



Analyzing Perception-Distortion Tradeoff Using Enhanced Perceptual Super-Resolution Network

Subeesh Vasu^(✉), Nimisha Thekke Madam, and A. N. Rajagopalan

Indian Institute of Technology, Madras, Chennai, India
subeeshvasu@gmail.com

Abstract. Convolutional neural network (CNN) based methods have recently achieved great success for image super-resolution (SR). However, most deep CNN based SR models attempt to improve distortion measures (e.g. PSNR, SSIM, IFC, VIF) while resulting in poor quantified perceptual quality (e.g. human opinion score, no-reference quality measures such as NIQE). Few works have attempted to improve the perceptual quality at the cost of performance reduction in distortion measures. A very recent study has revealed that distortion and perceptual quality are at odds with each other and there is always a trade-off between the two. Often the restoration algorithms that are superior in terms of perceptual quality, are inferior in terms of distortion measures. Our work attempts to analyze the trade-off between distortion and perceptual quality for the problem of single image SR. To this end, we use the well-known SR architecture- enhanced deep super-resolution (EDSR) network and show that it can be adapted to achieve better perceptual quality for a specific range of the distortion measure. While the original network of EDSR was trained to minimize the error defined based on per-pixel accuracy alone, we train our network using a generative adversarial network framework with EDSR as the generator module. Our proposed network, called enhanced perceptual super-resolution network (EPSR), is trained with a combination of mean squared error loss, perceptual loss, and adversarial loss. Our experiments reveal that EPSR achieves the state-of-the-art trade-off between distortion and perceptual quality while the existing methods perform well in either of these measures alone.

Keywords: Super-resolution · Deep learning · Perceptual quality
GAN

1 Introduction

The problem of single image super-resolution (SISR) has attracted much attention and progress in recent years. The primary objective of SISR algorithms is to recover the high-resolution (HR) image from a given single low-resolution (LR) image. By definition, SISR is an ill-posed problem as no unique solution exists for

a given LR image. The same LR image can be obtained by down-sampling a large number of different HR images. The ill-posedness of SISR becomes particularly pronounced when the scaling factor increases. Deep learning approaches attempt to solve this ill-posed problem by learning a mapping between the LR and its corresponding HR image in a direct or indirect manner. Recent works on deep neural networks based SISR have shown significant performance improvement in terms of peak signal-to-noise ratio (PSNR).

SISR with deep networks gained momentum with the primal work of Chao et al. [12]. While [12] used a 3 layer convolutional neural network (CNN), the subsequent works used deeper network architectures [23, 24] and new techniques to improve the restoration accuracy [20, 31] and computational complexity [13, 40]. Despite significant progress in both reconstruction accuracy and speed, a majority of the existing works are still far away from reconstructing realistic textures. This is mainly because of the fact that these works are aimed at improving distortion scores such as PSNR and structural similarity index (SSIM) by optimizing pixel-wise computed error measures such as mean squared error (MSE). In the context of SISR, the optimal MSE estimator returns the mean of many possible solutions [28, 39] which often leads to blurry, overly smooth, and unnatural appearance in the output, especially at the information-rich regions.

Previous studies [27, 46] revealed that pixel-wise computed error measures correlate poorly with human perception of image quality. Considering the fact that, the behavior of optimization-based SR methods are strongly influenced by the choice of objective function, one should be able to obtain high-quality images by picking the best suited objective function for the task at hand. This is the main motivation behind the recent works on SISR [22, 28, 34, 39] that came up with new ways to improve the perceptual quality of reconstructed images.

A detailed analysis conducted by [5] showed that distortion and perceptual quality are at odds with each other and there is always a trade-off between the two. As observed in [5], the restoration algorithms that are superior in terms of perceptual quality, are often inferior in terms of distortion measures. They came up with a new methodology for evaluating image restoration methods which can be used to better reveal this trade-off. They have proposed to map SR methods onto a perception-distortion plane and choose the SR method which yields the lowest perceptual score for a given range of distortion measure as the best performing method for that range. They have also suggested that adversarial loss can be used to achieve the desired trade-off for the specific application in mind. Though the work in [5] concluded that the existing SISR works perform well in either of these metrics, the possibility to achieve better trade-off in different regions of the perception-distortion plane was left unexplored.

In this work, we analyze the perception-distortion trade-off that can be achieved by the well-known SISR architecture- enhanced deep super-resolution (EDSR) network [31]. In our analysis, we limit our focus to SISR by a factor of 4 for LR images distorted by the bicubic down-sampling operator. Selection of EDSR was motivated by the fact that it is one of the state-of-the-art network architecture in terms of the distortion measure for SISR. Since the original work

of EDSR proposed in [31] is aimed at improving distortion measure alone, the perceptual quality achieved by EDSR is poor as pointed out by [5]. We train EDSR network using a combination of loss functions that can improve distortion measures as well as perceptual quality. Motivated by the observations in [5, 22, 28, 39], we use a combination of MSE loss, perceptual (VGG) loss, and adversarial loss to train EDSR. Use of adversarial loss to improve perceptual quality allowed our approach to traverse different regions in the perception-distortion plane with ease. We name our approach as enhanced perceptual super-resolution network (EPSR). Our experiments reveal that EPSR can be used to achieve the state-of-the-art trade-off between distortion measure and perceptual quality corresponding to three different regions in the perception-distortion plane.

Our main contributions are summarized below.

- We expand the scope of EDSR and show that it can be adapted to improve the perceptual quality by compromising on distortion measures.
- Our proposed approach achieves the state-of-the-art perception-distortion trade-off results corresponding to different regions in the perception-distortion plane.

2 Related Works

Though there exist extensive literature studies on multi-image SR [6, 14, 38], here we limit our discussions to SISR works alone. An overview of recent image SR methods can be found in [37, 47]. Early approaches on SISR used sampling theory based interpolation techniques [2, 29, 50] to recover the lost details. While these algorithms can be very fast, they cannot recover details and realistic textures. Majority of the recent works aim to establish a complex mapping between LR and HR image pairs. The works in [15, 16] were some of the early approaches to learn such a complex mapping using example-pairs of LR and HR training patches. In [18], the presence of patch redundancies across scales within an image was exploited to generate more realistic textures. This idea was further extended by [21] wherein self-dictionaries were constructed using self-similar patches that are related through small transformations and shape variations. The convolutional sparse coding framework in [19] process the whole image and exploits the consistency of neighboring patches to yield better image reconstruction.

To generate edge-preserving realistic textures, [42] employed a learning-based approach driven by a gradient profile prior. [30] tried to capture the patch redundancy across different scales using a multi-scale dictionary. HR images from the web with similar contents were used with-in a structure-aware matching criterion to super-resolve landmark images in [48]. The class of neighbor embedding approaches [3, 8, 17, 44, 45] aim to find similar looking LR training patches from a low dimensional manifold and then combine their corresponding HR patches for resolution enhancement. The overfitting tendency of neighborhood approaches was pointed out by [25] while also formulating a more generic approach using

kernel ridge regression. The work in [9] learned a multitude of patch-specific regressors and proposed to use the most appropriate regressors during testing.

Recently, deep neural networks based SR algorithms showed dramatic performance improvements in SISR. Preliminary attempts to deep-learning based SISR appeared in [11, 12] (SRCNN) wherein a 3 layer network was employed to learn the mapping between the desired HR image and its bicubic up-sampled LR image. This was followed by deeper network architectures [23, 24] promising performance improvement over SRCNN. [23] proposed to use residual-learning and gradient clipping with a high-learning rate, whereas [24] relied on a deep recursive layer architecture. The works in [13, 40] revealed that SR networks can be trained to learn feature representations at the LR dimension itself thereby allowing to use LR images as a direct input rather than using an interpolated image as the input. This improvisation led to significant reduction in computations while maintaining the model capacity and performance gain. To map from the LR feature maps to the final HR image, these works used upsampling modules at the very end of the network. For upsampling, [13] used a deconv layer whereas [40] employed an efficient sub-pixel convolution layer. The work in [28] came up with a deeper architecture made of residual blocks for LR feature learning, called SRResNet. The well-known architecture of EDSR [31] is built as a modification to SRResNet while using an improvised form of the residual block. They have employed a deeper network architecture with more number of feature units as compared to SRResNet to become the winners of NTIRE2017 [43]. The work in [20] proposed a deep back-projection network (DBPN) to achieve performance improvement over [43] for the distortion measure based SISR. It should be noted that all the above-mentioned deep-learning based works have attempted to improve the performance in terms of distortion measures by training loss functions computed in the form of pixel-wise error measures.

Of particular relevance for our paper are the works that have attempted to use loss functions that can better approximate perceptual similarity ensuring recovery of more convincing HR images. The works along this line includes [7, 10, 22, 28, 34, 39]. Both [7] and [22] attempted to use an error function derived from the features extracted from a pre-trained VGG network instead of low-level pixel-wise error measures [41]. More specifically, they used the Euclidean distance between feature maps extracted from the VGG19 network (called VGG loss) as the loss function that was found to give more visually appealing results as opposed to using the MSE loss computed at the pixel-space. SRGAN proposed in [28] was the first attempt to use a GAN-based network which optimizes for the so-called adversarial loss to improve the perceptual quality in SISR. While [28] used a combination of MSE, VGG, and perceptual loss, the work in [39] used an additional texture matching loss to generate more realistic textures. [34] employed contextual loss to replace the perceptual loss for improved perceptual quality. [10] proposed to combine the high-frequency information of a GAN based method and the content information of an MSE loss based method to obtain achieve the desired balance between distortion and perceptual quality.

3 Method

An LR image I_{LR} can be related to its corresponding HR counterpart (I_{HR}) as

$$I_{LR} = d_{\alpha}(I_{HR}) \tag{1}$$

where d_{α} refers to the degradation operator which when acts on I_{HR} results in I_{LR} and $\alpha (>1)$ is the scaling factor. Though the degrading factors involved in d_{α} can be a combination of blur, decimation, or noise, in this work, we assume d_{α} to represent a bicubic downsampling operation with a single scale factor of 4. The task of SISR is to find an approximate inverse $f \approx d^{-1}$ to yield an HR image estimate I_{est} from I_{LR} . This problem is highly ill-posed as there exists a large number of possible image estimates I_{est} for which the degradation relation ($d_{\alpha}(I_{est}) = I_{LR}$) holds true.

Majority of the deep-learning approaches attempt to find f by minimizing the MSE loss between the network output and the ground truth image ($\|I_{est} - I_{HR}\|_2^2$). While such a scheme can give excellent results in terms of distortion measures, the resulting images are often blurry and lack high-frequency textures. Previous works on perceptual SR have shown that this limitation can be overcome by employing the loss functions that favor perceptually pleasing results. However, such perceptual improvements result in the reduction of distortion measures. The objective of our work is to experimentally find the perception-distortion trade-off for the state-of-the-art SISR architecture of EDSR.

Next, we will explain the details of our approach, including the network architecture, loss functions, and the methodology that we adopted to find the best possible trade-off corresponding to the network architecture of EDSR.

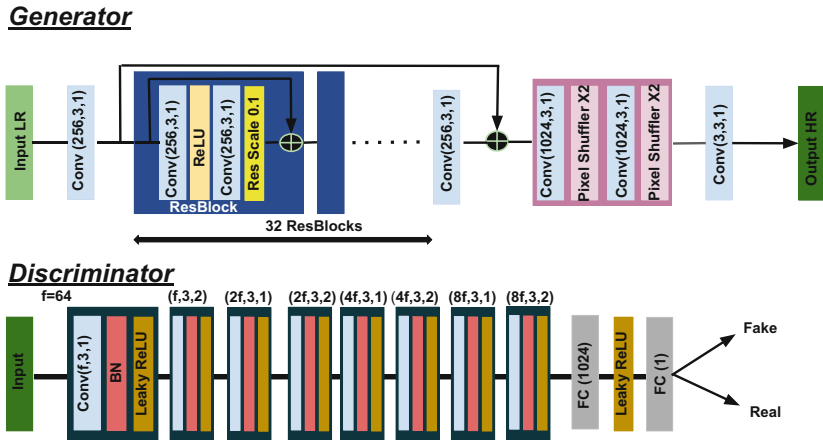


Fig. 1. Network architecture of EPSR.

3.1 Network Architecture

The complete architecture of the SR network used in our work is shown in Fig. 1. Our network consists of EDSR acting as the generator module and a CNN based classifier acting as a discriminator module. In the diagram shown in Fig. 1, $\mathbf{conv}(\mathbf{n}, \mathbf{k}, \mathbf{s})$ refers to a convolution layer with \mathbf{n} number of $\mathbf{k} \times \mathbf{k}$ filters performing convolution by a stride factor of \mathbf{s} . For simplicity we use the notation $(\mathbf{n}, \mathbf{k}, \mathbf{s})$ instead of $\mathbf{conv}(\mathbf{n}, \mathbf{k}, \mathbf{s})$ in the diagram corresponding to the discriminator. EDSR is built based on a modified form of residual block wherein no batch normalization units are used. An additional residual scaling layer (multiplication by a constant scale factor of 0.1) is inserted onto each residual block to numerically stabilize the training procedure in the absence of batch-normalization. This kind of a modified form of the residual block has allowed the work in [31] to employ a deeper network architecture with more number of feature units in each layer to improve the performance over the SRResNet architecture of [28]. The generator module comprises of 32 (modified form of) residual blocks (refer Fig. 1 for more details). The LR images are directly provided to the network as inputs. To increase the resolution by a factor of 4, residual blocks are followed by two pixel shuffler units each of which increases the spatial resolution by a factor of 2. As shown in Fig. 1, the discriminator that we used is a 10 layer network trained to produce a single output 0/1 depending on the input data which can be $I_{\text{est}}/I_{\text{HR}}$. The network consists of a number of convolution layers followed by fully connected layers which map from an image to a single output value.

3.2 Training and Loss Functions

We used the following form of loss function to train the network.

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{VGG}} + \lambda_2 \mathcal{L}_E + \lambda_3 \mathcal{L}_{\text{adv}} \quad (2)$$

where \mathcal{L} is the total loss function used for training the generator network. \mathcal{L}_E is the MSE between the network output and ground truth HR image given by

$$\mathcal{L}_E = \|I_{\text{est}} - I_{\text{HR}}\|_2^2 \quad (3)$$

\mathcal{L}_{VGG} is the perceptual loss [22] computed using the VGG network [41] as

$$\mathcal{L}_{\text{VGG}} = \|\phi(I_{\text{est}}) - \phi(I_{\text{HR}})\|_2^2 \quad (4)$$

where ϕ refers to VGG feature layers. Previous studies on perceptual SISR [28, 39] have shown that the use of perceptual loss \mathcal{L}_{VGG} can provide further boost in the detail enhancement if used along with adversarial loss. Following this line, we also use \mathcal{L}_{VGG} to provide an additional support to the adversarial loss for improving the perceptual quality. Similar to the work in [28], we used VGG54 as the feature extraction layer (i.e, the feature maps obtained by the 4th convolution (after activation) and before the 5th max-pooling layer). \mathcal{L}_{adv} is the adversarial loss derived from the discriminator network and is given by

$$\mathcal{L}_{\text{adv}} = -\log D(G(I_{\text{LR}})) \quad (5)$$

where $G(\cdot)$ and $D(\cdot)$ indicates the network outputs from the generator and discriminator respectively. λ_1 , λ_2 , and λ_3 are non-negative scale factors that can be varied to control the perception-distortion trade-off.

Motivated from the observation that GANs can provide a principled way to approach the perception-distortion bound [5], we train our network using different values of λ_2 and λ_3 (refer Table 1) to achieve the best possible perception-distortion trade-off using EPSR. The training of EPSR is done similar to that of [28]. The generator network is trained to learn a mapping from input image I_{LR} to an approximate estimate of the HR image I_{est} by optimizing the loss function \mathcal{L} . Simultaneously, the discriminative network D is trained to distinguish between real images I_{HR} from the training dataset and generated image estimates of the network $G(I_{LR})$. To train the discriminator we minimize the loss function.

$$\mathcal{L}_D = -\log(D(I_{HR})) - \log(1 - D(G(I_{LR}))) \quad (6)$$

During training, the discriminator was updated twice followed by a single generator update. Also, to train the network with different values of λ_2 and λ_3 , we initialized the model weights of generator using pre-trained weights of EDSR (obtained by training EDSR with $\lambda_1 = \lambda_3 = 0$).

Table 1. Parameter settings used for training BNet and EPSR to obtain results corresponding to Region 1, 2, and 3. BNet (refer Sect. 4) is a baseline network used for performance comparison.

Network model \rightarrow	BNet			EPSR		
	λ_1	λ_2	λ_3	λ_1	λ_2	λ_3
Region 1 (RMSE \leq 11.5)	1	0.1	0.4	1	.05	0.4
Region 2 (11.5 < RMSE \leq 12.5)	1	0.05	0.4	1	0.02	0.4
Region 3 (12.5 < RMSE \leq 16)	1	0.0005	0.6	1	0.0005	0.6

3.3 Implementation Details

To train our network, we used the first 800 images of DIV2K dataset [1]. The HR images were bicubically down-sampled by a factor of 4 to create the input LR images for training. We followed a patch-wise training wherein the patch-size of the network output was set to 192. We used ADAM [26] optimizer with a momentum of 0.9 and a batch size of 4. The network was trained for 300 epochs and the learning rate was initially set to $5e-5$ which was reduced by a factor of 0.5 after 150 epochs. We used pre-trained VGGNet weights to enforce the effect of perceptual loss. Our implementation was done in PyTorch and was built on top of the official PyTorch implementation of [31] which was available online. The code was run on TITAN-X Pascal GPU. It took around 45 h to complete the training of one single network. On an average, during testing, to super-resolve an input image of size 100×100 , EPSR takes around 0.5 s.

4 Evaluation

To evaluate the performance, we follow a procedure similar to that of “The PIRM challenge on perceptual super-resolution” (PIRM-SR) [4, 5]. The evaluation is done in a perceptual-quality aware manner [5], and not based solely on the basis of distortion measures. To this end, we divide the perception-distortion plane [5] into three regions defined by thresholds on the RMSE of the SR outputs. The thresholds used for the three regions are mentioned in Table 1.

Table 2. Results on public benchmark test data and PIRM-self validation data for existing distortion measure specific methods and our methods corresponding to region 1 (BNet₁ and EPSR₁). Bold indicates the best performance in Region 1 and italic indicates the second best.

Dataset	Scores	Bicubic	SRCNN [12]	EDSR [31]	DBPN [20]	BNet ₁	EPSR ₁
PIRM-self	RMSE	13.2923	12.0194	10.8934	10.9779	11.4956	11.4924
	PSNR	26.5006	27.5258	28.5754	28.4927	27.9752	27.9852
	SSIM	0.6980	0.7429	0.7808	0.7773	0.7511	0.7508
	PI	6.805	5.8247	5.0399	5.2043	<i>4.1492</i>	2.9459
Set5	PSNR	28.4164	30.5314	32.4034	32.3337	31.4505	31.6954
	SSIM	0.8096	0.8630	0.8960	0.8949	0.8739	0.8751
	PI	7.323	7.0858	5.8366	6.107	<i>5.4136</i>	4.8087
Set14	PSNR	25.6675	26.7191	27.4193	28.1266	27.0541	27.0123
	SSIM	0.6921	0.7316	0.7543	0.7686	0.7342	0.7315
	PI	6.968	6.0189	5.2942	5.5723	<i>4.4824</i>	3.7101
BSD100	PSNR	26.2128	26.7564	27.0088	27.0145	26.8711	26.7497
	SSIM	0.6839	0.7198	0.7396	0.7364	0.71782	0.7133
	PI	6.9485	5.9707	5.36	5.5362	<i>4.6416</i>	3.5503
Urban100	PSNR	22.7809	23.5834	24.5753	24.4825	24.1029	24.3012
	SSIM	0.6477	0.6984	0.7517	0.7460	0.72199	0.7302
	PI	6.8796	5.8414	5.0395	5.1944	<i>4.2223</i>	3.8994

We used perceptual index (PI) to quantify the perceptual quality. PI is computed by combining the quality measures of Ma-score [32] and NIQE [36] as follows

$$\text{PI} = 1/2((10 - \text{Ma-score}) + \text{NIQE}) \quad (7)$$

Note that, a lower PI indicates better perceptual quality. The algorithm with the best perceptual score (or equivalently lowest PI) in each region is treated as the one with most visually pleasing results corresponding to that particular region. This approach of region-wise comparison quantifies the accuracy and perceptual quality of algorithms jointly, and will, therefore, enable a fair comparison of perceptual-driven methods alongside algorithms that target PSNR maximization.

Table 3. Results on public benchmark test data and PIRM-self for existing perceptual quality specific methods and our proposed methods corresponding to Region 2 and Region 3 (EPSR₂ and EPSR₃). Bold indicates the best performance in Region 2 (and Region 3) and italic indicates the second best.

Dataset	Scores	ENet [39]	CX [34]	BNet ₂	BNet ₃	EPSR ₂	EPSR ₃
PIRM-self	RMSE	15.9853	15.2477	12.4709	15.6292	12.4094	15.3586
	PSNR	25.0642	25.4051	27.1789	25.2845	27.342	25.4541
	SSIM	0.6463	0.6744	0.7184	0.6560	0.72744	0.6655
	PI	2.6876	<i>2.131</i>	<i>2.4795</i>	2.2354	2.3881	2.0688
Set5	PSNR	28.5641	29.1017	30.7637	28.6764	31.2168	29.5757
	SSIM	0.80819	0.82982	0.85485	0.80948	0.8630	0.8388
	PI	2.9261	3.2947	4.0003	<i>3.2223</i>	<i>4.1123</i>	3.2571
Set14	PSNR	25.7521	25.2265	26.5242	25.2487	26.6068	25.5238
	SSIM	0.67953	0.67606	0.7104	0.6595	0.71342	0.6848
	PI	3.014	2.759	<i>3.1706</i>	2.6473	3.0246	<i>2.6982</i>
BSD100	PSNR	25.3764	24.2868	26.1619	24.7761	26.2819	24.9753
	SSIM	0.64268	0.6396	0.6826	0.6217	0.69054	0.64503
	PI	2.9297	<i>2.2501</i>	<i>2.801</i>	2.3674	2.7458	2.199
Urban100	PSNR	23.6771	22.8444	23.5657	22.0168	23.9985	22.7959
	SSIM	0.69775	0.6748	0.6934	0.6454	0.71798	0.66631
	PI	3.4679	3.3894	<i>3.6345</i>	3.2721	3.6236	<i>3.3316</i>

To have an idea about the performance level of EPSR, we compare it with that of the trade-off values achieved by a baseline network formed by our-self. We call our baseline network as BNet and is a simplified form of EPSR. Unlike EPSR, the generator of BNet has no residual scaling. BNet uses 32 number of residual blocks and 64 filters in each layer of the residual block. BNet is equivalent to the network in [28] (SRGAN) except for the fact that [28] use batch normalization units in the generator whereas BNet does not.

To perform a region-wise comparison, we train both BNet and EPSR with a different set of weights for MSE loss and adversarial loss. The weights for the best trade-off was empirically found for each region (refer to Table 1 for details). In the following comparisons, BNet₁ (/BNet₂/BNet₃) and EPSR₁ (/EPSR₂/EPSR₃) refers to the best model weights (i.e., the ones with the lowest PI) obtained for Region 1 (/2/3) corresponding to BNet and EPSR respectively. We perform the region-wise performance comparisons with the most relevant methods on distortion measure (bicubic interpolation, [12, 20, 31]) as well as perceptual quality [34, 39]. Since the code of SRGAN [28] was not available, an equivalent comparison is done using BNet. We could not compare with the other perceptual SR methods [7, 10, 22], as the source codes for them were not available.

Evaluation is done on the public benchmark data sets of Set5 [3], Set14 [49], BSD100 [33], Urban100 [21] and the self-validation data from PIRM-SR

(PIRM-self) [4]. Since PIRM-self contains 100 images with an equal distribution of scenes and quality, it can be treated as the most suited dataset for perceptual quality-based evaluation. Consequently, we use the average MSE values computed over PIRM-self to define the three regions in the perception-distortion plane.

4.1 Quantitative Results

To quantitatively compare the performance, we report the values of PSNR, SSIM, and PI. The results corresponding to [31] is obtained using the model weights of EDSR obtained through our own training. Also, the values that we have obtained for the existing methods on distortion measure is slightly different

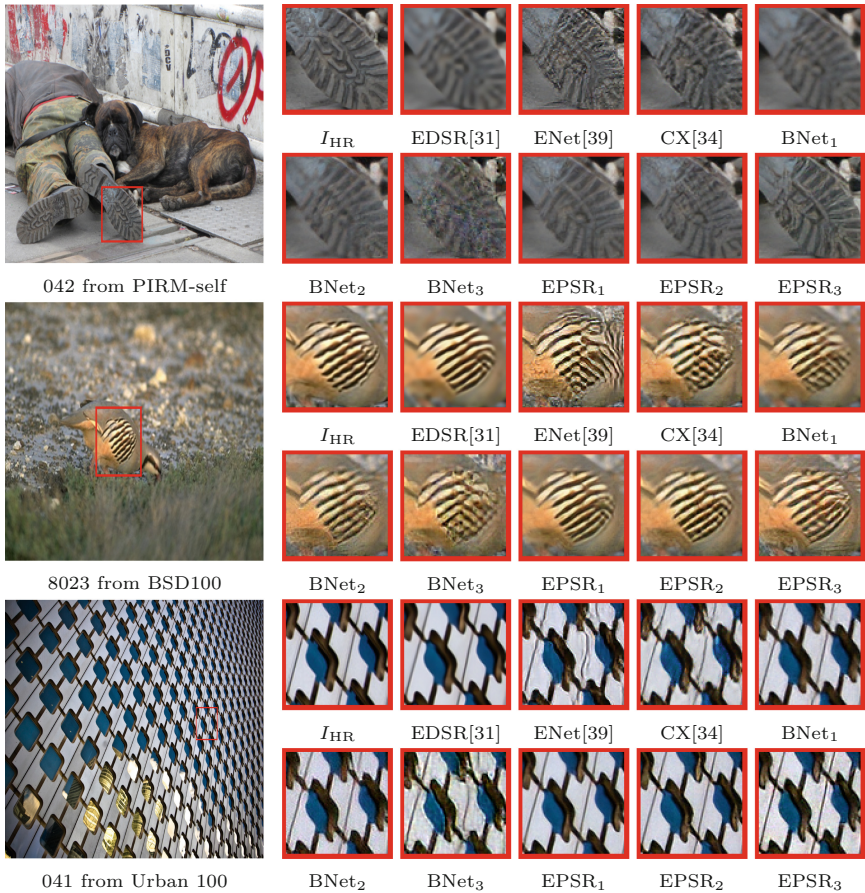


Fig. 2. Qualitative comparison of our models with other works on x4 super-resolution. The image examples are taken from datasets of PIRM-self (Row 1), BSD100 (Row 2), and Urban100 (Row 3).

as compared to the ones reported in the original papers. This could be due to the difference in the way we have computed the scores. All the scores reported in this paper are computed on the y-channel after removing a 4-pixel border.

Table 2 lists the quantitative comparison of distortion measure based methods with that of BNet₁ and EPSR₁.¹ As is evident from Table 2, EPSR performs the best and achieve the lowest PI in Region 1 and BNet turns out to be the second best. Both BNet₁ and EPSR₁ is able to deliver low PI values (i.e., better perceptual quality) while maintaining much better distortion measures (RMSE, PSNR, and SSIM) as compared to bicubic interpolation and SRCNN. A careful inspection of the distortion measure based method reveals that the perceptual quality improves as the PSNR increases, however, the relative improvement is very narrow. Differently, a comparison between EDSR and EPSR₁ shows that the use of adversarial loss has helped EPSR₁ to achieve significant improvement in perceptual quality but while subjected to reduction in distortion measures.

Table 3 lists the quantitative comparison of perceptual-SISR methods with that of BNet and EPSR corresponding to Region 2 and 3. It should be noted that, among all the datasets that we have compared, Set5, Set14, and Urban 100 are not the ideal ones for perceptual quality comparisons. Because Set5 and Set14 have only a small number of images whereas Urban100 covers only the images of urban scenes. Both, BSD 100 and PIRM-self covers wide-variety of scenes and can be treated as an ideal collection of natural images of different kinds. Comparisons over BSD 100 and PIRM-self in Table 3 underscore the superior perceptual quality improvement achieved by EPSR. In other datasets, the method which has the lowest PI varies. In Set5, ENet [39] performs best in Region 3, whereas BNet₂ performs best in Region 2. In Set14 and Urban 100, the best performing methods are CX [34], BNet, and EPSR with only a comparable performance difference between each other.

Considering all regions together, one can see that, EPSR achieves the best perceptual scores, with CX [34] being second best. By comparing BNet and EPSR scores across different regions we can notice the trade-off between the PI and RMSE. When we allowed having more distortion (i.e., higher RMSE), both BNet and EPSR are able to yield significant improvement in perceptual quality. Note that the generator network of BNet is inferior to that of EPSR in terms of distortion measures. This allows EPSR to achieve better perceptual quality than BNet for a fixed level of distortion. We believe the following as the primary reason for such an effect. To improve the perceptual quality, a network needs to generate more realistic textures resulting in an increase of the content deviation from the ground truth image. Therefore, for a given distortion range, a generator network which is superior in terms of distortion-measure is more likely to generate results with the best perceptual quality when trained using a GAN framework.

¹ Bicubic and SRCNN correspond to Region 2 since their RMSE values are above 11.5.

4.2 Qualitative Results

For qualitative comparisons, we show a few examples from the standard benchmark datasets. In all the cases, we also show the ground truth (GT) images to get an idea about the content distortions introduced by the perceptual SR methods and also to visualize the extent to which the distortion measure based methods can reveal the lost details. Figures 2 and 3 shows visual comparisons of seven examples in total. Examples in Figs. 2 and 3 clearly shows that, though ENet [39] is able to achieve a significant level of detail enhancement, the texture details added by the network is often very different from the ground-truth. Also, ENet [39] appears to add strong noise components while attempting to do detail enhancement. In comparison to ENet [39], the presence of noise and unrealistic texture is less for the case of CX [35] while maintaining a comparable level of detail enhancement. Contrarily, EPSR₃ is able to generate realistic textures that are faithful to both the GT image and the outputs from distortion-based methods.

The presence of spurious noise components in ENet [39] outputs can be seen in the first example of Fig. 2 as well as the first and second example of Fig. 3. For all these examples, BNet₃ also resulted in a very similar noise disturbance. However, EPSR₃ was able to generate visually pleasing realistic textures in the output. Second and third examples in Fig. 2 corresponds to failure case of ENet [39], CX [35], and BNet₃ wherein all of them resulted in texture patterns that are very different from the GT, whereas EPSR₃ has succeeded in generating outputs that are more faithful to the GT image. The fourth example of Fig. 3 shows the detail-preservation ability of EPSR as compared to the other perceptual methods. While EPSR₃ succeeded in reconstructing the seal whiskers to a great extent, both BNet and ENet [39] failed to do so.

In all the examples, the inadequacy of distortion based methods for reconstructing detailed textures is clearly evident. While outputs from both bicubic and SRCNN is affected by heavy blur, EDSR and DBPN output images with a minimal level of blur. The perceptual SR methods, on the other hand, generates detailed structures that are not necessarily consistent with the GT image. Among all the perceptual SR methods, EPSR performs the most convincing detail enhancement and is the one which generates detail enhanced outputs that are closest to the GT image. As indicated by the quantitative evaluation, EPSR₁ achieves significant perceptual quality improvement over EDSR while incurring only minimal distortion as compared to EDSR. This effect is predominantly visible in the first example of Fig. 2 and first two examples from Fig. 3. As is evident from the Visual comparison of images from EDSR and EPSR reveals the progressive detail recovery that can be achieved by EPSR while moving across different regions in the perception-distortion plane. A very similar observation can also be made by comparing the images corresponding to BNet too. The source code of our method can be downloaded from https://github.com/subeeshvasu/2018_subeesh_epsr_eccvw.

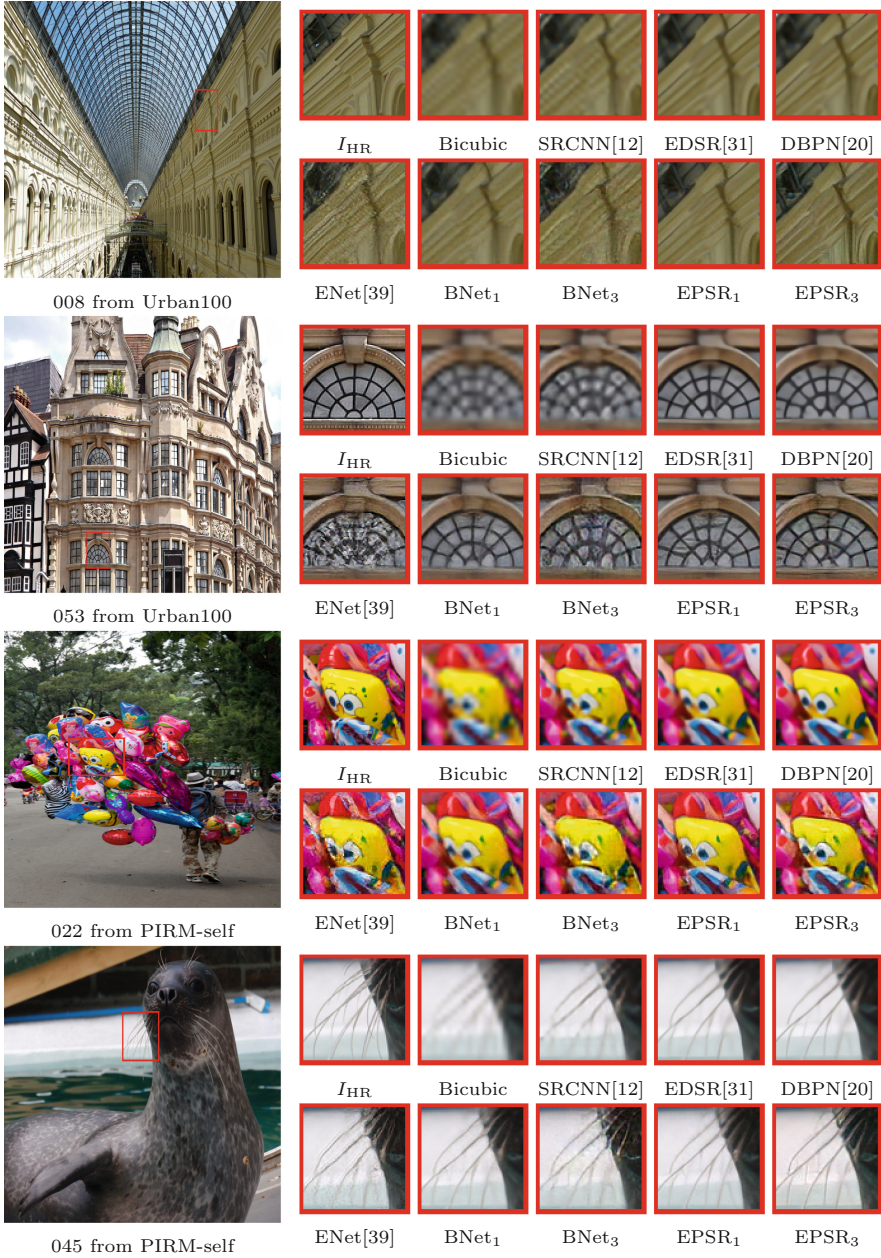


Fig. 3. Qualitative comparison of our models with other works on x4 super-resolution. Examples are taken from datasets of Urban100 (Rows 1–2) and PIRM-self (Rows 3–4).

4.3 Trade-Off Comparison with BNet

To analyze the impact of the generator module in achieving the trade-off, we perform a trade-off comparison between BNet and EPSR. Note that BNet uses a generator which is inferior to that of EPSR in terms of the distortion measures. Therefore, we expect to obtain a better perception-distortion trade-off using EPSR. Figure 4(a) is a plot corresponding to the trade-off comparison between BNet and EPSR, wherein we have used the network model weights corresponding to different parameter settings that span different regions in the perception-distortion plane. To generate the plot in Fig. 4, we use the PI and RMSE values computed based on the PIRM-self dataset. To obtain model weights corresponding to different trade-off points, we have trained BNet and EPSR with different parameter settings and chose a number of network weights that yields the lowest PI values over a certain range of RMSE. It is evident from Fig. 4 that EPSR is able to deliver a much better trade-off as compared to BNet as expected.

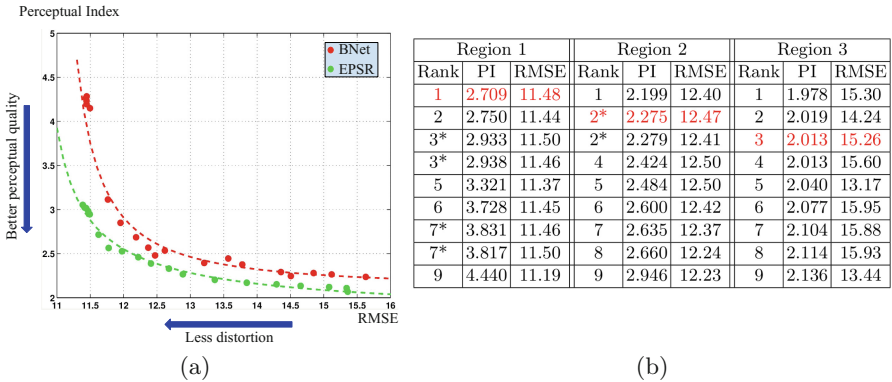


Fig. 4. (a) Perception-distortion trade-off between BNet and EPSR. For both methods, we plot the values corresponding to 19 model weights which span different regions on the perception-distortion plane and the corresponding curves that best fit these values. (b) Performance comparison of top 9 methods from PIRM-SR challenge [4]. Methods are ranked based on the PI and RMSE values corresponding to the test data of PIRM-SR. The entries from our approach are highlighted in red. Methods with a marginal difference in PI and RMSE values share the same rank and are indicated with a*. (Color figure online)

4.4 PIRM Challenge on Perceptual Super-Resolution

We have used our proposed model EPSR to participate in the PIRM-SR challenge [4] wherein the objective was to compare and rank perceptual SISR methods for an SR factor of 4. In order to rank each method, the perception-distortion plane was divided into three regions defined by thresholds on the RMSE. In each region, the winning algorithm is selected as the one that achieves the best perceptual quality. We have used parameter-tuned variants of EPSR to obtain the

results corresponding to all three regions. The RMSE range used to define the three regions and the parameter settings that we have used to generate the results corresponding to the three regions are mentioned in Table 1. Our method was ranked 1, 2, and 3 in region 1, 2, and 3 respectively as shown in Fig. 4(b).

5 Conclusions

We proposed an extension to the state-of-the-art EDSR network by using it within a GAN framework. The proposed approach, EPSR, scales well in different regions of the perception-distortion plane and achieves superior perceptual scores when compared in a region-wise manner to other existing works. The performance improvement achieved by our approach is a cumulative result of the following factors: state-of-the-art SR network (EDSR) as the generator module, careful selection of loss function weights, and initialization of GAN training with the pretrained weights of EDSR. Our analysis of the perception-distortion trade-off between BNet and EPSR signal the possibility to further boost the trade-off by adopting another generator module that yields better distortion measures.

References

1. Agustsson, E., Timofte, R.: NTIRE 2017 challenge on single image super-resolution: dataset and study. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, vol. 3, p. 2 (2017)
2. Allebach, J., Wong, P.W.: Edge-directed interpolation. In: Proceedings of the International Conference on Image Processing, vol. 3, pp. 707–710. IEEE (1996)
3. Bevilacqua, M., Roumy, A., Guillemot, C., Alberi-Morel, M.L.: Low-complexity single-image super-resolution based on nonnegative neighbor embedding (2012)
4. Blau, Y., Mechrez, R., Timofte, R., Michaeli, T., Zelnik-Manor, L.: 2018 PIRM challenge on perceptual image super-resolution. arXiv preprint [arXiv:1809.07517](https://arxiv.org/abs/1809.07517) (2018)
5. Blau, Y., Michaeli, T.: The perception-distortion tradeoff. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2018
6. Borman, S., Stevenson, R.L.: Super-resolution from image sequences—a review. In: Proceedings of 1998 Midwest Symposium on Circuits and Systems, pp. 374–378. IEEE (1998)
7. Bruna, J., Sprechmann, P., LeCun, Y.: Super-resolution with deep convolutional sufficient statistics. arXiv preprint [arXiv:1511.05666](https://arxiv.org/abs/1511.05666) (2015)
8. Chang, H., Yeung, D.Y., Xiong, Y.: Super-resolution through neighbor embedding. In: Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2004, vol. 1, pp. I-I. IEEE (2004)
9. Dai, D., Timofte, R., Van Gool, L.: Jointly optimized regressors for image super-resolution. In: Computer Graphics Forum, vol. 34, pp. 95–104. Wiley Online Library (2015)
10. Deng, X.: Enhancing image quality via style transfer for single image super-resolution. IEEE Signal Process. Lett. **25**(4), 571–575 (2018)

11. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8692, pp. 184–199. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10593-2_13
12. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(2), 295–307 (2016)
13. Dong, C., Loy, C.C., Tang, X.: Accelerating the Super-resolution convolutional neural network. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 391–407. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_25
14. Farsiu, S., Robinson, M.D., Elad, M., Milanfar, P.: Fast and robust multiframe super resolution. *IEEE Trans. Image Process.* **13**(10), 1327–1344 (2004)
15. Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. *IEEE Comput. Graph. Appl.* **22**(2), 56–65 (2002)
16. Freeman, W.T., Pasztor, E.C., Carmichael, O.T.: Learning low-level vision. *Int. J. Comput. Vis.* **40**(1), 25–47 (2000)
17. Gao, X., Zhang, K., Tao, D., Li, X.: Image super-resolution with sparse neighbor embedding. *IEEE Trans. Image Process.* **21**(7), 3194–3205 (2012)
18. Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: 2009 IEEE 12th International Conference on Computer Vision, pp. 349–356. IEEE (2009)
19. Gu, S., Zuo, W., Xie, Q., Meng, D., Feng, X., Zhang, L.: Convolutional sparse coding for image super-resolution. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1823–1831 (2015)
20. Haris, M., Shakhnarovich, G., Ukita, N.: Deep backprojection networks for super-resolution. In: Conference on Computer Vision and Pattern Recognition (2018)
21. Huang, J.B., Singh, A., Ahuja, N.: Single image super-resolution from transformed self-exemplars. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5197–5206 (2015)
22. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
23. Kim, J., Kwon Lee, J., Mu Lee, K.: Accurate image super-resolution using very deep convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1646–1654 (2016)
24. Kim, J., Kwon Lee, J., Mu Lee, K.: Deeply-recursive convolutional network for image super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1637–1645 (2016)
25. Kim, K.I., Kwon, Y.: Single-image super-resolution using sparse regression and natural image prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 1127–1133 (2010)
26. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. *arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)* (2014)
27. Laparra, V., Ballé, J., Berardino, A., Simoncelli, E.P.: Perceptual image quality assessment using a normalized Laplacian pyramid. *Electron. Imaging* **2016**(16), 1–6 (2016)
28. Ledig, C., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: CVPR, vol. 2, p. 4 (2017)
29. Li, X., Orchard, M.T.: New edge-directed interpolation. *IEEE Trans. Image Process.* **10**(10), 1521–1527 (2001)

30. Li, X., Tao, D., Gao, X., Zhang, K.: Multi-scale dictionary for single image super-resolution. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1114–1121. IEEE (2012)
31. Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M.: Enhanced deep residual networks for single image super-resolution. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, vol. 1, p. 4 (2017)
32. Ma, C., Yang, C.Y., Yang, X., Yang, M.H.: Learning a no-reference quality metric for single-image super-resolution. *Comput. Vis. Image Underst.* **158**, 1–16 (2017)
33. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: Proceedings of Eighth IEEE International Conference on Computer Vision, ICCV 2001, vol. 2, pp. 416–423. IEEE (2001)
34. Mechrez, R., Talmi, I., Shama, F., Zelnik-Manor, L.: Learning to maintain natural image statistics. arXiv preprint [arXiv:1803.04626](https://arxiv.org/abs/1803.04626) (2018)
35. Mechrez, R., Talmi, I., Zelnik-Manor, L.: The contextual loss for image transformation with non-aligned data. arXiv preprint [arXiv:1803.02077](https://arxiv.org/abs/1803.02077) (2018)
36. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. *IEEE Signal Process. Lett.* **20**(3), 209–212 (2013)
37. Nasrollahi, K., Moeslund, T.B.: Super-resolution: a comprehensive survey. *Mach. Vis. Appl.* **25**(6), 1423–1468 (2014)
38. Park, S.C., Park, M.K., Kang, M.G.: Super-resolution image reconstruction: a technical overview. *IEEE Signal Process. Mag.* **20**(3), 21–36 (2003)
39. Sajjadi, M.S., Schölkopf, B., Hirsch, M.: EnhanceNet: single image super-resolution through automated texture synthesis. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 4501–4510. IEEE (2017)
40. Shi, W., et al.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1874–1883 (2016)
41. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (ICLR) (2015)
42. Tai, Y.W., Liu, S., Brown, M.S., Lin, S.: Super resolution using edge prior and single image detail synthesis (2010)
43. Timofte, R., et al.: NTIRE 2017 challenge on single image super-resolution: methods and results. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1110–1121. IEEE (2017)
44. Timofte, R., De Smet, V., Van Gool, L.: Anchored neighborhood regression for fast example-based super-resolution. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1920–1927 (2013)
45. Timofte, R., De Smet, V., Van Gool, L.: A+: adjusted anchored neighborhood regression for fast super-resolution. In: Cremers, D., Reid, I., Saito, H., Yang, M.-H. (eds.) ACCV 2014. LNCS, vol. 9006, pp. 111–126. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-16817-3_8
46. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
47. Yang, C.-Y., Ma, C., Yang, M.-H.: Single-image super-resolution: a benchmark. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8692, pp. 372–386. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10593-2_25

48. Yue, H., Sun, X., Yang, J., Wu, F.: Landmark image super-resolution by retrieving web images. *IEEE Trans. Image Process.* **22**(12), 4865–4878 (2013)
49. Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparse-representations. In: Boissonnat, J.-D., et al. (eds.) *Curves and Surfaces 2010*. LNCS, vol. 6920, pp. 711–730. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-27413-8_47
50. Zhang, L., Wu, X.: An edge-guided image interpolation algorithm via directional filtering and data fusion. *IEEE Trans. Image Process.* **15**(8), 2226–2238 (2006)