to most of the image. We hypothesise that YOCO applied to `Horizontal flip` ruins the structure/consistency of iconic object images. For more complex images (either more objects or more background components) such as ImageNet, YOCO can be applied to `Horizontal flip`, as the structural information for the foreground object can be kept in more cases, while the image diversity regarding background/object location is improved. Also, the random `resized crop` pre-processing plays an important role, as it may change a non-iconic object image to a iconic object image and vice versa.

To verify our hypothesis, we evaluate on Oxford Flower102 (Nilsback & Zisserman, 2008) (fine-grained flower classification dataset containing iconic flower images) and 10% ImageNet. Results are presented in Figure 3. For iconic object images, YOCO is better than image-level augmentation only when a wider crop range is employed while YOCO always performs better on non-iconic object images. Thereby we conclude that for most augmentations, YOCO can outperform image-level augmentation, but for `Horizontal flip` YOCO should be employed depending on the properties and crop ratio of the dataset.

7.3. Why Should We Employ YOCO?

In the real-world environment, it is impossible to guarantee complete images as the source of inputs. That said, for a robust vision system applied in real-world practice, having the ability to recognize partial patterns is a must. In

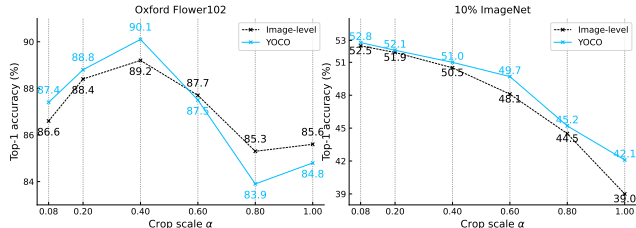


Figure 3. When to employ YOCO? Results of Horizontal flip performed with image-level augmentation and YOCO on Oxford Flower102 (left) and 10% ImageNet (right). We test multiple crop scale (α , 1).

common practice, neural networks mostly gain such ability from the random resized crop. Applying YOCO to multiple augmentations further equips neural networks to have stronger partial recognition abilities, which also result in better holistic level recognition abilities for the neural networks.

YOCO is a booster for training more robust vision systems that can be applied in varying complex real scenarios. YOCO also serves as a free plug-and-play for various vision tasks, where no extra computation costs are brought, no parameters need to be manually selected, and no complicated implementation is required.

7.4. Does YOCO Help Other Vision Tasks?

YOCO boosts data augmentation in multiple high-level vision tasks, trained from scratch, or transferred from contrastive pre-training. We thus study whether YOCO helps other vision tasks, where we employ image deraining (Zamir et al., 2021; Han et al., 2021a) and image super-resolution (Yoo et al., 2020), two representative low-level vision tasks, to further verify the generality of YOCO.

Image deraining: We include two augmentations, Horizontal flip and Vertical flip, along with YOCO, to the training process of MPRNet (Zamir et al., 2021). Besides the augmentations, MPRNet is trained with the identical dataset (synthetic rain dataset (Zamir et al., 2021) with 13712 pairs) and training settings. Results are presented in Table 9. Our gain over the baseline is **0.73dB** on the Rain100L test set.

Methods	Test100	Rain100H	Rain100L	Test2800	Test1200	Average
MPRNet	30.27	30.41	36.40	33.64	32.91	32.73
+ YOCO	30.33	30.53	37.13	33.64	32.84	32.89

Table 9. **Image deraining.** PSNR results of YOCO applied to image deraining tasks. We choose the state-of-the-art MPRNet as our baseline. Results are further improved with YOCO.

Image super-resolution: Following CutBlur (Yoo et al., 2020), we evaluate YOCO applied to mixture of augmen-

tations (Cutout, CutMix, Mixup, CutMixup, RGB perm, Blend, and CutBlur) on two common datasets, DIV2K (Agustsson & Timofte, 2017) and RealSR (Cai et al., 2019). All experiments are conducted using the CARN (Ahn et al., 2018) model, where the training settings follows CutBlur but without pre-training on DIV2K in $\times 2$ scale. In the synthetic training setting, models trained on the DIV2K dataset are evaluated on multiple test sets. Table 10 presents the results, where YOCO consistently outperforms image-level augmentation regardless of test sets in the synthetic setting, and shows on par performance in the realistic setting.

Aug	DIV2K	Set14	Synthetic Urban100	Manga109	Realistic RealSR
No aug	28.83	28.49	25.82	30.11	28.78
+ Image-level	28.83	28.48	25.82	30.08	28.99
+ YOCO	28.84	28.50	25.83	30.11	28.99

Table 10. **Super-resolution.** Quantitative comparison (PSNR) on super-resolution (scale $\times 4$) task in both synthetic and realistic settings. We evaluate the performance of image-level augmentation and YOCO applied to mixture of augmentations (7 in total).

We believe that our results on low-level vision further validate the generality and robustness of YOCO.

8. Limitation

Though thorough empirical results show the superiority and generality of YOCO, the theoretical foundation can be further developed for shedding light on the interpretability of YOCO. Additionally, applying YOCO can be harmful to certain classes in the supervised classification task, as the effect of data augmentation is class dependent (Balestriero et al., 2022). To achieve superior classification performance, designing a class-specific strategy of applying YOCO is desirable.

9. Conclusion

We propose YOCO to boost augmentations, where YOCO has shown positive results in a variety of computer vision tasks and datasets. The competitiveness of our minimalist YOCO suggests that performing data augmentations can be a core component in training neural networks. We hope our study will attract the community’s attention in revisiting how to perform data augmentations.

Acknowledgement

We would like to thank the anonymous reviewers for their constructive feedback.

References

- Agustsson, E. and Timofte, R. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 126–135, 2017.
- Ahn, N., Kang, B., and Sohn, K.-A. Fast, accurate, and lightweight super-resolution with cascading residual network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 252–268, 2018.
- Balestriero, R., Bottou, L., and LeCun, Y. The effects of regularization and data augmentation are class dependent. *arXiv preprint arXiv:2204.03632*, 2022.
- Bowles, C., Chen, L., Guerrero, R., Bentley, P., Gunn, R., Hammers, A., Dickie, D. A., Hernández, M. V., Wardlaw, J., and Rueckert, D. Gan augmentation: Augmenting training data using generative adversarial networks. *arXiv preprint arXiv:1810.10863*, 2018.
- Brendel, W. and Bethge, M. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. *arXiv preprint arXiv:1904.00760*, 2019.
- Cai, J., Zeng, H., Yong, H., Cao, Z., and Zhang, L. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3086–3095, 2019.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.
- Chen, P., Liu, S., and Jia, J. Jigsaw clustering for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11526–11535, 2021.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020a.
- Chen, X. and He, K. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758, 2021.
- Chen, X., Fan, H., Girshick, R., and He, K. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.
- Chollet, F. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.
- Csurka, G., Dance, C., Fan, L., Willamowski, J., and Bray, C. Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, volume 1, pp. 1–2. Prague, 2004.
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.
- DeVries, T. and Taylor, G. W. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Dwibedi, D., Misra, I., and Hebert, M. Cut, paste and learn: Surprisingly easy synthesis for instance detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 1301–1310, 2017.
- Efros, A. A. and Leung, T. K. Texture synthesis by non-parametric sampling. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pp. 1033–1038. IEEE, 1999.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88 (2):303–338, 2010.
- Georgakis, G., Mousavian, A., Berg, A. C., and Kosecka, J. Synthesizing training data for object detection in indoor scenes. *Robotics: Science and Systems (RSS)*, 2017.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 580–587, 2014.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Han, J., Li, W., Fang, P., Sun, C., Hong, J., Armin, M. A., Petersson, L., and Li, H. Blind image decomposition. In *arXiv preprint arXiv:2108.11364*, 2021a.

- Han, J., Shoeiby, M., Petersson, L., and Armin, M. A. Dual contrastive learning for unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2021b.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016a.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016b.
- He, K., Gkioxari, G., Dollár, P., and Girshick, R. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Hendrycks, D., Mazeika, M., and Dietterich, T. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018.
- Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., and Song, D. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15262–15271, 2021a.
- Hendrycks, D., Zou, A., Mazeika, M., Tang, L., Song, D., and Steinhardt, J. Pixmix: Dreamlike pictures comprehensively improve safety measures. *arXiv preprint arXiv:2112.05135*, 2021b.
- Huang, G., Sun, Y., Liu, Z., Sedra, D., and Weinberger, K. Q. Deep networks with stochastic depth. In *European conference on computer vision*, pp. 646–661. Springer, 2016.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Kervrann, C. and Boulanger, J. Optimal spatial adaptation for patch-based image denoising. *IEEE Transactions on Image Processing*, 15(10):2866–2878, 2006.
- Kim, H. Torchattacks: A pytorch repository for adversarial attacks. *arXiv preprint arXiv:2010.01950*, 2020.
- Kim, J.-H., Choo, W., Jeong, H., and Song, H. O. Co-mixup: Saliency guided joint mixup with supermodular diversity. *ICLR*, 2021.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012.
- Lazebnik, S., Schmid, C., and Ponce, J. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pp. 2169–2178. IEEE, 2006.
- Lin, S., Yu, T., Feng, R., Li, X., Jin, X., and Chen, Z. Local patch autoaugmt with multi-agent collaboration. *arXiv preprint arXiv:2103.11099*, 2021.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10012–10022, 2021.
- Long, J., Shelhamer, E., and Darrell, T. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
- Lopes, R. G., Yin, D., Poole, B., Gilmer, J., and Cubuk, E. D. Improving robustness without sacrificing accuracy with patch gaussian augmentation. *arXiv preprint arXiv:1906.02611*, 2019.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *ICLR*, 2018.
- Nguyen, K. and O’Connor, B. Posterior calibration and exploratory analysis for natural language processing models. *arXiv preprint arXiv:1508.05154*, 2015.
- Nilsback, M.-E. and Zisserman, A. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pp. 722–729. IEEE, 2008.
- Park, T., Efros, A. A., Zhang, R., and Zhu, J.-Y. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, pp. 319–345. Springer, 2020.

- Qin, Y., Zhang, C., Chen, T., Lakshminarayanan, B., Beutel, A., and Wang, X. Understanding and improving robustness of vision transformers through patch-based negative augmentation. *arXiv preprint arXiv:2110.07858*, 2021.
- Ren, S., He, K., Girshick, R., and Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252, 2015.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.
- Shocher, A., Cohen, N., and Irani, M. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3118–3126, 2018.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Sivic, J. and Zisserman, A. Video google: A text retrieval approach to object matching in videos. In *Computer Vision, IEEE International Conference on*, volume 3, pp. 1470–1470. IEEE Computer Society, 2003.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Upchurch, P., Gardner, J., Pleiss, G., Pless, R., Snavely, N., Bala, K., and Weinberger, K. Deep feature interpolation for image content changes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7064–7073, 2017.
- Wang, Y., Huang, G., Song, S., Pan, X., Xia, Y., and Wu, C. Regularizing deep networks with semantic data augmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Wightman, R. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. *CVPR*, 2017.
- Yoo, J., Ahn, N., and Sohn, K.-A. Rethinking data augmentation for image super-resolution: A comprehensive analysis and a new strategy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8375–8384, 2020.
- You, Y., Gitman, I., and Ginsburg, B. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6023–6032, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In *BMVC*, 2016.
- Zamir, S. W., Arora, A., Khan, S., Hayat, M., Khan, F. S., Yang, M.-H., and Shao, L. Multi-stage progressive image restoration. In *CVPR*, 2021.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Zhong, Z., Zheng, L., Kang, G., Li, S., and Yang, Y. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 13001–13008, 2020.

A. Augmentations

For all studied augmentations, we list a brief description and implementation details here. For YOCO, we do not modify the setting of augmentation operations (probability, magnitude, *etc.*) unless specified. We employ the Pytorch/Torchvision implementation for most augmentations.

A.1. Geometric transformation

Horizontal flip. Flip the image in the horizontal direction. We apply `Horizontal flip` with a probability of 0.5 as our default training setting.

Vertical flip. Flip the image in the vertical direction. We apply `Vertical flip` with a probability of 0.5.

A.2. Photometric transformation

Color jitter. Adjust the brightness, contrast, saturation, and hue of the image. We set $brightness = 0.4$, $contrast = 0.4$, $saturation = 0.4$, and $hue = 0.1$. The probability of being applied is 0.5.

Gaussian blur. Blur the image with `Gaussian blur`. We set $sigma = (0.1, 2.0)$, kernel size as approximately 10% of the image size, that is, for CIFAR images (32×32 resolution), a kernel size of 3 for image-level augmentation and kernel sizes of either 1 or 3 for YOCO. We apply `Gaussian blur` with a probability of 0.5.

A.3. Information dropping

Random erasing. Select a rectangular region in the image and erases its pixels. For CIFAR, we employ $scale = (0.02, 0.4)$, $ratio = (0.3, 3.3)$, and $value = 0$. For ImageNet, we change $scale$ from $(0.02, 0.4)$ to $(0.02, 0.33)$. `Random erasing` is applied with a probability of 0.5.

Cutout. Drops the pixels of a square region in the image. The mask size for `Cutout` is set to 25% of the image size (8×8 for CIFAR, 56×56 for ImageNet) and the location for dropping out is uniformly sampled. The dropped pixel is filled with 0 pixel. The probability of being applied is 0.5.

A.4. Search-based

AutoAug. Consists of 25 sub-policies (augmentations) searched from a pre-defined search space. Sub-policies span a wide range of augmentations. We employ the policy of CIFAR-10 and the policy of ImageNet for training on CIFAR datasets and ImageNet, respectively.

RandAug. Consists of searched augmentations from a reduced search space, where the magnitude strength of augmentations is controllable. RandAug can be applied uniformly across different tasks. For both CIFAR and ImageNet, we leverage an identical setting: 2 augmentation operations and a magnitude of 9.

A.5. Mix-based

Mixup. Produce an element-wise convex combination of two images. The label of training images changes with the mixing ratio and mixing images. We employ `Mixup` with a probability of 1, and a mixing ratio randomly sampled from $Beta(1, 1)$ for both CIFAR and ImageNet. For YOCO, the mixing image is identical for each piece, but with different mixing ratios. This way produces a slightly better performance than involving two mixing images. When YOCO is applied, the ImageNet *best* top-1 accuracy is **77.81** for one mixing image while being **77.72** for two mixing images.

CutMix. Combine `Mixup` and `Cutout`, randomly crops a patch from the mixing image, and pastes it into the corresponding position of the training image. Similar to `Mixup`, we employ `CutMix` with a probability of 1, and draw the mixing ratio from $Beta(1, 1)$ for both the CIFAR and ImageNet. We also find one mixing image outperforms two mixing images. For PreResNet18 on CIFAR-100 dataset, the top-1 accuracy of one mixing image and two mixing images are **80.70** and **79.73**, respectively.

For group of augmentations, the setting of each augmentation is identical to that augmentation applied individually.

B. Classification

Standard deviation. Table 11 reports the standard deviation of applying YOCO, and not applying YOCO, on every single model and every single dataset in classification tasks. Overall, image-level augmentation is more stable than YOCO on CIFAR-100. However, on CIFAR-10 and ImageNet, the stability of YOCO is on par to Image-level augmentation.

Models	CIFAR-10		CIFAR-100		ImageNet	
	YOCO	Image-level	YOCO	Image-level	YOCO	Image-level
PreResNet18	0.77	0.77	1.99	1.71	-	-
Xception	0.47	0.59	2.02	1.73	-	-
DenseNet121	0.65	0.55	1.84	1.72	-	-
ResNext50	0.55	0.67	2.02	1.77	-	-
WRN-28-10	0.54	0.49	1.92	1.40	-	-
ViT	0.74	0.70	2.03	1.67	-	-
Swin	0.70	0.62	1.97	1.67	-	-
Resnet50	-	-	-	-	0.26/0.29	0.23/0.33

Table 11. Standard deviation of YOCO and Image-level augmentation on every single model. For ImageNet, we report both *best* result (left) and *last* result (right).

C. Contrastive Learning

C.1. Implementation details

Augmentations. We follow the augmentations used in MoCo v2 and SimSiam. The augmentation procedure includes: Random resized crop with scale in $[0.2, 1.0]$, Color jitter with *brightness* = 0.4, *contrast* = 0.4, *saturation* = 0.4, *hue* = 0.1, and being applied with a probability of 0.5, Grayscale with a probability of 0.8, Guaussian blur with *sigma* = (0.1, 2.0) and is applied with a probability of 0.5. For convenience, we move Horizontal flip to the end, which is applied with a probability of 0.5.

Pre-training. We follow the training setting of MoCo v2 and SimSiam.

MoCo v2 is trained for 200 epochs with SGD as the optimizer. We use a weight decay of 0.0001 and a momentum of 0.9. The batch size is 256 and is distributed over 8 GPUs, and the initial learning rate is 0.03. The learning rate is multiplied by 0.1 at 120 and 160 epochs. The temperature is set to 0.2.

SimSiam is trained for 100 epochs only. The optimizer is SGD with a momentum of 0.9. The weight decay is 0.0001. We train with 8 GPUs and a batch size of 256. The initial learning rate is scaled to 0.05 with a cosine decay schedule.

Linear protocol. The linear protocol training setting follows MoCo v2 and SimSiam.

For MoCo v2, we train the linear classifier for 100 epochs. The initial learning rate is 30 and is multiplied by 0.1 at epochs 60 and 80. SGD optimizer is with a momentum of 0.9. The weight decay is 0. The linear classifier is trained with 8 GPUs and a batch size of 256.

SimSiam employs LARS optimizer (You et al., 2017). The linear classifier layer is trained for 90 epochs, and the initial learning rate is 0.2, which is scheduled with cosine decay. The weight decay is 0 and the momentum is 0.9. The batch size is 4096 and the training is distributed over 8 GPUs.

Fine-tuning on ImageNet. We follow the semi-supervised fine-tuning setting used in Jigsaw Clustering (Chen et al., 2021).

Pre-trained weights are fine-tuned for 100 epochs using SGD with a momentum of 0.9 as the optimizer. The batch size is 256 and we employ 4 GPUs for training. We do not apply any weight decay. The learning rate is multiplied by 0.1 at 30, 60, and 90 epochs. For 10% label fraction, we set the trunk learning rate to 0.01 while the last layer learning rate is set to 0.2. For 1% label fraction, the learning rate for the trunk is 0.02 and is 5 for the final layer.

D. Visualization

Augmented images. We show uncurated random samples on ImageNet. Each quadruplet presents the original image, an augmented image produced by image-level augmentation, and two augmented images created by YOCO. Visualizations are presented in Figure 4.

CAM visualization. We compute Grad-CAM (Selvaraju et al., 2017) for ResNet50 trained with YOCO and image-level

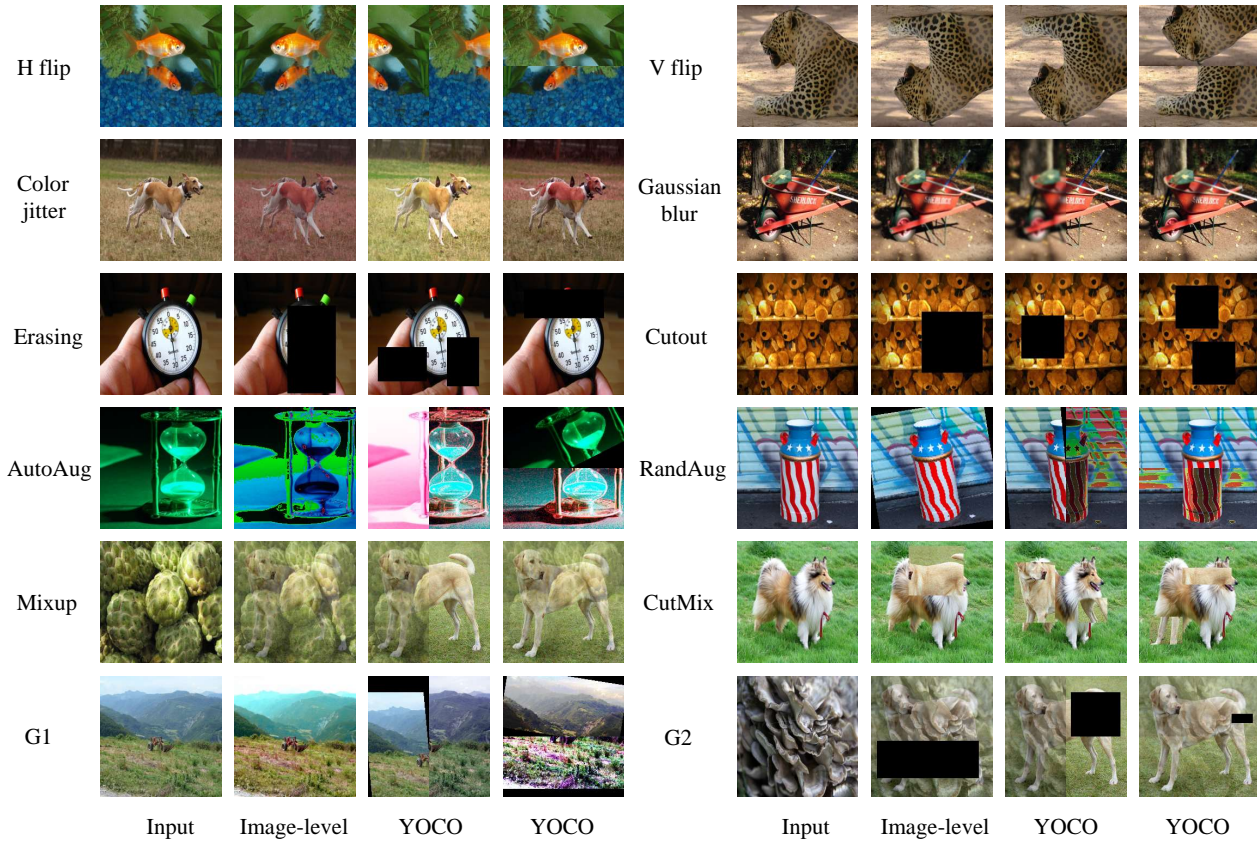


Figure 4. **Augmented images.** Randomly selected samples on ImageNet validation images. For each quadruplet, we show the original input image, augmented image from image-level augmentation, and two images from different cut dimensions produced by YOCO.

augmentation. We take `Horizontal flip` as an example. Visualizations are presented in Figure 5. For these samples, when only partial information is available (Banjo, Anemone fish), YOCO tends to focus more on that partial region instead of larger spatially distributed regions. YOCO also attends to focus on multiple structurally comprehensive regions, as shown in Hare, Entlebucher, and Yawl, where YOCO noticed two objects while image-level pays attention to a certain one. However, YOCO sometimes fails to comprehensively localize class-related cues, as presented in Desktop computer.

Differently classified images. To better understand the difference between image-level augmentation and YOCO. We present two collections containing randomly selected images that are differently classified by classifiers trained with YOCO and image-level augmentation. All images are taken from ImageNet validation set and images with private information are not presented. YOCO and image-level augmentation are both applied to `Horizontal flip`. Visualizations are presented in Figure 6. The set on the left side (image-level wrong, YOCO correct) contains more images that are with partial information only.

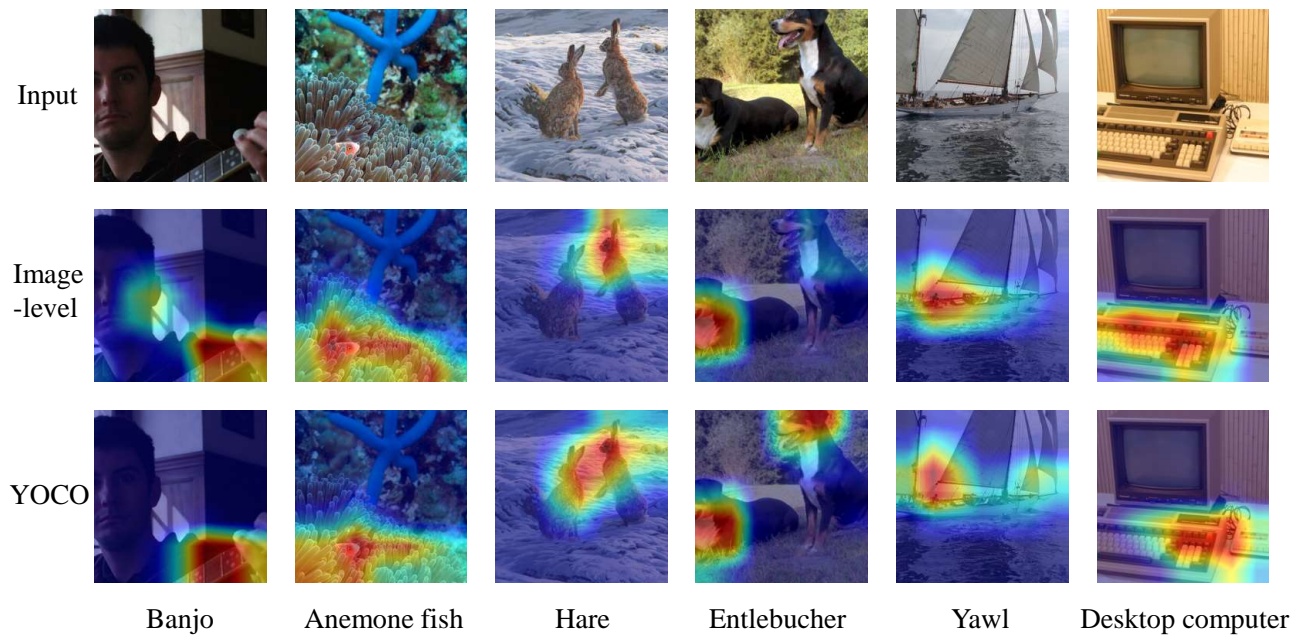


Figure 5. **Grad-CAM visualization.** We present the Grad-CAM results of models that are trained with image-level augmentation and YOCO.

You Only Cut Once: Boosting Data Augmentation with a Single Cut

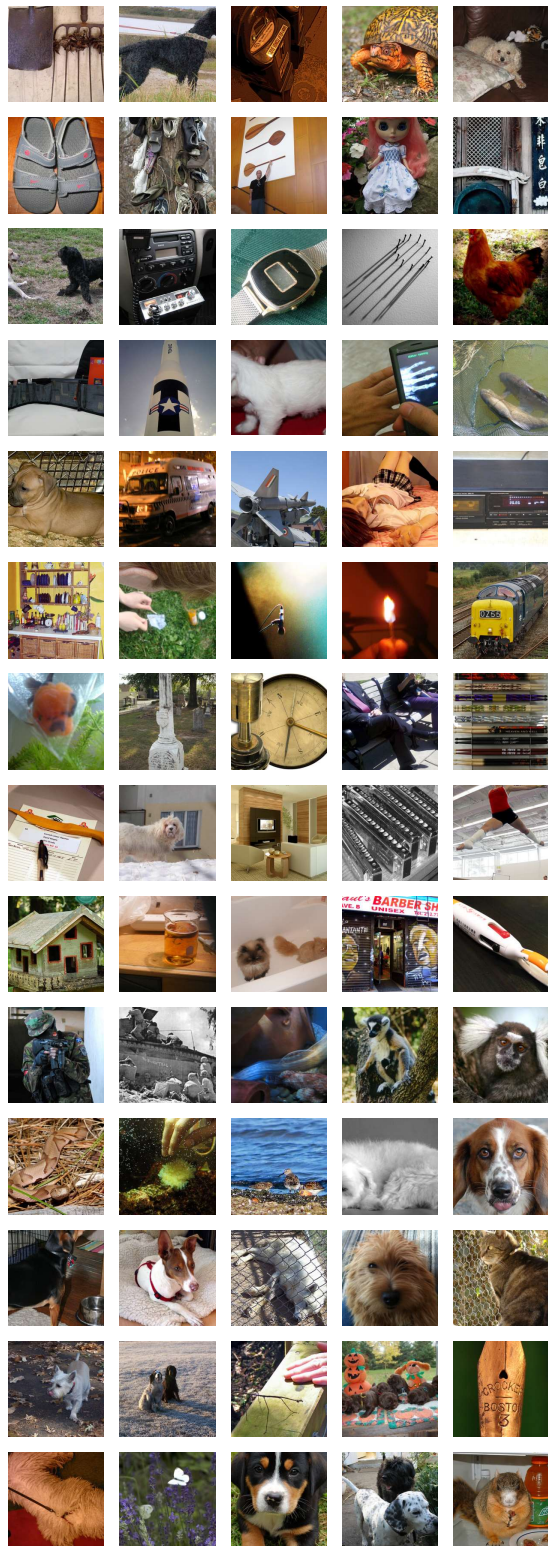
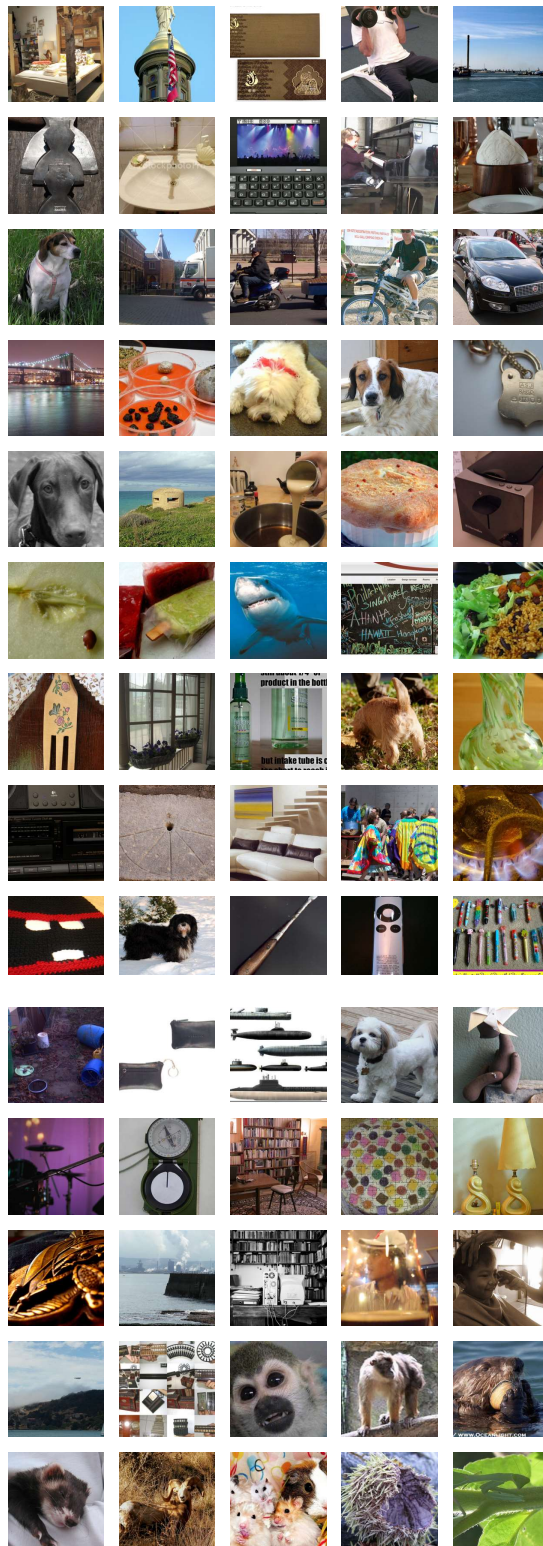


Image-level wrong, YOCO correct



YOCO wrong, image-level correct

Figure 6. Differently classified images. Uncurated ImageNet validation images that are differently classified by models trained with YOCO and image-level augmentation.