

Beyond Image Super-Resolution for Image Recognition with Task-Driven Perceptual Loss

Jaeha Kim¹ Junghun Oh¹ Kyoung Mu Lee^{1,2}

¹Dept. of ECE&ASRI, ²IPAI, Seoul National University, Korea

jhkim97s2@gmail.com, {dh6dh, kyoungmu}@snu.ac.kr

Abstract

In real-world scenarios, image recognition tasks, such as semantic segmentation and object detection, often pose greater challenges due to the lack of information available within low-resolution (LR) content. Image super-resolution (SR) is one of the promising solutions for addressing the challenges. However, due to the ill-posed property of SR, it is challenging for typical SR methods to restore task-relevant high-frequency contents, which may dilute the advantage of utilizing the SR method. Therefore, in this paper, we propose *Super-Resolution for Image Recognition (SR4IR)* that effectively guides the generation of SR images beneficial to achieving satisfactory image recognition performance when processing LR images. The critical component of our SR4IR is the task-driven perceptual (TDP) loss that enables the SR network to acquire task-specific knowledge from a network tailored for a specific task. Moreover, we propose a cross-quality patch mix and an alternate training framework that significantly enhances the efficacy of the TDP loss by addressing potential problems when employing the TDP loss. Through extensive experiments, we demonstrate that our SR4IR achieves outstanding task performance by generating SR images useful for a specific image recognition task, including semantic segmentation, object detection, and image classification. The implementation code is available at <https://github.com/JaehaKim97/SR4IR>.

1. Introduction

After the pioneering AlexNet [31], numerous deep neural network-based methods have achieved superior performance in various image recognition tasks, such as image classification [13, 20, 23, 24, 40, 53, 56], object detection [37, 39, 47, 48], and semantic segmentation [6, 42, 64, 72]. In real-world applications, such as medical image recognition and surveillance systems, a network for a task, which we refer to as a *task network*, may encounter input images containing low-resolution (LR) contents due to

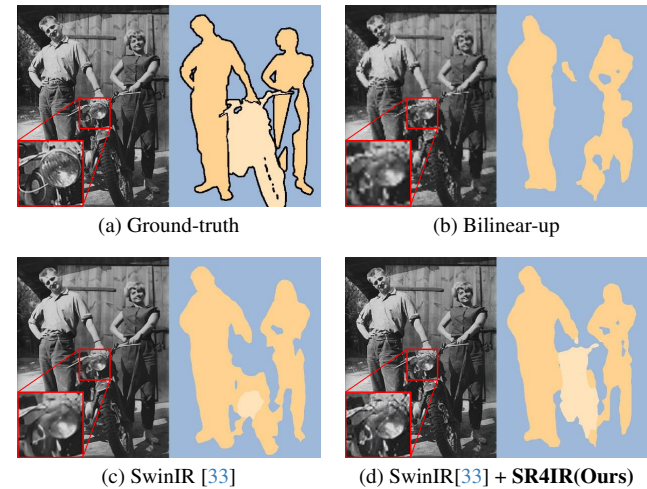


Figure 1. **Visualizations and comparisons of the results on semantic segmentation task with PASCAL VOC dataset [14].** (a) The ground-truth image and the class label map. (b) The bilinear-upscaled image and the predicted result when trained on bilinear-upscaled images. (c) The SR result of the SwinIR [33] and the predicted result when trained on the SR images. (d) the SR result of the SwinIR and the predicted result when trained by the proposed SR4IR framework. We use DeepLabV3 [7] as a task network and the downsampling scale factor is x4.

various factors, including limitations of cameras or storage capabilities and the presence of tiny objects. In this circumstance, achieving satisfactory task performance is limited by the lack of high-frequency components crucial for a task.

In this paper, we address these challenges by employing single image super-resolution (SR) methods, as SR can serve as a prominent pre-processing tool to restore missing high-frequency details in LR images. Conventionally, SR aims to restore high-resolution (HR) images mainly via pixel-wise loss [18, 29, 33, 35, 70, 74] or perceptual loss [5, 15, 27]. However, due to the ill-posed nature of SR, these approaches may not ensure the reconstruction of features crucial for achieving better performance for a specific task, thereby leading to marginal performance improvement from SR. Thus, in this paper, our primary goal is to develop a comprehensive framework, named *Super-Resolution for*

Image Recognition (SR4IR), wherein the SR and task network interplay harmoniously to generate SR images, especially focusing on incorporating task-relevant features, which are exploited to enhance task performance.

To attain our goal, we focus on leveraging the concept of perceptual loss, as it is devised to encourage SR images to mimic the HR counterparts within a feature space. Typically, the perceptual loss is computed using VGGNet [53] pre-trained on ImageNet dataset [49]. However, we note that employing perceptual loss in the conventional manner results in modest improvements in task performance. Instead, we propose a Task-Driven Perceptual (TDP) loss that enables the emulation of HR images within the feature space of a task network. By doing so, our TDP loss facilitates the restoration of high-frequency details, especially beneficial for enhancing task performance.

The effectiveness of our TDP loss is significantly influenced by the knowledge that the task network acquires. For high-level vision tasks, it is widely acknowledged that a network can fall into learning non-intrinsic features, leading to biased feature representations [16]. A network possessing biased knowledge on a task can diminish the efficacy of the TDP loss, as it may fail to deliver meaningful signals to the SR network for the restoration of other valuable features. To tackle the problem, we introduce a novel data augmentation strategy for training the task network, called Cross-Quality patch Mix (CQMix), which blends SR and HR images at the patch level to randomly eliminate high-frequency components. Consequently, CQMix can prevent a task network from learning shortcut features, thereby allowing us to fully harness the benefits of TDP loss.

In addition to calculating perceptual loss using a task network, we introduce further variations to the perceptual loss by applying it to a task network undergoing training instead of a pre-trained one. This strategy substantially improves the effectiveness of our TDP loss, since a pre-trained (then fixed) task network cannot learn to produce meaningful features from SR images. To this end, we constitute the SR4IR as an alternate training framework that trains the SR and task network alternately throughout the training procedure. In the first phase, we train the SR network using the TDP loss with the temporarily frozen task network. Subsequently, we train only the task network with task-specific loss using data augmented by the CQMix.

To evaluate the proposed methods, we have examined the integration of SR with three representative image recognition tasks: semantic segmentation, object detection, and image classification. For each task, we evaluate our method on two representative SR architectures, EDSR [35] and SwinIR [33], with scale factors of $\times 4$ and $\times 8$. Our extensive experiments clearly demonstrate that our SR4IR significantly improves task performance while achieving perceptually appealing SR results compared to baseline methods,

as shown in Figure 1. These results indicate that our framework can be generally adopted in various applications. We summarize our contributions as follows:

- To our knowledge, we are the first to introduce a comprehensive SR framework that addresses challenges posed by LR contents across various image recognition tasks.
- We propose the task-driven perceptual (TDP) loss that facilitates learning to restore task-related features acquired by a task network, enhancing task performance.
- We propose the cross-quality patch mix and the alternate training framework to address the potential problems of TDP loss, further enhancing the efficacy of TDP loss.
- The extensive experiments demonstrate that our SR4IR is generally applicable across a wide range of image recognition tasks to enhance task performance while generating perceptually satisfying SR images.

2. Related Work

Single Image Super-Resolution. In the era of deep learning, deep neural network-based methods have risen to prominence in the field of SR. After the introduction of the simple deep-learning-based SR network SRCNN [12], VDSR [29] brought about innovation in network architecture design by stacking deeper convolutional layers with the inclusion of residual connections. EDSR [35] and SRResNet [32] brought the Residual Block concept into SR networks, and subsequent methods such as RDN [74], DBPN [18] further evolved the Dense Residual Block, resulting in high-performance network structures. SAN [11], HAN [44], and RCAN [73] enhance the SR performance by incorporating channel attention mechanisms, enabling the exploitation of useful channel-wise features. Recently, transformer-based methods such as SwinIR [33], Restormer [67], and Uformer [62] have demonstrated remarkable SR performance by harnessing the powerful capabilities of Vision Transformers [13].

On the other hand, perceptual super-resolution (SR) methods constitute an alternative approach dedicated to restoring photorealistic SR results. Johnson *et al.* [27] introduced perceptual loss, which minimizes the distance in the feature space and has shown its effectiveness in reconstructing high-frequency components in SR. Zhang *et al.* [71] illustrated that the perceptual quality of images is not directly correlated with the peak signal-to-noise ratio (PSNR), then introduced the perceptual metric known as LPIPS. Subsequent to the development of adversarial loss [17], GAN-based SR methods [32, 34, 45, 51, 54, 55, 59, 60] have been successful in achieving realistic SR results. More recently, diffusion-based SR methods [8, 28, 43, 50, 61, 63, 65] have further pushed the boundaries by leveraging a diffusion-based model [21] to produce even more realistic SR results.

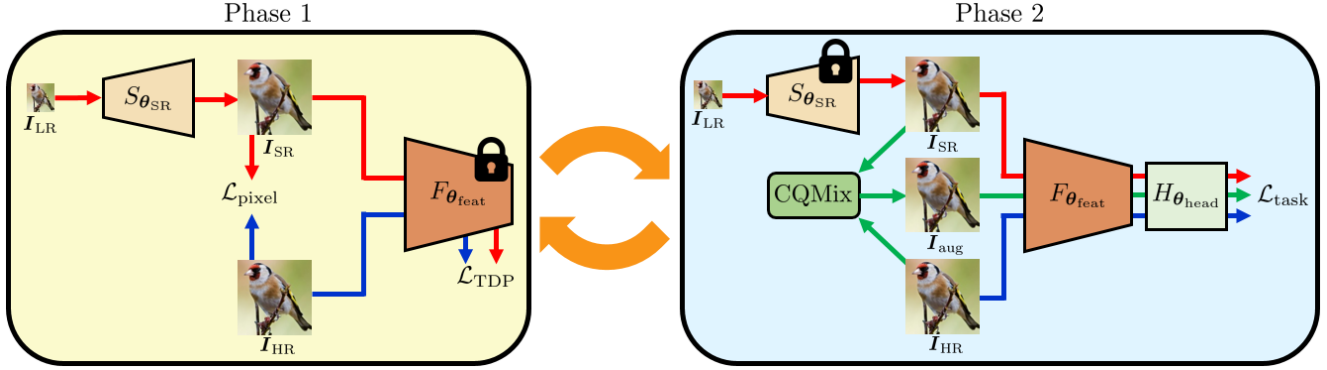


Figure 2. **Overview of the proposed SR4IR framework.** Our SR4IR framework consists of two training phases, where SR and task networks are alternately trained. During the first phase, SR4IR updates the SR network using the TDP loss, which is introduced in Section 3.1, while the task network is temporarily frozen. In the second phase, SR4IR updates the task network using the proposed data augmentation strategy called CQMix, which is introduced in Section 3.2, while the SR network is temporarily frozen.

Image restoration for high-level vision tasks. There exist pioneering recent works [9, 10, 38, 41, 52, 75] focused on designing image restoration techniques that offer advantages in high-level computer vision tasks. Zhou *et al.* [76] introduced a weight map that uses pixel loss primarily in high-frequency regions, showing its effectiveness in image classification. TDSR [19] presented an object detection-driven SR method that updates the SR network using a combination of reconstruction and detection losses and provides a comprehensive study on balancing the joint loss. Wang *et al.* [58] proposed adding an SR branch to the segmentation branch, enabling the effective utilization of HR representation. Bai *et al.* [2, 3] have demonstrated that incorporating adversarial loss through a multitask generative adversarial network can further enhance the performance of low-quality object detection. In contrast to previous approaches, our method successfully restores task-specific high-frequency components by leveraging the concept of perceptual loss. Moreover, we are the first to introduce a generalizable task-combined SR framework that achieves remarkable performance across a range of image recognition tasks.

3. Proposed method

Overview. Figure 2 shows the overview of the proposed Super-Resolution for Image Recognition (SR4IR) framework. The SR4IR comprises two distinct training phases that are alternately executed to train either the super-resolution (SR) or the task network. In the first phase, we train the SR network using the Task-Driven Perceptual (TDP) loss, which is designed to guide the SR network to restore high-frequency details relevant to the task (Section 3.1). In the second phase, we train the task network using the novel Cross-Quality patch Mix (CQMix) augmentation method, which prevents the task network from learning biased features, further enhancing the effectiveness of the TDP loss (Section 3.2). Lastly, we present the effect of alternate training with respect to TDP loss and summarize our final training objectives in each training phase (Section 3.3).

Notation. Let $I_{HR} \in \mathbb{I}_{HR}$ and $I_{LR} \in \mathbb{I}_{LR}$ denote the high-resolution (HR) and low-resolution (LR) images, paired with their corresponding task labels $\mathbf{y} \in \mathbb{Y}$ (e.g., class labels in image classification task). $S_{\theta_{SR}} : \mathbb{I}_{LR} \rightarrow \mathbb{I}_{SR}$ denotes the SR network parameterized by θ_{SR} where \mathbb{I}_{SR} is the super-resolved image set. Given the SR images I_{SR} , we predict the corresponding labels using the task network: $H_{\theta_{head}} \circ F_{\theta_{feat}} : \mathbb{I}_{SR} \rightarrow \mathbb{Y}$, where $F_{\theta_{feat}}$ and $H_{\theta_{head}}$ indicate the feature extractor and the task-specific head module (e.g., the last fully-connected layer for image classification tasks) parameterized by θ_{feat} and θ_{head} , respectively.

Problem Definition. The goal of SR4IR is to produce photo-realistic SR images from I_{LR} that can benefit the subsequent image recognition task, which is formally given by:

$$\min_{\theta_{SR}, \theta_{feat}, \theta_{head}} \mathcal{L}_{SR}(S_{\theta_{SR}}(I_{LR}), I_{HR}) + \mathcal{L}_{task}(H_{\theta_{head}} \circ F_{\theta_{feat}} \circ S_{\theta_{SR}}(I_{LR}), \mathbf{y}), \quad (1)$$

where the first and second terms are used to evaluate the quality of the super-resolved results (e.g., LPIPS [71]) and the performance for the target high-level vision task (e.g., classification accuracy), respectively.

3.1. Task-Driven Perceptual Loss

High-resolution (HR) images contain high-frequency information that is crucial to solving high-level vision tasks. However, due to the ill-posed nature of the SR, such task-relevant features are hard to restore using traditional pixel-wise reconstruction loss. Even with the use of perceptual loss [5, 15, 27], designed to produce realistic SR results, there is no guarantee of successfully reconstructing useful features for a specific task. To address this challenge, we propose task-driven perceptual (TDP) loss that effectively guides the SR network to produce images containing high-frequency details relevant to the task. Specifically, the TDP loss aims to maximize the similarity between HR and SR

images within the latent space of the feature extractor of the task network $F_{\theta_{\text{feat}}}$, whose outputs are substantially relevant to the task. We define the TDP loss as follows:

$$\mathcal{L}_{\text{TDP}} = \|F_{\theta_{\text{feat}}}(I_{\text{SR}}) - F_{\theta_{\text{feat}}}(I_{\text{HR}})\|_1. \quad (2)$$

3.2. Cross-Quality Patch Mix

The effectiveness of TDP loss is inherently related to the training procedure of the task network. For high-level vision tasks, it is widely known that a network is susceptible to shortcut learning [16], which means that a network can rely on biased and non-intrinsic features in training data, leading to poor generalization ability. In particular, learning biased features can pose challenges in employing the TDP loss, as biased features learned by the task network can hinder the SR network from acquiring comprehensive features present in HR images through the TDP loss. For instance, if a task network exclusively learns ‘beak’ features to classify bird species, the TDP loss may struggle to effectively guide SR images to restore other key high-frequency contents, such as ‘wings’ or ‘feet’. This limitation could hinder the overall task performance boosting that could be attained through the SR method.

To address this issue, we introduce a novel data augmentation strategy called **Cross-Quality patch Mix (CQMix)** when training the task network in the second phase. CQMix is designed to enhance the effectiveness of the TDP loss by encouraging the task network to use diverse high-frequency content as clues. Specifically, CQMix augments an input image by randomly choosing either HR or SR patches for each grid region, as shown in Figure 3, where the number of patches is hyper-parameter. By doing so, CQMix creates an image in which high-frequency components are randomly erased at a patch level, effectively preventing the task network from learning specific high-frequency features as shortcuts. Moreover, it is noteworthy that our CQMix is generally applicable to various vision tasks, unlike other data augmentation strategies, such as Mix-Up [69] or Cut-Mix [66], as long as HR and LR pairs are available.

3.3. Alternate Training Framework

Conventional perceptual loss calculates the similarity between SR and HR images within the feature spaces of a *fixed* network, such as fixed VGG [53] pre-trained on ImageNet dataset [49]. However, this approach can be problematic for TDP loss, which aims to restore task-relevant features from HR images. To effectively utilize the TDP loss to enhance task performance, the task network must be capable of extracting task-relevant features from SR images. However, since the fixed network is not originally trained on I_{SR} , which keeps changing during training, it cannot effectively




Figure 3. **Concept of CQMix.** The black and white region in the figure represents 0 and 1 region in the binary mask M .

extract task-relevant features from I_{SR} , thus reducing the effectiveness of the TDP loss. To resolve this domain gap issue, we propose an alternative training framework, as illustrated in Figure 2, which enables us to employ the feature space of the *on-training* feature extractor of the task network $F_{\theta_{\text{feat}}}$, for calculating the TDP loss. By doing so, we can further boost the efficacy of our TDP loss by providing a more meaningful guide to the SR network to acquire task-related knowledge.

Specifically, in the first phase, we train SR network $S_{\theta_{\text{SR}}}$ with pixel-wise reconstruction loss $\mathcal{L}_{\text{pixel}}$ and the proposed TDP loss, while keeping the feature extractor $F_{\theta_{\text{feat}}}$ frozen. In the second phase, we train the task network $H_{\theta_{\text{head}}} \circ F_{\theta_{\text{feat}}}$ with a task-specific loss $\mathcal{L}_{\text{task}}$, e.g., cross-entropy loss for image classification, while freezing the SR network. We use three types of images in this phase: HR, SR, and augmented images by our CQMix. In summary, the objectives of our SR4IR in each phase are defined as follows:

$$\begin{aligned} & \min_{\theta_{\text{SR}}} \mathcal{L}_{\text{pixel}}(S_{\theta_{\text{SR}}}(I_{\text{LR}}), I_{\text{HR}}) + \mathcal{L}_{\text{TDP}}, & \text{in phase 1,} \\ & \min_{\theta_{\text{feat}}, \theta_{\text{head}}} \mathcal{L}_{\text{task}}(H_{\theta_{\text{head}}} \circ F_{\theta_{\text{feat}}}(I_{\text{cat}}), \mathbf{y}), & \text{in phase 2,} \end{aligned} \quad (3)$$

where I_{cat} is the concatenation of I_{SR} , I_{HR} , and I_{aug} , where the I_{aug} are the images augmented by the CQMix.

4. Experiments

Datasets. We select representative datasets and corresponding evaluation metrics for various image recognition tasks, including semantic segmentation, object detection, and image classification. For semantic segmentation and object detection, we use the PASCAL VOC2012 dataset [14]. For the image classification task, we use the Stanford Cars [30] and CUB-200-2011 [57] datasets. Following the conventional LR datasets [1, 4, 68], we construct the corresponding LR datasets applying bicubic downsampling with scale factors x4 and x8.

Evaluation. We denote the task network $H_{\theta_{\text{head}}} \circ F_{\theta_{\text{feat}}}$ and the SR network $S_{\theta_{\text{SR}}}$ as T and S , respectively, for conve-

nience. Then, the available training scenarios when training both T and S concurrently are as follows:

- $\mathbb{I}_{HR} \rightarrow T$: Train the T with \mathcal{L}_{task} using $(\mathbf{I}_{HR}, \mathbf{y})$. This can be regarded as an oracle in our problem.
- $\mathbb{I}_{LR} \rightarrow T$: Train the T with \mathcal{L}_{task} using $(\mathbf{B}(\mathbf{I}_{LR}), \mathbf{y})$, where \mathbf{B} is bilinear upsampling operator.
- $S \rightarrow T$: Train the S with \mathcal{L}_{pixel} using $(\mathbf{I}_{LR}, \mathbf{I}_{HR})$ first. Then train T with \mathcal{L}_{task} using $(S(\mathbf{I}_{LR}), \mathbf{y})$, while freezing S , as in previous work [76].
- $T \rightarrow S$: Train T with \mathcal{L}_{task} using $(\mathbf{I}_{HR}, \mathbf{y})$ first. Then, train S with $\mathcal{L}_{pixel} + \mathcal{L}_{task}$ using $(\mathbf{I}_{LR}, \mathbf{I}_{HR}, \mathbf{y})$, while freezing T , as in previous work [19, 38, 41].
- $S + T$: Jointly train both S and T with $\mathcal{L}_{pixel} + \mathcal{L}_{task}$ using $(\mathbf{I}_{LR}, \mathbf{I}_{HR}, \mathbf{y})$.

We evaluate the proposed method, named **Super-Resolution for Image Recognition (SR4IR)**, by comparing it with the above methods. For evaluation metrics, we use mean Intersection over Union (mIoU) for semantic segmentation, mean Average Precision (mAP) for object detection, and Top-1 accuracy for image classification. The mAP score is calculated via the COCO [36] mAP evaluator, which uses 10 IoU thresholds from 0.5 to 0.95 with a step size of 0.05. Furthermore, we use LPIPS [71] and PSNR to measure the quality of the SR images.

Network architecture. For the SR network, we employ two architectures: EDSR-baseline [35] and SwinIR [33]. We initialize the SR network from network weights pre-trained on DIV2K [1], widely used for SR training. Regarding the task network, we select representative network architectures tailored to each specific task. Specifically, we adopt DeepLabV3 [7] for semantic segmentation, Faster-RCNN [48] for object detection, and ResNet18 [20] for image classification. The DeepLabV3 and Faster-RCNN are based on MobileNetV3 [22] backbone. The backbone networks are initialized using ImageNet [49] pre-trained weights provided by PyTorch [46] for easier optimization.

Implementations. We utilize AdamW [26] and SGD optimizer for the SR and task network. The learning rates are set to 10^{-4} for the SR network, 2×10^{-2} for DeepLabV3 and Faster-RCNN, and 3×10^{-2} for ResNet18, as these values have been shown to produce the best results for their respective tasks. The batch sizes are set to 16 for both semantic segmentation and object detection and 128 for image classification. We apply cosine annealing [25] scheduling to all optimizers, and all experiments are conducted under PyTorch [46] with eight NVIDIA RTX A6000 GPUs. For the CQMix, the number of patches is set to 64 for both semantic segmentation and object detection and 16 for image classification, regarding the image resolution for each dataset.

4.1. Quantitative results

Semantic segmentation. Table 1 shows the quantitative results on the semantic segmentation task. On the x4 SR scale, our SR4IR significantly boosts the mIoU scores compared to the case trained on the bilinear upsampled images ($\mathbb{I}_{LR} \rightarrow T$), by +2.1 and +3.6, for EDSR-baseline and SwinIR, respectively. Remarkably, our method with SwinIR can achieve a mIoU score of 62.2, which is comparable to the 63.3 mIoU score achieved by the oracle ($\mathbb{I}_{HR} \rightarrow T$), demonstrating the effectiveness of combining the SR method when dealing with LR images. Furthermore, SR4IR surpasses the mIoU scores of all SR combined baselines, $S \rightarrow T$, $T \rightarrow S$, and $S + T$, with a gain of +0.5 and +2.7 compared to the second-best baselines, in the results of the EDSR-baseline with SR scales of x4 and x8, respectively. This validates that our method successfully restores the high-frequency details that are beneficial for the subsequent task. In particular, compared to $\mathbb{I}_{LR} \rightarrow T$, the performance improvement of the proposed SR4IR becomes even more significant as the SR scale increases from x4 to x8, the performance gain increases from +2.1 to +5.7, for the EDSR-baseline case. This indicates that our method is even more effective when the LR images lose a substantial amount of essential high-frequency details relevant to the task, demonstrating the effectiveness of SR4IR.

Furthermore, the proposed SR4IR achieves the best LPIPS score in all cases, with a substantial margin of +12% gain over the second-best baseline for SwinIR on a scale of x4. These results demonstrate that our SR results are not only beneficial from a task perspective but are also capable of producing visually appealing results. Noticeably, separate and sequential training of SR and task networks ($S \rightarrow T$) achieves the best PSNR values but a much lower mIoU score than our SR4IR. The results validate that restoring high-frequency information relevant to the task is much more critical to improving task performance than minimizing the pixel-wise error from ground-truth images.

Moreover, it is noteworthy that training jointly both SR and the task network ($S + T$) performs even worse than sequential training of the SR and task networks ($S \rightarrow T$). This indicates that simultaneously updating both the SR and task networks diminishes the restoring power of the SR due to joint optimization, highlighting the advantages of our alternate training framework.

Object detection. Table 2 presents the numerical results of the object detection. The proposed SR4IR consistently achieves the best mAP score compared to all SR combined baselines. Specifically, for the case of x8 SR scale with the SwinIR model, our SR4IR outperforms the task network trained on bilinear upsampled images ($\mathbb{I}_{LR} \rightarrow T$) and the second-best SR combined baselines ($S \rightarrow T$) by substantial margins of +8.9 and +4.1, respectively.

Methods	SR Method	x4			x8		
		mIoU \uparrow	LPIPS \downarrow	PSNR \uparrow	mIoU \uparrow	LPIPS \downarrow	PSNR \uparrow
$\mathbb{I}_{HR} \rightarrow T$	Oracle (\mathbf{I}_{HR})	63.3	0.000	+ inf	63.3	0.000	+ inf
$\mathbb{I}_{LR} \rightarrow T$	Bilinear up	58.6	0.300	26.16	49.3	0.476	22.54
$S \rightarrow T$	EDSR- baseline [35]	<u>60.2</u>	<u>0.241</u>	28.94	<u>52.3</u>	<u>0.407</u>	24.41
$T \rightarrow S$		56.6	0.352	26.83	49.2	0.461	23.14
$S + T$		58.0	0.318	27.64	50.0	0.480	23.58
SR4IR (Ours)		60.7	0.220	<u>28.21</u>	55.0	0.380	<u>23.91</u>
$S \rightarrow T$	SwinIR [33]	61.4	<u>0.221</u>	29.57	53.9	<u>0.387</u>	24.75
$T \rightarrow S$		53.2	0.376	26.04	48.7	0.466	23.06
$S + T$		57.5	0.341	27.55	50.0	0.480	23.63
SR4IR (Ours)		62.2	0.194	<u>28.83</u>	56.5	0.355	<u>24.23</u>

Table 1. **Performance on the semantic segmentation.** We adopt DeepLabV3 [7] with the MobileNetV3 [22] backbone as the task network T and PASCAL VOC2012 [14] as an evaluation dataset. The **bold** and underlined mean the best and second best, respectively.

Methods	SR Method	x4			x8		
		mAP \uparrow	LPIPS \downarrow	PSNR \uparrow	mAP \uparrow	LPIPS \downarrow	PSNR \uparrow
$\mathbb{I}_{HR} \rightarrow T$	Oracle (\mathbf{I}_{HR})	36.7	0.000	+ inf	36.7	0.000	+ inf
$\mathbb{I}_{LR} \rightarrow T$	Bilinear up	27.5	0.411	22.77	18.9	0.559	20.29
$S \rightarrow T$	EDSR- baseline [35]	<u>30.3</u>	0.336	24.62	<u>21.9</u>	0.494	21.62
$T \rightarrow S$		27.2	<u>0.317</u>	23.37	15.5	<u>0.476</u>	20.48
$S + T$		28.3	0.354	23.71	20.3	<u>0.506</u>	<u>20.98</u>
SR4IR (Ours)		32.4	0.275	<u>23.94</u>	25.5	0.416	20.86
$S \rightarrow T$	SwinIR [33]	<u>31.5</u>	0.309	25.13	<u>23.7</u>	0.474	21.85
$T \rightarrow S$		29.7	<u>0.284</u>	23.59	18.9	<u>0.449</u>	20.59
$S + T$		31.4	0.331	23.94	21.9	<u>0.489</u>	<u>21.11</u>
SR4IR (Ours)		33.9	0.239	<u>24.38</u>	27.8	0.382	21.07

Table 2. **Performance on the object detection.** We adopt Faster-RCNN [48] with the MobileNetV3 backbone [22] as the task network T and PASCAL VOC2012 [14] as an evaluation dataset. The **bold** and underlined mean the best and second best, respectively.

Methods	SR Method	StanfordCars [30]				CUB-200-2011 [57]			
		x4		x8		x4		x8	
		Top-1 Acc. \uparrow (%)	LPIPS \downarrow	Top-1 Acc. \uparrow (%)	LPIPS \downarrow	Top-1 Acc. \uparrow (%)	LPIPS \downarrow	Top-1 Acc. \uparrow (%)	LPIPS \downarrow
$\mathbb{I}_{HR} \rightarrow T$	Oracle (\mathbf{I}_{HR})	86.4	0.000	86.4	0.000	78.3	0.000	78.3	0.000
$\mathbb{I}_{LR} \rightarrow T$	Bilinear up	78.7	0.304	61.0	0.478	70.8	0.298	58.2	0.465
$S \rightarrow T$	EDSR- baseline [35]	<u>82.7</u>	<u>0.176</u>	<u>68.3</u>	<u>0.328</u>	<u>71.1</u>	0.216	<u>59.8</u>	<u>0.365</u>
$T \rightarrow S$		80.1	0.190	56.9	0.448	70.5	<u>0.208</u>	56.5	0.401
$S + T$		80.8	0.196	65.2	0.378	70.8	0.226	58.6	0.380
SR4IR (Ours)		83.3	0.158	71.4	0.326	72.8	0.179	62.6	0.326
$S \rightarrow T$	SwinIR [33]	<u>84.0</u>	<u>0.159</u>	<u>71.4</u>	<u>0.308</u>	72.8	0.201	<u>61.2</u>	<u>0.349</u>
$T \rightarrow S$		82.0	0.162	61.0	0.397	72.1	<u>0.190</u>	58.7	0.377
$S + T$		83.4	0.179	66.6	0.355	<u>72.9</u>	0.212	61.2	0.364
SR4IR (Ours)		85.1	0.136	73.9	0.292	74.7	0.157	65.3	0.301

Table 3. **Performance on the image classification.** We adopt ResNet18 [20] with a fully-connected layer as the task network T . The **bold** and underlined mean the best and second best, respectively.

Image classification. Table 3 shows the quantitative results of the image classification task. Our SR4IR significantly improves image classification performance, remarkably achieving +12.9% improvement on the x8 scale StanfordCars dataset and a +7.1% enhancement on the x8 scale CUB-200-2011 dataset when combined with the SwinIR model, compared to the task network trained on bilinear up-sampled images ($\mathbb{I}_{LR} \rightarrow T$). Furthermore, we note that our method consistently improves performance across various datasets, demonstrating its strong generalization capability.

4.2. Qualitative results

In this section, we provide visual results of the SR with their corresponding task performance. Specifically, we compare

our SR4IR with two baselines, $\mathbb{I}_{LR} \rightarrow T$ and $S \rightarrow T$, where the latter achieved the second-best quantitative results. For a more extensive visual comparison with all baselines, please refer to our supplementary materials.

In the context of semantic segmentation, as shown in Figure 4, the proposed SR4IR successfully generates a segmentation map that closely resembles the ground truth. In contrast, the baselines $\mathbb{I}_{LR} \rightarrow T$ and $S \rightarrow T$ produce incorrect estimations on ‘husky’. Regarding object detection, illustrated in Figure 5, SR4IR enables the object detection network to accurately identify two people who are closely situated, while $\mathbb{I}_{LR} \rightarrow T$ and $S \rightarrow T$ fail to distinguish and predict them as a single person. In image classification, as shown in Figure 6, the proposed SR4IR empowers the sub-



Figure 4. **Visualization of images and semantic segmentation results on PASCAL VOC dataset.** We present the restored images and the corresponding predicted segmentation maps, respectively. For (b) and (c), we use the SwinIR model with an SR scale factor of x4.




Figure 5. **Visualization of object detection results on PASCAL VOC dataset.** The red box with orange annotation means the predicted object bounding box with the corresponding prediction. For (b) and (c), we use the SwinIR model with an SR scale factor of x4.

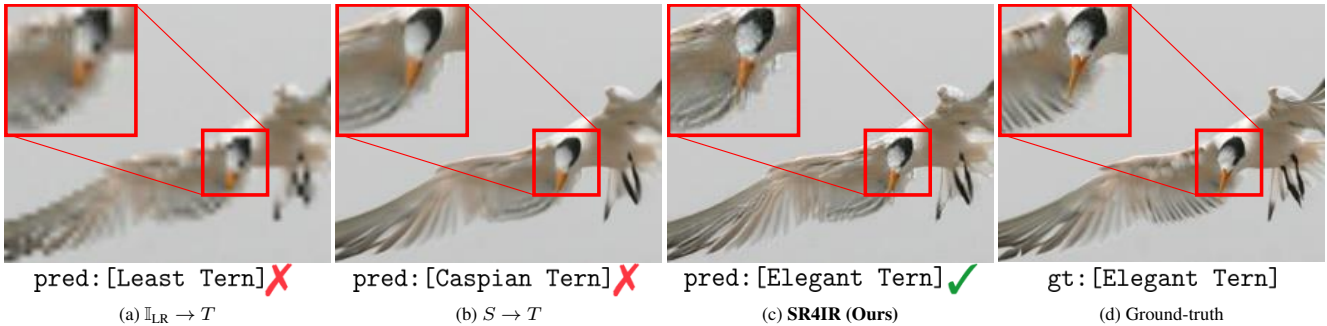


Figure 6. **Visualization of images and image classification results on CUB-200-2011 dataset.** We present the restored images and the corresponding caption. The caption below the image represents the predicted image classification results, which is indicated by a checkmark if the prediction is correct. For (b) and (c), we use the SwinIR model with an SR scale factor of x4.

sequent task network to make the correct label predictions, while the other methods fail to do so. Furthermore, it should be noted that the visualized images produced by our SR4IR achieve the most visually pleasing results.

4.3. Ablation studies

In the ablation studies, we analyze the effectiveness of each component of our method and further report the comparison with previous related works. We report the results of semantic segmentation on the x8 scale PASCAL VOC dataset and image classification on the x8 scale StanfordCars dataset, using the EDSR-baseline network. Please refer to our supplementary materials for more diverse ablation studies.

Effectiveness of the TDP loss. Table 4 shows the task performance of our SR4IR, comparing the cases with and without the proposed TDP loss. Without the TDP loss, the performance of SR4IR decreases significantly, with a decrease of -1.7 mIoU score in semantic segmentation and a decrease in precision of -1.5% in image classification. The

Methods	Segmentation		Classification	
	mIoU $_{\uparrow}$	LPIPS $_{\downarrow}$	Top-1 Acc. $_{\uparrow}$ (%)	LPIPS $_{\downarrow}$
SR4IR <i>without</i> TDP loss	53.3	0.407	69.9	0.327
SR4IR (Ours)	55.0	0.380	71.4	0.326

Table 4. **Effectiveness of the proposed TDP loss.** We evaluate the impact of TDP loss by including or excluding \mathcal{L}_{TDP} in phase 1 of equation (3) on the performance of SR4IR.

results demonstrate that the TDP loss is crucial for a significant boost in task performance in our SR4IR, showing that restoring the task-relevant high-frequency details is essential in enhancing the performance of the subsequent task network. Furthermore, the absence of TDP loss results in a significant deterioration in the LPIPS score with a -7% decrease in segmentation. These results validate that our TDP loss not only enhances task performance but also contributes to generating perceptually realistic SR results.

Effectiveness of the CQMix. To verify the effectiveness of the proposed CQMix, we conduct experiments on training the task network with various types of images, as shown in Table 5. When we train the task network with only I_{SR} ,

the performance is poor, indicating that the effectiveness of the TDP loss decreases when the task network is not trained to properly utilize the high-frequency information present in HR images. However, training the task network solely on I_{HR} also yielded subpar results, as the task network is not trained to handle I_{SR} . Combining both I_{SR} and I_{HR} performs better compared to individual methods. Finally, adding augmented images through our CQMix further enhances the performance, with gains of +0.3 in segmentation and +1.1% in classification, respectively, highlighting the substantial benefits of preventing shortcut learning.

Effectiveness of the alternate training framework. To validate the effectiveness of the proposed alternate training framework, we compare our methods with three perceptual-loss-related approaches, denoted as (A), (B) and (C). In (A), we substitute the proposed TDP loss to the conventional perceptual loss, which uses a fixed VGG network pre-trained on ImageNet. In (B), rather than using the on-training network proposed in our SR4IR, we employ TDP loss using the fixed task network pre-trained on the corresponding target task, *i.e.*, $I_{HR} \rightarrow T$. In (C), rather than adopting alternate training, we apply the TDP loss with a trainable on-training task network within a joint training, where both the SR and task networks are trained jointly.

Table 6 presents performance results between (A), (B), (C), and our method. Compared to the performance of SR4IR without TDP loss, which achieves a 53.3 mIoU score and 69.9% accuracy as presented in Table 4, the method (A) achieves similar or even worse performance, demonstrating that adopting conventional perceptual loss does not help to boost the task performance. The method (B) also results in marginal performance gains, as the fixed and pre-trained task network is not trained to handle the SR results. Notably, the method (C) fails to learn because optimizing the TDP loss leads to undesirable local optima when the task network is trainable, which results in forcing the task network to produce a constant output regardless of varied inputs. In contrast, our SR4IR achieves significant performance improvement compared to the above approaches, highlighting the importance of adopting an alternate training framework to fully leverage the power of the TDP loss.

Comparison with previous methods. To highlight the advantages of our SR4IR, we include a quantitative comparison with two related previous methods, TDSR [19] and SOD-MTGAN [3], which are task-driven restoration methods. The results are presented in Table 7. Due to the absence of a publicly available training code, we report the results of our re-implemented version, which is elaborated in the supplementary material. In comparison to TDSR and SOD-MTGAN, our SR4IR achieves significantly superior performance, underlining the effectiveness of leveraging perceptual loss when compared to the case of training without perceptual loss (TDSR) and adversarial loss (SOD-MTGAN).

Training image types	Segmentation (mIoU $_{\uparrow}$)	Classification (Top-1 Acc. % $_{\uparrow}$)
I_{SR}	51.7	66.7
I_{HR}	53.3	59.3
$I_{SR} + I_{HR}$	54.7	70.3
$I_{SR} + I_{HR} + I_{aug}$ (Ours)	55.0	71.4

Table 5. **Effectiveness of the proposed CQMix.** We compare the performance of SR4IR by varying the training images I_{cat} in phase 2 of equation (3).

Methods	Feature extractor for perceptual loss	Segmentation (mIoU $_{\uparrow}$)	Classification (Top-1 Acc. % $_{\uparrow}$)
(A)	Pre-trained on ImageNet	53.6	69.4
(B)	Pre-trained on corresponding task	53.6	70.3
(C)	On-training (joint)	3.5	3.7
Ours	On-training (alternate)	55.0	71.4

Table 6. **Effectiveness of the proposed alternate training framework.** We compare our SR4IR with three perceptual-loss-related approaches, by varying the feature extractor $F_{\theta_{feat}}$ used for the TDP loss in equation (2).

Methods	Segmentation (mIoU $_{\uparrow}$)	Classification (Top-1 Acc. % $_{\uparrow}$)
TDSR \dagger [19]	49.8	56.4
SOD-MTGAN \dagger [3]	51.6	68.1
SR4IR (Ours)	55.0	71.4

Table 7. **Comparison with related previous methods.** The symbol \dagger denotes the results of our re-implementation version.

5. Conclusion

In this paper, we tackle the practical scenarios for image recognition tasks where low-resolution (LR) content prevails when processing the images. We tackle the challenge by employing image super-resolution (SR) techniques. Considering the ill-posed nature of SR, we focus on developing methods that encourage an SR network to reconstruct task-related features beneficial for a specific task. In particular, we propose task-driven perceptual (TDP) loss that effectively guides the restoration of task-related high-frequency contents. Furthermore, our cross-quality patch mix (CQMix) and alternate training framework help us fully harness the advantages of our TDP loss by preventing a task network from learning biased features and eliminating the domain gap problem. Extensive experimental results demonstrate that our **Super-Resolution for Image Recognition (SR4IR)** can be generally and effectively applicable to various image recognition tasks, such as semantic segmentation, object detection, and image classification.

Acknowledgments. This work was supported in part by the IITP grants [No. 2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), No.2021-0-02068, and No.2023-0-00156], the NRF grant [No.2021M3A9E4080782] funded by the Korean government (MSIT).

References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *CVPR Workshops*, 2017. 4, 5
- [2] Yancheng Bai, Yongqiang Zhang, Mingli Ding, and Bernard Ghanem. Finding tiny faces in the wild with generative adversarial network. In *CVPR*, 2018. 3
- [3] Yancheng Bai, Yongqiang Zhang, Mingli Ding, and Bernard Ghanem. Sod-mtgan: Small object detection via multi-task generative adversarial network. In *ECCV*, 2018. 3, 8
- [4] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie-Line Alberi Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *BMVC*, 2012. 4
- [5] Joan Bruna, Pablo Sprechmann, and Yann LeCun. Super-resolution with deep convolutional sufficient statistics. In *ICLR*, 2016. 1, 3
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 2016. 1
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv*, 2017. 1, 5, 6
- [8] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *ICCV*, 2021. 2
- [9] Ziteng Cui, Yingying Zhu, Lin Gu, Guo-Jun Qi, Xiaoxiao Li, Renrui Zhang, Zenghui Zhang, and Tatsuya Harada. Exploring resolution and degradation clues as self-supervised signal for low quality object detection. In *ECCV*, 2022. 3
- [10] Dengxin Dai, Yujian Wang, Yuhua Chen, and Luc Van Gool. Is image super-resolution helpful for other vision tasks? In *WACV*, 2016. 3
- [11] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *CVPR*, 2019. 2
- [12] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE TPAMI*, 2016. 2
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 1, 2
- [14] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 1, 4, 6
- [15] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. In *NIPS*, 2015. 1, 3
- [16] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2020. 2, 4
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [18] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Deep back-projection networks for super-resolution. In *CVPR*, 2018. 1, 2
- [19] Muhammad Haris, Shakhnarovich Greg, and Ukita Norimichi. Task-driven super resolution: Object detection in low-resolution images. In *ICONIP*, 2021. 3, 5, 8
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 5, 6
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NIPS*, 2020. 2
- [22] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Efficientnet: Rethinking model scaling for convolutional neural networks. *ICCV*, 2019. 5, 6
- [23] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv*, 2017. 1
- [24] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 1
- [25] Loshchilov I. and Hutter F. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 5
- [26] Loshchilov Ilya and Hutter Frank. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [27] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 1, 2, 3
- [28] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *NIPS*, 2022. 2
- [29] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, 2016. 1, 2
- [30] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV workshops*, 2013. 4, 6
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [32] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. 2
- [33] Jingyun Liang, Jie Zhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image restoration using swin transformer. In *ICCV*, 2021. 1, 2, 5, 6

- [34] Jie Liang, Hui Zeng, and Lei Zhang. Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In *CVPR, 2022*. 2
- [35] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPR Workshops, 2017*. 1, 2, 5, 6
- [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV, 2014*. 5
- [37] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV, 2017*. 1
- [38] Ding Liu, Bihan Wen, Xianming Liu, Zhangyang Wang, and Thomas S Huang. When image denoising meets high-level vision tasks: A deep learning approach. In *IJCAI, 2018*. 3, 5
- [39] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. Ssd: Single shot multibox detector. In *ECCV, 2016*. 1
- [40] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CVPR, 2022*. 1
- [41] Zhuang Liu, Hung-Ju Wang, Tinghui Zhou, Zhiqiang Shen, Bingyi Kang, Evan Shelhamer, and Trevor Darrell. Exploring simple and transferable recognition-aware image processing. *IEEE TPAMI, 2022*. 3, 5
- [42] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR, 2015*. 1
- [43] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Image restoration with mean-reverting stochastic differential equations. In *ICML, 2023*. 2
- [44] Ben Niu, Weilei Wen, Wenqi Ren, Xiangde Zhang, Lianping Yang, Shuzhen Wang, Kaihao Zhang, Xiaochun Cao, and Haifeng Shen. Single image super-resolution via a holistic attention network. In *ECCV, 2020*. 2
- [45] JoonKyu Park, Sanghyun Son, and Kyoung Mu Lee. Content-aware local gan for photo-realistic super-resolution. In *ICCV, 2023*. 2
- [46] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS Workshops, 2017*. 5
- [47] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR, 2016*. 1
- [48] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS, 2015*. 1, 5, 6
- [49] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV, 2015*. 2, 4, 5
- [50] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI, 2022*. 2
- [51] Mehdi SM Sajjadi, Bernhard Scholkopf, and Michael Hirsch. Enhancenet: Single image super-resolution through automated texture synthesis. In *ICCV, 2017*. 2
- [52] Jacob Shermeyer and Adam Van Etten. The effects of super-resolution on object detection performance in satellite imagery. In *CVPR Workshops, 2019*. 3
- [53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR, 2015*. 1, 2, 4
- [54] Jae Woong Soh, Gu Yong Park, Junho Jo, and Nam Ik Cho. Natural and realistic single image super-resolution with explicit natural manifold discrimination. In *CVPR, 2019*. 2
- [55] Sanghyun Son, Jaeha Kim, Wei-Sheng Lai, Ming-Hsuan Yang, and Kyoung Mu Lee. Toward real-world super-resolution via adaptive downsampling models. *IEEE TPAMI, 2021*. 2
- [56] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *ICML, 2019*. 1
- [57] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 4, 6
- [58] Li Wang, Dong Li, Yousong Zhu, Lu Tian, and Yi Shan. Dual super-resolution learning for semantic segmentation. In *CVPR, 2020*. 3
- [59] 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 *ECCV Workshops, 2018*. 2
- [60] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *ICCV workshop, 2021*. 2
- [61] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *ICLR, 2023*. 2
- [62] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *CVPR, 2022*. 2
- [63] Bin Xia, Yulun Zhang, Shiyin Wang, Yitong Wang, Xinglong Wu, Yapeng Tian, Wenming Yang, and Luc Van Gool. Diffir: Efficient diffusion model for image restoration. In *ICCV, 2023*. 2
- [64] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. 2021. 1
- [65] Fanghua Yu, Jinjin Gu, Zheyuan Li, Jinfan Hu, Xiangtao Kong, Xintao Wang, Jingwen He, Yu Qiao, and Chao Dong. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild. In *CVPR, 2024*. 2
- [66] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV, 2019*. 4

- [67] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022. 2
- [68] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International Conference on Curves and Surfaces*, 2010. 4
- [69] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. In *ICLR*, 2017. 4
- [70] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *CVPR*, 2017. 1
- [71] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 2, 3, 5
- [72] Wenqiang Zhang, Zilong Huang, Guozhong Luo, Tao Chen, Xinggang Wang, Wenyu Liu, Gang Yu, and Chunhua Shen. Topformer: Token pyramid transformer for mobile semantic segmentation. In *CVPR*, 2022. 1
- [73] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, 2018. 2
- [74] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *CVPR*, 2018. 1, 2
- [75] Xiaotong Zhao, Wei Li, Yifan Zhang, and Zhiyong Feng. Residual super-resolution single shot network for low-resolution object detection. *IEEE Access*, 6:47780–47793, 2018. 3
- [76] Liguozhou, Guang Chen, Mingyue Feng, and Alois Knoll. Improving low-resolution image classification by super-resolution with enhancing high-frequency content. In *ICPR*, 2021. 3, 5