

Learning to Resize Images for Computer Vision Tasks

Hossein Talebi and Peyman Milanfar
 Google Research

Abstract

For all the ways convolutional neural nets have revolutionized computer vision in recent years, one important aspect has received surprisingly little attention: the effect of image size on the accuracy of tasks being trained for. Typically, to be efficient, the input images are resized to a relatively small spatial resolution (e.g. 224×224), and both training and inference are carried out at this resolution. The actual mechanism for this re-scaling has been an afterthought: Namely, off-the-shelf image resizers such as bilinear and bicubic are commonly used in most machine learning software frameworks. But do these resizers limit the on-task performance of the trained networks? The answer is yes. Indeed, we show that the typical linear resizer can be replaced with learned resizers that can substantially improve performance. Importantly, while the classical resizers typically result in better perceptual quality of the downsampled images, our proposed learned resizers do not necessarily give better visual quality, but instead improve task performance.

Our learned image resizer is jointly trained with a baseline vision model. This learned CNN-based resizer creates machine friendly visual manipulations that lead to a consistent improvement of the end task metric over the baseline model. Specifically, here we focus on the classification task with the ImageNet dataset [26], and experiment with four different models to learn resizers adapted to each model. Moreover, we show that the proposed resizer can also be useful for fine-tuning the classification baselines for other vision tasks. To this end, we experiment with three different baselines to develop image quality assessment (IQA) models on the AVA dataset [24].

1. Introduction

The emergence of deep neural networks along with large scale image datasets has led to major breakthroughs in machine visual recognition. Images in such datasets are typically obtained from the web and as a result have gone through various capture pipelines and post-processing steps. Besides these generally unknown processing operations, ef-




Figure 1. Our proposed framework for joint learning of the image resizer and recognition models.

Task	Model	Top-1 Error ↓	
		Bilinear Resizer	Proposed Resizer
Classification	Inception-v2 [34]	26.7%	24.0%
	DenseNet-121 [9]	33.1%	29.8%
	ResNet-50 [8]	24.7%	23.0%
	MobileNet-v2 [28]	29.5%	28.4%
IQA		PLCC ↑	
		Bicubic Resizer	Proposed Resizer
IQA	Inception-v2 [34]	0.662	0.686
	DenseNet-121 [9]	0.662	0.683
	EfficientNet-b0 [38]	0.642	0.671

Table 1. Summary of our results for image classification on ImageNet [26], and image quality assessment (IQA) on the AVA dataset [24].

ficient training of visual recognition CNNs requires additional image augmentations such as spatial resizing.

Image down-scaling is the most commonly used pre-processing module in classification models. The main reasons for spatial resizing are: (1) mini-batch learning through gradient descent requires the same spatial resolution for all images in a batch, (2) memory limitations prohibit training CNNs at high resolutions, and (3) large image sizes lead to slower training and inference. Given a fixed memory budget, there is a trade-off between the memory occupied by the spatial resolution and the batch size. This trade-off can have a significant impact on the accuracy of recognition CNNs [27, 39, 10, 12].

Currently rudimentary resizing methods such as nearest neighbor, bilinear, and bicubic are among the top adopted image resizers visual recognition systems. These resizers

are fast, and can be flexibly integrated into the train and test frameworks. However, these methods were developed decades before deep learning became the mainstream solution for visual recognition tasks, and hence are not optimized or adequate for machine perception.

Recent research on recognition-aware image processing has shown promising results on improving accuracy of classification models and simultaneously preserving the perceptual quality [20, 29]. This class of methods keep the classification model fixed, and only train the enhancement module. Meanwhile, there has been some effort on joint learning of both the pre-processor and the recognition model [2, 44, 19, 30, 18, 16, 45]. These algorithms set up training frameworks with hybrid losses that allow for learning better enhancement and recognition, concurrently. In practice, however, a recognition pre-processing operation such as resizing should not be optimized for better perceptual quality, because the end goal is for the recognition network to produce accurate results, not for the intermediate image to “look good” to a human observer.

In this paper we propose a novel image resizer that is jointly trained with classification models (see Figure 1), and is specifically designed to improve classification performance (see Table 1). To summarize our contributions:

- We couple our resizer with various classification models and show that it effectively adapts to each model and consistently improves over the baseline image classifier.
- The proposed resizer is not constrained by any pixel or perceptual loss, therefore our results present machine adaptive visual effects that differ from conventional image processing and super-resolution results.
- The proposed resizer model allows for down-scaling images at arbitrary scaling factors, hence we can conveniently search for the most optimal resolution for an underlying task.
- We expand the application of the proposed resizer to image quality assessment (IQA) and show that it successfully adapts to this task.

We purport the proposed method is the first pre-processing model that is specifically developed for vision tasks, and that aims to replace off-the-shelf resizers. Perhaps remote machine learning inference can be mentioned as a distinguishing application of the proposed resizer. In remote inference, transferring full resolution Mega pixel images from clients to servers imposes a significant latency on the system. While down-scaling by off-the-shelf resizers on the client side may reduce the latency issue, it may also negatively impact the recognition performance. The




Figure 2. Example of the proposed resizer trained for image classification on the ImageNet dataset [26]. The baseline classification model is Inception-v2 [34], and it is jointly trained with the resizer model shown in Fig. 3. The resized image fits the classification task better than the existing pre-processing resizers such as bilinear and bicubic.

proposed resizer, therefore, can be an alternative to the off-the-shelf resizers to effectively reduce the expected drop in the recognition performance.

Next, we briefly survey the research works related to this paper. Then, in Section 3 the proposed resizer model is discussed in detail. In Section 4 our results are presented, and finally we conclude in Section 5.

2. Related Work

There is ample reason to believe, given the past literature, that an optimized pre-processing module can improve performance of computer vision systems. For example, concurrent optimization of object recognition and enhancement tasks dates back to Zeiler et al. [43], where they used deconvolutional networks to learn robust features for image synthesis and analysis. Namboodiri et al. [25] show that enhancement algorithms such as super-resolution should be evaluated by classification driven metrics. More recently, a super-resolution algorithm [7] trained using an object detection loss showed superior results compared to the conventional super-resolution algorithms. Gondal et al. [40] studied the impact of super-resolution on recognition tasks by super-resolving down-sampled ImageNet [26] images and comparing classification accuracy for various up-scaling methods. They concluded that the choice of resizing method has a significant impact on the performance. More recently, Singh et al. [31] introduced a method for up-scaling very small images to improve face and digit recognition.

The impact of compression on recognition has also been studied in [21, 33, 35]. Luo et al. [21] show that JPEG quantization coefficients can be optimized to obtain lower bit-rates and at the same time preserve perceptual quality and recognition performance. Recently, several pre-editing methods for more efficient compression without sacrifice of the classification accuracy have been introduced [33, 36]. Typically these methods rely on some kind

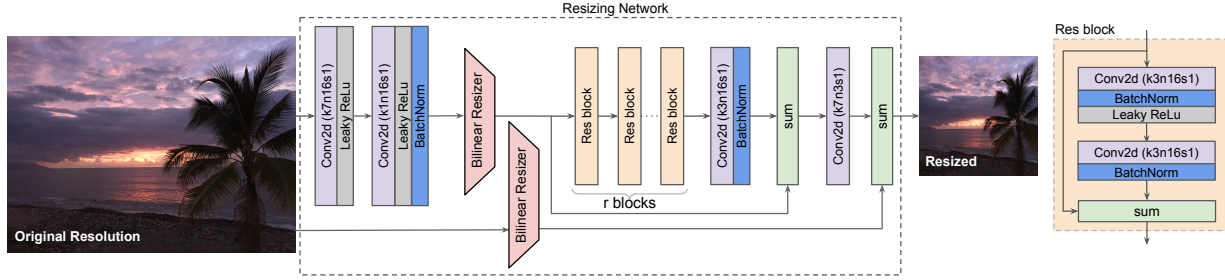


Figure 3. Our proposed CNN model for resizing images. The bilinear feature resizer allows for arbitrary up/down-scaling factors.

of rate-distortion-accuracy optimization.

Sharma et al. [30] train a generic enhancement model through joint learning with the classification network. This enhancement model does not alter the spatial resolution, and is trained via an L_2 loss on the enhanced image added to a cross-entropy classification loss. Denosing is yet another enhancement operation that has been successfully employed to improve the perceptual quality and recognition accuracy [5, 42]. The approach of Diamond et al. [5] focuses on low light imaging scenarios and shows that the best image processing algorithms for computer vision tasks are different from existing methods developed to produce visually pleasing photos. Also, the algorithm of Li et al. [15] shows that, dehazing can boost object detection and recognition performance on natural hazy images.

Recently Liu et al. [20] proposed an approach to improve machine interpretability of images by optimizing the recognition loss on the image processing network. They study super-resolution, denoising, and JPEG-deblocking as pre-processing operations, and show that the recognition performance gain can transfer when evaluated on different architectures and tasks. Our approach differs with [20] in that (1) we exclusively focus on improving the recognition performance, regardless of the perceptual quality, and (2) our pre-processing resizer is jointly trained with the recognition model, and consequently better adapts to the recognition architecture. This means the proposed model is not constrained to learn specific enhancements (e.g. denoising or deblurring), but rather it freely learns some unique machine friendly effects that result in a recognition gain (see Figure 2). This characteristic suits our model for applications where visually pleasing images are not the end goal.

3. Proposed Framework

In this section we introduce our resizer model, and discuss how we deploy it for training and testing image classification and IQA models.

Our resizer model is designed to be easily trainable, so it can be plugged into various learning frameworks and tasks. Also, it handles any arbitrary scaling factor, including up and down-scaling. This allows us to explore the resolution vs batch size trade-off, and as a result find the optimal resolution for the task in hand. In terms of perfor-

Filters \ Blocks	Blocks			
	$r = 1$	$r = 2$	$r = 3$	$r = 4$
$n=16$	11.87	16.48	21.08	25.69
$n=32$	38.08	56.51	74.94	93.37

Table 2. Number of parameters in the proposed resizer model are given in thousands. The number of residual blocks and the number of convolutional filters in Figure 3 are varied in this table.

mance, ideally, the net gain obtained by such adaptive resizing should surpass the extra computational complexity that the resizer adds to the system. These constraints make it almost impossible to use the existing super-resolution models [17, 14, 41, 46, 4, 23, 3]. On the other hand, image re-scaling methods such as bilinear and bicubic are, on their own, not trainable, and hence not suitable for this task. To this end, we design a model that satisfies these criteria.

3.1. Resizer Model

Our proposed resizer architecture is shown in Figure 3. Perhaps the most important characteristics of this model are (1) the bilinear feature resizing, and (2) the skip connection that accommodates combining the bilinearly resized image and the CNN features. The former factor allows for incorporation of features computed at original resolution into the model. Also, the skip connection accommodates for an easier learning process, because the resizer model can directly pass the bilinearly resized image into the baseline task. Note that the bilinear feature resizer shown in Figure 3 acts as a feed-forward bottleneck (down-scaling), but in principle it can also act as an inverse bottleneck as well (up-scaling). It is worth noting that unlike typical encoder-decoder architectures [22], the proposed architecture allows for resizing an image to any target size and aspect ratio. It is also important to highlight that performance of the learned resizer is barely dependent on the bilinear resizer choice, meaning it can be safely replaced with other off-the-shelf methods such as bicubic or Lanczos.

The residual blocks used in our model are inspired by [6, 14]. There are r identical residual blocks in our model and in our experiments we set $r = 1$ or 2. All intermediate convolutional layers have $n = 16$ kernels of size 3×3 . The first and the last layers consist of 7×7 kernels. The larger kernel size in the first layer allows for a 7×7 receptive field on the original image resolution. We also use batch normalization layers [11] and LeakyReLU activations with

a 0.2 negative slope coefficient.

The proposed resizer model is relatively lightweight and does not add a significant number of trainable parameters to the baseline task. The number of trainable parameters for various configurations of the CNN are shown in Table 2. These CNNs are significantly smaller than a baseline model such as ResNet-50 [8] which has about 23 million parameters. Performance of these configurations are compared in Section 4.4, where we show that even the lightest configuration with $n = 16, r = 1$ is quite effective.

3.2. Learning Losses

The resizer is jointly trained with the baseline model loss. Since our objective is to learn an optimal resizer for a baseline vision task, we do not apply any loss or regularization constraint on the resized image. A summary of the tasks explored in this paper is shown in Table 3.

3.2.1 Image Classification

The classification models are trained with the cross-entropy loss. More specifically, the loss is computed on the final logits with a Sigmoid layer. The ImageNet [26] classification challenge consists of 1000 object classes, hence, the final logits layer represents 1000 predicted classes. We also use the label-smoothing regularization proposed by Szegedy et al. [34]. The recognition loss can be expressed as

$$L_{\text{recog}} = - \sum_{k=1}^K \log(p_k) q'_k \quad (1)$$

where p and q' are the predictions and the smoothed labels, respectively, and K denotes the total number of classes. The smoothed label of an image with a ground truth label y is computed as $q'_k = (1 - \epsilon)\delta_{k,y} + \epsilon/K$, where $\delta_{k,y}$ is 1 when $k = y$, and 0 otherwise. We keep ϵ fixed as 0.1. The label regularization prevents the largest logit from dominating the other logits, leading to a less confident model and less overfitting.

3.2.2 Image Quality Assessment (IQA)

Our quality assessment models are trained through regression loss. Each image in the AVA dataset [24] has a histogram of human ratings, with scores ranging from 1 to 10. Following the recent work in [37], we use the Earth Mover's Distance (EMD) as our training loss. More specifically, the last layer of the baseline model is modified to have 10 logits, with a Softmax layer. The EMD loss can be represented as

$$L_{\text{quality}} = \left(\frac{1}{K} \sum_{k=1}^K |\text{CDF}(p_k) - \text{CDF}(q_k)|^d \right)^{1/d} \quad (2)$$

where $\text{CDF}(\cdot)$ is the cumulative distribution function. In our implementation we found $d = 2$ to be the most effective. Also, note that K is equal to 10 for the AVA dataset [24]. The EMD loss accommodates learning the distribution of human ratings. This has proven to be more effective than regressing to the mean ratings.

4. Experiments

Table 3 shows a summary of the specifics of our experiments. First, we train the baseline models on each dataset without the proposed resizer. For these cases we use the bilinear and the bicubic methods. These models are used as benchmarks to measure performance of the learned resizer. We also use these baselines to initialize the classification and IQA CNNs. For each baseline model and task, a separate resizer CNN is jointly trained with the baseline model. The resizer weights are randomly initialized.

To showcase the impact of the proposed resizer, we train the baseline model at various image resolutions with and without the resizer (shown in Table 4 and Table 5). More specifically, since the baseline models can be trained at any resolution, we vary the input size from the default 224×224 size to a larger 448×448 resolution. We use mini-batch learning, therefore the input image dimensions must be equal in each batch. To achieve this, and also allow the proposed resizer to train at higher resolutions, images in ImageNet and AVA are first resized by the bilinear or the bicubic method to a fixed resolution. The resizer's input resolution is always kept greater than or equal to its output resolution. Also, we apply the same resizing configuration at train and test time.

Feeding higher resolution images to CNNs leads to higher usage of computational resources. In principle, this extra computation should be justified by a corresponding boost in performance. This applies to all our experimental models with and without the learned resizer. To provide a fair comparison from the computational standpoint, in our experiments the floating point operations per second (FLOPS) are also reported (shown in Table 4 and Table 5).

We use Tensorflow [1] to train our networks with stochastic gradient descent with 4 NVIDIA V100 GPUs. Throughout our experiments we used the momentum optimizer [32] with a decay of 0.9. We used a learning rate of 0.05 when training from scratch, and 0.005 for fine-tuning. The learning rate is decayed every two epochs using an exponential rate of 0.94. Next, we discuss our results.

4.1. Classification

We select four baseline models to jointly train with the image resizer. We present our results for the ImageNet dataset in Table 4. We experiment with various resolutions and adjust the batch size to avoid exceeding our memory limits. The reported top-k error is the percentage of

Task	Benchmark Data	Resizer Initialization	Baseline Initialization	Baseline Models	Training Loss
Classification	ImageNet [26]	Random	Pre-trained	Inception-v2 [34], DenseNet-121 [9], ResNet-50 [8], MobileNet-v2 [28]	Cross-entropy
IQA	AVA [24]	Random	Pre-trained	EfficientNet-b0 [38], Inception-v2 [34], DenseNet-121 [9]	EMD

Table 3. Summary of the tasks implemented in this paper. The learned resizer shown in Figure 3 is jointly trained with each baseline model.

Classification Model	Resizer	Resizer's Input Resolution	Resizer's Output Resolution	Batch Size	Top-1 Error ↓	Top-5 Error ↓	Total FLOPS (Billions)
Inception-v2 [34]	Proposed	224×224	224×224	128	24.1%	7.5%	5.07
	Proposed	256×256	224×224	96	24.0%	7.4%	5.15
	Proposed	320×320	224×224	64	24.2%	7.3%	5.35
	Proposed	368×368	224×224	48	24.5%	7.4%	5.52
	Proposed	448×448	224×224	32	25.3%	7.9%	5.86
	Bilinear	original	224×224	128	26.7%	8.7%	3.88
	Bilinear	original	256×256	96	27.3%	9.1%	5.07
	Bilinear	original	320×320	64	27.3%	9.1%	7.92
	Bilinear	original	368×368	48	29.6%	10.4%	10.6
	Bilinear	original	448×448	32	30.6%	11.2%	15.52
DenseNet-121 [9]	Proposed	224×224	224×224	128	31.1%	11.6%	6.86
	Proposed	256×256	224×224	96	31.0%	11.4%	6.95
	Proposed	320×320	224×224	64	30.7%	11.1%	7.14
	Proposed	368×368	224×224	48	30.2%	10.9%	7.31
	Proposed	448×448	224×224	32	29.8%	10.8%	7.65
	Bilinear	original	224×224	128	33.1%	12.8%	5.67
	Bilinear	original	256×256	96	30.9%	11.7%	7.41
	Bilinear	original	320×320	64	29.9%	10.8%	11.57
	Bilinear	original	368×368	48	29.7%	10.7%	15.26
	Bilinear	original	448×448	32	31.5%	12.0%	22.68
ResNet-50 [8]	Proposed	224×224	224×224	128	23.7%	7.0%	8.16
	Proposed	256×256	224×224	96	23.8%	7.0%	8.24
	Proposed	320×320	224×224	64	23.4%	6.8%	8.43
	Proposed	368×368	224×224	48	23.0%	6.7%	8.61
	Proposed	448×448	224×224	32	23.7%	6.9%	8.95
	Bilinear	original	224×224	128	24.7%	7.5%	6.97
	Bilinear	original	256×256	96	23.5%	6.9%	9.10
	Bilinear	original	320×320	64	22.5%	6.3%	14.21
	Bilinear	original	368×368	48	22.1%	6.0%	19.17
	Bilinear	original	448×448	32	21.9%	5.8%	27.85
MobileNet-v2 [28]	Proposed	224×224	224×224	128	29.1%	10.1%	1.79
	Proposed	256×256	224×224	96	29.0%	10.1%	1.87
	Proposed	320×320	224×224	64	28.7%	9.9%	2.07
	Proposed	368×368	224×224	48	28.4%	9.8%	2.24
	Proposed	448×448	224×224	32	28.5%	9.8%	2.58
	Bilinear	original	224×224	128	29.5%	10.4%	0.60
	Bilinear	original	256×256	96	28.7%	9.6%	0.78
	Bilinear	original	320×320	64	27.2%	9.0%	1.23
	Bilinear	original	368×368	48	26.6%	8.6%	1.66
	Bilinear	original	448×448	32	26.1%	8.3%	2.40

Table 4. Classification errors on the ImageNet [26] validation set using various models. Each row represents a model trained with a different resizing configuration. The highlighted results represent the best performances among all models with 224×224 input resolution. Note that as the input resolution increases, the batch size is reduced to avoid memory consumption issues. Also, images are resized to fix resolutions before feeding them to the proposed resizer (shown under resizer's input resolution) to accommodate for mini-batch gradient descent.

time that the classifier does not return the correct class in the top k highest probability scores. We call the model trained with bilinear resizer and output resizing resolution 224×224 the default baseline. The highlighted results represent the best performances among models with 224×224 resolution. As can be seen, networks trained with the proposed resizer show an overall improvement over the default baseline. Comparing to the default baseline, DenseNet-121 and MobileNet-v2 baselines show the largest and smallest gains, respectively. Also, it is worth mentioning that for the Inception-v2, DenseNet-121, and ResNet-50 models, the proposed resizer performs better than the bilinear resizer with comparable FLOPS. However, training the MobileNet-

v2 model with bilinear resizer at higher resolutions is more effective than using the learned resizer with similar FLOPS.

Table 4 also shows that with or without the proposed resizer, increasing the input resolution benefits the performance of DenseNet-121, ResNet-50, and MobileNet-v2. The Inception-v2 model is an exception as it gains the most performance boost from training with larger batch sizes. It is worth noting that training resizers with equal input and output resolutions also results in improvement over the default baselines. In most cases, however, the best performance is obtained when the resizer's input is larger than its output.

We also present some examples to visually compare the




Figure 4. Examples of the proposed learned resizer trained together with various classification models on ImageNet [26]. The resizers lead to improved recognition performances.




Figure 5. Examples of the proposed learned resizer trained with various classification models on ImageNet [26]. The resizers lead to improved recognition performances.

trained resizers in Figure 4 and Figure 5. Perhaps the common feature among these results is the boost of the high frequency details. Interestingly these effects tend to make the classification model more effective. Aside from the MobileNet-v2 results, the other models tend to create overly sharpened results. This may intuitively explain the low performance gain obtained for MobileNet-v2. Overall, these effects do not meet the perceptual bar for human vision, but they surely improve the machine vision task.

4.2. Quality Assessment

We use 3 different baseline models to train with the AVA dataset [24]. The baseline models are initialized from pre-trained weights on ImageNet [26], and fine-tuned on the AVA dataset. Note that the resizer weights are initialized randomly. In this set of experiments we use the bicubic resizer as our baseline method. Our results are presented in Table 5. We measure the performance by reporting correlation between the mean ground truth score and mean predicted score. To this end, we use the Pearson linear correlation coefficient (PLCC), and Spearman rank correlation coefficient (SRCC). As can be seen, there is a consistent improvement over the baseline models. Also, it is worth noting that for Inception-v2 and DenseNet-121 models, the proposed resizer performs better than the bicubic resizer with comparable FLOPS. At higher FLOPS, EfficientNet seems to be a more challenging baseline for the learned resizer.

Examples of the trained resizers are shown in Figure 6.

The residual images show the difference between the bicubic and the learned resizer. As can be seen, the residual image for the Inception and the DenseNet models represent a lot of fine grain details. On the other hand, the EfficientNet resizer shows a strong color shift and modest detail manipulations.

4.3. Generalization

In this section generalization of the resizer model is discussed. To this end, we first jointly fine-tune the learned resizer with a target baseline that is different from the resizer’s default baseline. Then, we measure performance of the target baseline on the underlying task. We observed that fine-tuning on about 4 epochs of the training data suffices to adapt the resizer to the target model. This validation is a reasonable indicator of how well the trained resizers generalize to various architectures. Our results for classification and IQA are presented in Table 6 and Table 7. Each column shows an initialization checkpoint for the resizer model, and each row indicates a target baseline. These results show that a resizer trained for one baseline can be effectively used to develop a resizer for another baseline with minimal fine-tuning. In some cases such as the DenseNet and MobileNet models the fine-tuned resizers actually surpass the classification performance obtained by random initialization (see Table 6). The same observation holds true for the EfficientNet model in the IQA application (see Table 7). These improvements are perhaps because of the transfer learning ef-

Baseline Model	Resizer	Resizer's Input Resolution	Resizer's Output Resolution	Batch Size	PLCC \uparrow	SRCC \uparrow	Total FLOPS (Billions)
Inception-v2 [34]	Proposed	224 \times 224	224 \times 224	128	0.673	0.653	5.07
	Proposed	256 \times 256	224 \times 224	96	0.674	0.655	5.15
	Proposed	320 \times 320	224 \times 224	64	0.686	0.663	5.35
	Proposed	368 \times 368	224 \times 224	48	0.677	0.652	5.52
	Proposed	448 \times 448	224 \times 224	32	0.677	0.651	5.86
	Bicubic	original	224 \times 224	128	0.662	0.643	3.88
	Bicubic	original	256 \times 256	96	0.672	0.652	5.07
	Bicubic	original	320 \times 320	64	0.688	0.664	7.92
DenseNet-121 [9]	Proposed	224 \times 224	224 \times 224	128	0.672	0.644	6.86
	Proposed	256 \times 256	224 \times 224	96	0.672	0.645	6.95
	Proposed	320 \times 320	224 \times 224	64	0.683	0.655	7.14
	Proposed	368 \times 368	224 \times 224	48	0.675	0.644	7.31
	Proposed	448 \times 448	224 \times 224	32	0.673	0.642	7.65
	Bicubic	original	224 \times 224	128	0.662	0.636	5.67
	Bicubic	original	256 \times 256	96	0.672	0.644	7.41
	Bicubic	original	320 \times 320	64	0.694	0.666	11.57
EfficientNet-b0 [38]	Proposed	224 \times 224	224 \times 224	128	0.646	0.626	1.93
	Proposed	256 \times 256	224 \times 224	96	0.650	0.629	2.01
	Proposed	320 \times 320	224 \times 224	64	0.671	0.651	2.20
	Proposed	368 \times 368	224 \times 224	48	0.654	0.632	2.38
	Proposed	448 \times 448	224 \times 224	32	0.644	0.616	2.72
	Bicubic	original	224 \times 224	128	0.642	0.620	0.74
	Bicubic	original	256 \times 256	96	0.659	0.637	0.97
	Bicubic	original	320 \times 320	64	0.674	0.652	1.51
	Bicubic	original	368 \times 368	48	0.678	0.655	2.05
	Bicubic	original	448 \times 448	32	0.673	0.648	2.96

Table 5. IQA on the AVA dataset [24] with various models. Each row represents a model trained with a different resizing configuration. Performance of each model is quantified by the Pearson and Spearman correlations of the predicted and ground truth mean scores. The highlighted results represent the best performances among all models with 224 \times 224 input resolution. Note that as the input resolution increases, the batch size is reduced to avoid memory consumption issues. Also, images are resized to fix resolutions before feeding them to the proposed resizer (shown under resizer’s input resolution) to accommodate for mini-batch gradient descent.




Figure 6. Examples of the proposed learned resizer trained with various IQA models on the AVA dataset [24]. (c), (e), and (f) are results from trained resizers with respective base models. (d), (f), and (h) represent the difference between bicubic and learned resizers.

fect.

We also tried the above cross-model validation without fine-tuning, however, performance mostly degraded. This is likely because (1) the proposed resizer is exclusively trained for one baseline model, and (2) no intermediate pixel loss is used during training.

For cross-dataset validation we train the resizer on Ima-

geNet, and test it with a classifier trained on the CIFAR-10 benchmark [13]. With the addition of the learned resizer to the classification task on CIFAR-10, the top-1 error of ResNet-50 (6.9%) baseline drops by 0.4%. We observed a similar favorable trend in other baseline models.

We also compare our performance with [20] in Table 9. We followed the authors instructions in [20] to generate

Initial Target	Top-1 Error ↓				Top-5 Error ↓			
	Inception-v2	DenseNet-121	ResNet-50	MobileNet-v2	Inception-v2	DenseNet-121	ResNet-50	MobileNet-v2
Inception-v2	24.5%	24.6%	24.5%	24.6%	7.4%	7.5%	7.5%	7.4%
DenseNet-121	29.7%	30.2%	29.7%	30.1%	10.7%	10.9%	10.6%	10.9%
ResNet-50	22.9%	23.0%	23.0%	23.1%	6.5%	6.4%	6.7%	6.7%
MobileNet-v2	28.0%	28.2%	28.0%	28.4%	9.6%	9.6%	9.6%	9.8%

Table 6. Generalization of the resizer models for image classification [26]. The learned resizers are trained with the initial baseline and then jointly fine-tuned with the target baseline model. The resizer’s input and output resolutions are 368×368 and 224×224 , respectively.

Initial Target	PLCC ↑			SRCC ↑		
	Inception-v2	DenseNet-121	EfficientNet-b0	Inception-v2	DenseNet-121	EfficientNet-b0
Inception-v2 [34]	0.677	0.672	0.670	0.652	0.649	0.649
DenseNet-121 [9]	0.672	0.675	0.671	0.645	0.644	0.642
EfficientNet-b0 [38]	0.660	0.653	0.654	0.636	0.632	0.632

Table 7. Generalization of the resizer models for IQA [24]. The learned resizers are trained with the initial baseline and then jointly fine-tuned with the target baseline model. The resizer’s input and output resolutions are 368×368 and 224×224 , respectively.

Task	Model	Top-1 Error ↓				Top-5 Error ↓			
		$r = 1$ $n = 16$	$r = 2$ $n = 16$	$r = 1$ $n = 32$	$r = 2$ $n = 32$	$r = 1$ $n = 16$	$r = 2$ $n = 16$	$r = 1$ $n = 32$	$r = 2$ $n = 32$
Classification	Inception-v2 [34]	24.5%	25.5%	25.6%	26.1%	7.4%	8.0%	7.9%	8.3%
	DenseNet-121 [9]	30.2%	29.8%	29.9%	29.8%	10.9%	10.8%	10.8%	10.8%
	ResNet-50 [8]	23.0%	23.4%	23.4%	23.3%	6.7%	6.6%	6.7%	6.8%
	MobileNet-v2 [28]	28.4%	28.5%	28.4%	28.3%	9.8%	9.6%	9.7%	9.7%
IQA	Inception-v2 [34]	0.677	0.677	0.675	0.676	0.652	0.654	0.643	0.643
	DenseNet-121 [9]	0.675	0.677	0.670	0.671	0.644	0.645	0.629	0.630
	EfficientNet-b0 [38]	0.654	0.652	0.646	0.648	0.632	0.630	0.625	0.628

Table 8. Effect of the resizer model parameters on the classification [26], and image quality assessment (IQA) [24]. Parameters r and n denote the number of residual blocks and convolutional filters, respectively. These parameters are presented in Figure 3. The resizer’s input and output resolutions are 368×368 and 224×224 , respectively.

Model	Pre-processor	Down-scaled ($4\times$)	Noise	JPEG
ResNet-50	[20]	31.8%	29.1%	34.9%
	ours	31.1%	28.6%	34.5%

Table 9. Top-1 classification error on distorted and then enhanced ImageNet images. Results from [20] are their best performing models.

the distorted images with Gaussian noise (standard deviation 0.1), spatial down-scaling ($4\times$), and JPEG compression with quality factor 10. We trained the resizer to increase the resolution by a factor of 4 for the down-scaled inputs, and kept the resolution unchanged for the other distortions. As can be noted from Table 9, the proposed resizer outperforms [20] consistently. Also, it is worth pointing out that our pre-processor model obtains the largest margin of improvement for down-scaled images.

4.4. Ablation

In this section effects of our design choices in the resizer model are discussed. We vary the number of residual blocks r , and the number of filters n (see Figure 3), and report the performance of the jointly trained baseline models. Note that so far in the experimental results we have employed CNN resizers with the default configuration in which $r = 1$ and $n = 16$.

Our results for the classification and IQA tasks with various configurations are presented in Table 8. In the classifi-

cation task, as the resizer model gets bigger, the DenseNet and the MobileNet baselines show modest improvements over the default configuration. However, Inception and ResNet do not benefit from larger number of parameters in the resizer. A similar trend can be observed in the IQA task.

Perhaps one of the reasons for the non-growing performance of the larger resizer models is the lowered batch size. Note that given limited memory, larger resizers have to be trained with smaller batch sizes. This factor may inadvertently limit the observed performance gain.

5. Conclusions

We presented a framework for learning pre-processing effects that boost the performance of image recognition models. We focused on image resizing, and did not apply an intermediate pixel or perceptual loss on the resized images, hence the results are exclusively optimized for machine vision tasks. Our experiments show that task-optimized deep vision models can benefit from replacing traditional image resizers with learned resizers. We believe that customized pre-processing algorithms for machine vision tasks have not been studied extensively, and given the impact shown in this paper, there is significant room for research in this area. As part of future work we will extend our model to other vision tasks.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016. 4
- [2] Yancheng Bai, Yongqiang Zhang, Mingli Ding, and Bernard Ghanem. Finding tiny faces in the wild with generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–30, 2018. 2
- [3] David Berthelot, Peyman Milanfar, and Ian Goodfellow. Creating high resolution images with a latent adversarial generator. *arXiv preprint arXiv:2003.02365*, 2020. 3
- [4] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3086–3095, 2019. 3
- [5] Steven Diamond, Vincent Sitzmann, Stephen Boyd, Gordon Wetzstein, and Felix Heide. Dirty pixels: Optimizing image classification architectures for raw sensor data. *arXiv preprint arXiv:1701.06487*, 2017. 3
- [6] Sam Gross and Michael Wilber. Training and investigating residual nets. *Facebook AI Research*, 6, 2016. 3
- [7] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Task-driven super resolution: Object detection in low-resolution images. *arXiv preprint arXiv:1803.11316*, 2018. 2
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 4, 5, 8
- [9] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 1, 5, 7, 8
- [10] Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Chen, HyounJoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. In *Advances in neural information processing systems*, pages 103–112, 2019. 1
- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 3
- [12] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 1
- [13] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 7
- [14] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 3
- [15] Boyi Li, Xiulian Peng, Zhangyang Wang, Jizheng Xu, and Dan Feng. Aod-net: All-in-one dehazing network. In *Proceedings of the IEEE international conference on computer vision*, pages 4770–4778, 2017. 3
- [16] Boyi Li, Xiulian Peng, Zhangyang Wang, Jizheng Xu, and Dan Feng. End-to-end united video dehazing and detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 2
- [17] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 3
- [18] Ding Liu, Bihan Wen, Xianming Liu, Zhangyang Wang, and Thomas S Huang. When image denoising meets high-level vision tasks: A deep learning approach. *arXiv preprint arXiv:1706.04284*, 2017. 2
- [19] Feng Liu, Ronghang Zhu, Dan Zeng, Qijun Zhao, and Xiaoming Liu. Disentangling features in 3d face shapes for joint face reconstruction and recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5216–5225, 2018. 2
- [20] Zhuang Liu, Tinghui Zhou, Hung-Ju Wang, Zhiqiang Shen, Bingyi Kang, Evan Shelhamer, and Trevor Darrell. Transferable recognition-aware image processing. *arXiv preprint arXiv:1910.09185*, 2019. 2, 3, 7, 8
- [21] Xiyang Luo, Hossein Talebi, Feng Yang, Michael Elad, and Peyman Milanfar. The rate-distortion-accuracy tradeoff: Jpeg case study. *arXiv preprint arXiv:2008.00605*, 2020. 2
- [22] Xiao-Jiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. *arXiv preprint arXiv:1603.09056*, 2016. 3
- [23] Peyman Milanfar. *Super-resolution imaging*. CRC press, 2017. 3
- [24] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2408–2415. IEEE, 2012. 1, 4, 5, 6, 7, 8
- [25] Vinay P Namboodiri, Vincent De Smet, and Luc Van Gool. Systematic evaluation of super-resolution using classification. In *2011 Visual Communications and Image Processing (VCIP)*, pages 1–4. IEEE, 2011. 2
- [26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 1, 2, 4, 5, 6, 8
- [27] Carl F Sabottke and Bradley M Spieler. The effect of image resolution on deep learning in radiography. *Radiology: Artificial Intelligence*, 2(1):e190015, 2020. 1
- [28] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the*

- IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 1, 5, 8
- [29] Walter Scheirer, Rosaura VidalMata, Sreya Banerjee, Brandon RichardWebster, Michael Albright, Pedro Davalos, Scott McCloskey, Ben Miller, Asongu Tambo, Sushobhan Ghosh, et al. Bridging the gap between computational photography and visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 2
- [30] Vivek Sharma, Ali Diba, Davy Neven, Michael S Brown, Luc Van Gool, and Rainer Stiefelhagen. Classification-driven dynamic image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4033–4041, 2018. 2, 3
- [31] Maneet Singh, Shruti Nagpal, Richa Singh, and Mayank Vatsa. Dual directed capsule network for very low resolution image recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 340–349, 2019. 2
- [32] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147. PMLR, 2013. 4
- [33] Satoshi Suzuki, Motohiro Takagi, Kazuya Hayase, Takayuki Onishi, and Atsushi Shimizu. Image pre-transformation for recognition-aware image compression. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2686–2690. IEEE, 2019. 2
- [34] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 1, 2, 4, 5, 7, 8
- [35] Khalid Tahboub, David Güera, Amy R Reibman, and Edward J Delp. Quality-adaptive deep learning for pedestrian detection. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 4187–4191. IEEE, 2017. 2
- [36] Hossein Talebi, Damien Kelly, Xiyang Luo, Ignacio Garcia Dorado, Feng Yang, Peyman Milanfar, and Michael Elad. Better compression with deep pre-editing. *arXiv preprint arXiv:2002.00113*, 2020. 2
- [37] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018. 4
- [38] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019. 1, 5, 7, 8
- [39] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. In *Advances in neural information processing systems*, pages 8252–8262, 2019. 1
- [40] Muhammad Waleed Gondal, Bernhard Scholkopf, and Michael Hirsch. The unreasonable effectiveness of texture transfer for single image super-resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 2
- [41] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 3
- [42] Jonghwa Yim and Kyung-Ah Sohn. Enhancing the performance of convolutional neural networks on quality degraded datasets. In *2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8. IEEE, 2017. 3
- [43] Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pages 2528–2535. IEEE, 2010. 2
- [44] Kaipeng Zhang, Zhanpeng Zhang, Chia-Wen Cheng, Winston H Hsu, Yu Qiao, Wei Liu, and Tong Zhang. Super-identity convolutional neural network for face hallucination. In *Proceedings of the European conference on computer vision (ECCV)*, pages 183–198, 2018. 2
- [45] Chen Zhao and Bernard Ghanem. Thumbnet: One thumbnail image contains all you need for recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1506–1514, 2020. 2
- [46] Ruofan Zhou and Sabine Susstrunk. Kernel modeling super-resolution on real low-resolution images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2433–2443, 2019. 3