

On the Effectiveness of Spectral Discriminators for Perceptual Quality Improvement

Xin Luo Yunan Zhu Shunxin Xu Dong Liu
 University of Science and Technology of China, Hefei, China
 {xinluo, zhuyun, sxu}@mail.ustc.edu.cn, dongeliu@ustc.edu.cn
<https://github.com/Luciennnnnn/DualFormer>


Abstract

Several recent studies advocate the use of spectral discriminators, which evaluate the Fourier spectra of images for generative modeling. However, the effectiveness of the spectral discriminators is not well interpreted yet. We tackle this issue by examining the spectral discriminators in the context of perceptual image super-resolution (i.e., GAN-based SR), as SR image quality is susceptible to spectral changes. Our analyses reveal that the spectral discriminator indeed performs better than the ordinary (a.k.a. spatial) discriminator in identifying the differences in the high-frequency range; however, the spatial discriminator holds an advantage in the low-frequency range. Thus, we suggest that the spectral and spatial discriminators shall be used simultaneously. Moreover, we improve the spectral discriminators by first calculating the patch-wise Fourier spectrum and then aggregating the spectra by Transformer. We verify the effectiveness of the proposed method twofold. On the one hand, thanks to the additional spectral discriminator, our obtained SR images have their spectra better aligned to those of the real images, which leads to a better PD trade-off. On the other hand, our ensembled discriminator predicts the perceptual quality more accurately, as evidenced in the no-reference image quality assessment task.


1. Introduction

Generative Adversarial Networks (GANs) [19] have gained widespread adoption in low-level vision tasks, e.g., image super-resolution (SR) [31], which improves the perceptual quality of reconstructed images considerably. From

This work was supported by the Natural Science Foundation of China under Grants 62022075 and 62036005, and by the Fundamental Research Funds for the Central Universities under Grant WK3490000006. This work was also supported by the advanced computing resources provided by the Supercomputing Center of USTC. (Corresponding author: Dong Liu.)



(a) Spectra Statistics



(b) Ground Truth (c) Low-Quality (d) Super-Resolved

Figure 1: Denote the reduced spectrum $\tilde{S}(r)$ as the azimuthal average over the spectrum in normalized polar coordinates $r \in [0, 1]$. (a) shows spectra statistics $\mathbb{E}[\tilde{S}(r)]$ of Real-ESRNet [56] and Real-ESRGAN [56] on the DIV2K [1] dataset. It illustrates that Real-ESRGAN improves perceptual quality by producing more high-frequency information. (b), (c) and (d) present the average scores of Real-ESRGAN's discriminator w.r.t. three types of images. Given second-order degradation model [56] to generate low-resolution input, (c) are bicubic upsampled from low-resolution input, and (d) are generated by the generator, which takes low-resolution images as input. The extremely high score on (c) indicates that spatial discriminators may excessively favor high-frequency information, even if it is noise or artifacts.

the perspective of PD tradeoff [2], GANs improve perceptual quality by measuring the difference between image distributions (perception) rather than signal differences (distortion). In addition, Zhu *et al.* [73] found that trained discriminators can perform no-reference image quality assessment (NR-IQA). Therefore, the key to improving the perceptual quality in GAN-based SR is to enhance the discriminator's ability to measure the differences in image distribu-

tions, which can predict image quality more accurately.

Recently, many studies indicate that images produced by GAN have difficulty matching the spectral distribution of real data, especially in the high-frequency range [11, 13, 28, 12, 55, 3, 51]. Spectral discriminators [5, 26, 51] that utilize the Fourier spectrum as input have been shown to alleviate this problem. Nevertheless, the superiority of spectral discriminators over spatial discriminators remains unclear. In essence, the spectral discriminator measures the difference between the spectral distribution of real images and generated images. Thus, to understand the effectiveness of the spectral discriminator, a spectral viewpoint is indispensable.

In this paper, we investigate the spectral discriminators in the framework of GAN-based SR, as the quality of the super-resolved images is susceptible to spectral change¹. Before examining the spectral discriminator, we evaluate how the spatial discriminator behaves in SR. Fig. 1a shows the spectra statistics of Real-ESRNet [56] (distortion-oriented) and Real-ESRGAN [56] on the DIV2K [1] dataset, which illustrates that the introduction of the discriminator encourages the generator to produce more high-frequency components. In other words, the generator improves the perceptual quality by matching the spectra of real images². Moreover, we estimate the average score of Real-ESRGAN’s spatial discriminator on real, low-quality, and generated images (see Fig. 1b-Fig. 1d); despite the presence of noise and artifacts, low-quality images obtained the highest scores, indicating that the spatial discriminator may favor high-frequency data, even if it is incorrect. A natural question is whether the spectral discriminator’s effectiveness is related to its ability to compensate for this high-frequency flaw.

In an effort to answer the question, we further analyzed the spectra and observed that the spectra could be divided into three ranges, as revealed in Fig. 1a. We establish that it is a common phenomenon in image super-resolution algorithms from the frequency perspective of the PD trade-off, which encourages us to analyze the spatial/spectral discriminator by examining its robustness [49, 54, 47, 52] under frequency perturbations [54, 47, 52] within these three ranges. We observe that the spectral discriminator indeed performs better than the spatial discriminator in identifying the differences in the *high-frequency* range; however, the spatial discriminator has an advantage in the *low-frequency* range. Therefore, the spatial and spectral discriminators are complementary and are better to be used in combination.

Moreover, taking into account that the previous spectral discriminator [5, 26, 51, 14, 30] is an MLP with 1D

¹It is well known that downsampling leads to loss or aliasing of frequency components, so SR image quality is susceptible to spectral change.

²As the Fourier transform is linear and invertible, the difference between the distributions in the spatial domain and frequency domain is equivalent. Therefore, the closer the distributions are in the frequency domain, the closer they are in the spatial domain.

reduced spectrum as input, resulting in a loss of input information, we propose spectral Transformer, which calculate patch-wise Fourier transform and then aggregate spectra with Transformer [9]. Combined with spatial Transformer, our Dual Transformer (**DualFormer**) encourage the results of generator achieve better spectra alignment to real images, leading to an improved PD tradeoff. Additionally, we conducted NR-IQA tasks to demonstrate the alignment of our method with human perception. Our method achieved better performance on natural distortion datasets such as LIVE-itW [18] and KonIQ-10K [22], as well as on image-processing distortion datasets such as PIPAL [25], compared to the previous baseline [73].

2. Related Work

The frequency perspective of GAN. Many studies have examined GAN from a frequency perspective and have reached a consensus that there exist spectral discrepancies between the generator outputs and real images. While most studies attribute these discrepancies to the architecture of the generator [3, 11, 13, 15, 24, 28, 51], other works point to the discriminator [5, 26, 15, 51]. Concerning the generator, the frequency bias of the upsampling operations was revealed. Specifically, interpolation methods such as bilinear interpolation and nearest interpolation are resistant to generating high frequencies [3, 11, 51], whereas zero insertion between pixels introduces excessive high-frequency artifacts [3, 11, 51, 45]. Regarding the discriminator, Schwarz *et al.* [51] discovered that the spatial discriminator struggles with the low magnitude of the high-frequency content. They experimentally found that the spectral discriminator [26, 5] can reduce the spectral discrepancies, but the misalignment in the high-frequency is not fully addressed. Notably, image fidelity was improved by replacing the input of the spectral discriminator from the reduced spectrum to the full spectrum [51]. However, since the spectral discriminator is an MLP, which does not scale to the real-world setting, we solve this problem by introducing per-patch Fourier Transform, which enables compatibility between Transformer [9, 53] and frequency domain data.

GAN based image super-resolution. As one of the most important problems in image restoration [4, 63, 33, 64, 10, 65, 66], image super-resolution has seen rapid development [8, 33, 7, 72, 20, 35, 60, 29, 71, 31, 46] since the introduction of SRCNN [8]. Progress has been made on pixel-wise image reconstruction metrics, such as peak signal-to-noise ratio (PSNR). However, these metrics have been shown to have poor correlation with human perception [31, 2]. Thus, GAN [19] is leveraged to improve perceptual quality by matching reconstructed and real image distribution [31]. Following this, a lot of research have improved architecture [58, 56, 57] or training methods [50, 58, 69, 56, 41, 48, 40, 14]. Like our method, many

studies [14, 17] also utilize an additional spectral discriminator. However, they are focused on loss design for efficient SR, and we are driven by the effectiveness problem of the spectral discriminators. Besides answering the above question, we further improve the architecture of the spectral discriminator.

Opinion-unaware NR-IQA. Opinion-unaware NR-IQA (OU NR-IQA) aims to estimate the perceptual quality of images without access to human-labeled data. There are few studies on deep learning-based OU NR-IQA [37, 38, 73, 21, 43], and most of them are built on the idea of learning from ranking [16]. To illustrate, Liu *et al.* [37, 38] train a siamese Network [6] to rank images using synthetic distortions for which relative image quality is known. The most relevant work to us is Recycled [73], where the authors propose that a discriminator trained with adversarial loss can predict perceptual quality, and they show that the discriminator of ESRGAN [58] can act as a IQA predictor.

3. The Frequency Perspective of PD tradeoff

Fig. 1a shows that the spectra can be roughly divided into three ranges. We argue that it is a common phenomenon of the image SR algorithm through the lens of the frequency perspective of the PD tradeoff.

PD tradeoff [2] claims an image restoration (*e.g.*, SR) algorithm can be potentially improved only in terms of its perceptual quality or its distortion, one at the expense of the other. Specifically, the loss of an image restoration algorithm can be formulated as

$$\ell_{\text{gen}} = \lambda_1 \ell_{\text{distortion}} + \lambda_2 \ell_{\text{perception}}, \quad (1)$$

where $\ell_{\text{distortion}}$ is usually the L_1 loss measuring the per-pixel difference between original and reconstructed images, and $\ell_{\text{perception}}$ is the adversarial loss that measures how close the generator distribution and the real distribution are.

Given the spectral properties of the natural images, the spectra can be divided into three ranges from the frequency perspective of the PD tradeoff:

Low-frequency Range. As the power spectrum of natural images decays exponentially [51], the natural image mainly comprise low-frequency components. Therefore, the generator will have a high priority to recover low-frequency components since any deviation in this range would cause both high distortion and low perceptual quality. From Fig. 1a, both Real-ESRGAN and Real-ESRNet accurately match real spectra in this spectrum range.

Middle-frequency Range. The distortion metrics are dominated by the difference in the low-frequency range. As demonstrated in Fig. 1a, a generator without the adversarial loss, *e.g.*, Real-ESRNet, will sacrifice the ability to generate higher frequencies to focus on the low-frequency component, resulting in blurred output images. Conversely,

	Low	Middle	High
Real-ESRNet [56]	45.35	45.95	6.06
Real-ESRGAN [56]	56.10	37.82	4.79

Table 1: **Magnitude RMSE for each of the three frequency ranges.** The discriminator mitigates distortion in the middle and high-frequency ranges.

equipped with the adversarial loss, Real-ESRGAN is motivated to match real spectra in this range. Additionally, Tab. 1 shows Real-ESRGAN reaches higher distortion in the low-frequency range and lower distortion in higher frequency ranges compared to Real-ESRNet. In other words, higher perceptual quality is related to lower distortion in middle and high-frequency ranges.

High-frequency Range. The discriminator encourages the generator to generate more components in middle and high-frequency ranges. However, there are subtle yet critical differences between these two ranges. Specifically, the power spectrum of natural images drops suddenly at the highest frequency, resulting in a negligible impact of the distortion term in this range. This is clearly demonstrated by the catastrophic component loss of Real-ESRNet in the high-frequency range of Fig. 1a. Consequently, this range is primarily governed by the perception term. As shown in Fig. 1a, Real-ESRGAN generates random high-frequency details in this range guided by the adversarial loss.

4. Method and Analyses

To compare the spatial and spectral discriminators, a unified architecture that works for both domains is indispensable. But so far, such an architecture does not exist, making it unfair to compare different architectures. To solve this problem, we present our Spectral Transformer in this section, which uses the Transformer in the spectral domain by applying a per-patch Fourier Transform. Then, we analyze the differences between spatial and spectral discriminators from a frequency perspective. Furthermore, we investigate how patch size affects the Spectral Transformer and propose the Dual Transformer as a new discriminator.

4.1. Spectral Transformer

Several studies have demonstrated that it is possible to improve learning spectral statistics by adding a discriminator on the Fourier spectrum [51, 5, 26, 14, 30]. However, these studies take the reduced spectrum as input because their spectral discriminator adopts the MLP architecture. CNNs do not apply to spectrum data [51], and if the 2D Fourier spectrum is directly used as input to the MLP, the time and space complexity would be unaffordable. Schwarz *et al.* [51] showed that replacing the reduced spectrum with the full Fourier transform indeed improves the image fidelity in the low-resolution case. Thus, there is an urgent need for a new architecture that can process full

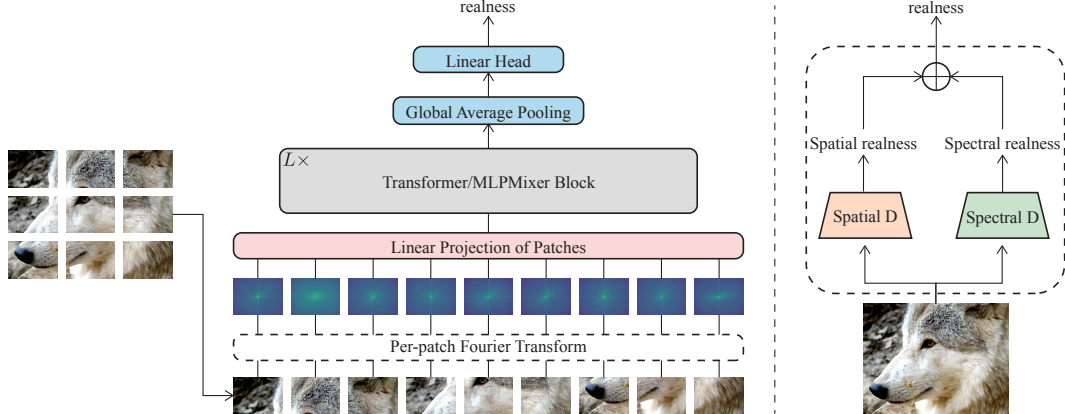


Figure 2: **Spectral Transformer**. After splitting the image into fixed-size patches, we apply patch-wise Fourier transform on each of them and aggregate the spectra by Transformer. The overall realness is the average of spatial realness and spectral realness.

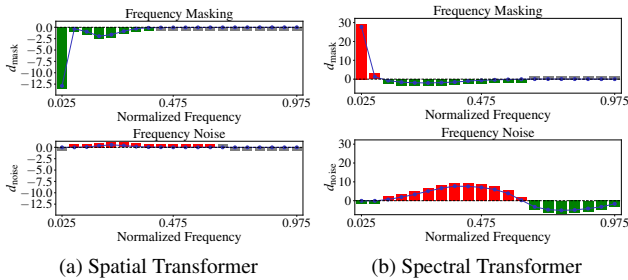


Figure 3: **Robustness behavior of discriminators under frequency perturbations**. The **positive** and **negative** values indicate the degree to which the discriminator considers the input image to be real or fake compared to the original image. Negligible differences are shown in **gray**. In summary, the spatial discriminator is an expert at discriminating low-frequency masking, while the spectral discriminator is better at discriminating high-frequency noise. Moreover, there is a tradeoff between the discriminator’s ability to determine frequency masking and noise in a specific range.

spectrum input efficiently and effectively.

We solve this problem with our Spectral Transformer (SpecFormer), which harnesses the Transformer by using a per-patch Fourier Transform. First, the image is cropped into patches in the spatial domain according to spatial continuity. Then, we apply Fourier Transform to each patch separately. Finally, these transformed patches are arranged in the original spatial order as input to Transformer [9] (we particularly call it Spatial Transformer, dubbed SpatFormer, to distinguish from our SpecFormer); see Fig. 2 for an intuitive explanation. Our approach has two benefits. Firstly, SpecFormer takes 2D full spectra, which are powerful and scalable. Secondly, the consistent architecture facilitates our analysis of the differences between the spatial and spectral discriminator in Sec. 4.2.

4.2. On the Robustness of Discriminator

The natural division of the spectra encourages us to analyze the discriminator by examining its robustness [49, 54,

47, 52] under frequency perturbations [54, 47] in the three spectra ranges. Specifically, we evaluate the robustness of the discriminator in the three ranges under frequency masking [54] and noise [47]. Here we focus on the SpatFormer and SpecFormer, they are same Transformer architecture with or without per-patch Fourier Transform. In addition, we have also studied the behavior of other discriminators and examined how the scaling factory of SR affects the three-range behavior of the discriminator, which further validates the rationality of our analysis method. Please refer to our supplementary for dedicated details.

Denote \mathbf{I}_{mask} and $\mathbf{I}_{\text{noise}}$ as the image altered by frequency masking and noise. They are defined as follows:

$$\begin{aligned} \mathbf{I}_{\text{mask}} &= \mathcal{F}^{-1} \left(\mathcal{F}(\mathbf{I}) \odot \overline{\mathbf{M}}_f(l, r) \right), \\ \mathbf{I}_{\text{noise}} &= \mathbf{I} + \mathcal{F}^{-1} \left(\mathcal{F}(\delta) \odot \mathbf{M}_f(l, r) \right), \end{aligned} \quad (2)$$

where \mathbf{I} is the clean image, δ is random noise. $\mathcal{F}(\cdot)$ and $\mathcal{F}^{-1}(\cdot)$ are Fourier transform and inverse Fourier transform, respectively. $\mathbf{M}_f(l, r)$ represents the frequency mask, which is 1 in the frequency interval $[l, r]$ and 0 elsewhere. $\overline{\mathbf{M}}_f$ is an element-wise logical inversion of \mathbf{M}_f . Namely, for components within the radius interval $[l, r]$ on the spectrum, we remove them or add noise to them. Then, we monitor the relative difference in the score between altered and real images, $d_{\text{mask}} = \mathbb{E}[D(\mathbf{I}_{\text{mask}})] - \mathbb{E}[D(\mathbf{I})]$ and $d_{\text{noise}} = \mathbb{E}[D(\mathbf{I}_{\text{noise}})] - \mathbb{E}[D(\mathbf{I})]$. These values indicate how true the discriminator considers it to be compared to the real image. For the purposes of the following analyses, it should be noted that frequency noise is a kind of small difference in a frequency range, while frequency mask corresponds to large deficiency in that range.

We expect the discriminator to always output negative values. However, as shown in Fig. 3, both the spatial and spectral discriminators fail in some ranges. In particular, the spectral discriminator exhibits a behavior mirror to that of the generator, characterized by three distinct ranges. This phenomenon can be understood from the frequency per-

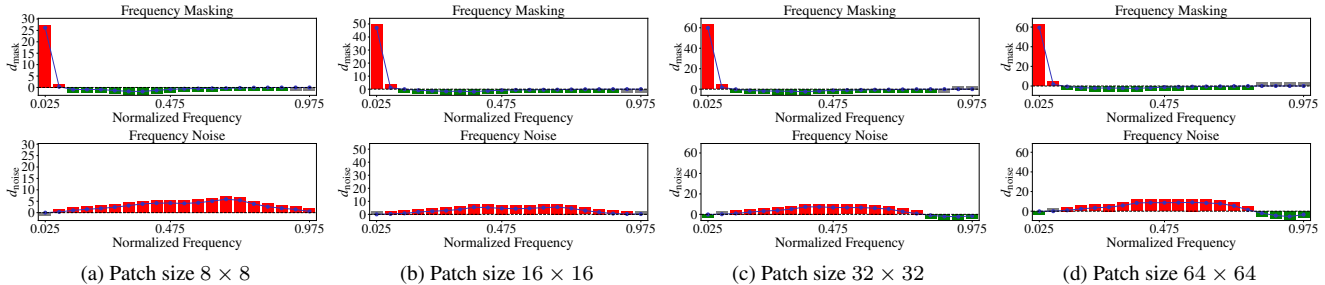


Figure 4: **Influence of patch size on the robustness behavior of spectral discriminators.** The ability that can discriminate high-frequency noises emerges at a relatively large patch size *i.e.* 32×32 . Consequently, for the Spectral Transformer, a larger patch size can be selected compared to the Spatial Transformer, achieving both effectiveness and computational efficiency.

spective of PD tradeoff. In the **low-frequency range**, the generated images closely resemble the ground truth images with only subtle differences. Consequently, the discriminator learns to discriminate subtle differences (the case of frequency noise), while it cannot discriminate frequency masking. Regarding the **middle-frequency range**, the generated images exhibit a deficiency state (lower magnitude than real images in frequency). Since compensating for this deficiency optimizes both the distortion and perceptual quality, the discriminator learns to discriminate this deficiency (corresponds to the frequency masking) in this range while not recognizing the frequency noise. In the **high-frequency range**, the distortion term barely affects the image, and the generator tricks the discriminator by generating high-frequency “noise.” Thus, the discriminator only learns to discriminate frequency noise. In conclusion, the discriminator learns based on what the generator produces, and due to the fact generator has different generation inclinations in each range, the discriminator also learns different abilities in each range. It is noteworthy that the spatial discriminator does not directly operate in the spectral domain, thereby we cannot analyze it using the same approach.

Finally, as shown in Fig. 3a-Fig. 3b, the application of the same architecture to spatial and frequency domains separately results in significant differences, which illustrates that *the spatial discriminator is an expert at discriminating deficiency in the low-frequency range, while the spectral discriminator is better at discriminating noise in the high-frequency range.* Thus, the spatial and spectral discriminators are complementary and should be used in combination.

4.3. How patch size affects Transformers

The choice of patch size has a crucial impact on the efficiency and effectiveness of Transformers. While small patch sizes may improve performance of SpatFormers, it results in lower efficiency [9]. However, this may not hold true for SpecFormers, since applying Fourier transform on small patches may not be meaningful. To verify this, we investigate the effect of patch size on the robustness behavior of SpecFormer. As shown in Fig. 4, the SpecFormer requires a patch size of at least 32×32 to effectively distin-

guish high-frequency noise. While we conducted a similar analysis for SpatFormer, its robustness behavior is not sensitive to patch size. In light of above observations, we can select a larger patch size for the SpecFormer, which yields better performance and is also more efficient.

4.4. Dual Transformer

Based on our previous analysis, we have found that the Spatial discriminator and Spectral discriminator possess complementary discriminative power. Therefore, we propose a novel discriminator called Dual Transformer (DualFormer) by combining the Spatial Transformer and Spectral Transformer. In particular, considering that the Spectral Transformer exhibits a preference for larger patch sizes, we use a Spectral Transformer with a minimum patch size of 32×32 in our work. For implementation, we utilize ViT [9] as the basis, adopt relative position encoding, and apply global average pooling for realness prediction. Our configuration is lightweight, with only 10 blocks and a hidden layer dimension of 96. We have tested heavier networks, but they did not exhibit any improvement in performance.

As we use two Transformers as discriminators, the efficiency of our method may become a concern. In fact, our discriminator is more efficient than commonly used VGG [58] and U-Net [56] (in terms of parameters, FLOPs and activations, see supplementary for details), since we use large patch sizes and small layer dimension.

5. Experimental Results

We validate the proposed method from two perspectives. Firstly, we verify whether our method can help the generator to better match spectra and achieve better SR performance. Secondly, we consider whether our method improves the discriminator’s ability to predict image quality from the perspective of OU NR-IQA [73].

5.1. Image Super-Resolution

Experimental Setup. To construct our model, we select ESRGAN [31] as the baseline and replace its discriminator with ours. The patch size of both SpatFormer and Spec-

Dataset	Metric	SRGAN [31]	RankSRGAN [69]	ESRGAN [58]	SPSR [40]	ESRGAN + LDL [34]	FFTGAN [14]	Ours
BSD100	↑PSNR	25.5544	25.5315	25.3726	25.7170	25.5685	25.7900	26.4826
	↑SSIM	0.6542	0.6518	0.6515	0.6658	0.6616	0.6580	0.6895
	↓LPIPS	0.1783	0.1772	0.1597	0.1558	0.1592	0.1580	0.1573
DIV2K	↑PSNR	28.1646	28.0916	28.0465	28.3978	28.2440	28.6300	29.3049
	↑SSIM	0.7745	0.0646	0.7669	0.7821	0.7758	0.7800	0.8023
	↓LPIPS	0.1257	0.1239	0.1142	0.1069	0.1133	0.1130	0.1030
Urban100	↑PSNR	24.4056	24.4233	24.6287	24.8393	24.4777	25.0500	25.6870
	↑SSIM	0.7298	0.7265	0.7411	0.7493	0.7389	0.7380	0.7727
	↓LPIPS	0.1439	0.1438	0.1240	0.1174	0.1243	0.1200	0.1147

Table 2: Quantitative comparison of GAN-based SR methods on $\times 4$ super-resolution.




Figure 5: Visual comparison of GAN-based SR methods on $\times 4$ super-resolution.

Former are set to 32×32 . The training settings are aligned with ESRGAN. We adopt DF2K [1, 36] as the training set. The training pairs are generated by bicubic downsampling with a scaling factor of 4. The HR patch size is set to 128, and the total batch size is set to 16. Networks are trained with 1000k using only L_1 loss and then 400k iterations with a combination of L_1 loss, perceptual loss [56], and GAN loss. The loss weights are kept same as ESRGAN.

Test sets and metrics. We evaluate our methods on three datasets: BSD100 [44], DIV2K validation set [1], and Urban100 [23]. We used PSNR and SSIM [59] as distortion metrics and LPIPS [68] for evaluating perceptual quality.

Comparison with state-of-the-arts. To validate the effectiveness of our method, besides ESRGAN, we compare it with several state-of-the-art (SOTA) GAN-based SR methods: RankSRGAN [69], SPSR [40], ESRGAN + LDL [34], FFTGAN (ESRGAN version) [14], where RankSRGAN is

constructed from SRGAN [31], other methods are based on ESRGAN. SPSR improves the architecture of the generator, while other studies, including ours, use the original generator of SRGAN/ESRGAN. For all methods except FFTGAN, we retrain the model using the code provided by the authors to keep the training settings the same as their basis, *i.e.*, SRGAN/ESRGAN. Since the code of FFTGAN is not publicly available, we take the results directly from their paper (notice that FFTGAN uses a large patch size of 256×256). Tab. 2 shows the quantitative comparisons. Overall, our method achieves the best distortion (PSNR/SSIM) and perceptual quality (LPIPS) compared to other SOTA methods on all benchmarks. Our discriminator improves the performance of ESRGAN significantly due to its better discriminating ability. Fig. 5 shows the visual comparison, demonstrating that our method produces cleaner results with fewer artifacts than other methods since the spectral discriminator




Figure 6: Visual comparisons of various combinations (generator/spatial discriminator/spectral discriminator) on $\times 4$ super-resolution.




Figure 7: Our method exhibits better spectra alignment.

Method	\uparrow PSNR	\uparrow SSIM	\downarrow LPIPS
ESRGAN	25.7327	0.7144	0.2161
ESRGAN + LDL	25.6998	0.7097	0.2215
SPSR	25.8033	0.7137	0.2160
RankSRGAN	25.7888	0.6953	0.2488
BebyGAN [32]	26.3378	0.7270	0.2077
Ours (ESRGAN version)	26.8768	0.7451	0.2039
Ours (BebyGAN version)	27.4823	0.7583	0.1993

Table 3: Comparison under hard gated degradation (DIV2K).

can identify high-frequency noise effectively.

Complex Degradation. In the preceding, we demonstrated the effectiveness of the spectral discriminator on bicubic degradation. However, as bicubic degradation is quite simple, ESRGAN can already match the real spectra well (see supplementary for details). To see the further potential of our method, we examine whether our discriminator can improve SRGAN models under a more complex degradation model (the hard gated degradation model [70]). In addition to ESRGAN, we also investigate whether our DualFormer can improve BebyGAN [32] in this case. From Tab. 3 it can be seen that our method once again shows superior performance, while achieving the best trade-off between distortion and perceptual quality.

Ablation study. We conduct ablation studies to investigate the role of the Spectral Transformer by selecting two representative networks, *i.e.*, RRDB [58, 56] and SwinIR [33], and evaluating the influence of various combinations of spatial and spectral discriminators. The experiments used a batch size of 32 and a patch size of 192, and were trained for 50,000 iterations under hard gated degra-

G	Spatial D	Spectral D	\uparrow PSNR	\uparrow SSIM	\downarrow LPIPS
RRDB	VGG	-	27.0408	0.7618	0.2359
RRDB	VGG	1D MLP [26]	26.4706	0.7448	0.2317
RRDB	VGG	SpecFormer	26.8322	0.7575	0.2271
RRDB	SpatFormer	-	26.9438	0.7553	0.2326
RRDB	SpatFormer	1D MLP	27.0769	0.7619	0.2493
RRDB	SpatFormer	SpecFormer	27.1460	0.7590	0.2284
RRDB	U-Net	-	26.9435	0.7585	0.2341
RRDB	U-Net	1D MLP	25.9467	0.7582	0.2358
RRDB	U-Net	SpecFormer	26.2829	0.7547	0.2258
RRDB	-	1D MLP	26.4098	0.7549	0.2331
RRDB	-	SpecFormer	26.4157	0.7544	0.2319
SwinIR	U-Net	-	27.2959	0.7593	0.2387
SwinIR	U-Net	1D MLP	26.5903	0.7611	0.2447
SwinIR	U-Net	SpecFormer	26.6090	0.7529	0.2216

Table 4: Quantitative comparison of various combinations on $\times 4$ super-resolution. Metrics are evaluated on DIV2K validation dataset.

ation. Tab. 4 and Fig. 6 showed that the Spectral Transformer helped improve the visual quality of the generators. Interestingly, when the Spatial Transformer served as the Spatial Discriminator, the Spectral Transformer not only improved visual quality but also helped reduce distortion. This phenomenon did not occur with other spatial discriminators. These findings suggest that the Spatial Transformer and Spectral Transformer complement each other well. Additionally, we plotted the spectrum profile of the RRDB with different discriminator combinations in Fig. 7, demonstrating that spectral discriminators promote high-frequency components and our SpecFormer aligns spectra better.

5.2. No-Reference Image Quality Assessment

Experimental Setup. Following RecycledD [73], we choose DF2K [1, 36] and OutdoorSceneTraining [57] as the training dataset. The generator is first trained with the L_1 loss for 500 iterations, and then the discriminator is introduced to train another 50k iterations. The HR patch size is set to 192, and the total batch size is set to 16. We report the result of the best-performing checkpoint for each method. For all experiments, the low-resolution images are generated by the simple gated degradation Model [70]. The

Method	LIVE-itW			KonIQ-10k		
	↑PLCC	↑SRCC	↑KRCC	↑PLCC	↑SRCC	↑KRCC
QAC [62]	0.2720	0.0457	0.0370	0.3719	0.3397	0.2302
LPSI [61]	0.2877	0.0834	0.0524	0.2066	0.2239	0.1504
IL-NIQE [67]	0.5039	0.4394	0.2985	0.5316	0.5057	0.3504
SNP-NIQE [39]	0.5201	0.4657	0.3162	0.6340	0.6285	0.4435
dipIQ [42]	0.3180	0.1774	0.1207	0.4429	0.2375	0.1594
RankIQ [37]	0.4528	0.4307	0.2945	0.5028	0.4983	0.3448
RecycleD [73]	-	-	-	0.6105	0.6020	0.4201
Ours	0.5068	0.4897	0.3312	0.6543	0.6321	0.4435

Table 5: Quantitative comparison on NR-IQA.

patch sizes of SpatFormer and SpecFormer are set to 8×8 and 64×64 respectively.

Test sets and metrics. We evaluate our spectral discriminator on three representative IQA datasets, *i.e.* LIVE-itW [18], KonIQ-10k [22] and PIPAL [25]. LIVE-itW and KonIQ-10k contain diverse, authentic distortions, while PIPAL constitutes with many image processing artifacts including the results of perceptual-oriented algorithms. We report three widely used metrics: Pearson linear correlation coefficient (PLCC), Spearman rank order correlation coefficient (SRCC), and Kendall rank order correlation coefficient [27] (KRCC).

Spatial D	Spectral D	Traditional SR	PSNR based SR	GAN based SR	Full
VGG [73]	-	0.3902	0.4630	0.0772	0.2842
VGG	1D MLP [26]	0.2692	0.3444	0.0488	0.2146
VGG	SpecFormer	0.3841	0.4609	0.0636	0.2889
U-Net	-	-0.0048	0.0354	0.0035	0.0162
U-Net	1D MLP	0.0144	0.0379	0.0189	0.0477
U-Net	SpecFormer	0.3700	0.4391	0.0598	0.2936
-	1D MLP	0.0368	0.0772	0.0301	0.0423
-	SpecFormer	0.3626	0.4346	0.0538	0.2883
SpatFormer	-	0.3689	0.4196	0.0640	0.2664
SpatFormer	1D MLP	0.2636	0.373	0.0364	0.2301
SpatFormer	SpecFormer	0.3578	0.4285	0.0647	0.2787

Table 6: Comparison (↑SRCC) on sub-types of PIPAL training set.

Comparison with state-of-the-arts. We compare our method to representative OU NR-IQA methods: QAC [62], LPSI [61], IL-NIQE [67], SNP-NIQE [39], dipIQ [42], RankIQ [37], and RecycleD [73]. Where QAC, LPSI, IL-NIQE, SNP-NIQE are conventional methods, dipIQ, RankIQ, RecycleD are neural network-based methods. Our method is directly based RecycleD, with an extra spectral discriminator. Notably, RecycleD uses sophisticated weighting strategy to get better results, we disable it for simplicity as our target is not performance but to validate the effectiveness of the spectral discriminator to predict perceptual quality. Tab. 5 shows the results, our method reached the best results on two datasets, which verified our ensemble discriminator could predict perceptual quality better than the spatial discriminator used in RecycleD.

Ablation study. We conducted experiments to investi-

Spatial D	Spectral D	↑PLCC	↑SRCC	↑KRCC
VGG [73]	-	0.6240	0.6034	0.4200
VGG	1D MLP [26]	0.5215	0.5013	0.3434
VGG	SpecFormer	0.6543	0.6321	0.4434
U-Net	-	0.4144	0.3458	0.2328
U-Net	1D MLP	0.3831	0.3021	0.2016
U-Net	SpecFormer	0.6459	0.6207	0.4355
-	1D MLP	0.2863	-0.2339	-0.1568
-	SpecFormer	0.6357	0.6094	0.4261
SpatFormer	-	0.6372	0.6124	0.4275
SpatFormer	1D MLP	0.575	0.5566	0.3847
SpatFormer	SpecFormer	0.6394	0.6220	0.4360

Table 7: Quantitative comparison of various combinations on NR-IQA. Metrics are evaluated on KonIQ-10k dataset.

gate the impact of SpecFormer on various spatial discriminators. As shown in Tab. 6, in most cases, using our SpecFormer achieved the best results. Specifically, when the spatial discriminator was VGG, SpecFormer improved the performance on the complete PIPAL dataset, but showed a slight decrease in performance on several SR-related subsets. Additionally, all combinations performed poorly in GAN-based SR, which is consistent with the results of Zhu *et al.* [73]. Results on the KonIQ-10K dataset are shown in Tab. 7, and our method achieved the best performance in all cases. Among them, VGG+SpecFormer achieved the best results, which may be because VGG can discriminate the widest range of frequency masking, while SpecFormer compensates for its inability to discriminate high-frequency noise (please refer to the supplementary for details).

6. Discussion and Limitations

With the introduction of an additional spectral discriminator, our obtained SR images have their spectra better aligned to those of the real images, thereby enhancing the overall perceptual quality of our methodology. Despite these benefits, we notice that better-aligned spectra may not always result in better perceptual quality. Further, the current method trains the spatial and spectral discriminator separately, increasing the computational and storage overhead. These issues indicate that there is still much room for improvement of the spectral discriminators.

7. Conclusion

This paper investigates the effectiveness of spectral discriminators. Our research reveals that spatial and spectral discriminators offer unique benefits, and they work best when being used together. We also introduce a per-patch Fourier Transform to improve the spectral discriminator’s structure. Our extensive experiments on image SR and NR-IQA confirm the effectiveness of our approach. Specifically, our method improves the alignment of spectra and the PD tradeoff in SR, as well as the IQA ability in NR-IQA.

References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, 2017. 1, 2, 6, 7
- [2] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *CVPR*, 2018. 1, 2, 3
- [3] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, and Ngai-Man Cheung. A closer look at fourier spectrum discrepancies for CNN-generated images detection. In *CVPR*, 2021. 2
- [4] Haoyu Chen, Jinjin Gu, Yihao Liu, Salma Abdel Magid, Chao Dong, Qiong Wang, Hanspeter Pfister, and Lei Zhu. Masked image training for generalizable deep image denoising. In *CVPR*, 2023. 2
- [5] Yuanqi Chen, Ge Li, Cece Jin, Shan Liu, and Thomas Li. SSD-GAN: Measuring the realness in the spatial and spectral domains. In *AAAI*, 2021. 2, 3
- [6] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. 3
- [7] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *CVPR*, 2019. 2
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE TPAMI*, 38(2):295–307, 2015. 2
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2, 4, 5
- [10] Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Burst image restoration and enhancement. In *CVPR*, 2022. 2
- [11] Ricard Durall, Margret Keuper, and Janis Keuper. Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions. In *CVPR*, 2020. 2
- [12] Tarik Dzanic, Karan Shah, and Freddie Witherden. Fourier spectrum discrepancies in deep network generated images. In *NeurIPS*, 2020. 2
- [13] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *ICML*, 2020. 2
- [14] Dario Fuoli, Luc Van Gool, and Radu Timofte. Fourier space losses for efficient perceptual image super-resolution. In *ICCV*, 2021. 2, 3, 6
- [15] Rinon Gal, Dana Cohen Hochberg, Amit Bermano, and Daniel Cohen-Or. SWAGAN: A style-based wavelet-driven generative model. *ACM TOG*, 40(4):1–11, 2021. 2
- [16] Fei Gao, Dacheng Tao, Xinbo Gao, and Xuelong Li. Learning to rank for blind image quality assessment. *IEEE TNNLS*, 26(10):2275–2290, 2015. 3
- [17] Anaïs Gastineau, Jean-François Aujol, Yannick Berthoumieu, and Christian Germain. Generative adversarial network for pansharpening with spectral and spatial discriminators. *IEEE TGARS*, 60:1–11, 2021. 3
- [18] Deepti Ghadiyaram and Alan C Bovik. Massive online crowdsourced study of subjective and objective picture quality. *IEEE TIP*, 25(1):372–387, 2015. 2, 8
- [19] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1, 2
- [20] Jinjin Gu and Chao Dong. Interpreting super-resolution networks with local attribution maps. In *CVPR*, 2021. 2
- [21] Jie Gu, Gaofeng Meng, Cheng Da, Shiming Xiang, and Chunhong Pan. No-reference image quality assessment with reinforcement recursive list-wise ranking. In *AAAI*, 2019. 3
- [22] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE TIP*, 29:4041–4056, 2020. 2, 8
- [23] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *CVPR*, 2015. 6
- [24] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *ICCV*, 2021. 2
- [25] Gu Jinjin, Cai Haoming, Chen Haoyu, Ye Xiaoxing, Jimmy S Ren, and Dong Chao. PIPAL: a large-scale image quality assessment dataset for perceptual image restoration. In *ECCV*, 2020. 2, 8
- [26] Steffen Jung and Margret Keuper. Spectral distribution aware image generation. In *AAAI*, 2021. 2, 3, 7, 8
- [27] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938. 8
- [28] Mahyar Khayatkhoei and Ahmed Elgammal. Spatial frequency bias in convolutional generative adversarial networks. In *AAAI*, 2022. 2
- [29] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, 2016. 2
- [30] Nahyun Kim, Donggon Jang, Sunhyeok Lee, Bomi Kim, and Dae-Shik Kim. Unsupervised image denoising with frequency domain knowledge. In *BMVC*, 2021. 2, 3
- [31] 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 *CVPR*, 2017. 1, 2, 5, 6
- [32] Wenbo Li, Kun Zhou, Lu Qi, Liying Lu, and Jiangbo Lu. Best-buddy GANs for highly detailed image super-resolution. In *AAAI*, 2022. 7
- [33] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image restoration using swin transformer. In *ICCV*, 2021. 2, 7
- [34] Jie Liang, Hui Zeng, and Lei Zhang. Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In *CVPR*, 2022. 6

- [35] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, 2017. 2
- [36] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, 2017. 6, 7
- [37] Xialei Liu, Joost Van De Weijer, and Andrew D Bagdanov. Rankiqa: Learning from rankings for no-reference image quality assessment. In *ICCV*, 2017. 3, 8
- [38] Xialei Liu, Joost Van De Weijer, and Andrew D Bagdanov. Exploiting unlabeled data in CNNs by self-supervised learning to rank. *IEEE TPAMI*, 41(8):1862–1878, 2019. 3
- [39] Yutao Liu, Ke Gu, Yongbing Zhang, Xiu Li, Guangtao Zhai, Debin Zhao, and Wen Gao. Unsupervised blind image quality evaluation via statistical measurements of structure, naturalness, and perception. *IEEE TCSVT*, 30(4):929–943, 2019. 8
- [40] Cheng Ma, Yongming Rao, Yean Cheng, Ce Chen, Jiwen Lu, and Jie Zhou. Structure-preserving super resolution with gradient guidance. In *CVPR*, 2020. 2, 6
- [41] Haichuan Ma, Dong Liu, and Feng Wu. Rectified wasserstein generative adversarial networks for perceptual image restoration. *IEEE TPAMI*, 2022. 2
- [42] Kede Ma, Wentao Liu, Tongliang Liu, Zhou Wang, and Dacheng Tao. dipIQ: Blind image quality assessment by learning-to-rank discriminable image pairs. *IEEE TIP*, 26(8):3951–3964, 2017. 8
- [43] Pavan C Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C Bovik. Image quality assessment using contrastive learning. *IEEE TIP*, 31:4149–4161, 2022. 3
- [44] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, 2001. 6
- [45] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 1(10):e3, 2016. 2
- [46] Zhihong Pan, Baopu Li, Dongliang He, Mingde Yao, Wenhao Wu, Tianwei Lin, Xin Li, and Errui Ding. Towards bidirectional arbitrary image rescaling: Joint optimization and cycle idempotence. In *CVPR*, 2022. 2
- [47] Namuk Park and Songkuk Kim. How do vision transformers work? In *ICLR*, 2022. 2, 4
- [48] Seong-Jin Park, Hyeongseok Son, Sunghyun Cho, Ki-Sang Hong, and Seungyong Lee. SRFeat: Single image super-resolution with feature discrimination. In *ECCV*, 2018. 2
- [49] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *ICML*, 2019. 2, 4
- [50] Mehdi SM Sajjadi, Bernhard Scholkopf, and Michael Hirsch. Enhancenet: Single image super-resolution through automated texture synthesis. In *ICCV*, 2017. 2
- [51] Katja Schwarz, Yiyi Liao, and Andreas Geiger. On the frequency bias of generative models. In *NeurIPS*, 2021. 2, 3
- [52] Yash Sharma, Gavin Weiguang Ding, and Marcus Brubaker. On the effectiveness of low frequency perturbations. In *IJ-CAI*, 2019. 2, 4
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2
- [54] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *CVPR*, 2020. 2, 4
- [55] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. CNN-generated images are surprisingly easy to spot... for now. In *CVPR*, 2020. 2
- [56] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-ESRGAN: Training real-world blind super-resolution with pure synthetic data. In *ICCVW*, 2021. 1, 2, 3, 5, 6, 7
- [57] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *CVPR*, 2018. 2, 7
- [58] 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 *ECCVW*, 2018. 2, 3, 5, 6, 7
- [59] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 6
- [60] Zhihao Wang, Jian Chen, and Steven CH Hoi. Deep learning for image super-resolution: A survey. *IEEE TPAMI*, 43(10):3365–3387, 2020. 2
- [61] Qingbo Wu, Zhou Wang, and Hongliang Li. A highly efficient method for blind image quality assessment. In *ICIP*, 2015. 8
- [62] Wufeng Xue, Lei Zhang, and Xuanqin Mou. Learning without human scores for blind image quality assessment. In *CVPR*, 2013. 8
- [63] Mingde Yao, Dongliang He, Xin Li, Fu Li, and Zhiwei Xiong. Towards interactive self-supervised denoising. *IEEE TCSVT*, 2023. 2
- [64] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022. 2
- [65] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. In *ECCV*, 2020. 2
- [66] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, 2021. 2
- [67] Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE TIP*, 24(8):2579–2591, 2015. 8
- [68] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [69] Wenlong Zhang, Yihao Liu, Chao Dong, and Yu Qiao. RankSRGAN: Generative adversarial networks with ranker for image super-resolution. In *ICCV*, 2019. 2, 6
- [70] Wenlong Zhang, Guangyuan Shi, Yihao Liu, Chao Dong, and Xiao-Ming Wu. A closer look at blind super-resolution:

Degradation models, baselines, and performance upper bounds. In *CVPRW*, 2022. [7](#)

- [71] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, 2018. [2](#)
- [72] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *CVPR*, 2018. [2](#)
- [73] Yunan Zhu, Haichuan Ma, Jialun Peng, Dong Liu, and Zhiwei Xiong. Recycling discriminator: Towards opinion-unaware image quality assessment using wasserstein gan. In *ACM MM*, 2021. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#)