

Simple and Efficient Unpaired Real-world Super-Resolution using Image Statistics

Kwangjin Yoon

SI Analytics

70 Yuseong-daero, Yuseong-gu, Daejeon 34047, Republic of Korea

yoon28@si-analytics.ai

Abstract

Learning super-resolution (SR) network without the paired low resolution (LR) and high resolution (HR) image is difficult because direct supervision through the corresponding HR counterpart is unavailable. Recently, many real-world SR researches take advantage of the unpaired image-to-image translation technique. That is, they used two or more generative adversarial networks (GANs), each of which translates images from one domain to another domain, e.g., translates images from the HR domain to the LR domain. However, it is not easy to stably learn such a translation with GANs using unpaired data. In this study, we present a simple and efficient method of training of real-world SR network. To stably train the network, we use statistics of an image patch, such as means and variances. Our real-world SR framework consists of two GANs, one for translating HR images to LR images (degradation task) and the other for translating LR to HR (SR task). We argue that the unpaired image translation using GANs can be learned efficiently with our proposed data sampling strategy, namely, variance matching. We test our method on the NTIRE 2020 real-world SR dataset. Our method outperforms the current state-of-the-art method in terms of the SSIM metric as well as produces comparable results on the LPIPS metric.

1. Introduction

Single image Super-Resolution (SR) is the task of increasing the resolution of a given image as well as sharpening its content by predicting the high-frequency component and the missing information. Many types of research [4, 20, 9, 16] were conducted on paired low-resolution (LR) and high-resolution (HR) images to train the neural network in a fully supervised manner. They artificially produced paired data by downscaling HR images with known kernels (e.g. bicubic) to generate corresponding LR images. Since

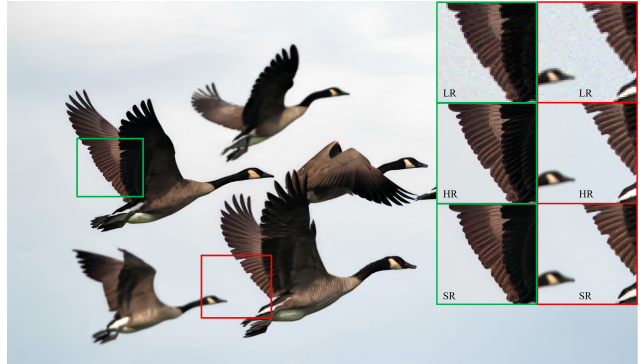


Figure 1. Result of super-resolved image ($\times 4$) by the proposed method.

there is a domain gap between downsampled images and real LR images, a network trained on downsampled images struggles to generalize to natural images [11]. On the other hand, some recent approaches tried to solve real-world SR problems by building real-world datasets where paired LR-HR images on the same scene are captured by adjusting the focal length of a digital camera [3, 19]. However, collecting such images is very expensive and time-consuming and requires extra conditions, namely non-moving objects in a scene.

Another stream of researches begins to focus on finding an actual degradation function, *i.e.* translating image domain from HR to real LR [2, 10, 17]. In other words, many real-world SR researches have recently emerged that take advantage of the unpaired image-to-image translation framework [21]. They used two or more generative adversarial networks (GANs), each of which translates images from one domain to another domain, *e.g.*, translating a given HR image to an LR image that is not distinguishable from a real LR image. However, it is difficult to learn such translations using GANs, particularly with unpaired data.

This study proposes a simple and efficient training method for real-world SR that takes unpaired data as the input. We show that our method is able to improve the SR

network by utilizing the statistics of an image patch. Specifically, the method restricts the amount of difference of the variance between LR and HR patches. By doing so, we can match the content density of those patches; hence a network can learn from unpaired patches that have a similar level of content.

We present the related work in section 2. Then, our method is described in detail in section 3. In section 4, we test our method on the NTIRE 2020 real-world SR dataset by performing $\times 4$ SR task for each dimension. Our method outperforms the current state-of-the-art method in terms of the SSIM metric as well as produces comparable results on the LPIPS [18] metric. Finally, we concluded the paper in section 5.

2. Related work

Recent super-resolution researches achieve strong performance on bicubic downsampled low-resolution images thanks to the convolutional neural network (CNN) [20, 9, 16]. Zhang *et al.* [20] proposed very deep residual channel attention network (RCAN), which uses a residual in residual structure to form a very deep network. Ledig *et al.* [9] used the residual connection mechanism to construct a deeper network with perceptual losses [8] and generative adversarial network (GAN) [6] which is called SRGAN. Wang *et al.* [16] introduced the enhanced SRGAN (ESRGAN), which improved the SRGAN by enhancing its architecture and perceptual loss. However, these SR models produce a poor result when a real-world LR image is inputted because they were trained on the paired data whose LR images were simply generated by downsampling HR counterparts with bicubic kernel. To address this issue, some researches [3, 19] collected the LR-HR pairs directly using particular camera instruments and introduced the real-world paired datasets.

Collecting the paired data requires an immense cost and extra conditions, such as a scene with non-moving objects. To overcome the problem, several real-world SR researches [2, 10, 17, 5] are proposed which use GANs to learn the conditional distribution of LR domain given an HR image. Bulat *et al.* [2] proposed a two-stage process which firstly trains a high-to-low GAN with unpaired HR and LR images, then the output of the network is inputted to low-to-high GAN, which is trained for super-resolution using paired LR and HR images. Lugmayr *et al.* [10] also used CycleGAN to translate HR images to LR images and consequently constructed the paired LR and HR dataset that is used for training the SR network. Yuan *et al.* [17] proposed a cycle-in-cycle network to simultaneously learn a degradation (high-to-low) network and SR network (low-to-high). In [5], Fritsche *et al.* proposed to separate the low and high image frequencies and treat them differently during training. The high-frequency information of an image is adver-

sarially trained with a discriminator, and the low-frequency information is learned with the $L1$ criterion. These methods use two or more GANs, each of which translates a set of images of one domain to the other domain. However, training GAN with unpaired data is prone to unstable. Our method can learn such image translation efficiently with simple image statistics.

In [7], Ji *et al.* proposed a novel degradation framework for real-world images by estimating various blur kernels as well as real noise distribution. In particular, they extracted noise patches from the real-world LR images whose variance is lower than a threshold. Then the extracted noise patch is added to downsampled HR image to imitate the real-world LR image. The assumption of their method is that the variance alone is enough to decouple noise and content. However, applying their method of generating LR images to other datasets needs far more effort than the learning-based methods since they engineered an image processing technique for a specific dataset. Furthermore, the existence of the threshold parameter needs sufficient prior knowledge about noise which is empirically set.

3. Method

In this section, we introduce the proposed method. A brief description of our unpaired SR framework and data flow is depicted in Figure 2. As shown in the figure, our network consists of two GANs. The detailed description of the architecture is presented in section 3.1.

Let x be an image in the LR domain X and y be an image in the HR domain Y . Then, we learn a mapping $F : Y \rightarrow X$ using a GAN such that the output $\hat{x} = F(y), y \in Y$, is indistinguishable from images $x \in X$, i.e. $\hat{x} \sim P_X$, where P_X is the distribution of X . In addition, we also learn another mapping $G : X \rightarrow Y$ with a GAN such that the output $\hat{y} = G(x), x \in X$, belongs to Y , i.e. $\hat{y} \sim P_Y$. Therefore, F degrades HR images into the real LR images, and G super-resolves the LR images. These two mappings are end-to-end trained with unpaired data. Since each of the two mappings is trained with the GAN framework, there are two adversarial discriminators, D_F and D_G for each generator, respectively.

In order to learn such mappings better, we also propose a simple and efficient training method which is described in section 3.3.

3.1. Network Architectures

In this section, we describe the architecture of our network which is depicted in Fig 2. Our network architectures are adopted from previous works [16, 21]. We used ESRGAN architecture [16] for generators. The network G consists of RRDBs [16] and two $\times 2$ upsampling layers, resulting in $\times 4$ SR network. Generator F , which translates images from domain Y to X (degradation task), also

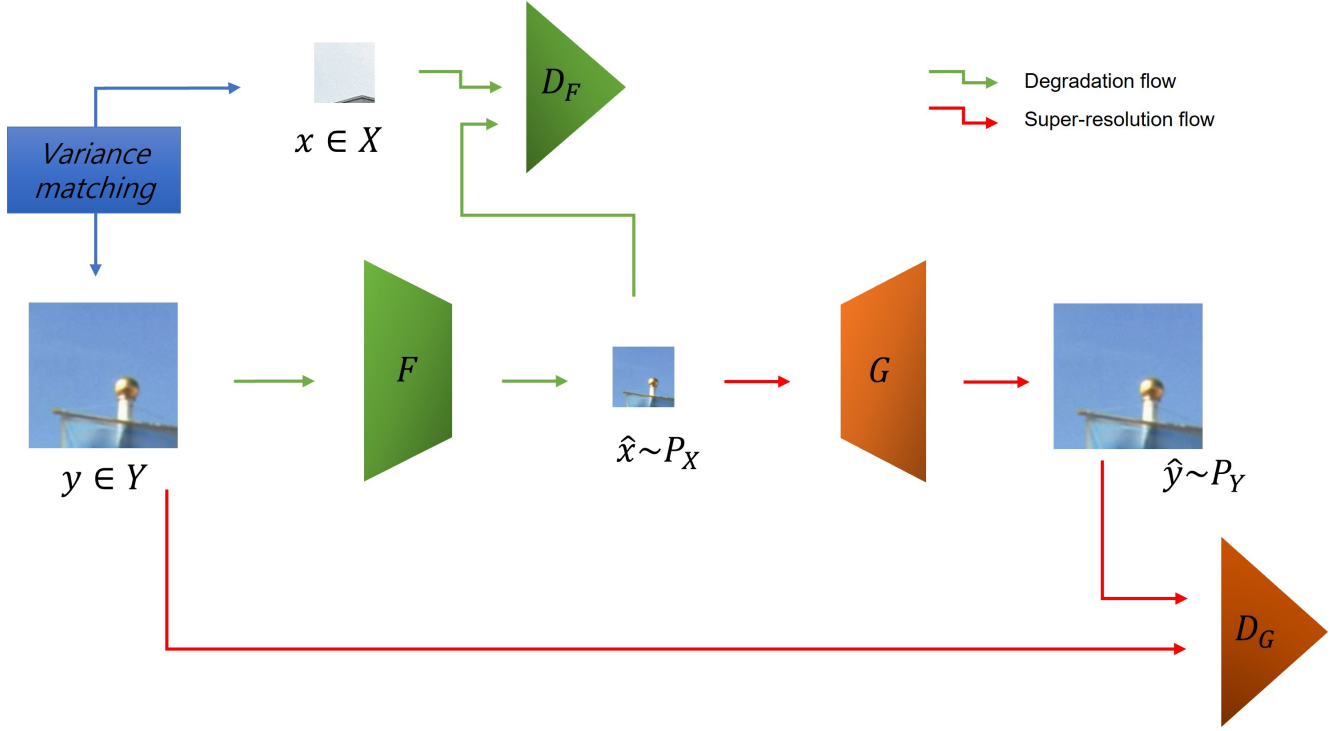


Figure 2. Framework and data flow of our method. F degrades HR images into the real LR images, and G super-resolves the LR images. *Variance matching* helps the training procedure of two GANs by constraining the variance difference between LR and HR patches so that they have a similar level of content. Detailed explanations are in section 3.3. For simplicity, the loss components are omitted.

adopts ESRGAN with modifications replacing upsampling layers with average pooling layers. For the discriminator networks, D_G and D_F , we use PatchGANs [21] which aim to classify whether 70×70 overlapping image patches are real or fake.

3.2. Losses

Our objective for both generators F and G contains the adversarial loss [6], the perceptual loss [8] and the feature matching loss [13]. The cycle consistency loss [21] is used for training G while generator F applies a content loss which is L_1 distance in the pixel domain. Therefore, the image translation cycle with our generators is only learned with the forward cycle consistency [21]. The objective of discriminators, D_G and D_F , is computed by LSGAN [12] discriminating real data from generated data.

Specifically, for the generator G , we optimize a loss:

$$L_G = \lambda_G^{adv} L_G^{adv} + \lambda_G^{cyc} L_G^{cyc} + \lambda_G^{per} L_G^{per} + \lambda_G^{fea} L_G^{fea}, \quad (1)$$

where L_G^{adv} is the adversarial loss [12]. L_G^{cyc} is the cycle consistency loss [21] which is a L_1 distance between y and $G(F(y))$. We used the improved perceptual loss (L_G^{per}) inspired by [16]. In addition, L_G^{fea} is the feature matching loss [13] that is a L_2 distance between feature vectors $D_G^f(y)$ and $D_G^f(G(\hat{x}))$, where D_G^f is an f -th intermediate layer of

the discriminator D_G . Finally, λ_G^{adv} , λ_G^{cyc} , λ_G^{per} and λ_G^{fea} are weight constants for corresponding losses.

For the generator F , the following loss is optimized:

$$L_F = \lambda_F^{adv} L_F^{adv} + \lambda_F^{con} L_F^{con} + \lambda_F^{per} L_F^{per} + \lambda_F^{fea} L_F^{fea}. \quad (2)$$

The adversarial loss (L_F^{adv}), perceptual loss (L_F^{per}) and feature matching loss (L_F^{fea}) are applied in a similar way to G . Furthermore, instead of the cycle consistency loss used in G , L_F^{con} is applied that is $\|B(y) - F(y)\|_1$. $B(\cdot)$ is a bicubic downsampling operation. Similarly, λ_F with the particular superscript is a weight constant for the corresponding loss.

3.3. Sampling with Image Statistics

We observed that learning the image translation task for SR with the unpaired dataset is difficult. We frequently observe that the procedure ends up in a trivial solution, *i.e.* produces an output that is merely a simple linear interpolation of a given input or fails to converge. We conjecture that the phenomenon is caused by random sampling from two independent variables, X , and Y . The unpaired input x and y are independent of each other, and hence each input contains completely unrelated contents and has varying degrees of image statics between them. To solve the problem, we propose to use image statistics when sampling LR and HR image patches. Specifically, we used the variance of the

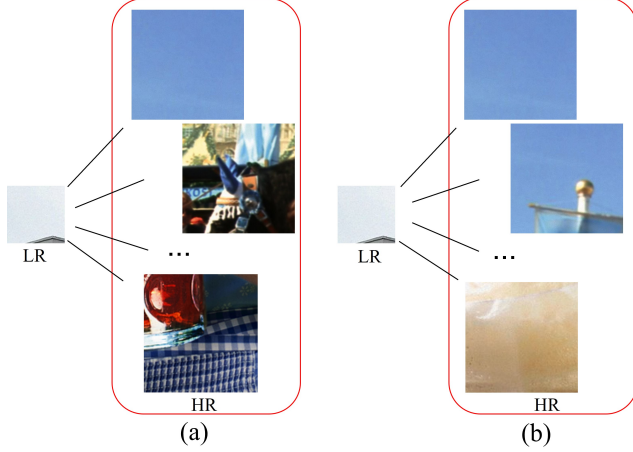


Figure 3. Exemplar diagram of our sampling method. An example LR patch has a low variance, where a large area is filled with the sky, and a roof of a building takes a tiny area at the bottom. Possible HR counterparts for the LR patch are shown in red rectangles. (a) independent sampling: There is no restriction for sampling an HR patch. (b) our method: an HR patch is sampled, of which variance is determined by the given LR patch.

image patch. Let σ_x^2 and σ_y^2 be the variance of LR patch x and HR patch y , respectively. After randomly sampling x , an HR patch y is retrieved that satisfying the following condition:

$$|\sigma_x^2 - \sigma_y^2| < \sigma_T^2 \quad (3)$$

where, σ_T^2 is a threshold value that controls the difference of variance between x and y . We call this *variance matching*.

Here, we used variance because the variance of an image patch is related to the content itself *e.g.* a patch with large variance has rich content [7]. By restricting the amount of difference of the variance between LR and HR patches, we can match the content density of those patches. Hence a network can learn from unpaired patches that have a similar level of content (Figure 3).

The variance matching at first helps the training of the degradation networks, F and D_F , by constraining the variance difference between HR and LR patches which are inputted F and D_F , respectively. So, they are trained with patches that have a similar level of content. Sequentially, the super-resolution networks (G , D_G) are also properly trained with \hat{x} that is indistinguishable from X (Figure 2). Again, our networks are trained in an end-to-end manner.

4. Experiments

We test our method on the NTIRE 2020 real-world SR dataset [11]. LR part of the training set is constructed by applying the degradation operation to the 2650 images of the Flickr2K [16] dataset. The degradation operation is undisclosed to the public. HR part of the training set is composed of the original 800 training images from DIV2K [15]

dataset. The validation set of the dataset consists of paired LR-HR images from the validation split of DIV2K, where the LR images are obtained by first downscaling the HR counterparts followed by the degradation. We did not use the test set of the dataset for the quantitative analysis since the ground truth of the set can not be accessed.

4.1. Implementation Details

The learning rates of G , F , D_G and D_F are initially set to 0.0001 and are halved at 100, 200, 400 and 700 epoch. For G , weight constants λ_G^{adv} , λ_G^{cyc} , λ_G^{per} and λ_G^{fea} are set to 0.3, 0.2, 0.5 and 20.0, respectively. For F λ_F^{adv} , λ_F^{con} , λ_F^{per} and λ_F^{fea} are set to 0.3, 0.5, 0.2 and 20.0, respectively. We used 64 for the variance threshold σ_T^2 . The size of input patch of G is 32^2 and that of F is 128^2 .

4.2. Sampling strategy

If we naively collect an HR patch after an LR patch is chosen, then it takes uneven times to sample an HR patch. This is because we have to search for an HR patch that satisfies the equation 3. We simply circumvent this issue and make the sampling time rather consistent. We first randomly collect a bunch of LR patches, namely N_{LR} patches, from an LR image. Similarly, N_{HR} HR patches are taken from an HR image. Then, the variance of each patch is calculated, and we finally take pairs that satisfy the equation 3 among $N_{LR} \times N_{HR}$ pairs. The law of large numbers inspires the idea of this simple sampling strategy. We used 30 for both N_{LR} and N_{HR} , which is larger than the batch size.

4.3. Quantitative Results

We compared our method with ESRGAN [16], ZSSR [14], K-ZSSR [1], CinC [17] and Ji *et al.* [7] in Table 1. Ours⁻ refers to the proposed method without the variance matching while Ours⁺ refers to the full implementation. Ji *et al.* took the first place in NTIRE 2020 Real-world SR challenge. Ours⁺ achieved the best SSIM score (\uparrow) which is 0.01 higher than Ji *et al.*, and the second best on the LPIPS [18] metric (\downarrow) which is 0.055 higher than the first ranker (Ji *et al.*).

It is noteworthy that, with the variance matching, the performance is considerably improved as Ours⁺ outperforms Ours⁻ at all three metrics in Table 1. This indicates that if cycle-GAN-like real-world SR methods [2, 10, 17] are equipped with the variance matching, their performance will be further improved. In addition, applying the variance matching to existing methods is very simple and effortless since it is only involved in the data sampling.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
ESRGAN [16]	19.06	0.2423	0.7552
ZSSR [14]	25.13	0.6268	0.6160
K-ZSSR [1]	18.46	0.3826	0.7307
CinC [17]	24.05	0.6583	0.4593
Ji <i>et al.</i> [7]	24.82	0.6619	0.2270
Ours ⁻	23.01	0.6389	0.3183
Ours ⁺	24.30	0.6731	0.2824

Table 1. Quantitative results on the validation set compared with ESRGAN, ZSSR, K-ZSSR, CinC and Ji *et al.* Ours⁻ refers to the proposed method without the variance matching while Ours⁺ refers to the full implementation.

σ_T^2	SSIM \uparrow	LPIPS \downarrow
N/A	0.6389	0.3183
576	0.6417	0.3021
256	0.6545	0.3115
100	0.6632	0.2989
64	0.6731	0.2824
36	0.6698	0.2822

Table 2. SSIM and LPIPS against the threshold σ_T^2 .

4.4. Ablation study

4.4.1 σ_T^2 of the variance matching

Here, we conducted an ablation study of the threshold value σ_T^2 . We reported the SSIM and LPIPS scores while we adjusted σ_T^2 in table 2, where N/A means that no variance matching was applied (Ours⁻). As shown in the table 2, the performance of SSIM was increased if we applied the variance matching. $\sigma_T^2 = 64$ produced the best result. It is until $\sigma_T^2 = 64$, the performance of SSIM increases as σ_T^2 decreases. However, if we further decrease σ_T^2 , the performance of SSIM is be dropped, *i.e.*, SSIM is dropped by 0.003 at $\sigma_T^2 = 36$ compared to $\sigma_T^2 = 64$.

For the LPIPS [18] metric, $\sigma_T^2 = 36$ shows the most improved result with 0.2822. However, it is worth noting that the performance gap between the best ($\sigma_T^2 = 36$) and second-best ($\sigma_T^2 = 64$) is only 0.0002, while the best one requires more time for sampling. In addition, the LPIPS has relatively fluctuated if compared with the SSIM.

4.4.2 Using mean as an image statistic

We also considered the mean of a patch for the image statistics. Together with the constraint in equation 3, the mean of an image patch is also used:

$$|\sigma_x^2 - \sigma_y^2| < \sigma_T^2 \wedge |\mu_x - \mu_y| < \mu_T, \quad (4)$$

where μ_x and μ_y are the mean of LR patch x and HR patch y , respectively. μ_T is a threshold value.

However, we found that using constraint 4 barely improved the performance. The results are shown in Table

2, where $\sigma_T^2 = 64$ is fixed and μ_T is decreased at interval of 100¹. It, on occasion, worsen the performance, *i.e.* $\mu_T = 155$. At $\mu_T = 55$, SSIM is very slightly increased by 0.0005. It is worth noting that the sampling time of $\mu_T = 55$ is much slower than not using it. It takes about $\times 2.5$ more time for an iteration due to the added constraint.

In addition, using the mean only for the image statistic was also examined. However, we did not observe any performance improvement. This is because that using only mean does not guarantee that the matched patches have a similar level of content. Therefore, we mainly use variance matching (Eq. 3).

μ_T	SSIM \uparrow
N/A	0.6731
155	0.6729
55	0.6736

Table 3. SSIM score against the threshold μ_T . σ_T^2 is fixed to 64.

4.5. Qualitative Results

In this section, we show the qualitative results. In Figure 4, the results of image translation using our method are shown. As shown in the left side of Figure 4, we compare the half of real LR image x and the half of generated LR image \hat{x} side-by-side. \hat{x} , the degraded result of y by $F(y)$, successfully mimics the real noises. Namely, it is hard to distinguish \hat{x} from x . On the right side of Figure 4, we also presented the side-by-side comparison between HR image y and super-resolved image \hat{y} . \hat{y} is the SR result generated from a real-world image x , *i.e.* $G(x)$.

In Figure 5, SR results of our method are depicted in the middle column. Our method produces natural texture, *e.g.* animal fur (the first row), as well as removes noise artifacts, *e.g.* sky (the second row). Lastly, visual comparisons among CinC [17], Ji *et al.* [7] and ours are presented in Figure 6.

5. Conclusion

In this study, we proposed a simple and efficient unpaired real-world SR method that utilized the variance of image patches so that it can match the content density between LR and HR pairs. An HR patch y is found after LR patch x is randomly sampled while it satisfies the variance matching condition. With this strategy, we stably train an unpaired SR network, which can produce real-world SR results with better performance. We test our method on the NTIRE 2020 real-world SR dataset. Our method outperforms the current state-of-the-art method in terms of the SSIM metric.

¹The pixel values are ranged in $[0, 255]$. Thus, $\mu_T = 255$ does not differ from N/A.



Figure 4. Qualitative comparisons between real LR image x and the degraded image $\hat{x} = F(y)$ (right), and between HR image y and the super-resolved image $\hat{y} = G(x)$ (left).

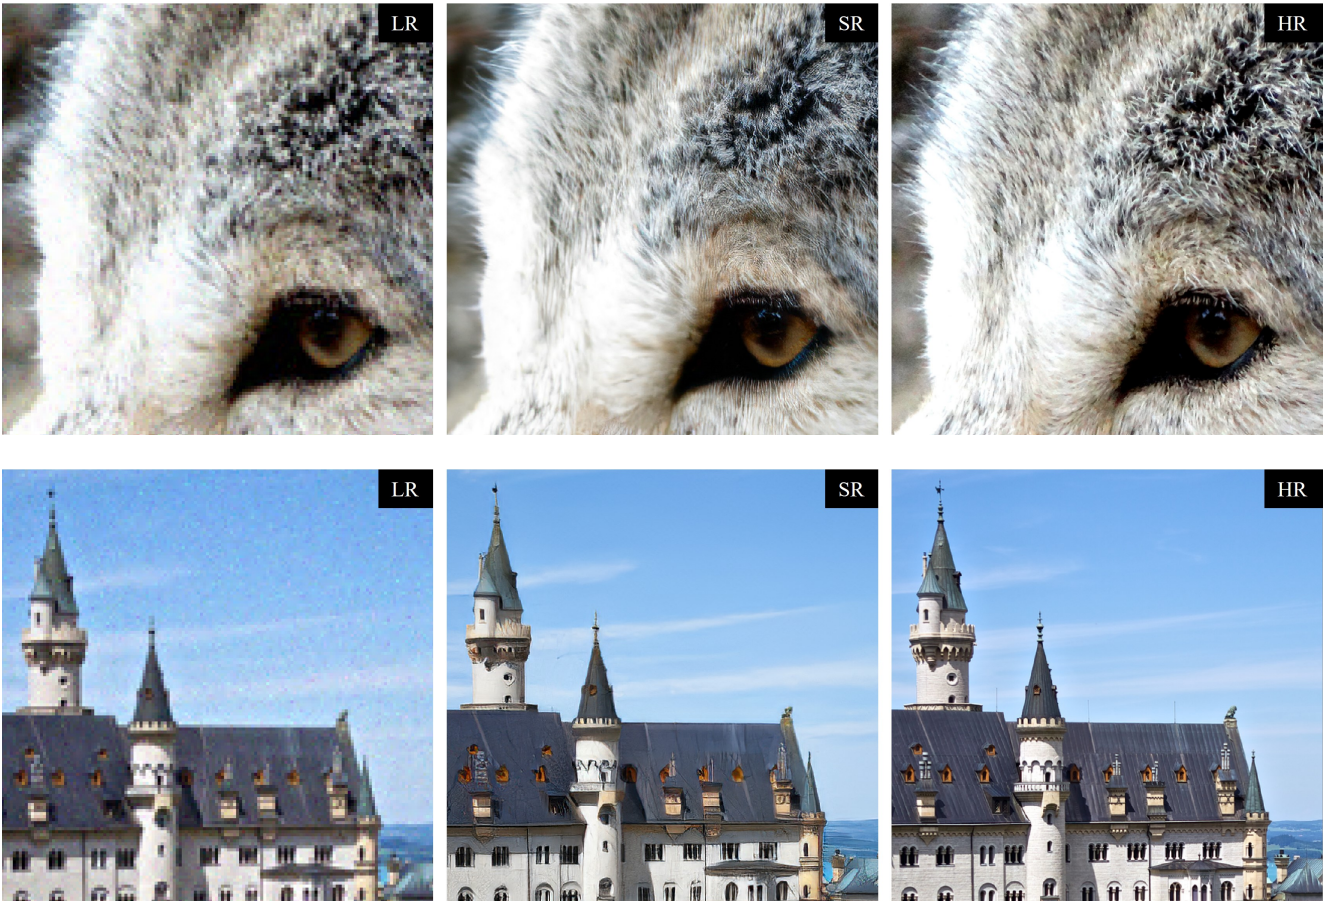


Figure 5. SR results of our method (middle column). LR (left column) and HR (right column) are also shown.

References

- [1] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. Blind super-resolution kernel estimation using an internal-gan. In *Advances in Neural Information Processing Systems*, pages 284–293, 2019. 4, 5
- [2] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do



Figure 6. Visualization comparison among CinC, Ji *et al.*, and our method with the variance matching. The red areas in LR images are zoomed in for the purpose of easier comparison.

- image degradation first. In *Proceedings of the European conference on computer vision (ECCV)*, pages 185–200, 2018. [1](#), [2](#), [4](#)
- [3] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. [1](#), [2](#)
- [4] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. [1](#)
- [5] Manuel Fritsche, Shuhang Gu, and Radu Timofte. Frequency separation for real-world super-resolution. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3599–3608. IEEE, 2019. [2](#)
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27:2672–2680, 2014. [2](#), [3](#)
- [7] Xiaozhong Ji, Yun Cao, Ying Tai, Chengjie Wang, Jilin Li, and Feiyue Huang. Real-world super-resolution via kernel estimation and noise injection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 466–467, 2020. [2](#), [4](#), [5](#)
- [8] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [2](#), [3](#)
- [9] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. [1](#), [2](#)
- [10] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Un-supervised learning for real-world super-resolution. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3408–3416. IEEE, 2019. [1](#), [2](#), [4](#)
- [11] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Ntire 2020 challenge on real-world image super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 494–495, 2020. [1](#), [4](#)
- [12] Xudong Mao, Q Li, H Xie, RYK Lau, and Z Wang. Least squares generative adversarial networks. cite. *arXiv preprint arxiv:1611.04076*, 4, 2016. [3](#)
- [13] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016. [3](#)

- [14] Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3118–3126, 2018. 4, 5
- [15] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 114–125, 2017. 4
- [16] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 1, 2, 3, 4, 5
- [17] Yuan Yuan, Siyuan Liu, Jiawei Zhang, Yongbing Zhang, Chao Dong, and Liang Lin. Unsupervised image super-resolution using cycle-in-cycle generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 701–710, 2018. 1, 2, 4, 5
- [18] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 2, 4, 5
- [19] Xuaner Zhang, Qifeng Chen, Ren Ng, and Vladlen Koltun. Zoom to learn, learn to zoom. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3762–3770, 2019. 1, 2
- [20] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 286–301, 2018. 1, 2
- [21] 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*, pages 2223–2232, 2017. 1, 2, 3