

JSI-GAN: GAN-Based Joint Super-Resolution and Inverse Tone-Mapping with Pixel-Wise Task-Specific Filters for UHD HDR Video

Soo Ye Kim,* Jihyong Oh,* Munchurl Kim

Korea Advanced Institute of Science and Technology
Republic of Korea

{sooyekim, jhoh94, mkimee}@kaist.ac.kr

Abstract

Joint learning of super-resolution (SR) and inverse tone-mapping (ITM) has been explored recently, to convert legacy low resolution (LR) standard dynamic range (SDR) videos to high resolution (HR) high dynamic range (HDR) videos for the growing need of UHD HDR TV/broadcasting applications. However, previous CNN-based methods directly reconstruct the HR HDR frames from LR SDR frames, and are only trained with a simple L2 loss. In this paper, we take a *divide-and-conquer* approach in designing a novel GAN-based joint SR-ITM network, called JSI-GAN, which is composed of three task-specific subnets: an image reconstruction subnet, a detail restoration (DR) subnet and a local contrast enhancement (LCE) subnet. We delicately design these subnets so that they are appropriately trained for the intended purpose, learning a pair of pixel-wise 1D separable filters via the DR subnet for detail restoration and a pixel-wise 2D local filter by the LCE subnet for contrast enhancement. Moreover, to train the JSI-GAN effectively, we propose a novel detail GAN loss alongside the conventional GAN loss, which helps enhancing both local details and contrasts to reconstruct high quality HR HDR results. When all subnets are jointly trained well, the predicted HR HDR results of higher quality are obtained with at least 0.41 dB gain in PSNR over those generated by the previous methods. The official Tensorflow code is available at <https://github.com/JihyongOh/JSI-GAN>.

Introduction

High dynamic range (HDR) videos can more realistically represent natural scenes, with higher bit depth per pixel and rich colors from a wider color gamut. They can be viewed on HDR TVs, which also tend to be UHD (Ultra High Definition) with very high resolution. However even with the latest UHD HDR TVs, the vast majority of the transmitted visual contents are still low resolution (LR), standard dynamic range (SDR) videos. There are also abundant existing legacy LR SDR videos, which brings about the need for appropriate conversion technologies that can generate high resolution (HR) HDR videos from LR SDR videos. This can

*Both authors contributed equally to this work.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

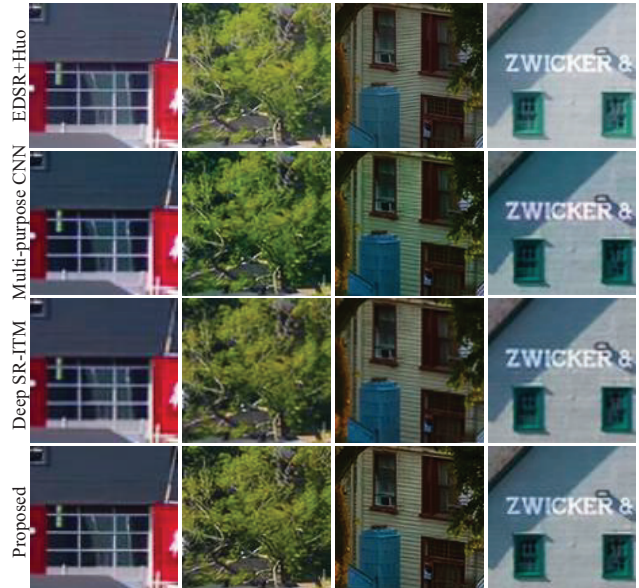


Figure 1: Qualitative comparison against other methods. Our method can reconstruct fine lines with realistic textures.

be achieved by joint super-resolution (SR) and inverse tone-mapping (ITM), where SR up-scales LR videos to HR, and ITM up-converts SDR videos to HDR.

In joint SR-ITM, it is important to restore details while up-scaling the image resolution, and to enhance local contrast while increasing the signal amplitudes. In this paper, we take a *divide-and-conquer* approach by dividing this joint problem into three sub-tasks: image reconstruction (IR), detail restoration (DR), and local contrast enhancement (LCE). A single subnet is dedicated for each of the tasks, but all subnets are jointly trained for the same goal of joint SR-ITM. To overcome the limitations of conventional shared-convolution filters over input channels in each layer, we design a pair of pixel-wise 1D separable filters in the DR subnet for detail restoration and a pixel-wise 2D local filter in the LCE subnet for contrast enhancement. Moreover, the 1D separable and 2D local filters are designed to be scalable

to up-scaling factors. This approach is inherently different from the approaches that directly produce the output HR HDR images. Furthermore, each input frame is divided into its base layer and a detail layer component by the guided filter (He, Sun, and Tang 2012). In our composite framework, the DR subnet, LCE subnet and the IR subnet are optimally combined to finally produce faithful HR HDR results.

Generative adversarial networks (GANs) have been widely applied in low level vision tasks, such as SR, where they tend to generate images with high subjective (perceptual) quality but low objective quality (e.g. PSNR, SSIM, etc.). For joint SR-ITM, direct generation of output images based on the conventional GAN-based methods can lead to unsatisfying results with lack of details and unnatural local contrasts, since simultaneously enhancing the local contrast and details while increasing both the bit-depth and the spatial resolution becomes a very challenging task in training GAN-based frameworks. Therefore, our GAN-based joint SR-ITM method, called JSI-GAN, incorporates a novel detail loss that enforces its generator to mimic the perceptually realistic details in the ground truth, and a feature-matching loss that helps mitigate the drop in objective performance by stabilizing the training process.

Our contributions are summarized as follows:

- We first propose a GAN framework for joint SR-ITM, called JSI-GAN, with a novel detail loss and a feature-matching loss that enable the restoration of realistic details and force stable training.
- The generator of JSI-GAN is designed to have task-specific subnets (DR/IR/LCE subnets) with pixel-wise 1D separable filters for local detail improvement and a 2D local filter for local contrast enhancement by considering local up-sampling operation given the up-scaling factor.
- The DR subnet focuses on the high frequency components to elaborately restore the details of the HR HDR output, while the LCE subnet effectively restores the local contrast by focusing on the base layer component of the LR SDR input.

Related Work

Pixel-Wise or Pixel-Aware Filter Learning

In convolution layers, the same convolution kernels are applied on all spatial positions of the input, and once trained, the same kernels are used for any input image. Therefore, to consider pixel-wise and sample-wise diversity, Brabandere *et al.* first introduced dynamic filter networks (De Brabandere *et al.* 2016) in video and stereo prediction, where position-specific filters are predicted through a CNN and applied as an inner product on each pixel position of the input image. Since a different filter is applied to each pixel, and different filters are predicted from different input images, they allow for sample-specific and position-specific filtering. This operation is called *dynamic local filtering*.

This concept was successfully utilized in other video-related tasks, such as frame interpolation (Niklaus, Mai, and Liu 2017b) and video SR (Jo *et al.* 2018). Niklaus *et al.*'s first idea (Niklaus, Mai, and Liu 2017a) was similar to that

of (De Brabandere *et al.* 2016) with 2D local filters being predicted, named as *adaptive convolution*. Their extended work (Niklaus, Mai, and Liu 2017b) with two 1D separable filters (horizontal/vertical) allowed to enlarge the final receptive field with the same number of parameters. For filter generation networks, the receptive field for the final output is solely defined by the size of the generated filter, implying that the depth or kernel sizes in the middle layers do not affect the final receptive field. Jo *et al.* used 3D-CNNs and added an up-sampling feature in generating the 2D filters (dynamic up-sampling filters), while incorporating a conventional residual network with direct reconstruction.

In our architecture, we design (i) the DR subnet with pixel-wise 1D separable horizontal and vertical filters to capture the *distinct details*; (ii) the LCE subnet with pixel-wise 2D local filters so that it can focus on the *local region contrast*. With the same number of filter parameters, 1D separable filters are coarse representations of a larger region, whereas 2D filters are finer representations of a local receptive area. In our framework, the DR subnet, LCE subnet and the IR subnet are combined to produce the final HR HDR images.

Generative Adversarial Networks

A GAN-based framework is typically composed of a generator and a discriminator, which are trained in an adversarial way, to force the generator to synthesize realistic images that are indistinguishable by the discriminator (Goodfellow *et al.* 2014). Many advanced techniques and variants of GANs have achieved significant performance improvements in various computer vision tasks, especially in image restoration, such as image dehazing (Qu *et al.* 2019), SR (Wang *et al.* 2018; Ledig *et al.* 2017), denoising (Chen *et al.* 2018a) and image enhancement (Chen *et al.* 2018b). These methods adopted various GAN-based frameworks to improve the perceptual quality for their individual purposes, and they generally train the main network as the generator, which is first trained with a pixel-wise norm-based loss, and is then fine-tuned by introducing an adversarial loss with the discriminator.

Motivated by the enhanced perceptual quality of GAN-based methods, we also design a GAN-based framework for joint SR-ITM. However, simply adopting conventional GANs for joint SR-ITM leads to difficulty in training the network for a more complex task of jointly improving the local variations of contrast and details along with up-converting to both higher bit depth per pixel and larger spatial resolution. To guide the generator in a more effective way to fool the discriminator and generate perceptually realistic HR HDR results with the GAN-based framework, we propose a novel detail GAN loss, which jointly optimizes a second discriminator for the detail components decomposed from the HR HDR prediction and the ground truth. We also employ the feature-matching loss that is measured in the feature space of the discriminator to reduce the drop in objective performance for a more stable training.

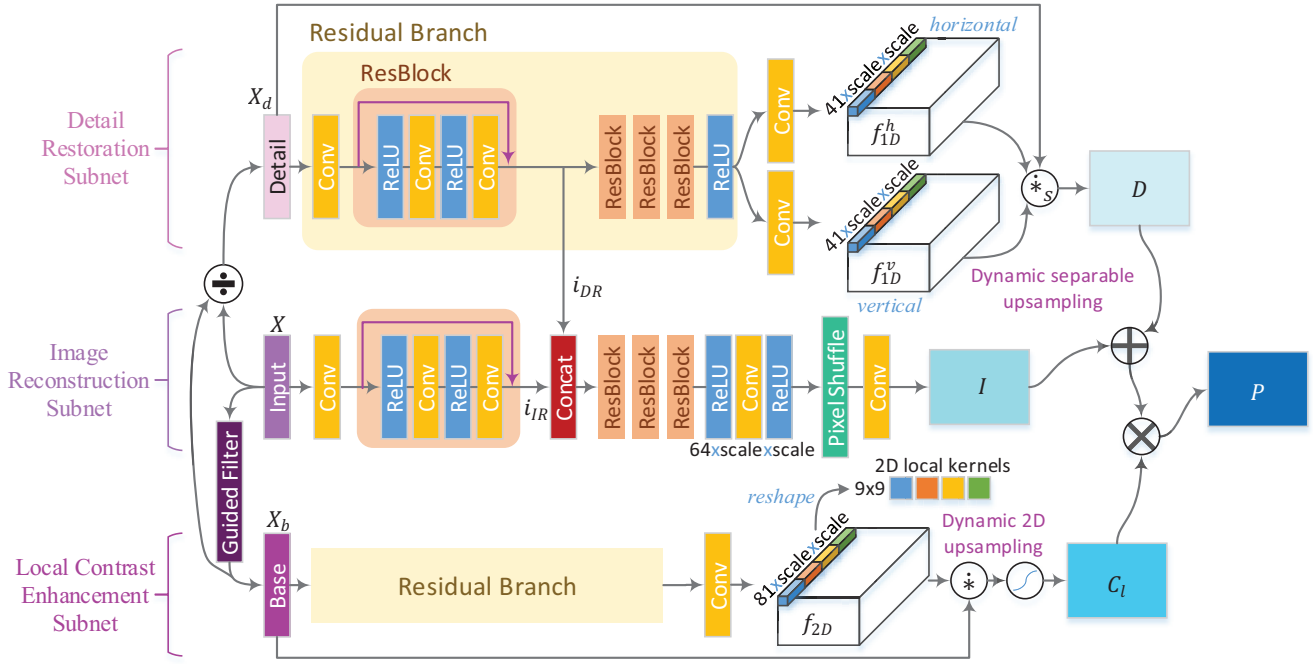


Figure 2: Network architecture of JSInet (Generator).

Joint SR-ITM

Deep-learning-based joint SR-ITM is a recent topic rising with the advent of UHD HDR TVs. The first CNN-based joint SR-ITM method (Kim and Kim 2018) took a multi-task learning approach and focused on the advantages of performing the individual SR and ITM tasks simultaneously, along with the joint SR-ITM task. The more recent method is the Deep SR-ITM (Kim, Oh, and Kim 2019), where the input is decomposed, and element-wise modulations are inserted to introduce spatially-variant operations. We also use the decomposed input for the DR subnet, but our architecture incorporates pixel-wise filters of two kinds: separable 1D kernels and 2D local kernels with up-sampling that are used to generate the final HR HDR output by filtering operations. Moreover, our network is a GAN-based framework, differing from the previous methods.

Proposed Method

We propose a GAN-based framework for joint SR-ITM, called JSI-GAN, where the generator, JSInet, is composed of three different task-specific subnets.

Network Architecture (Generator)

In joint SR-ITM, it is important to restore the high frequency details and enhance local contrast with the increase in image resolutions and pixel amplitudes to generate a high quality HR HDR image. Therefore, we design three subnets dedicated for each of these subtasks as a *divide-and-conquer* approach: the image reconstruction (IR) subnet reconstructs a coarse HR HDR image; the detail restoration (DR) subnet restores the details to be added on the coarse image; and the

local contrast enhancement (LCE) subnet generates a local contrast mask to boost the contrast in this image. A detailed structure of JSInet is depicted in Fig. 2.

Detail Restoration (DR) Subnet For the DR subnet, the input is the detail layer X_d , containing the high frequency components of the LR SDR input image X . X_d is given by,

$$X_d = X \oslash X_b, \quad (1)$$

where X_b is the guided filtered output of X , and \oslash denotes element-wise multiplication, as in (Kim, Oh, and Kim 2019). In our implementations, a small value 10^{-15} is added to the denominator to prevent X_d from diverging in case X_b approaches zero. X_d is utilized to generate horizontal and vertical 1D separable filters. Formally, a residual block (ResBlock) RB is defined as,

$$RB(x) = (Conv \circ RL \circ Conv \circ RL)(x) + x, \quad (2)$$

where x is the input to the ResBlock, $Conv$ is a convolution layer, and RL is a ReLU activation. Then, the horizontal 1D filter f_{1D}^h is obtained by,

$$f_{1D}^h = (Conv \circ RL \circ RB^4 \circ Conv)(X_d), \quad (3)$$

where RB^n denotes a serial cascade of n ResBlocks. The vertical 1D filter f_{1D}^v can be obtained in the same way as Eq. (3). As shown in Fig. 2, all layers except the last convolution layer are shared when producing f_{1D}^h and f_{1D}^v .

Each of the two last convolution layers consists of $41 \times scale \times scale$ output channels, where 41 is the length of the 1D separable kernel, each applied onto its corresponding grid location, and $scale \times scale$ takes into account the pixel shuffling operation for the up-scaling factor $scale$. Hence, this dynamic separable up-sampling operation ($*_s$) applies two 1D separable filters while producing

the spatially-upscaled output. Then, the final filtered output of the DR subnet is given by,

$$D = X_d \star_s (f_{1D}^v, f_{1D}^h). \quad (4)$$

The generated 1D kernels are position-specific, and also detail-specific, as different kernels are generated for different detail layers, unlike convolution filters that are fixed after training. In our implementations, f_{1D}^v was first applied to the detail layer via local filtering for each scale channel, followed by applying f_{1D}^h on its output. Finally, pixel shuffle was applied on the final filtered output with $scale \times scale$ channels for spatial up-scaling. With the same number of parameters, the 1D separable kernels represent a wider receptive field ($2k$ parameters), but as a coarse approximation compared to a 2D kernel (k^2 parameters). In our case, the separable 1D kernels of size $k = 41$ can be compared to the 2D kernel of size $k = 9$, with similar number of parameters.

Local Contrast Enhancement (LCE) Subnet The base layer X_b , obtained through guided filtering, is used as the input to the LCE subnet. The LCE subnet generates a 9×9 2D local filter at each pixel grid position. As in the DR subnet, it is also an up-sampling filter, and has $9 \times 9 \times scale \times scale$ output channels in the last layer. The 2D filter f_{2D} is then given by,

$$f_{2D} = (Conv \circ RL \circ RB^4 \circ Conv)(X_b). \quad (5)$$

With the 2D local filter predicted, dynamic 2D up-sampling operation (\star) is performed, and $2 \times sigmoid(x)$ is applied so that the the output lies in the range $[0, 2]$ with the middle value at $x = 0$ being 1. Formally, the final output C_l of the LCE subnet is given as,

$$C_l = 2 \times sigmoid(X_b \star f_{2D}). \quad (6)$$

As C_l is considered an LCE mask, and is element-wisely multiplied onto the sum of the two outputs from the IR and DR subnets, the JSInet converges better with the scaled sigmoid function. Without it, the pixel values of the initial predicted outputs (after multiplying with C_l) are too small, entailing much longer training time for the final HR HDR output of JSInet to reach an appropriate pixel range.

Image Reconstruction Subnet For the IR subnet, the LR SDR input X as itself is entered, to first produce the intermediate features i_{IR} as shown in Fig. 2, given by,

$$i_{IR} = (RB \circ Conv)(X). \quad (7)$$

Then, i_{IR} is concatenated with i_{DR} from the DR subnet, and the final output I of the IR subnet is directly generated (without local filtering) as,

$$I = (Conv \circ PS \circ RL \circ Conv \circ RL \circ RB^3)([i_{IR}, i_{DR}]), \quad (8)$$

where PS is a pixel-shuffle operator and $[x, y]$ is the concatenation of x and y in the channel direction.

Then, the final HR HDR prediction P is generated by adding the details (D) to I followed by multiplying the contrast mask (C_l) to the result, given by,

$$P = (I + D) \times C_l. \quad (9)$$

We provide an ablation study on the three subnets and demonstrate that they are acting according to their given tasks in the later sections of the paper.

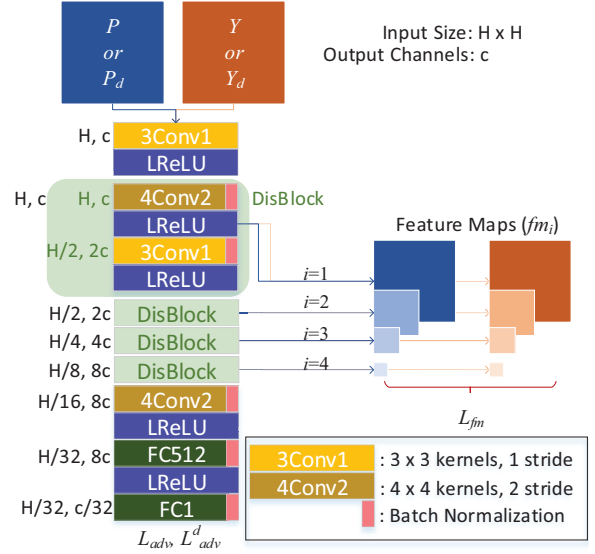


Figure 3: Discriminator architecture of JSI-GAN.

GAN-based Framework

Discriminator We employ a GAN-based framework, where the discriminator is designed as shown in Fig. 3 with spectral normalization (Miyato et al. 2018) for stable learning. The discriminator distinguishes the predicted HR HDR image (P) generated by the generator and the ground truth image (Y) alternatively during training. As shown in Fig. 3, when the input x (P or Y) is given to the discriminator D^f , the output is obtained by,

$$D^f(x) = (BN \circ FC1 \circ BN \circ FC512 \circ LRL \circ BN \circ 4Conv2 \circ DB^4 \circ LRL \circ 3Conv1)(x). \quad (10)$$

where BN is batch normalization (Ioffe and Szegedy 2015), LRL is a Leaky ReLU (LReLU) activation (Maas, Hannun, and Ng) with a slope size of 0.2, FCk is a fully-connected layer with k output channels, and $kConv_s$ denotes $k \times k$ kernels with a stride s . The input sizes ($H \times H$) and output channels (c) for each layer/block in Fig. 3. DB^n denotes n -times serially connected DisBlocks, where DB (DisBlock) is given as,

$$DB(x) = (LRL \circ BN \circ 3Conv1 \circ LRL \circ BN \circ 4Conv2)(x). \quad (11)$$

Adversarial Loss The RaHinge GAN loss (Jolicœur-Martineau 2019) is adopted as a basic adversarial loss for efficient training, which is given by

$$L_{adv}^D = \mathbb{E}_Y[\max(0, \tilde{Q}_{Y,P}^{(-)})] + \mathbb{E}_P[\max(0, \tilde{Q}_{P,Y}^{(+)})] \quad (12)$$

$$L_{adv}^G = \mathbb{E}_P[\max(0, \tilde{Q}_{P,Y}^{(-)})] + \mathbb{E}_Y[\max(0, \tilde{Q}_{Y,P}^{(+)})] \quad (13)$$

where L_{adv}^D and L_{adv}^G denote the RaHinge GAN losses for the discriminator D^f and the generator G , respectively, and

$\tilde{Q}_{P,Y}^{(\pm)} = 1 \pm \tilde{D}_{P,Y}$ with $\tilde{D}_{P,Y} = D_f(P) - \mathbb{E}_Y D_f(Y)$. It should be noted in Eq. (13) that L_{adv}^G contains both the Y and P , meaning that the generator is trained by gradients from both the Y and P during the adversarial training. We also use a feature-matching loss $L_{fm} = \sum_{i=1}^4 \|fm_i(Y) - fm_i(P)\|_2$, where the L2 loss is measured between feature maps $fm_i(\cdot)$ of Y and P , from the first LReLU output of the i -th DisBlock as shown in Fig. 3. However, simple utilization of the above two losses is insufficient to effectively train the generator for joint SR-ITM.

Detail GAN Loss We propose a novel detail GAN loss L_{adv}^d for the joint SR-ITM task, in order to enforce more stable training and generate visually pleasing HR HDR results. L_{adv}^d is calculated according to Eq. (12) and Eq. (13) by replacing Y with Y_d and P with P_d , where the subscript d denotes the detail layer component of the original image. For L_{adv}^d , we adopt a second discriminator (D_2) distinguished from the first discriminator (D_1), both of which have the same architecture as shown in Fig. 3, but D_2 alternatively takes two inputs P_d and Y_d , calculated by Eq. (1). L_{adv}^d not only guides the generator for a more stable training but also helps improve both local contrasts and details in the predicted HR HDR images.

Total Loss The total loss for our proposed GAN-based framework for joint SR-ITM is given by

$$L_{D_1} = L_{adv}^{D_1}, L_{D_2} = \lambda_d \cdot L_{adv}^{d,D_2}, \quad (14)$$

$$L_G = \lambda_{rec} \cdot \|Y - P\|_2 + \lambda_{adv} \cdot (L_{adv}^G + \lambda_d \cdot L_{adv}^{d,G}) + \lambda_{fm} \cdot (L_{fm} + \lambda_d \cdot L_{fm}^d), \quad (15)$$

where the superscript d means the loss for detail layer components (P_d, Y_d). We provide another ablation study on the losses L_{fm} and L^d , and especially show the effect of the newly proposed L_{adv}^d , in the later sections of the paper.

Experiment Results

Experiment Conditions

For training, the generator was first pre-trained based only upon the L2 loss with the initial learning rate of 10^{-4} , which is then decreased by a factor of 10 at epochs 200 and 225, of total 250 epochs, yielding the JSInet. Then it was fine-tuned based on a stable GAN-based framework with three losses (L_G, L_{D_1}, L_{D_2}) at the initial learning rate of 10^{-6} that linearly decays to zero from the 5-th epoch of total 10 epochs, which finally yields the JSI-GAN. For training, we used three Adam optimizers (Kingma and Ba 2015) for minimizing L_{D_1}, L_{D_2} and L_G , and all convolution weights were initialized by the Xavier method (Glorot and Bengio 2010). The generator G and the two discriminators (D_1, D_2) are trained alternatively by the three corresponding Adam optimizers. The weighting parameters for the losses were empirically set to $\lambda_{rec} = 1, \lambda_{adv} = 1, \lambda_{fm} = 0.5$ and $\lambda_d = 0.5$.

In the generator, the kernel size of the convolution filters were set to 3×3 with 64 output channels, except for the

Table 1: Ablation study on the subnets.

| Variant | (a) | (b) | (c) | (d) | (e) | (f) |
|---------|--------|--------|---------------|--------|--------|--------------|
| IR | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| DR | ✗ | ✓(-) | ✓ | ✗ | ✓(-) | ✓ |
| LCE | ✗ | ✗ | ✗ | ✓ | ✓ | ✓ |
| PSNR | 35.44 | 35.89 | 35.88 | 35.67 | 35.68 | 35.99 |
| SSIM | 0.9756 | 0.9762 | 0.9773 | 0.9761 | 0.9756 | 0.9768 |

*(-): Subtraction instead of division to obtain the detail layer.

**PSNR is measured in dB.

last layer that predicts local filters for the DR subnet and the LCE subnet, and the layer before pixel shuffle for the IR subnet. The structure of both D_1 and D_2 has the channel outputs of $c = 32$. The LR SDR patches of size 80×80 ($H = 80$) or 40×40 ($H = 40$) were cropped from 8-bit YUV frames in BT.709 (ITU-R 2002) for scale factor 2 and 4 respectively, and the ground truth (HR HDR) patches of size 160×160 ($H = 160$) were cropped from the corresponding 10-bit YUV frames in BT.2020 (ITU-R 2014) color container after PQ-OETF (SMPTE 2014), following the setting in previous work (Kim, Oh, and Kim 2019). For training and testing the JSI-GAN, we utilized the 4K HDR dataset used in (Kim, Oh, and Kim 2019) and adopted Tensorflow 1.13 in our implementations.

Performance of JSInet

We first investigate the performance of JSInet that is trained solely on the L2 loss *without* the GAN framework, by analyzing the efficacy of its three subnets.

Ablation Study of Subnets We first performed an ablation study on the three subnets in JSInet, by retraining different variants of the network. Table 1 shows the PSNR/SSIM performance of six combinations (variants) of the three subnets, where the IR subnet is essential for all cases. As shown in column (c) of Table 1, employing the DR subnet to the IR subnet brings 0.44 dB gain in PSNR, and additionally using the LCE subnet in column (f), further increases the PSNR by 0.11 dB, resulting in a total 0.55 dB gain over only using the IR subnet in column (a). It is also noted that the LCE subnet brings a higher performance gain when fused with the IR subnet (column (d)) with 0.23 dB gain, meaning that the LCE subnet is complementarily beneficial with the DR subnet.

We have also provided the experiment results for a different decomposition strategy, subtraction, on the input images instead of division in Eq. (1), with the sign (-) in Table 1. For the structure in column (b) and (c), there is minimal difference in PSNR, although division yields slightly better performance for SSIM. However, when the LCE subnet is jointly utilized, division outperforms subtraction by a larger margin of 0.31 dB in PSNR, since for column (e), as the local contrast map C_l (obtained from the base layer) is multiplied to the coarse image I with the added details D in Eq. (9), the subtraction operation is unmatched with the multiplication operation.

Some additional results not marked in Table 1 were the following. (i) If the concatenation of i_{DR} was removed,

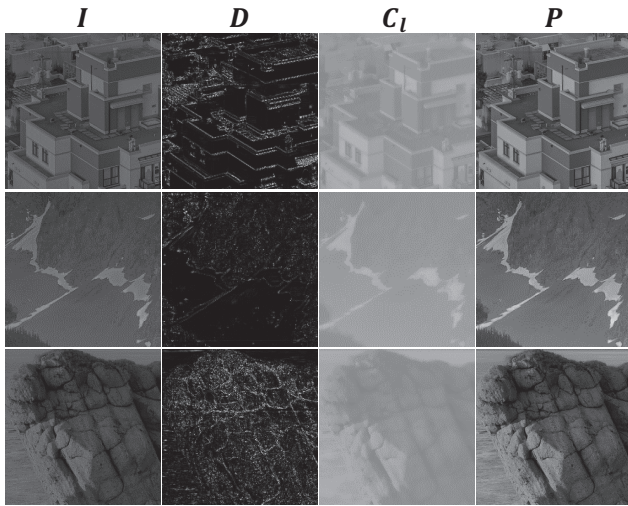


Figure 4: Visualization of the subnet predictions.

Table 2: Combinations of direct reconstruction and filter generation for the network with IR/DR subnets.

| IR subnet | direct | filter | direct |
|-----------|--------|--------|---------------|
| DR subnet | direct | direct | filter |
| PSNR (dB) | 35.73 | 34.14 | 35.88 |
| SSIM | 0.9765 | 0.9662 | 0.9773 |

there was a 0.02 dB drop in PSNR from the network in Table 1 (c). (ii) If feature maps i_{LCE} from the LCE subnet were additionally concatenated along with i_{IR} and i_{DR} , there was a 0.25 dB drop in PSNR from the network in Table 1 (f). Note that only i_{IR} and i_{DR} are stacked in the final generator architecture of our proposed JSInet in Fig. 2. (iii) If the image X was stacked with X_d for the input of the DR subnet, following the idea of providing guidance to the detail layer in (Kim, Oh, and Kim 2019), there was actually a minimal decrease, or no difference in PSNR, when experimented with the network in Table 1 (b) and Table 1 (c), respectively.

Visual Analysis To verify that each of the subnets is focusing on their dedicated tasks, the intermediate predictions (I, D, C_l) of the subnets and the final prediction P are visualized in Fig. 4. For visualization, $I, |D|, C_l$ and P were first linearly scaled to a maximum value of 8 bits/pixel, and $|D|$ was further scaled by 64 for better visualization. In Fig. 4, the added details D are invariant to the brightness context (1st and 2nd rows) and lighting conditions (3rd row), focusing only on the edges and texture as intended. The LCE mask C_l effectively modulates the local contrast, producing the final P with enhanced contrast.

Filter Generation Module To analyze the benefits of generating local filters, we performed an experiment with the combinations of direct reconstruction and kernel generation for a simple network with only the IR and DR subnet, as shown in Table 2. For the DR subnet, it is important to apply location-varying filters, as uniform convolution operations

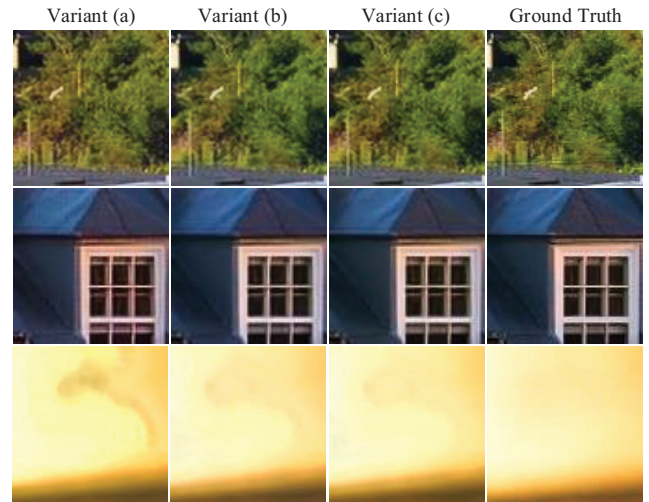


Figure 5: Effect of the weight parameters λ_{fm} and λ_d .

Table 3: Ablation study on the losses.

| Scale | $\times 2$ | | | $\times 4$ | | |
|----------------|------------|--------------|---------------|------------|---------------|--------------|
| Variant | (a) | (b) | (c) | (d) | (e) | (f) |
| λ_{fm} | \times | \checkmark | \checkmark | \times | \checkmark | \checkmark |
| λ_d | \times | \times | \checkmark | \times | \checkmark | \checkmark |
| PSNR (dB) | 33.73 | 35.63 | 35.73 | 32.11 | 33.48 | 33.50 |
| SSIM | 0.9643 | 0.9754 | 0.9763 | 0.9478 | 0.9577 | 0.9572 |

would blur the high frequency details. If the detail map D is directly generated, there is a performance drop of 0.15 dB in PSNR, as shown in Table 2. Moreover, if the filter generation module is used for the IR subnet, there is a drastic drop in PSNR performance compared to the proposed combination in the last column. For the LCE subnet, we employ a 2D filter generation module so that local information can be considered. If 1D separable kernels were predicted instead for the LCE subnet, there was a 0.07 dB drop in performance gain.

Performance of JSI-GAN

We also conducted an ablation study on the efficacy of various losses in terms of weighting parameters λ_{fm} and λ_d , to investigate their effects. Table 3 shows the average PSNR (dB) and SSIM performance of the JSI-GAN variants, each of which was trained via combinations with/without λ_{fm} and λ_d , for scales $\times 2$ and $\times 4$. If JSI-GAN is only trained with the basic adversarial loss where $\lambda_{rec} = 1, \lambda_{adv} = 1, \lambda_{fm} = 0$ and $\lambda_d = 0$, severe performance degradation is observed with at most 2 dB drop in PSNR (variants (a) and (d)). By comparing variants (b) to (a) and (e) to (d), additionally adopting the feature-matching loss L_{fm} between P and Y helps significantly in improving objective performance, with 1.9dB and 1.37dB gain in PSNR, respectively. Finally, our proposed detail GAN loss L_{adv}^d between P_d and Y_d allows for further improvements in the quantitative performance as shown in variants (c) and (f) of Table 3.

The effect of the losses is also shown qualitatively in Fig.

Table 4: Quantitative Comparison.

| Method | Scale | PSNR (dB) | SSIM | mPSNR (dB) | MS-SSIM | HDR-VDP (Q) | Inf. Time [s] | No. Parameters |
|------------------------------|-------|--------------|---------------|--------------|---------------|--------------|---------------|----------------|
| EDSR+Huo <i>et al.</i> | ×2 | 29.76 | 0.8934 | 31.81 | 0.9764 | 58.95 | - | - |
| EDSR+Eilertsen <i>et al.</i> | ×2 | 25.80 | 0.7586 | 28.22 | 0.9635 | 53.51 | - | - |
| Multi-purpose CNN | ×2 | 34.11 | 0.9671 | 36.38 | 0.9817 | 60.91 | 0.49 | 250K |
| Deep SR-ITM | ×2 | 35.58 | 0.9746 | 37.80 | 0.9839 | 61.39 | 5.02 | 2.50M |
| JSInet (w/o GAN) | ×2 | 35.99 | 0.9768 | 38.20 | 0.9843 | 60.58 | 1.72 | 1.45M |
| JSI-GAN (w/ GAN) | ×2 | 35.73 | 0.9763 | 37.96 | 0.9841 | 60.80 | 1.72 | 1.45M |
| EDSR+Huo <i>et al.</i> | ×4 | 28.90 | 0.8753 | 30.92 | 0.9693 | 55.59 | - | - |
| EDSR+Eilertsen <i>et al.</i> | ×4 | 26.54 | 0.7822 | 28.75 | 0.9631 | 53.88 | - | - |
| Multi-purpose CNN | ×4 | 33.10 | 0.9499 | 35.26 | 0.9758 | 56.41 | 0.34 | 283K |
| Deep SR-ITM | ×4 | 33.61 | 0.9561 | 35.73 | 0.9748 | 56.07 | 4.06 | 2.64M |
| JSInet (w/o GAN) | ×4 | 33.74 | 0.9598 | 35.93 | 0.9759 | 56.45 | 1.05 | 3.03M |
| JSI-GAN (w/ GAN) | ×4 | 33.50 | 0.9572 | 34.82 | 0.9743 | 56.41 | 1.05 | 3.03M |

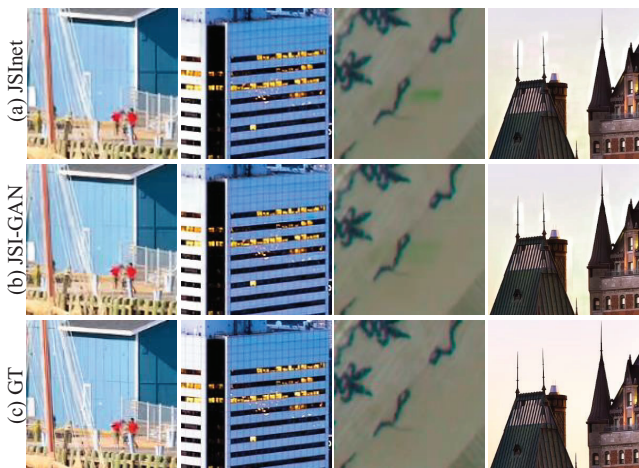


Figure 6: Visual comparison: JSInet and JSI-GAN.

5. Just simply adopting the basic GAN loss not only degrades the quantitative performance, but also severely deteriorates the visual qualities with the checkerboard artifacts and unnatural details/contrasts, as shown in the leftmost column in Fig. 5. Although the feature-matching loss helps the generator improve the overall visual quality when comparing the variant (b) to (a), the proposed detail-component-related losses ($L_{adv}^{d,D}$, $L_{adv}^{d,G}$, L_{fm}^d) additionally improve both visual qualities and objective performances comparing the variant (c) to the variants (a) and (b). As a result, the final JSI-GAN enables the restoration of realistic details with stable training while allowing for high objective quality of the HR HDR predictions.

Visual Comparison: JSInet and JSI-GAN

We also compare the visual results of the JSInet (without GAN) and JSI-GAN (with GAN). Although there is a slight drop in the performance of JSI-GAN in terms of PSNR (dB) in Table 4, there are some qualitative differences as shown in Fig. 6. Since JSInet is only trained via reconstruction loss (L_2) *without* the GAN framework, the visual results of JSInet shown in the 1st row of Fig. 6 tends to be blurred in

general. By comparing the row (b) to row (a) in 1st and 2nd column in Fig. 6, both local details and contrasts are boosted by utilizing proposed GAN-based framework. Furthermore, the ringing artifacts or color artifacts observed in the homogeneous regions (especially in background) of the results by JSInet, shown in 3rd and 4th columns of Fig. 6, are also enhanced by JSI-GAN.

Comparison with Other Methods

We compare our JSI-GAN with the two previous joint SR-ITM methods, the Multi-purpose CNN (Kim and Kim 2018) and Deep SR-ITM (Kim, Oh, and Kim 2019). The proposed JSI-GAN is trained on the same data as Deep SR-ITM, and the Multi-purpose CNN was re-trained on the same data. Additionally, the JSI-GAN is also compared with the cascades of an SR method, EDSR (Lim *et al.* 2017), and two ITM methods (Huo *et al.* 2014; Eilertsen *et al.* 2017). Note that Huo *et al.*'s method is not data-driven, and thus invariant to training data. Since Eilertsen *et al.*'s method was designed to be trained only for saturated pixel regions, it was not trained with our data that does not contain enough saturated pixel regions (unlike luminance-map-type data used in their dataset). The pre-trained model of EDSR as provided in the official code is used. The previous methods were compared following the experiment protocol on the ITM prediction pipeline and visualization as described in Deep SR-ITM (Kim, Oh, and Kim 2019).

Quantitative Comparison The quantitative comparison of the proposed JSInet and JSI-GAN against the previous methods is given in Table 4. The performance is measured using error-based metrics such as PSNR and mPSNR, and structural metrics such as SSIM and MS-SSIM, as well as HDR-VDP-2.2.1 (Mantiuk *et al.* 2011), which is able to measure the performance degradation in all luminance conditions. The CNN-based joint SR-ITM methods outperform the cascaded methods by a large margin, and our JSInet outperforms the other methods in all cases except for HDR-VDP for scale factor 2. The proposed JSI-GAN also shows good quantitative performance, thanks to the stable training that mitigates the drop in objective performance metrics.

Additionally, we provide the inference time in seconds and the number of parameters for the joint SR-ITM meth-

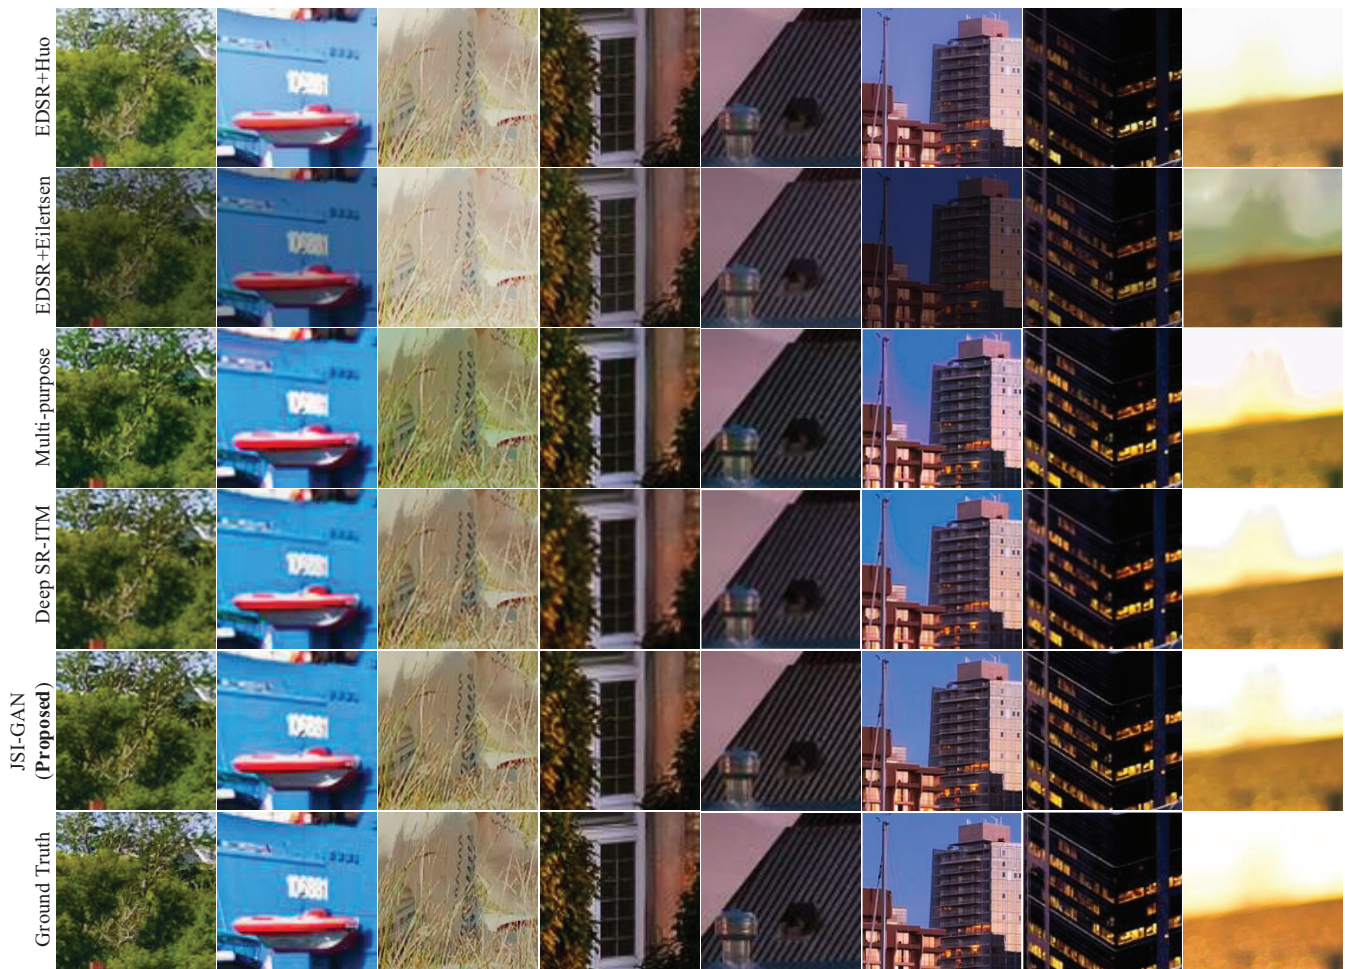


Figure 7: Qualitative Comparison.

ods, which are denoted as *Inf. Time* and *No. Parameters* in Table 4, when converting to 4K target resolution measured with an NVIDIA TITAN Xp GPU. Note that only the generator parameters are computed for JSI-GAN, and the number of parameters increases proportionally with the scale factor due to the parameters dedicated for the local filters.

Qualitative Comparison The qualitative comparison of the proposed JSI-GAN is given in Fig. 1 and Fig. 7. In Fig. 1, our method is able to reconstruct the fine lines on the window, produce more tree-like textures, and generate correct horizontal patterns on the wall. Likewise in Fig. 7, JSI-GAN generates *fine details* with *enhanced contrast* without artifacts in the homogeneous regions, thanks to the task-specific subnets and the detail component-related losses.

Conclusion

In this paper, we first proposed a GAN-based framework for joint SR-ITM, called JSI-GAN, where the generator (JSInet) jointly optimizes the three task-specific subnets designed with pixel-wise 1D separable filters for fine detail restoration and pixel-wise 2D local filters for contrast enhance-

ment. These subnets were carefully designed for their intended purposes to boost the output performance. Moreover, we also proposed a novel detail GAN loss alongside the conventional GAN loss, which helps enhancing both the local details and contrasts for generating high quality HR HDR reconstructions. We analyzed the efficacy of the subnet components and the weighting parameters for losses with intensive ablation studies. Our proposed JSI-GAN, which is applicable for directly converting LR SDR frames to HR HDR ones in real-world applications, achieves the state-of-the-art performance over the previous methods.

Acknowledgement

This work was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIT) (No. 2017-0-00419, Intelligent High Realistic Visual Processing for Smart Broadcasting Media).

References

- Chen, J.; Chen, J.; Chao, H.; and Yang, M. 2018a. Image blind denoising with generative adversarial network based noise modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3155–3164.
- Chen, Y.-S.; Wang, Y.-C.; Kao, M.-H.; and Chuang, Y.-Y. 2018b. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6306–6314.
- De Brabandere, B.; Jia, X.; Tuytelaars, T.; and Gool, L. V. 2016. Dynamic filter networks. In *Advances in Neural Information Processing Systems*, 667–675.
- Eilertsen, G.; Kronander, J.; Denes, G.; Mantiuk, R. K.; and Unger, J. 2017. Hdr image reconstruction from a single exposure using deep cnns. *ACM Transactions on Graphics* 36(6):178.
- Glorot, X., and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, 249–256.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- He, K.; Sun, J.; and Tang, X. 2012. Guided image filtering. *IEEE transactions on pattern analysis and machine intelligence* 35(6):1397–1409.
- Huo, Y.; Yang, F.; Dong, L.; and Brost, V. 2014. Physiological inverse tone mapping based on retina response. *The Visual Computer* 30(5):507–517.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning*, 448–456.
- ITU-R. 2002. Parameter values for the hdtv standards for production and international programme exchange. ITU-R Rec. BT.709-5. [Online]. Available: <http://www.itu.int/rec/R-REC-BT.709>.
- ITU-R. 2014. Parameter values for ultra-high definition television systems for production and international programme exchange. document ITU-R Rec. BT.2020-1. [Online]. Available: <http://www.itu.int/rec/R-REC-BT.2020>.
- Jo, Y.; Oh, S. W.; Kang, J.; and Kim, S. J. 2018. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3224–3232.
- Jolicoeur-Martineau, A. 2019. The relativistic discriminator: a key element missing from standard gan. In *International Conference on Learning Representations*.
- Kim, S. Y., and Kim, M. 2018. A multi-purpose convolutional neural network for simultaneous super-resolution and high dynamic range image reconstruction. In *Proceedings of the 14th Asian Conference on Computer Vision*, 379–394.
- Kim, S. Y.; Oh, J.; and Kim, M. 2019. Deep sr-itm: Joint learning of super-resolution and inverse tone-mapping for 4k uhd hdr applications. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4681–4690.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Lee, K. M. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 136–144.
- Maas, A. L.; Hannun, A. Y.; and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of the International Conference on Machine Learning*.
- Mantiuk, R.; Kim, K. J.; Rempel, A. G.; and Heidrich, W. 2011. A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics* 30(4):40.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*.
- Niklaus, S.; Mai, L.; and Liu, F. 2017a. Video frame interpolation via adaptive convolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 670–679.
- Niklaus, S.; Mai, L.; and Liu, F. 2017b. Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE International Conference on Computer Vision*, 261–270.
- Qu, Y.; Chen, Y.; Huang, J.; and Xie, Y. 2019. Enhanced pix2pix dehazing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8160–8168.
- SMPTE. 2014. High dynamic range electro-optical transfer function of mastering reference displays. ST2084:2014.
- Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Change Loy, C. 2018. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision Workshops*.