

URIE: Universal Image Enhancement for Visual Recognition in the Wild

Taeyoung Son¹ Juwon Kang¹ Namyup Kim¹ Sunghyun Cho² Suha Kwak²

¹Department of Computer Science and Engineering

²Graduate School of Artificial Intelligence

POSTECH, Pohang, Korea

<http://cvlab.postech.ac.kr/research/URIE/>

Abstract. Despite the great advances in visual recognition, it has been witnessed that recognition models trained on clean images of common datasets are not robust against distorted images in the real world. To tackle this issue, we present a Universal and Recognition-friendly Image Enhancement network, dubbed URIE, which is attached in front of existing recognition models and enhances distorted input to improve their performance without retraining them. URIE is universal in that it aims to handle various factors of image degradation and to be incorporated with any arbitrary recognition models. Also, it is recognition-friendly since it is optimized to improve the robustness of following recognition models, instead of perceptual quality of output image. Our experiments demonstrate that URIE can handle various and latent image distortions and improve the performance of existing models for five diverse recognition tasks where input images are degraded.

Keywords: Visual recognition, Image enhancement

1 Introduction

The development of deep learning has made the great advances in visual recognition. Especially, the recent advances in this field have been concentrated on accurate and comprehensive understanding of high-quality images taken in limited environments. However, images in the real world are often corrupted by various factors including adverse weather conditions, sensor noise, under- or over-exposure, motion blur, and compression artifact. Since most of existing recognition models do not consider these factors, their performance could be easily degraded when input image is corrupted [4,11,42], which could be fatal in safety-critical applications like autonomous driving.

A straightforward way to resolve this issue is to improve the quality of input image fed to recognition models by a preprocessing method. Image restoration seems suited to this purpose at first glance, but unfortunately, it is not in general for the following reasons. First, most of restoration methods assume that images are degraded by a single and known factor of corruption. They are thus hard to handle realistic recognition scenarios where images can be corrupted by various

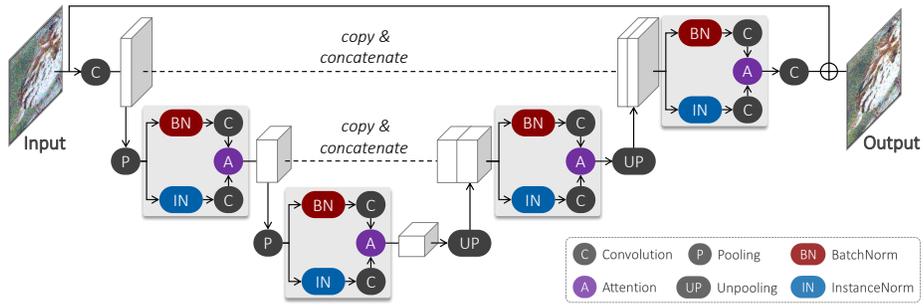


Fig. 1. Overall architecture of URIE. The Selective Enhancement Modules (SEM) are indicated by gray rectangles. Details of these modules are illustrated in Fig. 2.

factors whose types and degrees are hidden. Second, they are not trained for visual recognition but for human perception. Hence input images restored by them do not always guarantee performance improvement of recognition models as demonstrated in [4,16,24,35]. Finally, recent restoration networks tend to be heavy in computation, thus could make the entire system overly expensive when integrated with recognition models.

This paper addresses the problem through a novel image enhancement model dedicated to robust visual recognition in the wild. We call our model Universal and Recognition-friendly Image Enhancement network, dubbed URIE, since we keep the following properties in our mind when developing the model. (1) *Universally applicable*: Our model aims to be attached in front of any existing recognition models to improve their robustness in any situations without retraining them. To this end, it has to deal with various factors of degradation whose types and intensities are latent, and be applicable to diverse recognition tasks and models. (2) *Recognition friendly*: Our model should be optimized for improving the performance of following recognition models, instead of making images looking plausible. (3) *Computationally efficient*: The complexity of our model has to be low to minimize the computational burden it additionally imposes.

We come up with a new network architecture and its own training strategy to implement these properties. In particular, the architecture of URIE is devised to deal with various and latent distortions efficiently and effectively. To handle diverse image distortions, URIE runs a large number of different image enhancement procedures internally by a single feed-forward path. At the same time, it has the capability to select enhancement procedures appropriate for dealing with the latent type of input distortion. These features are given by Selective Enhancement Module (SEM), a basic building block of URIE. The architecture of URIE composed of SEMs is illustrated in Fig. 1. Each SEM provides two different steps towards image enhancement. Hence, by concatenating SEMs, the network will have an exponentially large number of enhancement procedures. Moreover, each SEM has an attention mechanism that assigns a larger attention weights to its enhancement step more appropriate to deal with the input distor-

tion than the other; this enables URIE to dynamically select useful enhancement procedures depending on the latent type of distortion.

In addition, our new training method enables URIE to be recognition-friendly as well as universal. During training, URIE is coupled to an image classification network pretrained on the original ImageNet dataset [3]. It is then trained in an end-to-end manner to improve the performance of the classifier on distorted images. As training data, we adopt ImageNet-C [11], in which images of the original ImageNet are degraded by 19 classes of distortion. This training strategy gives URIE the recognition-friendly property since the training objective is not the perceptual quality of output image but the recognition performance of the coupled classifier. In addition, we expect that training on such a large-scale dataset with a variety of corruptions allows our model to be less sensitive to the task and architecture of the following recognition model, analogously to ImageNet pretrained networks that have been used for various purposes [22,23,30] other than their original target.

The efficacy of URIE is evaluated on five different recognition tasks, in which recognition models are pretrained on clean images and tested with URIE on corrupted ones with no finetuning. Experimental results demonstrate that URIE improves recognition performance in the presence of image distortions including those unseen during training, and regardless of target tasks and recognition networks coupled with it. URIE is also compared with the state-of-the-art image restoration model [32], and outperforms it in terms of both recognition accuracy and computational efficiency. In summary, our contribution is four-fold:

- We introduce a new method towards robust visual recognition. Our method improves the robustness of existing recognition models without retraining them through a single and universal image enhancement module.
- We propose a novel network architecture for image enhancement that can handle various and latent types of image distortions effectively and efficiently.
- We present a new training strategy that enables our enhancement model to be universally applicable and recognition-friendly.
- In our experiments, our model improves the recognition performance of existing models substantially in the presence of various image distortions regardless of tasks and models it is tested on.

2 Related Work

2.1 Fragility of Visual Recognition Models

Fragility of deep neural networks for visual recognition has been studied to comprehend their robustness against challenging conditions that they often encounter in the real world. Diamond *et al.* [4] demonstrated that image classification networks are sensitive to noise and blur. Similarly, Pei *et al.* [24] showed that haze can degrade the performance of classification networks. On the other hand, Zendel *et al.* [42] defined various factors of image degradation appearing

frequently in autonomous driving scenarios, and analyzed the robustness of existing semantic segmentation networks to them. Hendrycks and Dietterich [11] studied performance degradation of image classifiers by simulated image distortions. These results indicate that existing recognition networks trained on clean images are fragile in the presence of image distortion. Geirhos *et al.* [7] showed that one way to resolve this issue is to finetune models to distorted images but such models are not well generalized when tested on other distortion types.

Our work is motivated by these observations. We address the above problem through an image enhancement model that deals with a variety of image distortions and is optimized to improve the robustness of recognition models.

2.2 Recognition of Distorted Images

Previous methods to robust recognition of distorted images can be grouped into two classes: *direct recognition* and *recognition-friendly image enhancement*.

Direct recognition of corrupted images. Classical methods in this category rely on visual features robust against image distortions like blur-invariant features [9,38]. Recent methods focus on developing new network architectures and training strategies dedicated to the robust recognition. Wang *et al.* [36] and Singh *et al.* [31] developed neural networks for classifying low resolution images, and Lee *et al.* [14] proposed a multi-task network for understanding road markings in rainy days and nighttime. Also, domain adaptation techniques have been widely adopted for corruption-aware training of recognition models. Considering different types of degradation as individual domains, they allow models for object detection and semantic segmentation to be adapted to adverse weather conditions [2,27,39] and nighttime images [26].

Recognition-friendly image enhancement. Models in this category are designed and trained to enhance input images so that it becomes more informative and task-optimized for following recognizers. Diamond *et al.* [4] proposed a differentiable image processing module learned jointly with a classifier for improved classification of noisy and blurred images. Similarly, Liu *et al.* [19] developed denoising networks for classification and semantic segmentation of noisy images. Gomez *et al.* [8] studied network architectures that enhance images taken in difficult illumination conditions for robust visual odometry. Sharma *et al.* [29] investigated a way to dynamically predict image processing filters that are applied to input images and make them more classification-friendly.

The aforementioned methods of both categories are limited in terms of universality as they handle a single degradation type and/or are optimized for only one recognition task. On the other hand, our goal is to build a single enhancement model that can cope with various types of image degradation and be attached to any architectures for any recognition tasks.

2.3 Image Restoration

Most of image restoration methods focus on a single type of image distortion (*e.g.*, denoising [43], deblurring [21], dehazing [15], deraining [40], and super-

resolution [5]), and consequently they are incapable of handling real world images where the type and intensity of distortion are hidden in general. A few network architectures for image restoration are designed to handle various types of degradation [18,33], but not at the same time as they have to be trained for each restoration objective. Recently, the above issue of conventional image restoration has been addressed by new architectures that can deal with various and latent factors of degradation [32,41]. However, they are not optimized for image recognition, thus may not guarantee improved performance of coupled visual recognition models. In addition, since they are computationally heavy, they will increase the complexity of overall system significantly if integrated.

3 URIE: Architecture and Training

Our final target in this paper is to improve the robustness of existing recognition models to diverse image distortions *without retraining them*. URIE is proposed to address this challenging problem by converting an input image degraded by latent distortion into another that can be recognized more reliably by any existing models. To this end, it is designed and trained to be universally applicable to various image distortions and recognition architectures, while being recognition-friendly for improved performance of the following recognition models.

Specifically, URIE is devised to derive many different enhancement procedures from a single feed-forward architecture, and at the same time, to dynamically select procedures useful to handle the input distortion. Thanks to these features, URIE can be universally applicable to a large variety of image distortions whose types and intensities are latent. This architecture is embodied by a series of basic building blocks, which we call SEMs. Furthermore, we provide a new training method that enables URIE to be recognition-friendly and universally applicable to any recognition models.

The remainder of this section presents details of SEM, discusses overall architecture and other design points of URIE, and describes our training strategy that makes URIE universally applicable and recognition-friendly.

3.1 Selective Enhancement Module

In URIE, an image enhancement procedure is represented as a sequence of convolution layers. As a basic building block of the network, SEM provides individual steps towards image enhancement (*i.e.*, convolution layers) so that a series of such steps forms an enhancement procedure. To be specific, a single SEM presents two diverse enhancement steps, and has the capability to estimate which one among them is more appropriate to deal with the input distortion. The detailed architecture of SEM is illustrated in Fig. 2.

The two enhancement steps in SEM have separate convolution layers associated with different normalization operations, Instance Normalization (IN) [34] and Batch Normalization (BN) [12]. We adopt these two different normalization operations to learn a pair of diverse and complementary enhancement functions.

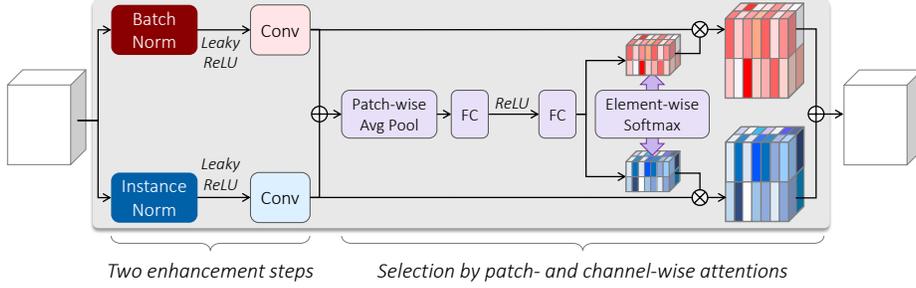


Fig. 2. Details of SEM. \oplus and \otimes indicate element-wise summation and multiplication between feature maps, respectively.

IN has been used in neural style transfer to normalize the style of image while keeping its major content by normalizing statistics per image independently. Regarding a distortion type as an image style, we expect the enhancement step with IN to moderate the effect of distortion for individual images. On the other hand, the enhancement step with BN maintains the distortion information disregarded by the other, and is trained by considering a variety of distorted images together within each minibatch. We thus expect that it is complementary to its IN counterpart and useful for understanding input distortion. Empirical justification of using IN and BN can be found in the supplementary material.

The two outputs of these enhancement steps are fed to the attention module that determines which one among them is more useful to deal with the input distortion for each channel and each local patch of the input image. Let $F_{\text{in}} \in \mathbb{R}^{w \times h \times C}$ and $F_{\text{bn}} \in \mathbb{R}^{w \times h \times C}$ be the output feature maps of the two enhancement steps with resolution of $w \times h$ and C channels. They are first aggregated by element-wise summation into a single feature map, then we divide it into the 4×4 regular grid and conduct average pooling on each cell so that feature vectors for 16 local patches are obtained. The local feature vectors are then concatenated into one large vector, which is in turn fed to the following fully connected layers. The final output of the fully connected layers are two score vectors $\mathbf{s}_{\text{in}} \in \mathbb{R}^{16C}$ and $\mathbf{s}_{\text{bn}} \in \mathbb{R}^{16C}$, each of whose elements indicates how useful a specific channel of F_{in} or F_{bn} is for handling the distortion in a specific local patch.

The score vectors \mathbf{s}_{in} and \mathbf{s}_{bn} are used to compute attention weights. They are first reshaped into score tensors $S_{\text{in}} \in \mathbb{R}^{4 \times 4 \times C}$ and $S_{\text{bn}} \in \mathbb{R}^{4 \times 4 \times C}$, then converted to attention weight tensors A_{in} and A_{bn} of the same dimension through the element-wise softmax operation:

$$A_{\text{in}}[i, j, c] = \frac{\exp S_{\text{in}}[i, j, c]}{\exp S_{\text{in}}[i, j, c] + \exp S_{\text{bn}}[i, j, c]}, \quad (1)$$

$$A_{\text{bn}}[i, j, c] = \frac{\exp S_{\text{bn}}[i, j, c]}{\exp S_{\text{in}}[i, j, c] + \exp S_{\text{bn}}[i, j, c]}. \quad (2)$$

The spatial dimension of the two attention tensors are enlarged to $w \times h$ by bilinear interpolation to be multiplied to the outputs of the two enhancement steps, F_{in} and F_{bn} . The final output of SEM, denoted by $Y \in \mathbb{R}^{w \times h \times C}$, is then computed by weighted summation of F_{in} and F_{bn} :

$$Y = (A'_{\text{in}} \otimes F_{\text{in}}) \oplus (A'_{\text{bn}} \otimes F_{\text{bn}}), \quad (3)$$

where A'_{in} and A'_{bn} are the resized attention tensors, \otimes indicates element-wise multiplication, and \oplus denotes element-wise summation.

Note that the attention tensors have both spatial (*i.e.*, patch) and channel dimensions. In Eq. (3), the channel-wise attention facilitates more flexible and diverse integration of the two enhancement steps since different channels can be processed by different enhancement steps; although a SEM provides only two enhancement steps, the number of their output combinations is in principle extremely large. On the other hand, the patch-wise attention allows to process different image regions with different enhancement processes.

3.2 Overall Architecture

The overall architecture of URIE is illustrated in Fig. 1. Its structure resembles U-Net [25], but is clearly distinct from the conventional U-Net architectures in that its all layers, except for the first and the last, are not ordinary convolution layers but SEMs introduced in Sec. 3.1. The reason for following the design principle of U-Net is two-fold. First, it effectively enlarges the receptive field so as to capture high-level semantic information. This is essential for URIE that performs recognition-aware image enhancement, unlike ordinary restoration models that do not necessarily need to understand image contents. Second, it is computationally efficient since it reduces the sizes of intermediate features using pooling operations. Reducing feature resolution may degrade the perceptual quality of output as a side effect, which however is not our concern since our objective is improved recognition performance of following recognition model.

The key components in this architecture are SEMs, which allow URIE to handle multiple and latent factors of image degradation efficiently and effectively. Since each SEM provides two enhancement steps and a sequence of such steps can be regarded as an individual enhancement procedure, the concatenation of SEMs leads to an exponentially large number of enhancement procedures, which can handle a large variety of distortions. Moreover, through the attentions drawn by SEMs, URIE can dynamically select enhancement procedures depending on

Table 1. Details of URIE. $W \times H$ denotes the input resolution.

	Output	Layer specification
(1)	$W \times H$	9×9 conv, 32, stride 1
(2)	$W/2 \times H/2$	2×2 max pool, stride 2
(3)	$W/2 \times H/2$	SEM[3×3 , 64, stride 1]
(4)	$W/4 \times H/4$	2×2 max pool, stride 2
(5)	$W/4 \times H/4$	SEM[3×3 , 64, stride 1]
(6)	$W/2 \times H/2$	$2 \times$ bilinear upsampling
(7)	$W/2 \times H/2$	concat[(3), (6)]
(8)	$W/2 \times H/2$	SEM[3×3 , 32, stride 1]
(9)	$W \times H$	$2 \times$ bilinear upsampling
(10)	$W \times H$	concat[(5), (9)]
(11)	$W \times H$	SEM[3×3 , 16, stride 1]
(12)	$W \times H$	3×3 conv, 3, stride 1

the latent type of input distortion. Since the selection is done for each local patch of the input image through the patch-wise attentions, URIE can handle different image regions with different enhancement procedures. This is more useful and realistic than applying a single enhancement function to the whole input image since individual regions of the input could be distorted by different factors in realistic scenarios, *e.g.*, images with locally different haze densities with respect to depths, and images partially over- and under-exposed at the same time.

The specification of URIE, including the size of convolution kernels and the number of channels per layer, is given in Tab. 1. In the table, SEM[$k \times k$, C , stride s] means that each convolution layer in the SEM consists of C number of $k \times k$ kernels with stride s . Note that we set the size of convolution kernels for the first layer especially large to capture rich information about image distortion and to enlarge the receptive field of our shallow network architecture.

3.3 Training Strategy

Our training scheme for URIE is based on two key ideas. (1) *Recognition-aware loss*: We employ a pretrained recognition network to provide URIE with supervisory signals that lead it to be recognition-friendly. Specifically, URIE is coupled with a recognition network that is pretrained on clean images and frozen during training. It is then trained to minimize a recognition loss of the pretrained model on distorted images in an end-to-end manner. (2) *Large-scale learning with distorted images*: We suggest learning URIE using a large-scale image dataset in the presence of diverse distortions so that it may learn enhancement functions insensitive to distortion types and recognition models. This is motivated by the fact that deep neural networks trained in large-scale datasets have been used for various purposes other than their original targets (*e.g.*, ImageNet pretrained classifiers used for other recognition tasks [22,23,30]).

In detail, we employ ResNet-50 [10] pretrained on the original ImageNet [3] as the coupled recognition model, and adopt the ImageNet-C dataset [11] as our training data. URIE is then trained on the ImageNet-C to minimize the cross-entropy loss of the pretrained model. The training dataset consists of images from the original ImageNet yet corrupted artificially by in total 19 different distortions with 5 intensity levels. We use only 15 distortion classes* among them to degrade training images, and call them *seen corruptions*. The other distortion classes** are kept as *unseen corruptions* for evaluating the generalization ability of our model in testing. Note that clean images from the original ImageNet are also used for training to prevent URIE from being biased to distortions.

3.4 Discussion

How our model works. To investigate how URIE works, we apply the model trained in Sec. 3.3 to random images sampled from other datasets and corrupted

* Gaussian noise, shot noise, impulse noise, defocus blur, glass blur, motion blur, zoom blur, snow, frost, fog, brightness, contrast, elastic transform, pixelation, jpeg.

** Speckle noise, Gaussian blur, spatter, saturation.

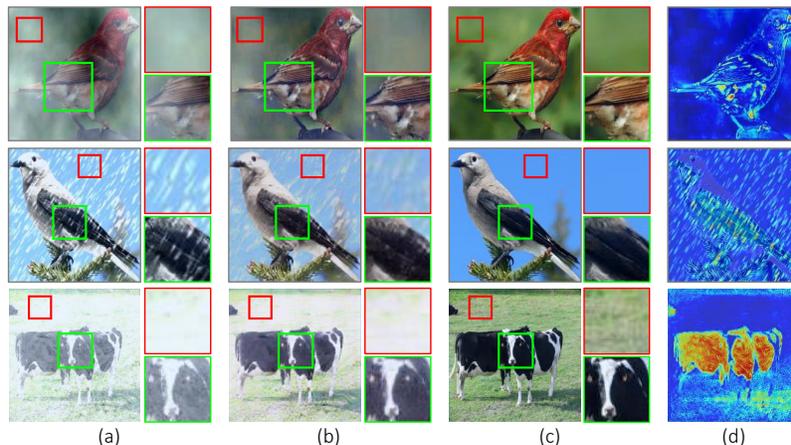


Fig. 3. Example outputs of URIE. (a) Distorted input images. (b) Outputs of URIE. (c) Ground-truth images. (d) Magnitudes of per-pixel intensity change by URIE.

by the same synthetic distortions with the ImageNet-C. Note that URIE is directly used as-is with no finetuning. Qualitative results of URIE on these images are given in Fig. 3. Interestingly, URIE does alleviate the effect of distortions to some degree although it is not trained explicitly for that purpose. More importantly, Fig. 3(b) and 3(d) demonstrate that it tends to manipulate salient local areas only (*e.g.*, object parts) and ignore remainders (*e.g.*, background). This behavior is a result of the recognition-aware learning in Sec. 3.3 that prompts the model to focus more on image regions important for recognition.

Comparison to SKNet. URIE shares a similar idea with the Selective Kernel Network (SKNet) [17] that adopts a self-attention mechanism to utilize intermediate feature maps selectively. However, the two models are clearly distinct since their overall architectures and training schemes are totally different due to their different objectives: image classification in SKNet, and image enhancement in URIE. Besides the above, their attention mechanisms are different in the following aspects. First, SKNet selects from convolution layers of different kernel sizes to implement dynamic receptive fields, but URIE selects from those with different normalization operations, IN and BN, to adapt dynamically to the latent input distortion. Second, SKNet draws only channel-wise attentions, but URIE takes both channel and spatial locations into account when drawing attentions to treat different image regions differently.

Comparison to OWAN. OWAN [32] is the current state of the art in image restoration. It is similar with URIE in that both models consider multiple and latent types of image distortion and use attentions to select appropriate image manipulation process dynamically. The major differences between URIE and OWAN lie in their techniques implementing the ideas. OWAN draws attentions on a number of convolution layers to select useful ones among them, thus has to compute individual feature maps of the layers, which is time consuming. On

the other hands, URIE selects from only two enhancement steps, but mix their outputs up in a patch- and channel-wise manner by drawing attention weights on individual channels and local patches. This scheme allows our model to enjoy an extremely large number of output combinations while keeping computational efficiency. Thanks to these differences, URIE is substantially more efficient than OWAN, *about 6 times faster* in terms of multiply-accumulate operations.

4 Experiments

In this section, we empirically verify that URIE improves performance of diverse recognition models for image classification, object detection, and semantic segmentation on real and artificially degraded images. For image restoration performance of URIE, refer to Tab. 5 of the supplementary material.

4.1 Training Configurations

Each training image is first resized to 256×256 , cropped to 224×224 , and flipped horizontally at random for data augmentation. It is then degraded on-the-fly by the synthetic distortion generator of the ImageNet-C [11], where the type and intensity of distortion are chosen at random. Note that only the 15 seen corruptions are used during training as described in Sec. 3.3. For optimization, we employ Adam [13] with learning rate 0.001, and decay the learning rate every 8 epochs by 10. Our model is trained for 30 epochs with mini-batches of size 112.

4.2 Experimental Configurations

For evaluation of our model, we employ widely used image recognition datasets, ImageNet [3], CUB [37], and PASCAL VOC [6], and generate three different versions of their test or validation sets: *Clean*—original uncorrupted images, *Seen*—images degraded by the 15 seen corruptions, and *Unseen*—images degraded by the 4 unseen corruptions. We did not generate the datasets on-the-fly but prior to evaluation so that they are fixed during testing. In addition, our model is evaluated on the collection of real haze images introduced in [24].

To verify the effectiveness of URIE, it is compared with two different image processing models. One of them is OWAN [32], the state-of-the-art image restoration model devised to deal with multiple and latent types of image distortion. OWAN is trained on the same data with ours for fair comparisons, but by using the original loss proposed in [32] since it is not straightforward to train this model with our recognition-aware loss on the large-scale dataset due to its heavy computational complexity (refer to Sec. 3.4). The other one is URIE-MSE, which is a restoration-oriented variant of URIE and trained using a Mean-Squared-Error loss instead of the recognition-aware loss. URIE and these two models are evaluated in terms of performance of coupled recognizers on the aforementioned datasets. Note that they are never finetuned to test datasets, and all the recognizers coupled with them are trained on the clean images and tested on the distorted images as-is.

Table 2. Classification accuracy on the ImageNet dataset. The numbers in parentheses indicate the differences from the baseline. V16, R50, and R101 denote VGG-16, ResNet-50, and ResNet-101, respectively.

	OWAN			URIE-MSE			URIE		
	Clean	Seen	Unseen	Clean	Seen	Unseen	Clean	Seen	Unseen
V16	69.8 (-0.2)	31.6 (+2.0)	39.8 (+0.8)	59.5 (-10.5)	32.3 (+2.7)	38.1 (-0.9)	67.1 (-2.9)	42.4 (+12.8)	44.8 (+5.8)
R50	74.4 (-0.4)	42.6 (+2.0)	50.3 (+1.2)	66.7 (-7.8)	44.5 (+3.9)	49.3 (+0.2)	72.9 (-1.6)	55.1 (+14.5)	56.5 (+7.4)
R101	75.9 (0.0)	47.4 (+1.5)	54.6 (+0.8)	68.8 (-7.1)	50.0 (+4.1)	54.0 (+0.2)	74.1 (-1.8)	57.8 (+11.9)	59.4 (+5.6)

Table 3. Classification accuracy on the CUB dataset. The numbers in parentheses indicate the differences from the baseline. V16, R50, and R101 denote VGG-16, ResNet-50, and ResNet-101, respectively.

	OWAN			URIE-MSE			URIE		
	Clean	Seen	Unseen	Clean	Seen	Unseen	Clean	Seen	Unseen
V16	79.1 (0.0)	50.4 (+4.2)	42.3 (+1.7)	68.5 (-11.6)	50.5 (+4.3)	39.6 (-1.0)	77.9 (-1.2)	58.8 (+12.6)	48.9 (+8.3)
R50	84.6 (0.0)	52.8 (+3.3)	48.2 (+1.2)	76.5 (-8.1)	58.4 (+8.9)	49.3 (+2.1)	83.7 (-0.9)	64.7 (+15.2)	54.9 (+7.7)
R101	84.9 (-0.1)	55.5 (+3.6)	48.5 (+1.8)	76.2 (-8.8)	59.9 (+8.0)	50.5 (+3.8)	84.1 (-0.9)	67.1 (+15.2)	56.6 (+9.9)

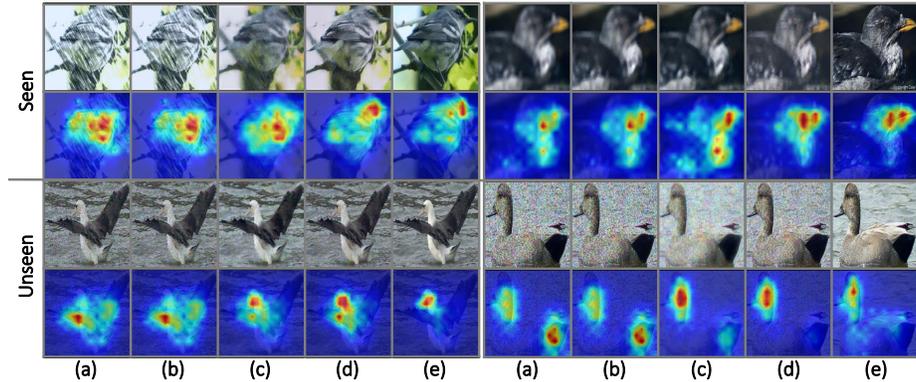


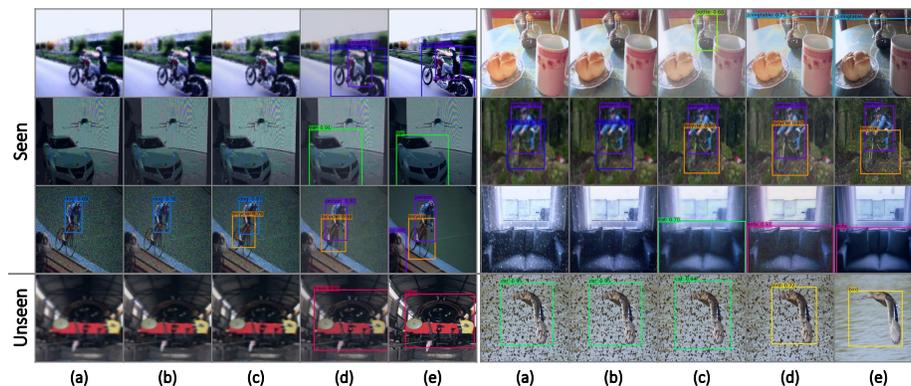
Fig. 4. Qualitative results on the CUB dataset. (a) Input distorted images. (b) OWAN. (c) URIE-MSE. (d) URIE. (e) Ground-truth images. For all images, their grad-CAMs drawn by the ResNet-50 classifier are presented alongside. Examples in the first row are degraded by seen corruptions and the others are by unseen corruptions.

4.3 Performance Evaluation

Classification on ImageNet. We first evaluate and compare the effectiveness of URIE and the other two models for classification on the ImageNet validation set. For evaluation, they are integrated with three different classification networks: VGG-16 [30], ResNet-50 [10] and ResNet-101 [10]. As can be seen in Tab. 2, the performance improvement by URIE is substantial in the case of seen corruptions and nontrivial for unseen corruptions as well, regardless of the classification network it is integrated with. These results suggest that URIE is universally applicable to diverse corruptions and to various recognition models. Also, the performance gap between URIE and the other two methods is significant, which verifies the advantage of our method. In particular, the difference

Table 4. Object detection performance of SSD 300 in mAP (%) on the VOC 2007 dataset. The numbers in parentheses indicate the differences from the baseline.

OWAN			URIE-MSE			URIE		
Clean	Seen	Unseen	Clean	Seen	Unseen	Clean	Seen	Unseen
77.4 (0.0)	51.3 (+1.7)	60.2 (+0.9)	76.5 (-0.9)	52.7 (+3.1)	59.9 (+0.6)	76.5 (-0.9)	59.4 (+9.8)	62.7 (+3.4)

**Fig. 5.** Qualitative results of SSD 300 on the VOC 2007 dataset. (a) Corrupted input. (b) OWAN. (c) URIE-MSE. (d) URIE. (e) Ground-truth. Examples in the top three rows are degraded by seen corruptions and the others are by unseen corruptions.

between URIE and URIE-MSE in performance shows the clear advantage of our recognition-aware training. Moreover, URIE-MSE is comparable to OWAN, although it is six times more efficient as discussed in Sec. 3.4. This demonstrates the advantage of the network architecture we proposed. Unfortunately, URIE marginally degrades the performance in the case of *clean*; we suspect that the architecture of URIE could be biased to distortions while OWAN has more skip connections and better preserves input consequently.

Fine-grained classification on CUB. Our model and the others are evaluated on the CUB validation set to examine if they work universally even in other domains where data distributions are different from that of the training data, *i.e.*, ImageNet. The evaluation setting in this dataset is the same with that of the ImageNet classification. As summarized in Tab. 3, all models turn out to be transferable to some degree in spite of the domain gap between the ImageNet and CUB datasets; we suspect that this is an effect of the large-scale training using the ImageNet-C. However, URIE demonstrates its superiority over the other models by outperforming them substantially. Fig. 3 presents qualitative results of the models and their grad-CAMs [28] computed by the ResNet-50 classifier using the ground-truth labels. In this figure, grad-CAMs computed from the results of URIE are closest to those of ground-truth images, which means that URIE best recovers image regions critical for correct classification.

Object detection and segmentation on PASCAL VOC. This time we evaluate the effectiveness of our model for object detection and semantic seg-

Table 5. Semantic segmentation performance of DeepLab v3 in mIoU (%) on the VOC 2012 dataset. The numbers in parentheses indicate the differences from the baseline.

OWAN			URIE-MSE			URIE		
Clean	Seen	Unseen	Clean	Seen	Unseen	Clean	Seen	Unseen
78.7 (-0.2)	58.2 (+3.3)	65.4(+2.1)	78.1 (-0.8)	60.1 (+5.2)	63.6 (+0.3)	78.6 (-0.3)	67.0 (+12.1)	67.2 (+3.9)

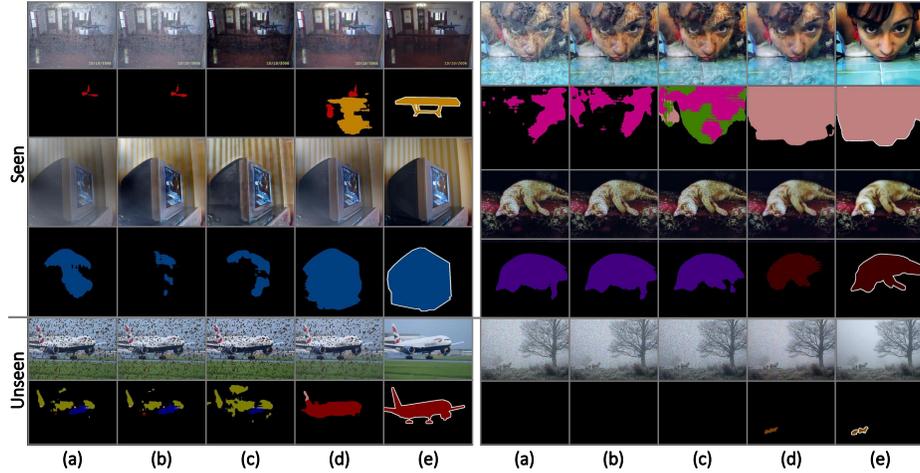


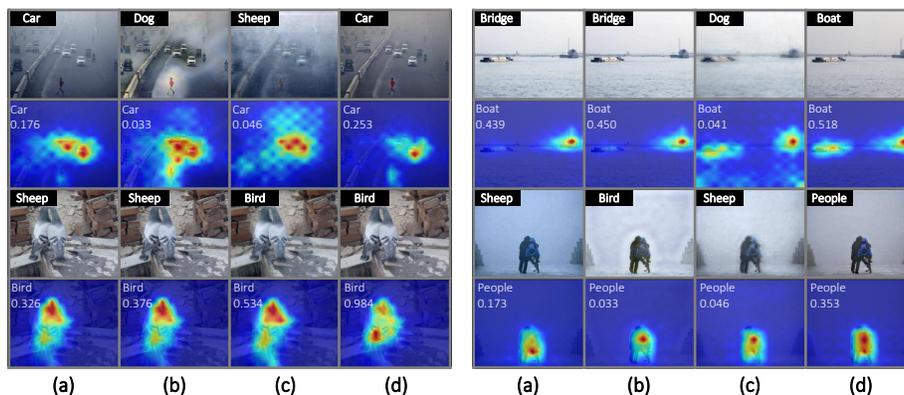
Fig. 6. Qualitative results of DeepLab v3 on the VOC 2012 dataset. (a) Corrupted input. (b) OWAN. (c) URIE-MSE. (d) URIE. (e) Ground-truth. Examples in the top two rows are degraded by seen corruptions and the others are by unseen corruptions.

mentation on the PASCAL VOC dataset to validate its universal applicability to totally different recognition tasks. We employ SSD 300 [20] for object detection and DeepLab v3 [1] for semantic segmentation as the recognition models. Following the convention, the object detection model is evaluated on the VOC 2007 test set while the segmentation model is tested on the VOC 2012 validation set. Tab. 4 and 5 summarize performance of the two recognition models, and Fig. 5 and 6 present their qualitative results when coupled with URIE and the two restoration models. In both of object detection and semantic segmentation, URIE improves recognition performance noticeably, and clearly outperforms the other two models for both seen and unseen corruptions.

Classification on Haze-20. Finally, the effectiveness of URIE is evaluated on the Haze-20 dataset [24], containing haze images of 20 object classes captured in real environments. This dataset is paired with another set, HazeClear-20, a collection of real haze-free images of the same 20 object categories. URIE and the two restoration models are applied to both of datasets and evaluated in terms of performance of a ResNet-50 classifier taking their results as input. This setting is challenging for the three models since they are trained only with synthetically distorted images. Moreover, Pei *et al.* [24] demonstrated by extensive experiments that existing dehazing algorithms do not help improve classification per-

Table 6. Accuracy of the ResNet-50 classifier on the Haze-20 and HazeClear-20 datasets. The numbers in parentheses indicate the differences from the baseline.

OWAN		URIE-MSE		URIE	
Haze	HazeClear	Haze	HazeClear	Haze	HazeClear
58.7 (+1.7)	97.8 (-0.1)	53.4 (-3.6)	94.4 (-3.5)	60.8 (+3.8)	97.6 (-0.3)

**Fig. 7.** Qualitative results on the Haze-20 dataset. (a) Corrupted input. (b) OWAN. (c) URIE-MSE. (d) URIE. Top-1 prediction of the ResNet-50 classifier together with its confidence score and grad-CAM are presented alongside per example.

formance on the Haze-20 dataset. Nevertheless, as summarized in Tab. 6, URIE is still able to improve the classification accuracy in the presence of real-world haze, and at the same time, outperforms the two restoration methods. Fig. 7 shows that URIE enhances areas around objects better than OWAN and URIE-MSE, and consequently allows the coupled classifier to predict the correct class label and draw attention most accurately.

5 Conclusion

This paper has presented a new image enhancement network dedicated to robust visual recognition in the presence of input distortion. Our network is universally applicable to various types of image distortions and many different recognition models. At the same time, it is recognition-friendly since it is optimized to improve robustness of recognition models taking its output as input. These features of our model are given by a novel network architecture and a training strategy we carefully designed. The advantages of our model have been verified throughout experiments in various settings. We however have found that our model marginally degrades performance on distortion-free images. Next on our agenda is to resolve this issue by revising the model architecture.

Acknowledgement: This work was supported by Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-IT1801-05.

References

1. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proc. European Conference on Computer Vision (ECCV) (2018) [13](#)
2. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [4](#)
3. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2009) [3](#), [8](#), [10](#)
4. Diamond, S., Sitzmann, V., Boyd, S.P., Wetzstein, G., Heide, F.: Dirty pixels: Optimizing image classification architectures for raw sensor data. arXiv preprint arXiv:1701.06487 (2017) [1](#), [2](#), [3](#), [4](#)
5. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **38**(2), 295–307 (2016) [5](#)
6. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The Pascal Visual Object Classes (VOC) Challenge. International Journal of Computer Vision (IJCV) (2010) [10](#)
7. Geirhos, R., Temme, C.R.M., Rauber, J., Schütt, H.H., Bethge, M., Wichmann, F.A.: Generalisation in humans and deep neural networks. In: Proc. Neural Information Processing Systems (NeurIPS) (2018) [4](#)
8. Gomez, R., Zhang, Z., González-Jiménez, J., Scaramuzza, D.: Learning-based image enhancement for visual odometry in challenging hdr environments. In: Proc. International Conference on Robotics and Automation (ICRA) (2018) [4](#)
9. Gopalan, R., Taheri, S., Turaga, P., Chellappa, R.: A blur-robust descriptor with applications to face recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **34**(6), 1220–1226 (June 2012) [4](#)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016) [8](#), [11](#)
11. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: Proc. International Conference on Learning Representations (ICLR) (2019) [1](#), [3](#), [4](#), [8](#), [10](#)
12. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: Proc. International Conference on Machine Learning (ICML) (2015) [5](#)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proc. International Conference on Learning Representations (ICLR) (2015) [10](#)
14. Lee, S., Kim, J., Yoon, J.S., Shin, S., Bailo, O., Kim, N., Lee, T.H., Hong, H.S., Han, S.H., Kweon, I.S.: Vpnet: Vanishing point guided network for lane and road marking detection and recognition. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2017) [4](#)
15. Li, B., Peng, X., Wang, Z., Xu, J., Feng, D.: Aod-net: All-in-one dehazing network. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2017) [4](#)
16. Li, S., Araujo, I.B., Ren, W., Wang, Z., Tokuda, E.K., Junior, R.H., Cesar-Junior, R., Zhang, J., Guo, X., Cao, X.: Single image deraining: A comprehensive benchmark analysis. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [2](#)

17. Li, X., Wang, W., Hu, X., Yang, J.: Selective kernel networks. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 9
18. Liu, D., Wen, B., Fan, Y., Loy, C.C., Huang, T.S.: Non-local recurrent network for image restoration. In: Proc. Neural Information Processing Systems (NeurIPS) (2018) 5
19. Liu, D., Wen, B., Liu, X., Wang, Z., Huang, T.S.: When image denoising meets high-level vision tasks: A deep learning approach. In: Proc. International Joint Conference on Artificial Intelligence (IJCAI) (2018) 4
20. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single shot multibox detector. In: Proc. European Conference on Computer Vision (ECCV) (2016) 13
21. Nah, S., Hyun Kim, T., Mu Lee, K.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 4
22. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2015) 3, 8
23. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Learning and transferring mid-level image representations using convolutional neural networks. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014) 3, 8
24. Pei, Y., Huang, Y., Zou, Q., Lu, Y., Wang, S.: Does haze removal help cnn-based image classification? In: Proc. European Conference on Computer Vision (ECCV) (2018) 2, 3, 10, 13
25. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) (2015) 7
26. Sakaridis, C., Dai, D., Gool, L.V.: Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2019) 4
27. Sakaridis, C., Dai, D., Hecker, S., Van Gool, L.: Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In: Proc. European Conference on Computer Vision (ECCV) (2018) 4
28. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2017) 12
29. Sharma, V., Diba, A., Neven, D., Brown, M.S., Van Gool, L., Stiefelhagen, R.: Classification-driven dynamic image enhancement. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) 4
30. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Proc. International Conference on Learning Representations (ICLR) (2015) 3, 8, 11
31. Singh, M., Nagpal, S., Singh, R., Vatsa, M.: Dual directed capsule network for very low resolution image recognition. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2019) 4
32. Suganuma, M., Liu, X., Okatani, T.: Attention-based adaptive selection of operations for image restoration in the presence of unknown combined distortions. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 3, 5, 9, 10
33. Tai, Y., Yang, J., Liu, X., Xu, C.: Memnet: A persistent memory network for image restoration. In: Proc. IEEE International Conference on Computer Vision (ICCV). pp. 4539–4547 (2017) 5

34. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022 (2016) 5
35. Vidal, R.G., Banerjee, S., Grm, K., Struc, V., Scheirer, W.J.: Ug²: a video benchmark for assessing the impact of image restoration and enhancement on automatic visual recognition. In: Proc. IEEE Winter Conference on Applications of Computer Vision (WACV) (2018) 2
36. Wang, Z., Chang, S., Yang, Y., Liu, D., Huang, T.S.: Studying very low resolution recognition using deep networks. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) 4
37. Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-UCSD Birds 200. Tech. Rep. CNS-TR-2010-001, California Institute of Technology (2010) 10
38. Wu, Y., Ling, H., Yu, J., Li, F., Mei, X., Cheng, E.: Blurred target tracking by blur-driven tracker. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2011) 4
39. Wu, Z., Suresh, K., Narayanan, P., Xu, H., Kwon, H., Wang, Z.: Delving into robust object detection from unmanned aerial vehicles: A deep nuisance disentanglement approach. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2019) 4
40. Yasarla, R., Patel, V.M.: Uncertainty guided multi-scale residual learning-using a cycle spinning cnn for single image de-raining. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 4
41. Yu, K., Dong, C., Lin, L., Change Loy, C.: Crafting a toolchain for image restoration by deep reinforcement learning. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) 5
42. Zendel, O., Honauer, K., Murschitz, M., Steininger, D., Fernandez Dominguez, G.: Wilddash - creating hazard-aware benchmarks. In: Proc. European Conference on Computer Vision (ECCV) (2018) 1, 3
43. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. IEEE Transactions on Image Processing (TIP) 26(7), 3142–3155 (July 2017) 4