

# Exploring Resolution and Degradation Clues as Self-supervised Signal for Low Quality Object Detection

Ziteng Cui<sup>1</sup>, Yingying Zhu<sup>2</sup>, Lin Gu<sup>3,4\*</sup>, Guo-Jun Qi<sup>5</sup>, Xiaoxiao Li<sup>6</sup>,  
Renrui Zhang<sup>7</sup>, Zenghui Zhang<sup>1</sup>, and Tatsuya Harada<sup>4,3</sup>

<sup>1</sup> Shanghai Jiao Tong University

<sup>2</sup> University of Texas at Arlington

<sup>3</sup> RIKEN AIP

<sup>4</sup> The University of Tokyo

<sup>5</sup> Laboratory for Machine Perception and Learning

<sup>6</sup> The University of British Columbia

<sup>7</sup> Shanghai AI Laboratory

**Abstract.** Image restoration algorithms such as super resolution (SR) are indispensable pre-processing modules for object detection in low quality images. Most of these algorithms assume the degradation is fixed and known a priori. However, in practical, either the real degradation or optimal up-sampling ratio rate is unknown or differs from assumption, leading to a deteriorating performance for both the pre-processing module and the consequent high-level task such as object detection. Here, we propose a novel self-supervised framework to detect objects in degraded low resolution images. We utilize the downsampling degradation as a kind of transformation for self-supervised signals to explore the equivariant representation against various resolutions and other degradation conditions. The Auto Encoding Resolution in Self-supervision (AERIS) framework could further take the advantage of advanced SR architectures with an arbitrary resolution restoring decoder to reconstruct the original correspondence from the degraded input image. Both the representation learning and object detection are optimized jointly in an end-to-end training fashion. The generic AERIS framework could be implemented on various mainstream object detection architectures with different backbones. The extensive experiments show that our method has achieved superior performance compared with existing methods when facing variant degradation situations. Code is available at [this link](#).

**Keywords:** Self-supervised Learning, Computational Photography, Object Detection

## 1 Introduction

High level vision tasks (*i.e.* image classification, object detection, and semantic segmentation) have witnessed great success thanks to the large scale dataset

---

\* Corresponding author.

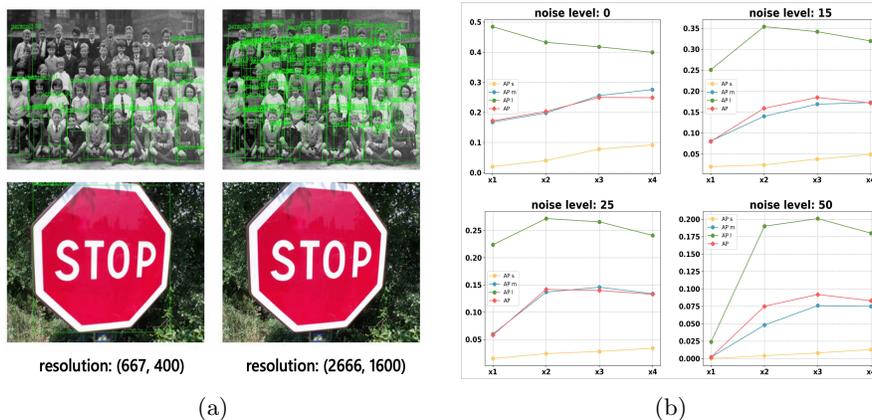


Fig. 1: Illustration of scale variance bottleneck. (a): Tiny people are well detected at high resolution while the large traffic sign is recognized in low resolution. (b): Detection result on down-sampled MS COCO [35] dataset (down scale rate: 4) with different noise level. Specifically, we up-scale the images with different ratio (2, 3, 4) before detection. X-axis is up-scale ratio and Y-axis is mAP result of CenterNet [64]. We also report the results on small, medium and large objects.

[10,35,14]. Images in these datasets are mainly captured by commercial cameras with higher resolution and signal-to-noise ratio (SNR). Trained and optimized on these high-quality images, high-level vision would suffer a performance drop on low resolution [9,19,53] or low quality images [38,52,45,1,8,41].

To improve the performance of vision algorithms on degraded low resolution images, Dai *et al.* [9] presented the first comprehensive study advocating pre-processing images with super resolution (SR) algorithms. Other high-level tasks like face recognition [66], face detection [2], image classification [55,38] and semantic segmentation [53], also benefit from the restoration module to extract more discriminate features.

Most existing enhancement methods, especially SR algorithms [3,60,59], assume target images are from a **known and fixed** degradation model [13,36]:

$$t(x) = (x \otimes k) \downarrow_s + n, \quad (1)$$

where  $t(x)$  and  $x$  denote the degraded low resolution (LR) image and original high resolution (HR) input respectively.  $k$  is the blur kernel while  $\downarrow_s$  is the down-sampling operation with ratio  $s$ .  $n$  is the additive noise. However, the performance of these enhancement algorithms would decline severely when the real degradation deviates from the assumption [18]. To make it worse, for machine perception tasks, as shown in Fig 1b, higher resolution does not necessarily guarantee a better performance in high level tasks. Like object detection, the optimal SR ratio varies across the images due to the scale variance bottleneck [48,49], there is a trade off that certain high level predictions are better handled at

lower resolution and others better processed at higher resolution. As illustrated in Fig 1a, though working well on individual tiny person at high resolution (2666,1600), the detection method ignores the large traffic sign. On the contrary, in low resolution (667,400) images, detecting network’s reception field could observe more global context for large structure, at the cost of sacrificing the small objects. Fig 1b also quantitatively demonstrates this bottleneck. The detection performance does not necessary increases with super resolution ratio, especially for large objects.

Instead of explicitly enhancing an input image with a fixed restoration module, we exploit the intrinsic equivariant representation against various resolutions and degradation. *I know who I was when I got up this morning, but I think I must have been changed several times since then.*<sup>1</sup>. Either being small enough to squeeze through the door or so big to shed a pool of tears, Alice should be encoded with a equivariant representation to show who she is in the world. Based on the encoded representation shown in Fig.2, we propose an end-to-end framework for object detection in low quality images. To capture the complex patterns of visual structures, we utilize groups of downsampling degradation transformations under different downsampling rate, noise and degradation kernel as the self-supervised signal [63,16].

During the training, we generate a degraded LR image  $t(x)$  from the original HR image  $x$  through a random degradation transformation  $t$ . As shown in Fig.2, to train the Encoder  $E$  to learn the degradation equivariant representation  $E(t(x))$ , we introduce an arbitrary-resolution restoration decoder (ARRD) decoder  $D_r$ . ARRD implicitly decodes  $t$  to reconstruct the original HR data  $x$  from the representation  $E(t(x))$  of various degraded LR image  $t(x)$ . If the self-supervised signal is reconstructed, the representation should capture the dynamics of how they change under different resolution and other degradation as much as possible [7,44,63]. The nature of reconstructing HR data also allows us to leverage the advance of the fast-growing SR research by directly using their successful architectures.

On the encoded representation  $E(t(x))$ , we further impose an object detection decoder  $D_o$  to supervise the encoder  $E$  to encode the image structure relevant to the consequent tasks. The object detection decoder  $D_o$  performs the detection task to get the object’s location and class. During inference, the target image is directly passed through the encoder  $E$  and object detection decoder  $D_o$  in Fig.2 for detection. Compared to pre-processing module based methods [46,19], our inference pipeline is more computation efficient as we avoid explicitly reconstructing the image details.

To cover the diverse degradation and resolutions, in real scenario, we generate degraded  $t(x)$  by randomly sampling a transformations  $t$  according to practical down-sampling degradation model [36,59]. As shown in Fig.2, the transformation  $t$  is characterised by down-sampling ratio  $s$ , blur kernel  $k$ , and noise level  $n$  in Eq.1.

Our contributions could be summarised as follows:

---

<sup>1</sup> Chapter 5, Alice in Wonderland

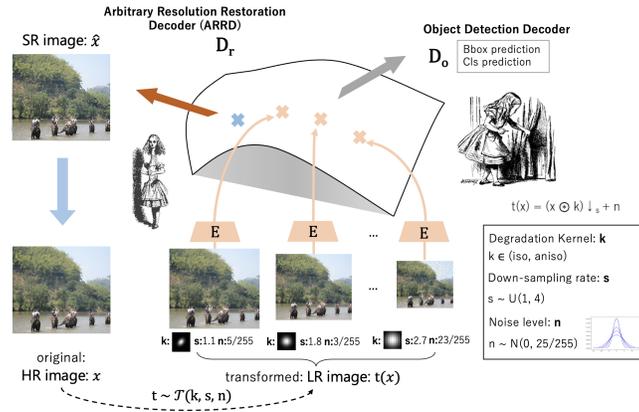


Fig. 2: A simple illustration of Auto Encoding Resolution in Self-supervision (AERIS). Encoder  $E$  encodes images transformed under various self-supervised signal into degradation equivariant representations (like Alice after drinking ”Drink me”). ARRD  $D_r$  implicitly decodes the self-supervised signal to reconstruct the original image/transformation while detection decoder  $D_o$  decodes the representations for object detection task.

- We propose a novel framework, Auto Encoding Resolution in Self-supervision (AERIS), to detect objects in degraded low resolution images by utilizing the resolution and degradation clues as self-supervised signal. Specifically, we learn the degradation equivariant representation that captures the dynamics of feature representations under diverse resolutions and degradation types. Our AERIS is generic and could be implemented on several mainstream object detection architecture.
- AERIS method takes the strength of advanced super resolution (SR) research by training an arbitrary resolution restoration decoder (ARRD) that reconstructs the high resolution details. Furthermore, by optimizing the representation learning and detection in a unified end-to-end training framework, the representation preserves the intrinsic visual structure that is discriminative for detection.
- We evaluate our method on mainstream public dataset MS COCO [35] and KITTI [15]. The experiment results also show that our method has achieved SOTA performance on several degradation conditions.

## 2 Related Works

### 2.1 Single Image Super Resolution

Image restoration algorithms are intuitive solutions to handle degradation, here we mainly introduce single image super-resolution (SISR), since the object detection

task is sensitive to resolution, and the other low-level vision tasks (*i.e.* denoise, deblur) also have connections with SISR task.

The very first CCN-based SR was proposed by Dong *et al.* [11] with a three-layer neural network. Then Kim *et al.* [29] extended the depth of network to 20 layers with gradient clipping and residual learning. Batch normalization is later identified to impose the negative effect on the SR reconstruction. By removing this layer, EDSR [34] achieves SOTA in 2017. After ESDR, better SR architectures are designed by integrating the successful deep learning techniques such as Laplacian pyramid structure [30], dense connection [51], back projection [20], transformer blocks [33] and so on. Besides designing sophisticated architecture, losses like perceptual loss [26] and adversarial loss [31] are also demonstrated to improve the SR reconstruction quality.

SR algorithms heavily rely on the assumption of degradation model and fixed resolution. Much efforts are spent to relax the constraint. Not limited to a unfixed up-sampling scale, like Hu *et al.* [24] first proposed Meta-SR to super-resolve images with arbitrary scale factor. After that, LIIF [6] utilize implicit function to solve this task, and FuncNet [40] further been proposed also for noise and blur condition.

To deploy SR for real scenarios, blind SR methods assume the degradation information is not known. One direction is to convert the problem into non-blind SR by provide prior degradation information [62] or initially estimate the degradation parameters [3]. However, the applied non-blind SR algorithm is very sensitive to the error of the degradation estimation. Gu *et al.* [18] then proposed to iterative correct the estimated degradation with an iterative kernel correction (IKC) method. Without explicitly estimating degradation parameters, Wang *et al.* [54] introduced a contrastive loss to design the degradation-aware SR network based on the learned representations. Recently, Zhang *et al.* [59] solved the general blind SISR by designing a practical model considering complex degradation. This model has been demonstrated to cover the degradation space of real images. Therefore, we adopt this practical model to synthesize various degraded LR images as the self-supervised signal to train our model.

## 2.2 Image Restoration for Machine Perception

There is sufficient evidence that degraded scene would give negative impact on high-level vision tasks [57,9,19,27]. As for resolution, Dai *et al.* [9] made the first analyze on improving several vision tasks with SR as pre-process. Wang *et al.* [55] analyzed the effectiveness of SR in image classification task while DSRL [53] improved the low-resolution semantic segmentation with an additional SR block. Shermeyer and Etten [46] evaluate the effectiveness of a SR pre-process step on aerial image object detection. Recently, Haris *et al.* [19] jointly optimise object detection loss along with SR sub-network [20] to improve detection performance.

Similarly, noise and blur’s effect on high-level vision have also been well-studied. Hendrycks *et al.* [22] evaluate image classification robustness under multiply degradation conditions including noise and blur. Kamann *et al.* [27] studied the impact of noise and blur on different semantic segmentation methods.

Liu *et al.* [38] combine a denoise network in classifier to improve classification’s performance under noisy condition. Very recently, Mohamed and Gabriel [45] analyse motion blur and propose several methods to improve detection performance on motion blurry images.

However, most of these existing works assume the degradation parameters such as the down-sampling ratio is known and fixed. Based on the degradation equivariant representation, our framework is robust to various degradation in real-world scenarios. Without an explicit restoration module or restoration step, we directly perform the detection on low-dimension encoded features that saves much computational burden.

### 3 Down-sampling Degradation Transformations

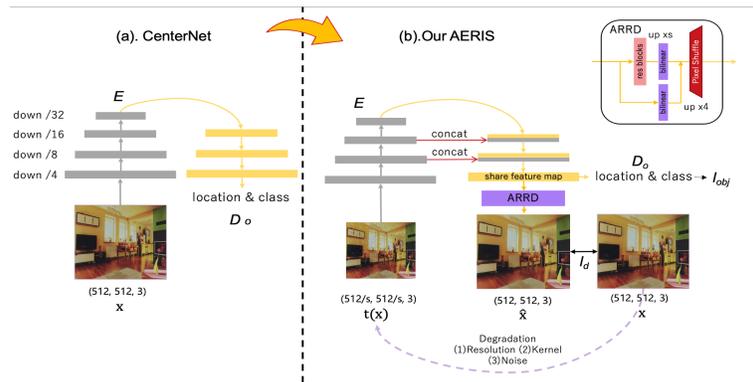


Fig. 3: An illustration of how our AERIS is implemented based on CenterNet [64]. Left is the original CenterNet while the right one is the architecture of AERIS.

In real scenarios, the image may be captured and processed in various ways. To cover these generally unknown operations, it is necessary to select a practical degradation model for the degradation transformation. This model would transform the high resolution (HR) image  $x$  to the low resolution counterpart  $t(x)$  with Eq. 1.

Early restoration methods assume a simple degradation model where LR is directly down-sampled from the HR images without or with simple noise. Instead of dealing with synthetic images, recent methods now focus on more realistic degradation models. For example, [4] directly train the model on the LR and HR images pair captured by the real camera system. USRNet [60] effectively handled the degradation models with different parameters such as scale factors by unfolding the model-based energy function. Here, we adopt the practical degradation model [13,36,3,60,59] that accounts for diverse degradation in real images.

**Convolution Operation:** Gaussian blur kernel is the most common kernel to blur the image [12,37,3]. Here, we choose two Gaussian degradation kernels  $k$ : isotropic Gaussian kernels  $k_{iso}$  and anisotropic Gaussian kernels  $k_{aniso}$  [3,62,59]. We also consider none degradation kernel as  $k_{none}$ . Following [59], the kernel size is uniformly sampled from  $\{7 \times 7, 9 \times 9, \dots, 21 \times 21\}$  and the  $k_{iso}$ 's width is uniformly chosen from (0.1, 2.4). For  $k_{aniso}$ , the kernel angle is uniformly chosen from  $(0, \pi)$  and the longer kernel width is uniformly chosen from (0.5, 6).

**Noise:** When dealing with the real-world scenarios, the Gaussian noise model is usually adopted to simulate the noises from camera sensor noise [28], low-light noise [56] to quantization noise [25] *etc.* Therefore, we adopt a zero-mean additive white Gaussian noise (AWGN) model  $n \sim N(0, \sigma)$  in Eq.1. The variance  $\sigma$  is randomly chosen from a uniform distribution  $U(0, 25/255)$  (*e.g.* 13.2/255).

**Down-sampling:** For the down-sampling process, the sampling ratio  $s$  is randomly chosen from uniform distribution  $U(1, 4)$  (*e.g.* 1.9) while the down-sampling methods is randomly chosen from nearest method  $d_{nearest}^s$ , bilinear method  $d_{bilinear}^s$  and bicubic method  $d_{bicubic}^s$ .

So the final down-sampling degradation transformation  $t$  would take from random blur kernel  $k$ , noise level  $n$  and down-sampling ratio  $s$  in the total distribution  $\mathcal{T}$ , as  $t \sim \mathcal{T}(k, n, s)$ .

## 4 Our Framework

Due to its concise structure, we mainly take one-stage object detector CenterNet [64] to illustrate how to implement our AERIS. More object detectors' implements and results please refer to supplementary. We think that our AERIS framework is a generic framework that could also be implemented on other popular object detectors [43,17] and various backbones [23,39,32,42].

### 4.1 CenterNet

CenterNet [64] is an efficient one-stage anchor-free object detector. We show its vanilla structure of [64] in Fig.3 (a). Input image is fed to the backbone (*i.e.* ResNet18 [21]) to extract /32 bottleneck feature, and then upsampled to a /4 feature map by three  $\times 2$  deconvolution blocks. This /4 feature map is passed to scaling blocks with three independent convolution blocks to generate the final feature maps. Based on this feature map, there are prediction heads conducting class-wise bbox center detection, bbox height and width regression, offset regression respectively. For more details, please refer to [64].

The CenterNet could be decomposed into an encoder-decoder style structure. Here, we denote the network backbone part (gray part in Fig.3.a) as encoder  $E$ . The object detection decoder  $D_o$ , comprised by three prediction heads (colored in orange in Fig.3 (a)), decodes the object information.

## 4.2 Architecture and Training Pipeline

Fig.3 (b) illustrates how to implement our AERIS based on CenterNet [64]. The detailed training procedure is given in Algo.1. When training AERIS, we first apply the degradation transformation  $t \sim \mathcal{T}(k, n, s)$  convert  $x$  to a random generated  $t(x)$ , covering the degradation space of real scenarios. In each batch  $B$ , the generated  $t(x) \in \mathbb{R}^{B \times 3 \times \frac{H}{s} \times \frac{W}{s}}$  should keep the same down-sampling rate  $s \sim (1, 4)$ . Then the transformed degraded LR image  $t(x)$  are sent to the encoder  $E$  to encode the degradation equivariant representation  $E(t(x)) \in \mathbb{R}^{B \times C \times \frac{H}{s*d} \times \frac{W}{s*d}}$ . Here the encoder  $E$  refers to the backbone in detection network, and  $d$  is the down-sample rate of backbone (*i.e.* 32 for ResNet [21]).

We then input degradation equivariant representation  $E(t(x))$  into upscaling blocks with three  $\times 2$  deconvolution blocks to generate the final feature map for object detection decoder  $D_o$  and arbitrary resolution restoration decoder (ARRD)  $D_r$ . As illustrated by the red arrows in Fig.3, we introduce the skip connection on  $/8$  and  $/16$  feature maps between the backbone encoder  $E$  and the deconvolution blocks. Fusing features from different scales is a common process in low-level vision tasks, and could enhance semantic information and contribute to the subsequent  $D_o$  and  $D_r$ .

We further regularise the representation encode  $E$  with our unique ARRD  $D_r$  that implicitly estimate the self-supervised signal  $t$  to recover HR image  $\hat{x}$ . Since the downsampling rate  $s$  is not a fixed integral number in the training stage, ARRD could deal with an arbitrary scale factor. ARRD  $D_r$  could force the encoder  $E$  to not only capture the dynamics of how images change under different transformations, but also extracts the complex patterns of visual structures. Since ARRD  $D_r$  aims to recover the original resolution of clean image  $x$  from  $E(t(x))$ , it could also support the object detection decoder  $D_o$  with more detailed features. Inspired by the learnable resizer model [50], we design this decoder with a residual bilinear model shown in Fig. 3, which ends up with a  $\times 4$  pixel shuffle layer [47]. ARRD is a light weight structure that uses fewer parameters (0.06M) compared to the backbone encoder (11.17M), upscaling blocks (3.61M) and detection decoder  $D_o$  (0.12M). The ARRD loss  $l_d$  is defined as an L1 loss between output image  $\hat{x}$  and ground truth image  $x$ :

$$l_d = |\hat{x} - x|_1 = |D_r(E(t(x))) - x|_1. \quad (2)$$

We adopt the three CenterNet prediction heads as the object detection decoder  $D_o$  to conduct detection on the final feature map generated by the upscaling block.

As shown in Algo.1, we optimise the total loss  $l_{total}$  including detection loss  $l_{obj}$  (*i.e.* classwise bbox center loss, bbox width and height loss, bbox offset loss for CenterNet [64]), and data restoration loss  $l_d$ :

$$l_{total} = l_{obj} + \lambda \cdot l_d, \quad (3)$$

where  $\lambda$  is the non-negative parameters for loss balancing. Which we set to 0.4 in CenterNet experiments and 0.8 in DETR experiments, more ablation details please refer to supplementary.

---

**Algorithm 1** AERIS Algorithm Pipeline

---

**(1). Data Generation:**

B: batch size, C: channel, H: image height, W: image width

**inputs:** HR image  $x = (B, C, H, W)$ , down-sample factor  $s \sim (1.0, 4.0)$ **outputs:** degraded LR image  $t(x) = (B, C, \frac{H}{s}, \frac{W}{s})$ **for**  $i$  in range( $B$ ): **do**

- (1). Convolution with blur kernel  $k$
- (2). Down-sampling with rate  $s$
- (3). Add noise  $n$

**end for****(2). Training:****inputs:** Degraded LR image  $t(x) = (B, C, \frac{H}{s}, \frac{W}{s})$ **outputs:** detection output, estimated SR image  $\hat{x}$ **encoding:** $t(x) \xrightarrow{E} E(t(x))$ **decoding:**data restoration decoding:  $\hat{x} = D_r(E(t(x)))$ detection decoding: detection results =  $D_o(E(t(x)))$ 

---

### 4.3 Inference Procedure

The inference procedure only involves encoder  $E$ , upscaling block and object detection decoder  $D_o$  as illustrated in Fig.3. Specifically, the encoder  $E$  encodes the input target image before  $D_o$  performs the detection. Compared to explicitly pre-processing image for high-level tasks [9,46,19], our AERIS saves much computational time as we avoid reconstructing HR details of data.

We could also reconstruct the HR image with our ARRD decoder  $D_r$ . Very interestingly, our restored images  $\hat{x}$  are more machine vision oriented and exhibit artifacts around the center of the object, as shown in Fig.4.

## 5 Experiments and Details

### 5.1 Datasets and Implementation Details

**Dataset.** We adopt two widely used object detection datasets MS COCO [35] and KITTI [15] for detection robustness evaluation. MS COCO contains  $\sim 118k$  images with bounding box annotation in 80 categories. We use COCO **train2017** set as train set and use COCO **val2017** set as normal condition evaluation set. Also COCO dataset differentiates the labels of different scale level’s objects, and give them specific evaluation metrics (small:  $AP_s$ , middle:  $AP_m$ , large:  $AP_l$ ), which could show us different degradation conditions’ influence on different scale objects, especially the down-sampling process  $s$ .

KITTI [15] is a popular small object detection dataset for autonomous driving. For KITTI dataset, we evaluate car class, and use KITTI **train** set as train set

Table 1: Comparison with SOTA restoration methods and different training strategies on **COCO-d** dataset. Here **CenterNet** with **ResNet-18** backbone and **Swin-T** backbone were adopted. The inference images with **higher resolution** are in **blue** background.

Test Set	Pre-process	Training Strategy	CenterNet (ResNet-18)				CenterNet (Swin-T)					
			AP	AP <sub>s</sub>	AP <sub>m</sub>	AP <sub>l</sub>	AP	AP <sub>s</sub>	AP <sub>m</sub>	AP <sub>l</sub>		
COCO-d	-	Detection	30.1	10.6	33.2	47.2	36.9	17.9	41.8	52.9		
			14.5	1.2	10.4	38.6	19.9	2.7	16.9	46.2		
			16.2	4.1	15.3	31.1	18.6	4.0	17.8	39.7		
			8.0	4.6	10.5	10.1	10.6	<b>5.7</b>	12.8	16.7		
			14.8	2.6	14.3	27.9	16.6	3.0	16.5	33.4		
			15.0	3.5	14.3	27.4	16.7	3.4	16.1	32.0		
			14.2	2.6	12.4	29.5	17.3	3.6	17.0	34.1		
			16.8	4.2	15.8	36.9	20.2	4.8	18.1	40.5		
			10.4	0.8	6.8	27.9	10.9	0.7	8.8	35.1		
			11.4	1.2	7.2	34.8	11.9	1.4	8.9	33.4		
				Deg $t$	17.6	2.3	15.4	41.9	20.9	3.1	20.3	47.6
				Deg $t + N$	17.9	2.5	15.9	42.5	21.0	3.0	20.4	48.2
				$D_r +$ Detection	17.7	<b>4.8</b>	15.8	41.0	21.4	5.6	19.6	46.3
				<b>AERIS</b>	<b>18.4</b>	2.7	<b>16.4</b>	<b>42.5</b>	<b>21.6</b>	3.2	<b>20.4</b>	<b>49.0</b>

and use KITTI val set as normal condition evaluation set, and show the AP rate for comparison.

**Implementation Details.** We build our framework based on the open-source object detection toolbox `mmdetection` [5]. Throughout the experiments, the backbone ResNet-18 [21] and Swin-T [39] are initialed with ImageNet [21] pre-train weights. We apply the data augmentation pipeline in `mmdetection` [5], specifically we adopt random crop, random flip and multi-scale test.

During training stage, all the models are trained on 4 Tesla V100 GPUs. Same as setting in [5], for AERIS-CenterNet training, the input image shape is resized to  $512 \times 512$ . The model has been trained for 140 epochs with SGD optimizer. Batch size is set to 16 per GPU. Momentum and weight decay are set to 0.9 and  $1e-4$ . Initial learning rate is 0.01 and warms up at first 500 iterations and would decays to one-tenth at 90 and 120 epoch.

**Comparison Methods.** To evaluate object detectors’ robustness under different conditions’ degradation, we separately set the multi-degradation evaluation and single-degradation evaluation (see Sec. 5.2 and Sec. 5.3 for details). We first compare our methods with SOTA image restoration methods: non-blind SR methods [31,20,33], blind SR methods [4,59], denoise methods [61,33,58]. Also

Table 2: Comparison with SOTA restoration methods and different training strategies on **KITTI-d** dataset. Here **CenterNet** with **ResNet-18** backbone is adopted. We also show the inference speed (FPS) in the table.

Methods	-	bicubic (x2)	bicubic (x4)	SRGAN [31] (x2)	DBPN [20] (x2)	BSRGAN [59] (x2)
AP	42.2	50.6	36.5	54.3	55.6	70.8
FPS	<b>87.4</b>	50.0	16.2	51.1	50.6	51.4

Methods	BM3D Restormer [58]	Deg $t$	Deg $t + N$	$D_r + \text{Detection}$	<b>AERIS</b>
AP	50.9	52.6	76.0	80.0	<b>80.5</b>
FPS	86.0	86.1	<b>87.4</b>	<b>87.4</b>	43.6

in Sec. 5.3, we separately using different type of restoration methods to handle different type of degradation, for degradation specific comparison.

On the other hand, we also compare with other training strategies of the network setting, to evaluate their robustness improvement on detection. As it shown in Table 1 and Table 3. “Deg  $t$ ” corresponds to train the detector with the LR images  $t(x)$  random generated from HR images  $x$ , with the random degradation transformations in Sec. 3. “Deg  $t + N$ ” means to mix the training data of HR image  $x$  and LR images  $t(x)$ . “ $D_r + \text{Detection}$ ” is the structure like [19] which joint optimize a pre-process SR block and following object detector.

For fairness, all comparison methods adopt the same data augmentation process and same training setting. In the testing stage, all the results are tested on a single RTX 6000 GPU. We compare the speed by reporting the frames per second (FPS) in the experiments, a simple illustration is shown in Table 2.

## 5.2 Multi-Degradation Evaluation

To evaluate object detectors’ robustness under diverse degradation of real-world images. Different from previous works [22,27,38,45] that only consider single degradation type at a time. Following the practical degradation model [13,36] in Eq.1, we design the experiments on images with multiply degradation conditions and down-scale ratios, to verify detection robustness in real-world diverse condition. We generate **COCO-d** dataset from original COCO **va12017** dataset and **KITTI-d** dataset from original KITTI **va1** dataset. We give per-image a random blur kernel (isotropic Gaussian kernel  $k_{iso}$ , anisotropic Gaussian kernel  $k_{aniso}$ ) and random noise level (AWGN noise with variance  $\sigma \sim U(0, 25/255)$ ). As for the resolutions, we down-sample per image in COCO **va12017** with a random rate  $s \sim U(1.0, 4.0)$ .

The experimental results are shown in Table 1 and Table 2, we add the up-scale ratio ( $\times 2$ ,  $\times 4$ ) after name of SR methods. We first give the detection results on original COCO **va12017** and **COCO-d** dataset. The object detector is easily affected and the performance suffers a large decrease on the multi-degradation condition. Table 1 also verifies up-scale higher resolution (either by interpolate or SR pre-process) improves the detection performance on small objects, but has a negative impact on the middle and large objects. Restoration methods would also

Table 3: Comparison with SOTA restoration methods and SSL methods on COCO val2017 with **noise**, **gaussian blur** condition and **low resolution** condition (down-sampling ratio: **2** and **4**). Here **CenterNet** with **ResNet-18** backbone was adopted. Higher resolution results are in **blue** background.

(a) Noise.					(b) Blur.						
Method	$\sigma$	(5, 50)			Method	$k$	Mix				
		15	25	50				26.7	25.8	23.9	23.1
-		22.8	26.8	23.8	15.4	-	26.7	25.8	23.9	23.1	
IRCNN [61]		22.6	26.8	24.2	16.8	EPLL [65]	27.8	26.8	25.4	25.2	
Swin-IR [33]		24.2	28.0	25.6	19.3	IRCNN [61]	26.7	26.9	24.1	22.8	
Restormer [58]		23.8	27.6	25.1	18.9	Deg $t + N$	28.8	27.5	27.6	27.8	
Deg $t + N$		24.3	27.6	25.0	18.3	$D_r + \text{Detection}$	28.5	27.7	27.5	27.3	
$D_r + \text{Detection}$		24.8	28.5	25.5	19.4	<b>AERIS</b>	<b>29.3</b>	<b>28.6</b>	<b>28.0</b>	<b>28.2</b>	
<b>AERIS</b>		<b>25.1</b>	<b>28.7</b>	<b>26.5</b>	<b>20.2</b>						
(c) Down-sampling (Ratio: 2).					(d) Down-sampling (Ratio: 4).						
Method	metric	AP	AP <sub>s</sub>	AP <sub>m</sub>	AP <sub>l</sub>	Method	metric	AP	AP <sub>s</sub>	AP <sub>m</sub>	AP <sub>l</sub>
		20.2	1.5	16.1	49.8			-	8.2	0.0	3.2
SRGAN [31] ( $\times 2$ )		24.0	6.2	25.4	39.8	DBPN [20] ( $\times 2$ )		14.8	1.0	9.9	39.7
DBPN [20] ( $\times 2$ )		25.1	7.3	27.1	41.9	DBPN [20] ( $\times 4$ )		12.2	1.7	11.7	23.4
Swin-IR [33] ( $\times 2$ )		25.4	7.6	27.0	42.4	Swin-IR [33] ( $\times 2$ )		15.2	1.1	10.1	39.9
$D_r + \text{Detection}$		26.0	8.4	26.2	46.5	Swin-IR [33] ( $\times 4$ )		12.8	1.8	12.2	23.4
<b>AERIS</b>		25.4	4.6	25.8	<b>49.6</b>	$D_r + \text{Detection}$		15.1	1.8	12.7	40.1
<b>AERIS</b> ( $\times 2$ )		<b>26.8</b>	<b>8.6</b>	<b>28.8</b>	45.2	<b>AERIS</b>		13.0	0.8	10.2	<b>42.6</b>
						<b>AERIS</b> ( $\times 2$ )		<b>15.8</b>	<b>2.0</b>	<b>13.2</b>	40.9

invalid if degradation types and down-sampling scales are diverse, among several restoration methods, real-world SR method BSRGAN [59] could get satisfactory results. Our AERIS model could get best performance in most of metrics, even with a lower input resolution, but the one limitation is that AERIS could not get best performance on small object.

### 5.3 Degradation Specific Evaluation

To further understand the advantages of AERIS, we design the experiments on single degradation conditions. We separately make experiments on noise, gaussian blur and low-resolution conditions. For noise and blur condition, in training stage of three SSL methods, we generate  $t(x)$  from  $x$  with noise  $n$  (variance  $\sigma \sim U(0, 50/255)$ ) and blur kernel  $k$  (same as Sec. 3). As for low

Table 4: Ablation results of CenterNet on **COCO-d** dataset.

CenterNet	+Deg $t$	+feature connect	+ARRD loc 1	+ARRD loc 2	+ARRD loc 3	mAP (COCO-d)
✓						14.5
✓	✓					16.8 (+2.3)
✓	✓	✓				17.5 (+3.0)
✓	✓			✓		17.7 (+3.2)
✓	✓	✓	✓			18.2 (+3.7)
✓	✓	✓		✓		18.4 (+3.9)
✓	✓	✓			✓	17.9 (+3.4)

resolution condition, we generate  $t(x)$  from  $x$  with down-sampling ratio  $s$ . Here we discuss three conditions as follow:

**Performance *w.r.t* Noise & Blur.** For noise’s affect on object detection, we process COCO **val2017** with random Gaussian noise  $n \sim N(0, \sigma)$ , we first random choose variance  $\sigma$  from uniform distribution  $U(5/255, 50/255)$  for mix noise level evaluation. Then we take three different noise level:  $\sigma = 15$ ,  $\sigma = 25$  and  $\sigma = 50$  for specific evaluation. We compare AERIS with SOTA denoise methods IRCNN [61], Swin-IR [33] and Restormer [58] and also compare with other training strategies. We report the average precision (AP) in Table 3a.

For blur’s affect on object detection, we first process COCO **val2017** with random isotropic Gaussian kernel  $k_{iso}$  and anisotropic Gaussian kernel  $k_{aniso}$  (probability both 0.5) as **Mix** evaluation. Then we specifically choose three degradation kernels for specific evaluation (see Table 3b). We compare with Gaussian deblur methods EPLL [65] and IRCNN [61] and other two other training strategies, then report AP value in Table 3b. Our AERIS gains **best** performance under various noise and blur conditions, among image restoration methods and other training strategies.

**Performance *w.r.t* Low Resolution.** To evaluate low resolution’s affect on detection task, we down-scale original COCO **val2017** with a fixed down-sampling ratio  $s$ . Here we set the down-sampling ratio to 2 and 4, and then compare with SOTA SISR methods SRGAN [31], DBPN [20], Swin-IR [33] and pre-stage method “ $D_r$  + Detection”. We report the total AP and different level objects’ detection performance ( $AP_s$ ,  $AP_m$ ,  $AP_l$ ) in Table 3c and Table 3d. We also make an additional experiments to up-scale input images with ratio 2 and then send into AERIS model as **AERIS** ( $\times 2$ ), for same resolution comparison with SR methods and pre-upsampling method.

#### 5.4 Ablation Study

To evaluate the location of ARRD on CenterNet [64], we separately make the ablation study to evaluate each part’s effectiveness, as shown in Table 4, “+Deg  $t$ ” refers to adding the degradation transformation in Sec. 3. “+feature connect” means the feature connection process in Fig. 3. “+ARRD” adds the decoder

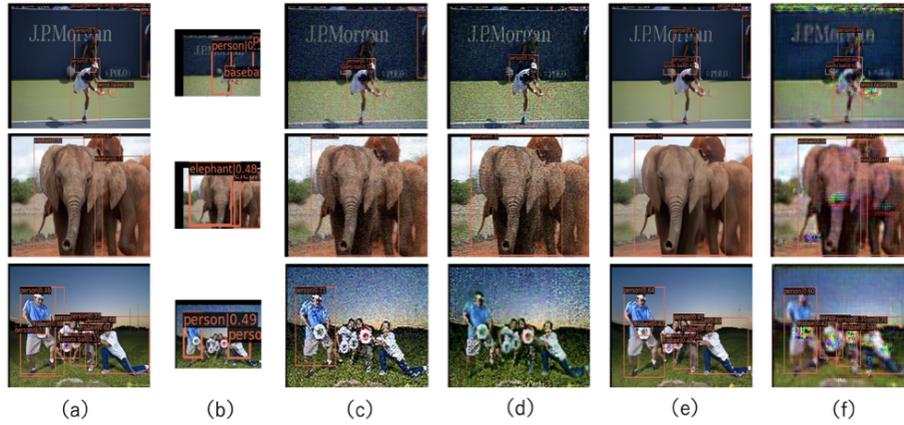


Fig. 4: Exemplar detection results on MS COCO 2017 dataset [35]. (a)/(b) is CenterNet trained on normal images and tested on normal/**COCO-d** dataset, (c)/(d)/(e) is CenterNet tested on the degraded image restored by individual SR algorithm SRGAN [31]/Real-SR [4]/BSRGan [59]. (f) is the detection result of our AERIS and we use the output of ARRD  $D_r$  as background images.

$D_r$  upon the network structure. We also evaluate adding ARRD on different location of the up-sampling blocks, “loc 1” means the shallow up-sampling layer, “loc 2” means middle up-sampling layer and “loc 3” means the final up-sampling layer (“loc 3” also connect with the detection decoder  $D_o$ ), as our finding, to implement ARRD on the middle layer could get best result.

## 6 Conclusion

In this paper, we propose a novel self-supervised framework, AERIS, to handle object detection for arbitrary degraded low resolution images. To capture the dynamics of feature representations under diverse resolution and degradation conditions, we propose a degradation equivariant representation that is generic and could be implemented on popular detection architectures. To further combine the strength of the existing progress on super resolution (SR), we also introduce an arbitrary-resolution restoration decoder that supervises the latent representation to preserve the visual structure. The extensive experiments demonstrate that our AERIS achieves SOTA results on two mainstream public datasets among different degradation conditions (resolution, noise and blur).

## 7 Acknowledgement

This work was supported by JST Moonshot R&D Grant Number JPMJMS2011 and JST ACT-X Grant Number JPMJAX190D, Japan.

## References

1. Affi, M., Brown, M.S.: What else can fool deep learning? addressing color constancy errors on deep neural network performance. In: International Conference on Computer Vision (ICCV) (2019)
2. Bai, Y., Zhang, Y., Ding, M., Ghanem, B.: Finding tiny faces in the wild with generative adversarial network. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21–30 (2018). <https://doi.org/10.1109/CVPR.2018.00010>
3. Bell-Kligler, S., Shocher, A., Irani, M.: Blind super-resolution kernel estimation using an internal-gan. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019), <https://proceedings.neurips.cc/paper/2019/file/5fd0b37cd7dbbb00f97ba6ce92bf5add-Paper.pdf>
4. Cai, J., Zeng, H., Yong, H., Cao, Z., Zhang, L.: Toward real-world single image super-resolution: A new benchmark and a new model. In: Proceedings of the IEEE International Conference on Computer Vision (2019)
5. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
6. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. arXiv preprint arXiv:2012.09161 (2020)
7. Cohen, T.S., Welling, M.: Group equivariant convolutional networks. In: Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48. p. 2990–2999. ICML'16, JMLR.org (2016)
8. Cui, Z., Qi, G.J., Gu, L., You, S., Zhang, Z., Harada, T.: Multitask aet with orthogonal tangent regularity for dark object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2553–2562 (October 2021)
9. Dai, D., Wang, Y., Chen, Y., Van Gool, L.: Is image super-resolution helpful for other vision tasks? In: IEEE Winter Conference on Applications of Computer Vision (WACV) (2016)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
11. Dong, C., Loy, C.C., He, K., *et.al*: Learning a deep convolutional network for image super-resolution. In: Computer Vision – ECCV 2014. pp. 184–199 (2014)
12. Efrat, N., Glasner, D., Apartsin, A., Nadler, B., Levin, A.: Accurate blur models vs. image priors in single image super-resolution. In: 2013 IEEE International Conference on Computer Vision. pp. 2832–2839 (2013). <https://doi.org/10.1109/ICCV.2013.352>
13. Elad, M., Feuer, A.: Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. IEEE Transactions on Image Processing **6**(12), 1646–1658 (1997). <https://doi.org/10.1109/83.650118>
14. Everingham, M., Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. Int. J. Comput. Vision **88**(2), 303–338 (Jun 2010). <https://doi.org/10.1007/s11263-009-0275-4>, <https://doi.org/10.1007/s11263-009-0275-4>

15. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)
16. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=S1v4N210->
17. Girshick, R.: Fast r-cnn. In: 2015 IEEE International Conference on Computer Vision (ICCV). pp. 1440–1448 (2015). <https://doi.org/10.1109/ICCV.2015.169>
18. Gu, J., Lu, H., Zuo, W., Dong, C.: Blind super-resolution with iterative kernel correction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
19. Haris, M., Shakhnarovich, G., Ukita, N.: Task-driven super resolution: Object detection in low-resolution images. In: Mantoro, T., Lee, M., Ayu, M.A., Wong, K.W., Hidayanto, A.N. (eds.) Neural Information Processing. pp. 387–395. Springer International Publishing, Cham (2021)
20. Haris, M., Shakhnarovich, G., Ukita, N.: Deep back-projection networks for super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016). <https://doi.org/10.1109/CVPR.2016.90>
22. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. Proceedings of the International Conference on Learning Representations (2019)
23. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. CoRR **abs/1704.04861** (2017), <http://arxiv.org/abs/1704.04861>
24. Hu, X., Mu, H., Zhang, X., Wang, Z., Tan, T., Sun, J.: Meta-sr: A magnification-arbitrary network for super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
25. Irie, K., McKinnon, A.E., Unsworth, K., Woodhead, I.M.: A technique for evaluation of ccd video-camera noise. IEEE Transactions on Circuits and Systems for Video Technology **18**(2), 280–284 (2008). <https://doi.org/10.1109/TCSVT.2007.913972>
26. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European Conference on Computer Vision (2016)
27. Kamann, C., Rother, C.: Benchmarking the robustness of semantic segmentation models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
28. Karaimer, H.C., Brown, M.S.: A software platform for manipulating the camera imaging pipeline. In: European Conference on Computer Vision (ECCV) (2016)
29. Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1646–1654 (2016). <https://doi.org/10.1109/CVPR.2016.182>
30. Lai, W.S., Huang, J.B., Ahuja, N., Yang, M.H.: Deep laplacian pyramid networks for fast and accurate super-resolution. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5835–5843 (2017). <https://doi.org/10.1109/CVPR.2017.618>

31. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., Shi, W.: Photo-realistic single image super-resolution using a generative adversarial network. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 105–114 (2017). <https://doi.org/10.1109/CVPR.2017.19>
32. Li, K., Wang, Y., Zhang, J., Gao, P., Song, G., Liu, Y., Li, H., Qiao, Y.: Uniformer: Unifying convolution and self-attention for visual recognition (2022)
33. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: IEEE International Conference on Computer Vision Workshops (2021)
34. Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M.: Enhanced deep residual networks for single image super-resolution. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 1132–1140 (2017). <https://doi.org/10.1109/CVPRW.2017.151>
35. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) Computer Vision – ECCV 2014. pp. 740–755. Springer International Publishing, Cham (2014)
36. Liu, C., Sun, D.: On bayesian adaptive video super resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(2), 346–360 (2014). <https://doi.org/10.1109/TPAMI.2013.127>
37. Liu, C., Sun, D.: On bayesian adaptive video super resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(2), 346–360 (2014). <https://doi.org/10.1109/TPAMI.2013.127>
38. Liu, D., Wen, B., Liu, X., Wang, Z., Huang, T.S.: When image denoising meets high-level vision tasks: A deep learning approach. In: IJCAI (2018)
39. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
40. Luo, F., Wu, X., Guo, Y.: Functional neural networks for parametric image restoration problems. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems* (2021), <https://openreview.net/forum?id=MMZ4djXrwbu>
41. Ma, X., Wang, Z., Zhan, Y., Zheng, Y., Wang, Z., Dai, D., Lin, C.W.: Both style and fog matter: Cumulative domain adaptation for semantic foggy scene understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18922–18931 (June 2022)
42. Mao, M., Peng, G., Zhang, R., Zheng, H., Ma, T., Peng, Y., Ding, E., Zhang, B., Han, S.: Dual-stream network for visual recognition. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems* (2021), <https://openreview.net/forum?id=Ajfd1JjeVKN>
43. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 779–788 (2016). <https://doi.org/10.1109/CVPR.2016.91>
44. Sabour, S., Frosst, N., Hinton, G.E.: Dynamic routing between capsules. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 3859–3869. NIPS’17, Curran Associates Inc., Red Hook, NY, USA (2017)
45. Sayed, M., Brostow, G.: Improved handling of motion blur in online object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1706–1716 (June 2021)

46. Shermeyer, J., Van Etten, A.: The effects of super-resolution on object detection performance in satellite imagery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2019)
47. Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1874–1883 (2016). <https://doi.org/10.1109/CVPR.2016.207>
48. Singh, B., Davis, L.S.: An analysis of scale invariance in object detection snip. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
49. Singh, B., Najibi, M., Davis, L.S.: Sniper: Efficient multi-scale training. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 31. Curran Associates, Inc. (2018), <https://proceedings.neurips.cc/paper/2018/file/166cee72e93a992007a89b39eb29628b-Paper.pdf>
50. Talebi, H., Milanfar, P.: Learning to resize images for computer vision tasks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 497–506 (October 2021)
51. Tong, T., Li, G., Liu, X., Gao, Q.: Image super-resolution using dense skip connections. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 4809–4817 (2017). <https://doi.org/10.1109/ICCV.2017.514>
52. Vasiljevic, I., Chakrabarti, A., Shakhnarovich, G.: Examining the impact of blur on recognition by convolutional networks (2017)
53. Wang, L., Li, D., Zhu, Y., Tian, L., Shan, Y.: Dual super-resolution learning for semantic segmentation. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3773–3782 (2020). <https://doi.org/10.1109/CVPR42600.2020.00383>
54. Wang, L., Wang, Y., Dong, X., Xu, Q., Yang, J., An, W., Guo, Y.: Unsupervised degradation representation learning for blind super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10581–10590 (June 2021)
55. Wang, Z., Chang, S., Yang, Y., Liu, D., Huang, T.S.: Studying very low resolution recognition using deep networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
56. Wei, K., Fu, Y., Yang, J., Huang, H.: A physics-based noise formation model for extreme low-light raw denoising. In: IEEE Conference on Computer Vision and Pattern Recognition (2020)
57. Yang, W., Yuan, Y., Ren, W., *et al.*: Advancing image understanding in poor visibility environments: A collective benchmark study. IEEE Transactions on Image Processing **29**, 5737–5752 (2020). <https://doi.org/10.1109/TIP.2020.2981922>
58. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: CVPR (2022)
59. Zhang, K., Liang, J., Van Gool, L., Timofte, R.: Designing a practical degradation model for deep blind image super-resolution. In: IEEE International Conference on Computer Vision (2021)
60. Zhang, K., Van Gool, L., Timofte, R.: Deep unfolding network for image super-resolution. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 3217–3226 (2020)

61. Zhang, K., Zuo, W., Gu, S., Zhang, L.: Learning deep cnn denoiser prior for image restoration. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 3929–3938 (2017)
62. Zhang, K., Zuo, W., Zhang, L.: Learning a single convolutional super-resolution network for multiple degradations. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 3262–3271 (2018)
63. Zhang, L., Qi, G.J., Wang, L., Luo, J.: Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2547–2555 (2019)
64. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. CoRR **abs/1904.07850** (2019), <http://arxiv.org/abs/1904.07850>
65. Zoran, D., Weiss, Y.: From learning models of natural image patches to whole image restoration. In: 2011 International Conference on Computer Vision. pp. 479–486 (2011). <https://doi.org/10.1109/ICCV.2011.6126278>
66. Zou, W.W.W., Yuen, P.C.: Very low resolution face recognition problem. IEEE Transactions on Image Processing **21**(1), 327–340 (2012). <https://doi.org/10.1109/TIP.2011.2162423>