

DIVIDE-AND-CONQUER ADVERSARIAL LEARNING FOR HIGH-RESOLUTION IMAGE ENHANCEMENT

Anonymous authors

Paper under double-blind review

ABSTRACT

This paper introduces a divide-and-conquer inspired adversarial learning (DACAL) approach for photo enhancement. The key idea is to decompose the photo enhancement process into hierarchically multiple sub-problems, which can be better conquered from bottom to up. On the top level, we propose a perception-based division to learn additive and multiplicative components, required to translate a low-quality image into its high-quality counterpart. On the intermediate level, we use a frequency-based division with generative adversarial network (GAN) to weakly supervise the photo enhancement process. On the lower level, we design a dimension-based division that enables the GAN model to better approximate the distribution distance on multiple independent one-dimensional data to train the GAN model. While considering all three hierarchies, we develop a multiscale training approach to optimize the image enhancement process, suitable for high-resolution images, in a weakly-supervised manner. Both quantitative and qualitative results clearly demonstrate that the proposed DACAL achieves the state-of-the-art performance for high-resolution image enhancement.

1 INTRODUCTION

Despite many mobile camera technological advances we have today, our captured images often still come with limited dynamic range, undesirable color rendition, and unsatisfactory texture sharpness. Among many possible causes, low-light environments and under/over-exposed regions usually introduce severe lack of texture details and low-dynamic range coverage, respectively. Another critical issue is the amplification (during the enhancement process) of noise in the dark and/or texture-less regions, where the enhancement may not even be necessary. Due to these issues, images acquired under different conditions, or different parts of a single image, may require separate enhancement operations. In this context, customarily used context/content agnostic enhancement methods often lead to a poor performance in the overall visual assessment. Therefore, comprehensive methods that improve the perceptual quality of images are in high demand.

In recent years, deep image enhancement methods have demonstrated their superiority for color enrichment, texture sharpening as well as contrast adjustment. A large amount of works like (Dong et al. (2014); Kim et al. (2016); Mao et al. (2016); Shi et al. (2016); Wang et al. (2018b); Timofte et al. (2018); Nah et al. (2017); Zhang et al. (2017b; 2016); Yu et al. (2018); Wu et al. (2018); Yang et al. (2018); Ruixing Wang & Jia (2019)) have made great success in one of such sub-tasks. However, in the real-world setting it is non-trivial to decompose the low-quality factors and treat them respectively from a low-quality images. Another family of approaches (e.g., Gharbi et al. (2017); Ignatov et al. (2017); Liu & Jung (2018); Huang et al. (2018); Chen et al. (2018)) learns the enhancement as some non-linear mapping between low and high-quality image pairs, in a supervised manner. These methods require the desired high-quality images to be well aligned with the low-quality ones. Such image pairs are collected either by manually enhancing the low-quality images (using professional artists) or by using a calibrated rig of low and high-end cameras, with a very small baseline. Unfortunately, the data collection for these setups is not convenient, mainly due to the expensive manual retouching process and the physical limitation for small baseline. Moreover, fully-supervised methods often lack flexibility towards new domain adaptation, thus requiring low-high paired acquisitions for every low-end camera.

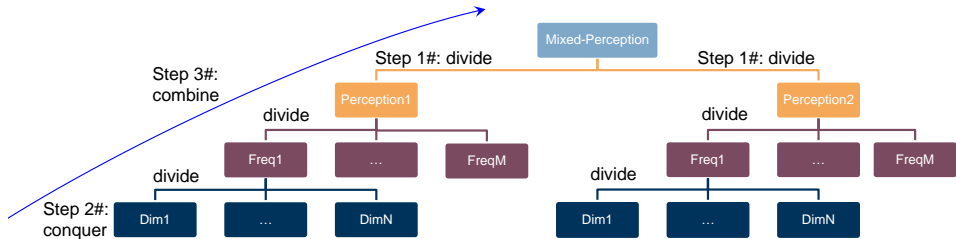


Figure 1: Conceptual illustration of the proposed adversarial learning approach of divide-and-conquer (DACAL). It decomposes the high-quality image distribution learning problem into perception-, frequency- and dimension-based sub-problems in a hierarchical structure. Following the division, the sub-problems are conquered and combined in a deep adversarial training fashion.

To address the limitations of fully-supervised methods, the weakly-supervised deep photo enhancement methods (e.g., Ignatov et al. (2018); Chen et al. (2018)) have emerged recently. These methods merely require a set of target good quality images, whose contents and viewpoints share some similarity with those of the low-quality image domain, without requiring image acquisitions for the same scenes. On the collected unpaired low- and high-quality images, such kind of methods typically apply generative adversarial networks (GANs), which alternately optimizes its two components (i.e., generator and discriminator) using a min-max objective. The discriminator is optimized to approximate the desired-quality image distribution, and the generator is guided by the discriminator to learn the enhancement map. While achieving some success, they still have drawbacks when enhancing mixed visual perceptions or treating high-resolution images. Without the implementation of a pixel-wise loss function in a supervised setup, it appears exceptionally difficult to simultaneously enhance multiple perceptual components (e.g., color, texture and illumination). Besides, due to the computational complexity and unstable adversarial training, the existing image enhancement methods treat high-resolution images using either the downscaled version or patch-wise processing strategies, neither of which are optimal. It is needless to mention that the downscaling loses images details, whereas, patch-wise enhancement does not maintain the spatial consistency over the whole image.

In this paper, inspired by the concept of the divide-and-conquer (DAC) paradigm, we aim at decomposing the original problem of high-quality image distribution learning into sub-problems, which can be better addressed from bottom to up. For this purpose, we propose a DAC inspired adversarial learning approach¹ to better deal with the high-resolution photo enhancement problem. As depicted in Figure 1, our basic learning process innovatively introduces the following three parts to the context of adversarial learning for high-resolution image enhancement.

- **Divide:** We hierarchically decompose the problem into three sub-problems: perception-based (additive and multiplicative maps), frequency-based (low and high frequencies), and dimension-based (multiple one dimensions) data distribution learning problems (from top to bottom). The third sub-problem that learns for each dimension separately is the basics for the original problem. In terms of implementation, these divisions are respectively carried out by: (a) perception-wise branching of the generator; (b) frequency-wise division of the discriminator’s inputs; and (c) dimension-wise division of the discriminator’s output.
- **Conquer:** To conquer the base sub-problems, we exploit an adaptive sliced Wasserstein GAN (AdaSWGAN) model. The basic idea of the AdaSWGAN is to first factorize high-dimensional distributions into their multiple one-dimensional marginal distributions, followed by the approximation of each one-dimensional distributions independently.
- **Combine:** Our model combines the solutions from bottom to top by, (a) aggregating the approximated one-dimensional distributions to obtain the frequency-based solutions; (b) averaging frequency-based solutions to obtain the perception-based solutions; (c) summing two perception-based solutions to learn the high-quality image distribution.

¹There exist some networks like (Romaniuk & Hall (1993); Ghosh et al. (2017); Nowak et al. (2018)) applying the paradigm of DAC to general neural networks. While we found two concurrent works (Kim et al. (2019); Lin et al. (2019)) introduce the DAC concept to adversarial networks, they merely study the decomposition of generator. In contrast, our approach applies the DAC concept to the whole adversarial learning process.

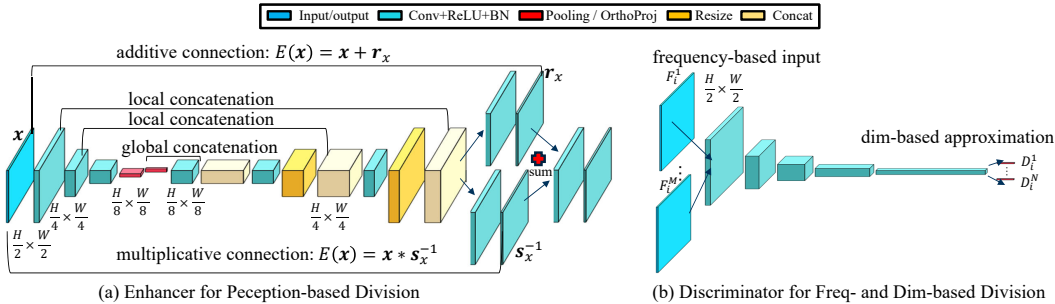


Figure 2: Illustration of the proposed network design for our adversarial learning approach of divide-and-conquer. Specially, we apply the perception-based division to the generator, the frequency-based one to the discriminator input, and the dimension-based one to the discriminator output.

2 PROPOSED METHOD

2.1 PERCEPTION-BASED DIVISION

Given an input low-quality image $\mathbf{x} \in X$, image enhancement seeks a mapping function $E : X \rightarrow Y$, such that $\mathbf{y} = E(\mathbf{x})$, where $\mathbf{y} \in Y$ is the desired-quality image. We suggest to divide the mapping function learning problem into two subproblems, which learn additive and multiplicative perceptions respectively. On one hand, learning additive map is encouraged by the success of residual learning on image colorization and super-resolution *et cetera*. They apply the additive mapping $E(\mathbf{x}) = \mathbf{x} + \mathbf{r}_x$ (where \mathbf{r}_x is the residual component) to learn additive differences from low-quality images to desired-quality ones. On the other hand, learning multiplicative map is inspired by existing works like (Huang et al. (2009); Bioucas-Dias & Figueiredo (2010); Aubert & Aujol (2008); Engel et al. (2016); Ruixing Wang & Jia (2019); Guo et al. (2017); Ghosh et al. (2017); Fu et al. (2016)). Such works apply the multiplicative mapping $E(\mathbf{x}) = \mathbf{x} * \mathbf{s}_x^{-1}$ (where \mathbf{s}_x is the multiplicative component) to address the problem of multiplicative noise removal or illumination improvement successfully. Accordingly, we introduce the perception-wise division to the generator so that the benefits of both additive and multiplicative map learning approaches can be fully exploited. Formally, the final map of the generator is

$$E(\mathbf{x}) = \mathbf{x} + \mathbf{r}_x + \beta(\mathbf{x} * \mathbf{s}_x^{-1}), \quad (1)$$

where β plays the trade-off between such two perception-based components² and the operator ‘*’ means element-wise multiplication.

To approximate the additive component \mathbf{r}_x and the multiplicative component \mathbf{s}_x , we design a two-stream U-Net (Ronneberger et al. (2015)) for the generator (i.e., photo enhancer). As shown in Figure 2, we follow the standard U-Net design to apply the skip connection across the encoder and decoder³. As skip connections concatenate local feature maps, they are good for reliable local feature extraction. In addition, we investigate that global information (e.g., scene category, object class or overall illumination) is useful for individual pixels to guide their local adjustment. Accordingly, we exploit a global concatenation operation to embed high-level information into decoded feature maps. While the encoder can transform input images into low-dimensional features, it does not always achieve a global feature (with a vector form) when the input images are of different sizes. To address this issue, we additionally apply an average pooling operation with a fully convolutional layer to extract the global feature vector. As a result, this design enables our model to enhance full-resolution images while it is trained on down-scaled images. On top of the decoder, we design two different branches to learn the additive and multiplicative components separately. During the early training, we optimize the additive and multiplicative mappings alternately. In the end of training, we aggregate the resulting approximations from these two branches, followed by a group of regular convolutional network layers.

²Our experiments were conducted with $\beta = 1$, which can be chosen differently for known sub-tasks priors.

³The encoder includes groups of convolution, SELU (Klambauer et al. (2017)) and batch normalization (Ioffe & Szegedy (2015)), while the decoder consists of convolution, resizing operation based layers.

The idea of our designed two-stream U-Net is related to Chen et al. (2018); Ruixing Wang & Jia (2019). For the additive component learning, Chen et al. (2018) introduces a global U-Net. Compared to our model, the global U-Net from Chen et al. (2018) can only accept images of a fixed low-resolution due to the hard design (a fixed filter for the fully-connected layer) for the global feature extraction. For example, when feeding $8 \times 8 \times 128$ feature map into the global feature extraction, it has to use 8×8 filters to obtain the global feature $1 \times 1 \times 128$. For different image sizes, the resulting embedded feature size will be different, say $4 \times 8 \times 128$. In such cases, the global feature extraction will fail. For the multiplicative component learning, Ruixing Wang & Jia (2019) designs a bilateral-filtering based network, while our enhancer adopts regular convolutions and resizing operations.

2.2 FREQUENCY-BASED DIVISION

For the supervised photo enhancement problem that requires corresponding ground-truth high-quality images for guidance, the common reconstruction loss can be used to train the proposed deep enhancer. In contrast, the weakly-supervised photo enhancement task merely has a set of good quality images for reference, and thus for this task we need to learn the map from the low-quality image domain to high-quality image domain. One of the most promising approaches is cyclic GAN models (Zhu et al. (2017)). Hence, we choose to adopt the cyclic GAN objective to train our developed enhancer:

$$\min_{E, \hat{E}} \max_{C, \hat{C}} \mathcal{L}_{\text{GAN}}(E, C) + \mathcal{L}_{\text{GAN}}(\hat{E}, \hat{C}) + \gamma_1 \mathcal{L}_{\text{cyc}}(E, \hat{E}) + \gamma_2 \mathcal{L}_{\text{id}}(E, \hat{E}), \quad (2)$$

where $E : X \rightarrow Y$ and $\hat{E} : Y \rightarrow X$ are used for forward and backward image translation, respectively. They share the same design using the proposed two-stream U-Net. Here, C and \hat{C} are the corresponding discriminators (or critics) which are used to guide the training of E and \hat{E} , respectively. The GAN loss $\mathcal{L}_{\text{GAN}}(E, C)$ and $\mathcal{L}_{\text{GAN}}(\hat{E}, \hat{C})$ are first suggested to use the adaptive Wasserstein GAN loss (Chen et al. (2018)) as,

$$\begin{aligned} \mathcal{L}_{\text{GAN}}^C(E, C) &= \mathbb{E}_{\mathbf{y} \sim P_y} [C(\mathbf{y})] - \mathbb{E}_{\hat{\mathbf{y}} \sim P_E} [C(E(\hat{\mathbf{y}}))] + \lambda \mathbb{E}_{\hat{\mathbf{y}} \sim P_{\hat{y}}} [\max(0, \|\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})\|_2 - 1)], \\ \mathcal{L}_{\text{GAN}}^E(E, C) &= \mathbb{E}_{\mathbf{x} \sim P_x} [C(E(\mathbf{x}))], \end{aligned} \quad (3)$$

where $\hat{\mathbf{y}}$ are random samples following the distribution $P_{\hat{y}}$, which is uniformly sampled along straight lines between pairs of points sampled from P_y and P_E , and $\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})$ is the gradient w.r.t $\hat{\mathbf{y}}$.

Standard cyclic GAN models suggest to use the cyclic consistency loss and identity mapping loss. We follow (Chen et al. (2018)) to define these two losses based on mean square error (MSE):

$$\begin{aligned} \mathcal{L}_{\text{cyc}}(E, \hat{E}) &= \mathbb{E}_{\mathbf{x} \sim P_x} [\|\hat{E}(E(\mathbf{x})) - \mathbf{x}\|_2] + \mathbb{E}_{\mathbf{y} \sim P_y} [\|E(\hat{E}(\mathbf{y})) - \mathbf{y}\|_2], \\ \mathcal{L}_{\text{id}}(E, \hat{E}) &= \mathbb{E}_{\mathbf{x} \sim P_x} [\|E(\mathbf{x}) - \mathbf{x}\|_2] + \mathbb{E}_{\mathbf{y} \sim P_y} [\|\hat{E}(\mathbf{y}) - \mathbf{y}\|_2]. \end{aligned} \quad (4)$$

To further decompose the perception-based distribution learning, we follow (Ignatov et al. (2017; 2018)) to apply frequency-based division to the discriminator C . Specifically, we adopt the separation of enhancing on low frequencies (e.g., main structures, major contents and colors), and optimizing on high frequencies (e.g., textures and small details). For this purpose, we replicate the discriminators to specialize on various frequencies of input images. Without loss of generality, we study two frequencies for discriminators: For low frequency data, we utilize the RGB image blurred by a Gaussian kernels, whereas, the grayscale component of the discriminator input images is used as the high frequency data. Both are equal in architecture and parameter values, and their GAN losses are averaged in the end. In this case, individual discriminators are encouraged to focus on an easier distribution approximation task, and thus become more reliable to supervise the training of the proposed photo enhancer.

2.3 DIMENSION-BASED DIVISION AND SOLUTION

While the frequency-based division for the discriminator’s input can reduce the complexity of training Wasserstein GAN to some extent, it still suffers from the unstable training issue when the raw images lie in a high dimensional space, where we can only select samples very sparsely during training. To address this issue, we suggest to further decompose the challenging estimation of a high-dimensional distribution into simpler estimation of multiple one-dimensional distributions, which is

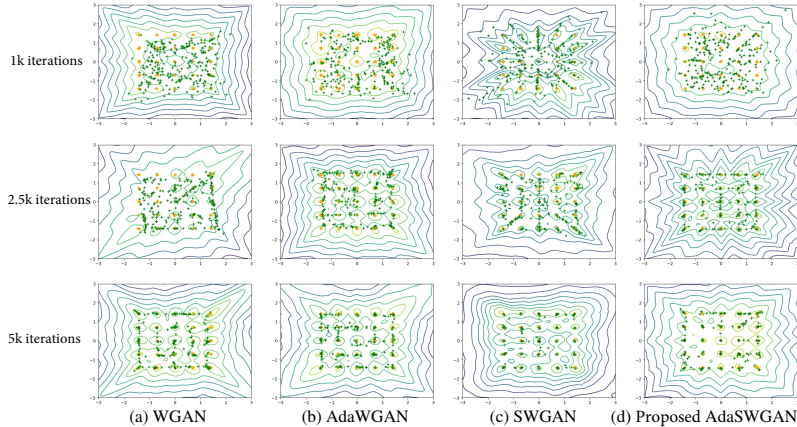


Figure 3: Results of WGAN (Gulrajani et al. (2017)), AdaWGAN (Chen et al. (2018)), SWGAN (Wu et al. (2019)) and the proposed AdaSWGAN (three rows indicate training for 1000, 2500 and 5000 iterations resp.) for the 25 Gaussians dataset. The orange points represent the target distribution, green points are the generated samples, and the curves indicate the value surfaces of the discriminators of compared models.

the underlying idea of sliced Wasserstein distance computation. To achieve this goal, we exploit an adaptive version of sliced Wasserstein GAN (AdaSWGAN) loss:

$$\min_E \max_C \int_{\theta \in \mathbb{S}^{n-1}} \left(\mathbb{E}_{\mathbf{y} \sim P_y} [C(\mathbf{y})] - \mathbb{E}_{\tilde{\mathbf{y}} \sim P_E} [C(E(\mathbf{x}))] \right) + \lambda \mathbb{E}_{\hat{\mathbf{y}} \sim P_{\hat{\mathbf{y}}}} [\max(0, \|\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})\|_2 - 1)], \quad (5)$$

where θ is the orthogonal projection matrix for the decomposition from high-dimensional data to independent one-dimensional ones under C , $\hat{\mathbf{y}}$ denotes random samples following the distribution $P_{\hat{\mathbf{X}}}$ which is sampled uniformly along straight lines between pairs of points sampled from P_X and P_G , and $\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})$ is the gradient with respect to $\hat{\mathbf{y}}$. As the parameter λ weights the gradient penalty, it can highly influence the trade-off between the original sliced Wasserstein distance and the Lipschitz constraint. To reach a good balance, we adopt an adaptive learning manner to update the hyperparameter by observing the moving average of gradients:

$$\overline{\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})} = \eta \overline{\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})} + (1 - \eta) \frac{\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})}{\lambda}, \quad (6)$$

where λ is the weight for gradient penalty, and η is the decay constant. When $\overline{\nabla_{\hat{\mathbf{y}}} C(\hat{\mathbf{y}})}$ is larger than a fixed upper bound τ , the relative influence of λ becomes too small and therefore the penalty is not strong enough to satisfy the Lipschitz constraint. In this case, the weight λ should be increased with a certain scale (in our experiments, we always set 2 times bigger). Otherwise, λ is decreased to half of the current value.

During training, the parameterized projection matrices θ should remain orthogonal. For this requirement we first initialize the parameters with random orthogonal matrices through QR decomposition, then update them on a Stiefel manifold during training. Particularly, we follow (Wu et al. (2019)) to optimize the orthogonal matrices θ on Stiefel manifolds.

Few variations of sliced Wasserstein GANs (Deshpande et al. (2018); Ishan Deshpande (2019); Wu et al. (2019)) are proposed in the recent years. Differently from (Deshpande et al. (2018); Ishan Deshpande (2019)), we use the parameterized projection matrix to optimize the sliced Wasserstein distance approximation. When compared to (Wu et al. (2019)), we adopt an adaptive weight scheme to adjust the penalty weight. To show the advantage of the proposed AdaSWGAN, we follow the sampling strategy of (Gulrajani et al. (2017)) to generate 25 Gaussian distributed data. On the toy dataset, we compare it against the four existing Wasserstein GAN models⁴. The results in Figure 3 show that the proposed AdaSWGAN has a better convergence to fit the toy data than the other compared methods.

⁴For fair comparison, same generator and discriminator architectures (of Gulrajani et al. (2017)) were used for all compared methods

3 MULTISCALE NETWORK AND TRAINING

To improve the enhancement on high-resolution images, we introduce a multiscale enhancer that is trained on images with various resolutions. This design allows for a wider range of learned receptive fields, from coarse to fine. The multiscale enhancer consists of various enhancers at multiple scales. Let’s take a two-scale case for example. The two enhancers are denoted as E_l, E_h , which enhance images at resolution of 512×256 and 1024×512 , respectively. The multiscaled design of enhancers is expected to effectively aggregate the visual features over various scales.

The high-scale enhancer E_h has the same structure as the low-scale enhancer E_l , which is shown in Figure 2. To transmit the low-scale information smoothly from the low-scale to the high-scale components, the high-scale enhancer also inherits the learned parameters from the low-scale enhancer. In addition, it repeats some more layers on the bottom and the top of the low-scale enhancer. Specially, it adds one more down-sampling component including convolution, SELU, and batch normalization at the bottom. One more up-sampling component is applied to the top. It contains an up-scale re-sizing and a convolution operation. Following the lower-scale enhancer design, we apply the same two-stream design to learn the additive and multiplicative components for the high-resolution image enhancement as well.

It is known that synthesizing high-resolution images yields significant challenges to the GAN training. For example, differentiating high-resolution real and generated images requires the discriminator to have a deeper architecture or convolutions with a larger receptive field. But such two solutions will often lead to overfitting and out of memory issues. Besides, due to the memory limit, we have to use smaller minibatches that will further harm training stability. Therefore, we also follow the design of multiscale enhancer to exploit multiscale discriminators for more reliable training guidance on different resolutions. Similarly, we insert some more layers into the bottom of the lower-scale discriminator to construct higher-scale discriminators for more reliable training guidance on higher-resolutions.

For more details on the design of the multiscale enhancer and discriminator, please refer to the appendix. We train the proposed multiscale enhancer with three stages. At the first stage, we train the lower-enhancer E_l using either reconstruction loss in the supervised case or using the proposed frequency- and dimension-based GAN loss for the weakly-supervised case. At the second stage, we tune the parameters of the newly added layers in the higher-scale enhancer E_h while fixing the parameters of the lower-scale enhancer E_l . Finally, we train the multiscale enhancer as a whole.

To overcome the unstable GAN training issue on high-resolution images, there exist various cascaded or multiscaled network designs like (Denton et al. (2015); Dosovitskiy & Brox (2016); Gatys et al. (2016); Johnson et al. (2016); Zhang et al. (2017a); Huang et al. (2017); Karras et al. (2017); Wang et al. (2018a)) that synthesize high-resolution images. However, such works rarely study the multiscaled network based on the U-Net design, which has shown its advantages over other networks for image enhancement (Chen et al. (2018); Huang et al. (2018)). In contrast, our multiscale enhancer is built upon a two-stream U-Net based model, which is tailored for better high-resolution image enhancement.

4 EXPERIMENT

4.1 SUPERVISED PHOTO ENHANCEMENT

This task aims at learning the enhancement mapping between low- and high-quality image pairs in a supervised fashion. We use the commonly-used MIT-Adobe FiveK (Bychkovsky et al. (2011)) dataset for this task. The dataset consists of 5,000 high-resolution (higher than 1K) images, each of which is retouched by five experts. For the evaluation, we use 2,250 images and their retouched versions from photographer C for training, and the remaining 498 images are used for testing. Distinct from most of existing works, we go for high-resolution image enhancement. To study this task, we compare our proposed adversarial learning approach of divide-and-conquer (DACAL) against state-of-the-art methods including White-Box (WB) (Hu et al. (2018)), Distort-and-Recover (DR) (Park et al. (2018)), DSLR-Photo Enhancement (DPED) (Ignatov et al. (2017)), and the supervised version of the Deep Photo Enhancement (DPE) method (Chen et al. (2018)). Except DPED, the rest three methods are originally designed for down-scaled image enhancement. In contrast, our model

Table 1: PSNR and SSIM results for the MIT-Adobe FiveK (Bychkovsky et al. (2011)) test images. Here, WB and DR indicate the White-Box and Distort-and-Recover methods, respectively. $DACAL_{l_1}$, $DACAL_{l_2}$, $DACAL_{l_3}$ and $DACAL_l$ represent the use of individual additive, individual multiplicative, multiplicative cascaded by additive, and our suggested parallel fusion (two-stream strategy), respectively. $DACAL_h$ is our higher-scale version. $PSNR_d/SSIM_d$ and $PSNR_f/SSIM_f$ indicate the results on downsampled images and full-resolution images, respectively.

	WB	DR	DPED	DPE	$DACAL_{l_1}$	$DACAL_{l_2}$	$DACAL_{l_3}$	$DACAL_l$	$DACAL_h$
$PSNR_d$	18.86	21.64	21.05	22.10	22.73	22.99	23.01	23.52	24.15
$PSNR_f$	19.09	21.52	20.86	21.65	22.43	22.69	23.02	23.56	24.07
$SSIM_d$	0.928	0.936	0.922	0.947	0.958	0.942	0.949	0.959	0.962
$SSIM_f$	0.920	0.922	0.916	0.894	0.948	0.942	0.940	0.954	0.956

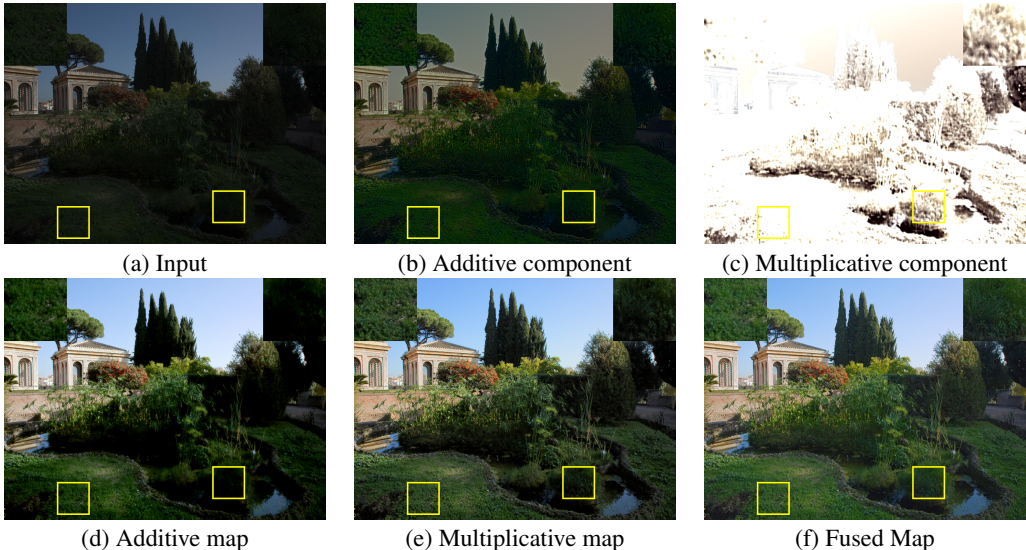


Figure 4: Visual results of additive, multiplicative components and their combination. It is best to zoom in on the high-resolution pictures.

can be trained on multiscaled images. For details on our training details⁵, please refer to the appendix. As the dataset has paired original/retouched images, we use both the Peak Signal-to-Noise Ratio (PSNR) and the multi-scale structural similarity index measure (SSIM)⁶ (Wang et al. (2003)) to evaluate the comparing methods.

Table 1 summarizes the PSNR and SSIM results on both downsampled images of size 512×256 and full-resolution ones. Compared to the state-of-the-art methods, our proposed DACAL achieves the highest PSNR and SSIM. Note that our DACAL achieves the highest PSNR ⁷ (**24.15 dB**) on the MIT-Adobe FiveK dataset. In particular, we discover the PSNR and SSIM values of some competing methods like DPE decline by a certain margin when enhancing full-resolution images. In contrast, our proposed DACAL has relatively stable PSNR and SSIM values in the high-res case. For a quantitative perceptual quality comparison, we employ 20 Amazon Mechanical Turk (MTurk) workers to compare 100 image pairs, which are from our DACAL and the competing methods respectively. As shown in Table 3, the highest preference score shows the superiority of our DACAL.

In addition, we conduct the ablation study on the proposed two-stream U-Net. Specially, we evaluate the separative use of additive and multiplicative streams denoted by $DACAL_{l_1}$ and $DACAL_{l_2}$, respectively. As shown in Table 1, they are outperformed by our two-stream version $DACAL_l$. We also evaluate the performance by cascading these two streams: additive cascaded by multiplicative

⁵The code will be released once the paper is published.

⁶For short, we denote the mutli-scale SSIM as SSIM throughout the paper.

⁷Recently the CVPR19 paper (Ruixing Wang & Jia (2019)) has corrected their model’s PSNR from 30.80 dB to 23.04 dB, on its GitHub webpage.

Table 2: PSNR and SSIM results for the DPED (Ignatov et al. (2017)) test 100×100 image patches. Here, l, f, d for DACAL represent the use of our proposed perception-based, frequency-based and dimension-based division respectively. $DACAL_h$ is our higher-scale version.

	WESPE	DPE	$DACAL_l$	$DACAL_{l+f}$	$DACAL_{l+f+d}$	$DACAL_h$
PSNR ₁₀₀	17.45	18.53	19.62	20.01	20.43	20.90
SSIM ₁₀₀	0.854	0.861	0.868	0.869	0.872	0.874

Table 3: Preference ratio (one vote is lost for DACAL vs. Input) from MTurk user study

DACAL vs.	Input	White-Box	Distort-and-Recover	DPED	WESPE	DPE
Adobe	395:4	389:11	303:97	365:35	N/A	370:30
DPED	384:16	N/A	N/A	N/A	384:16	344:56
Huawei-Flickr	366:34	N/A	N/A	N/A	341:59	360:40

stream, or multiplicative cascaded by additive stream. As we find the former case works better than the later one in terms of PSNR and SSIM. Here, we only report the results of the former one $DACAL_{l_3}$. The improvement of our suggested two-stream version further justifies its effectiveness. In Figure 4, we also study the visual results of the two components and their resulting maps. The results reflect that the additive one targets for better color and texture, and the multiplicative one targets on enhancing illumination. The result of final map shows the benefit of fusing additive and multiplicative streams.

4.2 WEAKLY-SUPERVISED PHOTO ENHANCEMENT

This task performs photo enhancement using one set of good-quality images instead of paired before/after images. For this task, we use the DPED dataset (Ignatov et al. (2017)) and our collected Huawei-Flickr dataset. The DPED dataset is composed of images from smartphone cameras (i.e., iPhone 3GS, BlackBerry Passport and Sony Xperia Z) paired with images of the same scenes captured by a DSLR camera (i.e., Canon 70D). We use the iPhone-DSLR data (5,614 iPhone images of size 2048×1536 , 5,902 DSLR images of size 3648×2432) to evaluate the competing methods on the full-resolution images in a weakly-supervised manner. As the dataset also contains 4,353 well-aligned 100×100 iPhone-DSLR image patch pairs, we use its testing image patches to compute the PSNR and SSIM values for quantitative evaluation. As for our Huawei-Flickr dataset, we collect 5,138 images of size 1920×1080 by Huawei Mate20 pro and use the downloaded 604 HDR Flickr images with around 1K resolution from (Chen et al. (2018)) for the target.

Table 2 reports the PSNR and SSIM results of the evaluated methods on the DPED dataset. When compared with the state-of-art weakly-supervised methods WESPE (Ignatov et al. (2018)) and DPE (Chen et al. (2018)), our DACAL reaches the highest PSNR (**20.90 dB**) and SSIM. The table also shows the improvement of the introduced frequency based GAN and sliced SWGAN loss. For a quantitative perceptual quality comparison on our DACAL, WESPE and DPE, we ask 20 MTurk works for user study on 120 image pairs for DPED and Huawei-Flickr. The results in Table 3 show that our method achieves much higher preference score than all competing methods in terms of the sampled human perception. The qualitative results from Figure 5 reflect that our DACAL has the best performance in terms of color rendering, texture sharpening, and illumination adjustment for the DPED iPhone images. In Figure 6, we present enhancement results for our Huawei-Flickr dataset. We can find that WESPE tends to make the image over-exposed, while DPE has serious texture issues due to its design for low-resolution image enhancement. By comparison, our DACAL performs most promisingly for color, texture, and light condition improvements.

5 CONCLUSION

This paper addresses mix-perception enhancement of images with a divide-and-conquer inspired adversarial learning approach, which divides and merges perception-based, frequency-based and dimension-based sub-problems. For high-resolution image enhancement, we suggest to progressively grow the proposed multiscaled model along various resolutions. Evaluations on both su-



Figure 5: High-resolution image enhancement results for the DPED iPhone-DSLR data. It is best to zoom in on the high-resolution pictures.

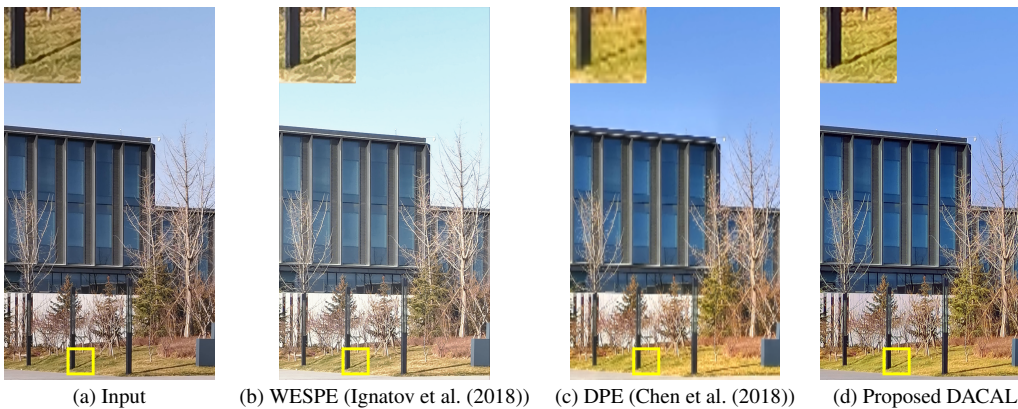


Figure 6: High-resolution image enhancement results for our Huawei-Flickr data. It is best to zoom in on the high-resolution pictures.

pervised and weakly-supervised image enhancement tasks demonstrate the clear superiority of our proposed enhancer against the competing methods.

REFERENCES

Gilles Aubert and Jean-Francois Aujol. A variational approach to removing multiplicative noise. *SIAM journal on applied mathematics*, 68(4):925–946, 2008.

- José M Bioucas-Dias and Mário AT Figueiredo. Multiplicative noise removal using variable splitting and constrained optimization. *IEEE Transactions on Image Processing*, 19(7):1720–1730, 2010.
- Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input/output image pairs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 97–104. IEEE, 2011.
- Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6306–6314, 2018.
- Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pp. 1486–1494, 2015.
- Ishan Deshpande, Ziyu Zhang, and Alexander Schwing. Generative modeling using the sliced Wasserstein distance. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pp. 184–199. Springer, 2014.
- Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems*, pp. 658–666, 2016.
- Jakob Engel, Vladyslav Usenko, and Daniel Cremers. A photometrically calibrated benchmark for monocular visual odometry. *arXiv preprint arXiv:1607.02555*, 2016.
- Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2782–2790, 2016.
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2414–2423, 2016.
- Michaël Gharbi, Jiawen Chen, Jonathan T Barron, Samuel W Hasinoff, and Frédo Durand. Deep bilateral learning for real-time image enhancement. *ACM Transactions on Graphics (TOG)*, 36(4):118, 2017.
- Dibya Ghosh, Avi Singh, Aravind Rajeswaran, Vikash Kumar, and Sergey Levine. Divide-and-conquer reinforcement learning. *arXiv preprint arXiv:1711.09874*, 2017.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pp. 5767–5777, 2017.
- Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Trans. Image Processing*, 26(2):982–993, 2017.
- Yuanming Hu, Hao He, Chenxi Xu, Baoyuan Wang, and Stephen Lin. Exposure: A white-box photo post-processing framework. *ACM Transactions on Graphics (TOG)*, 37(2):26, 2018.
- Jie Huang, Pengfei Zhu, Mingrui Geng, Jiewen Ran, Xingguang Zhou, Chen Xing, Pengfei Wan, and Xiangyang Ji. Range scaling global u-net for perceptual image enhancement on mobile devices. In *European Conference on Computer Vision*, pp. 230–242. Springer, 2018.
- Xun Huang, Yixuan Li, Omid Poursaeed, John Hopcroft, and Serge Belongie. Stacked generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5077–5086, 2017.
- Yu-Mei Huang, Michael K Ng, and You-Wei Wen. A new total variation method for multiplicative noise removal. *SIAM Journal on imaging sciences*, 2(1):20–40, 2009.

- Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3277–3285, 2017.
- Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Wespe: weakly supervised photo enhancer for digital cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 691–700, 2018.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Ruoyu Sun Ayis Pyrros Nasir Siddiqui Sanmi Koyejo Zhizhen Zhao David Forsyth Alexander Schwing Ishan Deshpande, Yuan-Ting Hu. Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2019.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pp. 694–711. Springer, 2016.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1646–1654, 2016.
- Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Jsi-gan: Gan-based joint super-resolution and inverse tone-mapping with pixel-wise task-specific filters for uhd hdr video. *arXiv preprint arXiv:1909.04391*, 2019.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *Advances in neural information processing systems*, pp. 971–980, 2017.
- Chieh Hubert Lin, Chia-Che Chang, Yu-Sheng Chen, Da-Cheng Juan, Wei Wei, and Hwann-Tzong Chen. Coco-gan: Generation by parts via conditional coordinating. *arXiv preprint arXiv:1904.00284*, 2019.
- Jie Liu and Cheolkon Jung. Multiple connected residual network for image enhancement on smart-phones. In *European Conference on Computer Vision*, pp. 182–196. Springer, 2018.
- Xiaojiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Advances in neural information processing systems*, pp. 2802–2810, 2016.
- Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3883–3891, 2017.
- Alex Nowak, David Folqué, and Joan Bruna. Divide and conquer networks. In *ICLR*, 2018.
- Jongchan Park, Joon-Young Lee, Donggeun Yoo, and In So Kweon. Distort-and-recover: Color enhancement using deep reinforcement learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5928–5936, 2018.
- Steve G Romaniuk and Lawrence O Hall. Divide and conquer neural networks. *Neural Networks*, 6(8):1105–1116, 1993.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015.
- Chi-Wing Fu Xiaoyong Shen Wei-Shi Zheng Ruixing Wang, Qing Zhang and Jiaya Jia. under-exposed photo enhancement using deep illumination estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2019.

- Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1874–1883, 2016.
- Radu Timofte, Shuhang Gu, Jiqing Wu, and Luc Van Gool. Ntire 2018 challenge on single image super-resolution: methods and results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 852–863, 2018.
- Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8798–8807, 2018a.
- Yifan Wang, Federico Perazzi, Brian McWilliams, Alexander Sorkine-Hornung, Olga Sorkine-Hornung, and Christopher Schroers. A fully progressive approach to single-image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 864–873, 2018b.
- Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pp. 1398–1402. Ieee, 2003.
- Jiqing Wu, Zhiwu Huang, Dinesh Acharya, Wen Li, Janine Thoma, Danda Pani Paudel, and Luc Van Gool. Sliced wasserstein generative models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2019.
- Shangzhe Wu, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Deep high dynamic range imaging with large foreground motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 117–132, 2018.
- Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, and Rynson WH Lau. Image correction via deep reciprocating hdr transformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1798–1807, 2018.
- Runsheng Yu, Wenyu Liu, Yasen Zhang, Zhi Qu, Deli Zhao, and Bo Zhang. Deepexposure: Learning to expose photos with asynchronously reinforced adversarial learning. In *Advances in Neural Information Processing Systems*, pp. 2149–2159, 2018.
- Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5907–5915, 2017a.
- Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017b.
- Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pp. 649–666. Springer, 2016.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2223–2232, 2017.

A APPENDIX

A.1 NETWORK ARCHITECTURE

In the paper, we use two-scale design for both enhancer and discriminator. Figure 7 illustrates the network architecture of our proposed method. Specially, Figure 7 (a) shows the network design of our multiscaled enhancer, consisting of 5 components: netG₀, netG₁, netG₂, netG₃ and netG₄.

```

===== net_name = _netG_0 =====
[ input][ 0] : ( 3, 512, 1024, 3)
[ conv][ 1] : ( 3, 512, 1024, 16)
[ selu][ 2] : ( 3, 512, 1024, 16)
[ bn][ 3] : ( 3, 512, 1024, 16)
[ conv][ 4] : ( 3, 256, 512, 3)
[ selu][ 5] : ( 3, 256, 512, 3)
[ bn][ 6] : ( 3, 256, 512, 3)
===== net_name = _netG_1 =====
[ input][ 6] : ( 3, 256, 512, 3)
[ conv][ 7] : ( 3, 256, 512, 16)
[ selu][ 8] : ( 3, 256, 512, 16)
[ bn][ 9] : ( 3, 256, 512, 16)
[ conv][ 10] : ( 3, 128, 256, 32)
[ selu][ 11] : ( 3, 128, 256, 32)
[ bn][ 12] : ( 3, 128, 256, 32)
[ conv][ 13] : ( 3, 64, 128, 64)
[ selu][ 14] : ( 3, 64, 128, 64)
[ bn][ 15] : ( 3, 64, 128, 64)
[ conv][ 16] : ( 3, 32, 64, 128)
[ selu][ 17] : ( 3, 32, 64, 128)
[ bn][ 18] : ( 3, 32, 64, 128)
===== net_name = _netG_2 =====
[ input][ 18] : ( 3, 32, 64, 128)
[ conv][ 19] : ( 3, 16, 32, 128)
[ selu][ 20] : ( 3, 16, 32, 128)
[ bn][ 21] : ( 3, 16, 32, 128)
[ conv][ 22] : ( 3, 8, 16, 128)
[ selu][ 23] : ( 3, 8, 16, 128)
[ bn][ 24] : ( 3, 8, 16, 128)
[ conv][ 25] : ( 3, 1, 9, 128)
[ selu][ 26] : ( 3, 1, 9, 128)
[ conv][ 27] : ( 3, 1, 9, 128)
===== net_name = _netG_3 =====
[ input][ 18] : ( 3, 32, 64, 128)
[ conv][ 28] : ( 3, 32, 64, 128)
[ g_concat][ 29] : ( 3, 32, 64, 256)

[ conv][ 30] : ( 3, 32, 64, 128)
[ selu][ 31] : ( 3, 32, 64, 128)
[ bn][ 32] : ( 3, 32, 64, 128)
[ conv][ 33] : ( 3, 32, 64, 128)
[ resize][ 34] : ( 3, 64, 128, 128)
[ concat][ 35] : ( 3, 64, 128, 192)

[ selu][ 36] : ( 3, 64, 128, 192)
[ bn][ 37] : ( 3, 64, 128, 192)
[ conv][ 38] : ( 3, 64, 128, 64)
[ resize][ 39] : ( 3, 128, 256, 64)
[ concat][ 40] : ( 3, 128, 256, 96)

[ selu][ 41] : ( 3, 128, 256, 96)
[ bn][ 42] : ( 3, 128, 256, 96)
[ conv][ 43] : ( 3, 128, 256, 32)
[ resize][ 44] : ( 3, 256, 512, 32)
[ concat][ 45] : ( 3, 256, 512, 48)

[ selu][ 46] : ( 3, 256, 512, 48)
[ bn][ 47] : ( 3, 256, 512, 48)
[ conv][ 48] : ( 3, 256, 512, 16)
[ selu][ 49] : ( 3, 256, 512, 16)
[ bn][ 50] : ( 3, 256, 512, 16)
[ conv_color_illu][ 51] : ( 3, 256, 512, 3)

===== net_name = _netG_4 =====
[ input][ 51] : ( 3, 256, 512, 3)
[ conv][ 52] : ( 3, 256, 512, 3)
[ resize][ 53] : ( 3, 512, 1024, 3)
[ concat][ 54] : ( 3, 512, 1024, 19)

[ selu][ 55] : ( 3, 512, 1024, 19)
[ bn][ 56] : ( 3, 512, 1024, 19)
[ conv][ 57] : ( 3, 512, 1024, 16)
[ selu][ 58] : ( 3, 512, 1024, 16)
[ bn][ 59] : ( 3, 512, 1024, 16)
[ conv_color_illu][ 60] : ( 3, 512, 1024, 3)

```

(a) Multi-scaled Enhancer

```

===== net_name = _netD_0 =====
[ input][ 0] : ( 3, 512, 1024, 3)
[ conv][ 1] : ( 3, 256, 512, 3)
[ lrelu][ 2] : ( 3, 256, 512, 3)
[ in][ 3] : ( 3, 256, 512, 3)
===== net_name = _netD_1 =====
[ input][ 3] : ( 3, 256, 512, 3)
[ conv][ 4] : ( 3, 256, 512, 16)
[ lrelu][ 5] : ( 3, 256, 512, 16)
[ in][ 6] : ( 3, 256, 512, 16)
[ conv][ 7] : ( 3, 128, 256, 32)
[ lrelu][ 8] : ( 3, 128, 256, 32)
[ in][ 9] : ( 3, 128, 256, 32)
[ conv][ 10] : ( 3, 64, 128, 64)
[ lrelu][ 11] : ( 3, 64, 128, 64)
[ in][ 12] : ( 3, 64, 128, 64)
[ conv][ 13] : ( 3, 32, 64, 128)
[ lrelu][ 14] : ( 3, 32, 64, 128)
[ in][ 15] : ( 3, 32, 64, 128)
[ conv][ 16] : ( 3, 16, 32, 128)
[ lrelu][ 17] : ( 3, 16, 32, 128)
[ in][ 18] : ( 3, 16, 32, 128)
[ conv][ 19] : ( 3, 16, 32, 128)
[ lrelu][ 20] : ( 3, 16, 32, 128)
[ in][ 21] : ( 3, 16, 32, 128)
[ conv][ 22] : ( 3, 1, 17, 1)

===== net_name = _netS_1 =====
[ input][ 0] : ( 3, 1, 17, 1)
[ d_swd][ 1] : ( 3, 17)
[ d_swd][ 2] : ( 3, 17)
[ d_swd][ 3] : ( 3, 17)
[ concat_swd][ 4] : ( 3, 34)

[ concat_swd][ 5] : ( 3, 51)

```

(b) Multi-scaled Discriminator

Figure 7: Illustration of the used multi-scaled enhancer (a) and multi-scaled discriminator (b) for our proposed method

Among them, netG_1 , netG_2 , netG_3 form the lower-scale enhancer. Therefore, the higher-scale enhancer inherits netG_1 , netG_2 , netG_3 from lower-scale enhancer, and introduces two new components that are netG_0 and netG_4 . Note that the layer named ‘conv-color-illu’ realizes the two-stream branches, which contains separate groups of convolutions operations, and finally aggregates them using the Eqn.1 in the major paper. Figure 7 (b) shows the structure of our multiscaled discriminator that contains netD_0 , netD_1 and netS_1 . The component netS_1 is designed for the computation of sliced Wasserstein distance. Analogously, netD_1 and netS_1 are from lower-scale discriminator. The higher-scale discriminators add one more component netD_0 .

A.2 TRAINING DETAILS

For the supervised image enhancement task, we train the proposed two-scaled enhancer on one single NVIDIA TITAN Xp GPU with 12GB GPU memory. We first train the lower-scaled enhancer over down-scaled (512×256) images for 20 epochs, and then train higher-scaled enhancer through tuning the new parameters of netG_0 and netG_4 while fixing the parameters of netG_1 , netG_2 , netG_3 over higher-resolution (1024×512) images for 20 epochs. Finally, we train the whole multi-scale enhancer for 20 epochs. For the weakly supervised image enhancement task, we train the two-scaled enhancer on two NVIDIA TITAN Xp GPUs, each of which has 12GB GPU memory. Similar to the supervised case, we use the same style to train the multi-scaled enhancer. For the training, we empirically set the hyperparameters $\lambda = 10$, $\eta = 0.99$, $\tau = 0.05$, $\gamma_1 = 10000$, $\gamma_2 = 1000$ from the proposed GAN model. Note that it is still feasible for us to train our enhancer on higher-resolution images if we have more computational resources.