Perception-Distortion Balanced ADMM Optimization for Single-Image Super-Resolution

Yuehan Zhang¹[0000-0002-5017-0097]</sup>, Bo Ji¹[0000-0002-5718-2074]</sup>, Jia Hao²[0000-0002-0909-9444]</sup>, and Angela Yao¹[0000-0001-7418-6141]</sup>

 ¹ National University of Singapore,
 ² HiSilicon Technologies, Shanghai {zyuehan,jibo,ayao}@comp.nus.edu.sg, hao.jia@huawei.com

Abstract. In image super-resolution, both pixel-wise accuracy and perceptual fidelity are desirable. However, most deep learning methods only achieve high performance in one aspect due to the perception-distortion trade-off, and works that successfully balance the trade-off rely on fusing results from separately trained models with ad-hoc post-processing. In this paper, we propose a novel super-resolution model with a lowfrequency constraint (LFc-SR), which balances the objective and perceptual quality through a single model and yields super-resolved images with high PSNR and perceptual scores. We further introduce an ADMMbased alternating optimization method for the non-trivial learning of the constrained model. Experiments showed that our method, without cumbersome post-processing procedures, achieved the state-of-the-art performance. The code is available at https://github.com/Yuehan717/PDASR.

Keywords: Image Super-Resolution · Perception-Distortion Trade-Off · Constrained Optimization

1 Introduction

Single image super-resolution (SISR) recovers a high-resolution (HR) image from a low-resolution (LR) input. There are two types of quality assessments for super-resolved images: objective quality, evaluated by PSNR and SSIM, and perceptual quality, based on metrics such as NRQM [25] and LPIPS [47]. SISR methods aiming at high objective quality [5, 14, 15, 23, 31, 33, 50, 51] achieve high PSNR values, but the results look blurry. Another line of research focuses on improving perceptual quality [19, 26, 29, 34, 44, 49], which produces HR images that are visually shaper, but have lower PSNR scores and unrealistic patterns (see comparisons in Fig. 1).

An ideally reconstructed image is similar to the ground truth HR image, with limited distortion *and* high perceptual quality, where distortion refers to a drop in objective quality. However, most deep learning methods can achieve high performance in only one of the two qualities. This perception-distortion (PD) tradeoff is rooted in the supervised training process of SISR methods. Since perception and distortion measurements are incoherent with each other, optimization in one

2 Y. Zhang et al.



(c) Perceptual-aimed(d) Balanced Quality

(e) SOTAs performance

direction naturally leads to sacrifice in the other [2, 9]. A recently emerging line of work seeks to balance or improve both the objective and perceptual quality of SISR images [6, 7, 42, 41, 44, 32, 36, 20, 43]. Some researchers have studied the trade-off problem through model optimization [41, 32, 36, 20, 43]. However, they have only focused on traversing the PD trade-off instead of improving both the objective and perceptual quality. Others bypassed the optimization incoherence and merged the models or model outputs trained separately for the two objectives via post-processing [6, 7, 42, 44]. Some methods [6, 7] achieved excellent results, but their post-processing can be computationally costly.

Previous studies have aligned image quality with different image frequency subbands. Low-frequency (LF) information captures the overall scene structure [35, 54], while high-frequency (HF) details are critical to achieving high perceptual quality [19, 54]. This observation has been leveraged to reconstruct details in single-image SR [37, 54] and suppress misalignment in real-world video SR [45]. However, it has not been used to explicitly balance the objective and perceptual quality in model optimization.

This paper present a new low-frequency constrained SISR model (LFc-SR) model that improves both objective and perceptual quality. The model has two stages: stage one focuses on super-resolution for objective quality, while stage two subsequently refines the image for perceptual quality (see Fig. 2). At our model's core is a similarity constraint that keeps the low-frequency subbands of both stages similar during optimization. The constraint allows the optimization of perceptual quality to be oriented towards the high-frequency bands while ensuring that the overall scene retains the effects of objective-focused learning.

Fig. 1: (a)-(d) are ground truth from Urban100[12] and ×4 super-resolved images. Note that (b) produced by NLSA [31] is blurry in detailing, while (c) from RankSRGAN [49] has unnatural patterns; our balanced approach in (d) mitigates both types of artifacts. (e) is performance of state-of-the-art ×4 SISR models on Urban100. Higher PSNR and NRQM scores indicate better objective and perceptual quality. Our method reaches state-of-the-art while being single-shot.

As such, our method yields images with high performance on both qualities in a one-shot manner.

Our key novelty is in our formulation of LFc-SR training as constrained multiobjective optimization. Enforcing constraints within deep learning is non-trivial with challenges to find satisfying solutions. To cooperate the low-frequency constraint, we designed a novel ADMM-based alternating optimization method. The alternating direction method of multipliers (ADMM) was originally introduced as a tool for convex optimization problems, decomposing the constrained optimization into solvable substeps [3]. However, it has also been widely applied to non-convex problems, including deep neural network training [13, 52, 22, 48, 38, 21]. While finding the optimum for non-convex problems is not guaranteed, ADMM often finds satisfying approximations. We refer to our optimization method as PD-ADMM, as it aims at mitigating the perception-distortion (PD) trade-off.

To summarize our contributions in this paper:

- 1. We propose LFc-SR, a novel SISR model trained with a low-frequency constraint, to balance the perception-distortion trade-off.
- 2. We propose a novel formulation of SISR as a constrained multi-objective optimization problem and show how to incorporate an ADMM-based method (PD-ADMM) for solutions that balance the perception-distortion trade-off.
- 3. Our LFc-SR model learned with the PD-ADMM optimization scheme yields HR images with high objective and perceptual quality. The proposed oneshot model requires significantly less computational expense than competing trade-off methods that rely on heavy post-processing.

2 Related Work

2.1 Objective Quality versus Perceptual Quality

SISR methods designed for objective quality are deep neural networks that use pixel-wise losses, such as L_1 and MSE. Starting with the CNN-based SRCNN [8], later variants leverage residual [14, 23, 19], dense [51, 40], and attention mechanisms [31, 33, 5]. These works strongly focus on architecture design and achieve high PSNR and SSIM scores. However, their results are blurry and evaluated poorly on perceptual metrics, especially at larger scaling factors like $\times 4$.

Many SISR methods often incorporate GANs [19, 44, 34, 29, 49, 26] and combine the adversarial loss with a content loss to improve the perceptual quality. Common content losses include the VGG [19] and contextual loss [29]. GANbased models yield sharper lines and more high-frequency details. They achieve high scores on perceptual measures such as LPIPS [47] and NRQM [25], attributed to implicit distribution learning. At the same time, these methods suffer from unrealistic artifacts resulting from adversarial training. 4 Y. Zhang et al.

2.2 Perception-Distortion Trade-off

Most state-of-the-art methods perform well either in objective or perceptual quality, but not in both. This trade-off phenomenon was explored by Blau *et al.* [2]; they attributed it to the incoherence between distortion and perception measurements. Some methods have tried to control the compromise between the two qualities. One solution was to fuse either two models [44] or model outputs [42,7,6] trained separately for objective and perceptual quality. The most notable of these post-processing works is WDST [7], which, like our work, set out to take separate frequency band considerations. A significant difference between our work and WDST is that WDST merges wavelet channels from the outputs of two separately trained networks, whereas our method is one-shot with a single model. Moreover, as WDST relies on a style-transfer [10] for the fusion, their inference procedure is computationally expensive and time-consuming, as the output channels are iteratively merged at test time. Our one-shot inference consumes significantly less computational resources and is much more efficient.

In addition to post-processing, another line of balancing approaches focuses on the training strategy. They balance the loss terms [41] or introduce a controlling factor as training inputs [32, 36, 43, 20]. For example, Vasu *et al.* [41] directly added MSE, VGG, and adversarial loss together and weighted them in different proportions. CFSNet [43] achieved a test-time transition between perceptual and objective quality by using an extra input factor α in their training strategy. These approaches create a smooth transition between the two qualities; but do not attempt to improve both of them. In contrast, our method focuses on the training strategy *and* improves both qualities.

3 Method

3.1 Revisiting the Multi-Objective SR Formulation

SISR recovers a high-resolution image $Y \in \mathbb{R}^{H \times W}$, given an input low-resolution image $X \in \mathbb{R}^{h \times w}$ via a neural network G, *i.e.* Y = G(X). Here, $h \times w$ and $H \times W$ are input and output image heights and widths respectively. As it is challenging to gather real-world pairs of images in low- and high-resolution, X is typically generated from a ground truth high-resolution image \hat{Y} via a downsampling function, *i.e.* $X = f_{\downarrow}(\hat{Y})$. Bicubic downsampling is a commonly used $f_{\downarrow}(\cdot)$, with standard downsampling factors of $\times 2$ or $\times 4$, *e.g.* $H = 2 \times h, W = 2 \times w$.

Considering both distortion and perception, a simple way to learn $G(\cdot)$ is by weighting different loss terms L_O and L_P with weights λ_O and λ_P^3 for objective and perceptual quality respectively, *i.e.*

$$L_G = \lambda_O \cdot L_O(Y) + \lambda_P \cdot L_P(Y). \tag{1}$$

However, this does not achieve a good PD trade-off because of the incoherence between perception and distortion losses [2]. Specifically, minimizing L_O

³ Throughout the paper, we will use O and P (either subscript or superscript) to denote objective- and perception-focused items respectively.

(*i.e.* with MSE, L_1) leads to low expectation of pixel-wise error, whereas L_P (*i.e.* adversarial loss combined with contextual loss) targets distribution-based learning for sharp results [2]. It is therefore common for methods seeking a good trade-off to learn two separate networks and fuse the results via post-processing, *i.e.*

$$Y = F(Y^O, Y^P), \quad \text{where} \quad Y^O = G_O(X), \text{ and } Y^P = G_P(X), \tag{2}$$

where G_O and G_P are two separately trained models and F is some fusion module. F can be quite complex and computationally expensive. For example, in Deng [6] and WDST [7], F is a style transfer module applied during inference.

Inspired by previous research in human perception [16, 54] and the direct link between low- and high-frequency information to objective and perceptual quality [7, 19], we take a frequency-banded approach. A naïve formulation is to separate the learning objectives for the low- and high-frequency bands, *i.e.*

$$Y = G(X), \text{ where } L_G = \lambda_O \cdot L_O(T_{LF}(Y)) + \lambda_P \cdot L_P(T_{HF}(Y)).$$
(3)

In this formulation, the network G is learned by applying a weighted loss on the low-pass and high-pass decomposed subbands of the output Y, based on the respective transformations T_{LF} and T_{HF} . Like previous works [7, 53, 24], we used the discrete wavelet transform with a Haar wavelet basis for decomposing the image into frequency subbands. Supervising a network based on the loss of Eq. (3) assumes that the different frequency bands can be optimized independently, which is not the case. A hard separation increases the difficulty of learning, especially for high-frequency bands $T_{HF}(Y)$, as L_P includes an adversarial loss and is, from our observations, unstable and prone to training collapse.

3.2 Low-Frequency Constrained SR (LFc-SR)

Due to the lack of independence, we contend that the optimization of each objective, be it objective or perceptual, must observe the entire band of image frequencies. Our solution to encourage stable learning of high-frequency information of Y while maintaining objective quality is to allow a separate (early stage) output Y' for the objective quality. Specifically, we partitioned the model G into two stages, where $G = \{\mathcal{G}_O, \mathcal{G}_P\}$ (see Fig. 2), *i.e.*

$$Y' = \mathcal{G}_O(X), \quad Y = \mathcal{G}_P(Y') \qquad s.t. \ \operatorname{T}_{LF}(Y) = \operatorname{T}_{LF}(Y'), \tag{4}$$

where $\mathcal{G}_O(\cdot)$ represents the first objective-focused stage of G and Y' is a superresolved image with high objective quality. $\mathcal{G}_P(\cdot)$ is the perception-focused stage targeting a further improvement on high-frequency detailing that takes Y' as input. To encourage the final output Y to maintain high objective quality, we placed a constraint that the low-frequency subbands of Y should be equal to those of Y', which in practice encourages the two to be similar.

The two stages accordingly have different loss functions. Outputs from \mathcal{G}_O are supervised by a L_1 loss:

$$L_O = L_1(Y', \hat{Y}) \tag{5}$$



Fig. 2: Proposed two-stage model with a low-frequency constraint. The first stage learns high-fidelity LF subbands to yield high objective quality. The second stage focuses on the distribution learning of HF subbands to improve perceptual quality for the overall image. Discrete Wavelet Transform (DWT) decreases the spatial size of each channel without information loss, and Inverse Discrete Wavelet Transform (iDWT) recovers the spatial resolution (see Sec. 4.1 for details). The two stages are trained with the low-frequency constraint to encourage the similarity of the LF subbands of the two outputs.

where \hat{Y} is the ground truth for the training input X. For \mathcal{G}_P , we used the contextual loss [29] together with adversarial training for high perceptual quality. The contextual loss approximates KL-divergence between the deep features of super-resolved and ground-truth images, and we refer to Mechrez *et al.* [30, 29] for further details. The overall loss function for \mathcal{G}_P is as follows:

$$L_P = \lambda_{CX} \cdot L_{CX}(Y, \hat{Y}) + \lambda_D \cdot L_1(f_{\downarrow}(Y), f_{\downarrow}(\hat{Y})) + \lambda_{Gen} \cdot L_{Gen}(Y), \quad (6)$$

where L_{CX} is the contextual loss and $f_{\downarrow}(\cdot)$ estimates the low-resolution counterparts of HR images. The L_1 loss here is, as introduced in Mechrez *et al.* [30], to boost the spatial structure similarity with the contextual loss, and L_{Gen} represents the generator loss used in adversarial training.

A straightforward way to implement the equivalence constraint in Eq. (4) is to apply a L_1 regularizer on the outputs of two stages, and the total loss is as following:

$$\lambda_O \cdot L_O + \lambda_P \cdot L_P + \lambda_r \cdot \|\mathbf{T}_{LF}(Y) - \mathbf{T}_{LF}(Y')\|_1.$$
(7)

Theoretically, regularization will also encourage LF similarity. However, minimizing Eq. (7) requires a difficult-to-achieve balance between L_O , L_P [4] and the selection of a good λ_r . To bypass this difficulty, we interpret Eq. (4) as a multi-objective optimization problem with a combinatorial constraint. ADMM is an established algorithm for solving such constrained optimizations and in the next section, we outline PD-ADMM, our ADMM-based optimization method.

3.3 Alternating Optimization with ADMM

We designed an alternating training method, PD-ADMM, based on the Alternating Direction Method of Multipliers (ADMM) [3] for the optimization of LFc-SR. Considering Eq. (4), finding solutions for the LFc-SR model is actually a constrained multi-objective optimization problem, *i.e.* there are two different objectives with the constraint on low-frequency coherency:

$$\min L_P + L_O \quad s.t. \quad \mathcal{T}_{LF}(Y) = \mathcal{T}_{LF}(Y'), \tag{8}$$

where L_P and L_O are loss functions in Sec. 3.2, and the constraint is the same as that in Eq. (4). Since the model parameters are variables in this optimization, we express Eq. (8) as functions of θ_P and θ_O , the parameters of \mathcal{G}_P and \mathcal{G}_O :

$$\min_{\theta_P,\theta_O} P(\theta_P) + O(\theta_O) \quad s.t. \ \mathrm{LF}_{\mathrm{p}}(\theta_P) = \mathrm{LF}_{\mathrm{O}}(\theta_O), \tag{9}$$

where $P(\theta_P)$ measures the perceptual error of the perceptual stage output, and $O(\theta_O)$ measures the objective error of the objective stage; $LF_P(\theta_P) = T_{LF}(\mathcal{G}_P(Y';\theta_P))$ and $LF_O(\theta_O) = T_{LF}(\mathcal{G}_O(X;\theta_O))$ represent the low-frequency subbands of outputs from the perceptual- and objective-focused stage.

The augmented Lagrangian function of Eq. (9) is:

$$\mathcal{L}(\theta_P, \theta_O, u) = P(\theta_P) + O(\theta_O) + u^T (\mathrm{LF}_{\mathrm{P}}(\theta_P) - \mathrm{LF}_{\mathrm{O}}(\theta_O)) + \frac{\rho}{2} \|\mathrm{LF}_{\mathrm{P}}(\theta_P) - \mathrm{LF}_{\mathrm{O}}(\theta_O)\|_2^2.$$
(10)

where u is the Lagrangian multiplier, and ρ is the penalty parameter. It is often easier to express the above function in the scaled form by defining $u = \rho s$ [48, 13,21], resulting in

$$\mathcal{L}(\theta_P, \theta_O, s) = P(\theta_P) + O(\theta_O) + \frac{\rho}{2} \left\| \mathrm{LF}_{\mathrm{P}}(\theta_P) - \mathrm{LF}_{\mathrm{O}}(\theta_O) + s \right\|_2^2 - \frac{\rho}{2} \left\| s \right\|_2^2.$$
(11)

ADMM solves Eq. (11) by a *decomposition* into three iterative sub-steps:

$$\theta_P^{k+1} = \operatorname*{argmin}_{\theta_P} \mathcal{L}(\theta_P, \theta_O^k, s^k)$$
(12a)

$$\theta_O^{k+1} = \operatorname*{argmin}_{\theta_O} \mathcal{L}(\theta_P^{k+1}, \theta_O, s^k)$$
(12b)

$$s^{k+1} = s^k + LF_P(\theta_P^{k+1}) - LF_O(\theta_O^{k+1})$$
 (12c)

The first two steps are equivalent to:

$$\underset{\theta_P}{\operatorname{argmin}} P(\theta_P) + \frac{\rho}{2} \left\| \operatorname{LF}_{\mathrm{P}}(\theta_P) - \operatorname{LF}_{\mathrm{O}}(\theta_O) + s \right\|_2^2$$
(13a)

$$\underset{\theta_O}{\operatorname{argmin}} O(\theta_O) + \frac{\rho}{2} \left\| \operatorname{LF}_{\mathrm{P}}(\theta_P) - \operatorname{LF}_{\mathrm{O}}(\theta_O) + s \right\|_2^2, \tag{13b}$$

where the first term of each minimization function is the loss function of the perception-focused stage or objective-focused stage, *i.e.* L_P or L_O discussed in Sec. 3.2. The second terms are special L_2 regularizers applied to variable s and the low-frequency subbands of the outputs from the two stages. Thus, we can solve these two sub-problems through a deep model optimization algorithm, such as Stochastic Gradient Descent. Based on above analysis through ADMM, our training algorithm for the LFc-SR model consists of three alternating steps and is concluded by dual variable updates in Eq. (12c) after optimizing the perceptual- and objective-focused stage through Eq. (13a) and Eq. (13b).

4 Experiments

4.1 Settings

Dataset & Evaluation: We used the DIV2K dataset [39] for training and validation, We evaluated on standard benchmarks: Set5 [1], Set14 [46], BSD100 [27], Urban100 [12] and Manga109 [28]. To measure objective quality, we computed the PNSR and SSIM on the image's Y channel with Matlab functions. For perceptual quality, we used the no-reference metric NRQM [25] and the full-reference metric LPIPS [47] for a comprehensive treatment. We provide visual comparisons in Fig. 4 and Fig. 5 to show the qualitative efficacy of our method.

Image Decomposition: We used the discrete wavelet transform (DWT) to decompose our image into frequency subbands. DWT is an invertible frequency transformation common in image processing [53, 24, 11]. It decomposes the image into four half-resolution channels, LL, HL, LH and HH, including low- (L) or high- (H) frequency subbands of height and width dimension. DWT can be applied on LL iteratively. We used the LL channel from a 1-level Haar DWT of the image as the LF information in Eq. (9). Other transformations are also feasible, and we experimented with Gaussian blur as an ablation. For downsampling function $f_{\downarrow}(\cdot)$ in Eq. (6), we also used DWT for convenience; we used the LL channel from a 2-level DWT for ×4 super-resolution.

Model Architecture: We tested our method for $\times 4$ super-resolution. For \mathcal{G}_O , we directly adopted the architecture of HAN [33], which finish upscaling the images into high resolution. \mathcal{G}_P starts with DWT and ends with inverse DWT (iDWT), saving training time by losslessly reducing the spatial resolution of processed feature maps. Between the two transformations are 15 Res-clique blocks designed by Zhong *et al.* [53]. For the architecture and loss calculation of discriminator for adversarial training, we adopted what was used in SRGAN [19]. More details are in the Supplementary.

Training Details: We took RGB patches of size 36×36 as inputs and trained the model with the ADAM optimizer [18] using the settings $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$ on a minibatch of 16. The two stages were pretrained with L_1 loss jointly for 400 epochs. The initial learning rate was 10^{-4} and decay rates were 0.1 every 100 epochs. We subsequently trained the model using the proposed PD-ADMM algorithm with a penalty factor $\rho = 10^{-4}$ and initialized s with 0. We trained the whole model for another 200 epochs with an initial learning rate of 5×10^{-5} and decreasing to one-eighth after 20, 50, 70, 100, and 140 epochs.

4.2 Comparison with State-of-the-Art Methods

Quantitative Comparison with Trade-off Methods: We first compare our method with existing methods aiming to balance objective and perceptual quality. These include CFSNet [43], G-MGBP [32], PESR [42], ESRGAN with network interpolation [44], Deng [6], and WDST [7]. We note, however, that both Deng and WDST rely on image style transfer [10] for post-processing, which requires much more computational power and impractical inference runtimes (see

detailed discussion on inference complexity in Sec. 4.2). Hence, a fair comparison with them is impossible. Nevertheless, as shown in Table 1, our method still reaches comparable performance to WDST and surpasses all other methods.

Table 1: Comparison with other PD trade-off state-of-the-art models aiming at both objective and perceptual quality on ×4 super-resolution. We color the best performance in red and second best in blue. We separate Deng [6] and WDST [7] as they rely on three orders of magnitude higher computational power and a fair comparison with other listed methods is impossible. For reference, we set the best performance of Deng and WDST that surpass all other methods to **bold**. Higher PSNR, SSIM, NRQM score and *lower* LPIPS score mean better performance.

Dateset	Metric	Deng [6]	WDST [7]	G-MGBP	[32] PESR [42]	ESRGAN [44]	CFSNet [43]	Ours
Set5 [1]	PSNR	31.14	31.46	30.87	30.76	31.11	31.00	31.79
	SSIM	0.8917	0.8929	0.8807	0.8915	0.8839	0.8894	0.8910
	NRQM	7.0022	7.5180	7.3155	7.1344	7.0724	7.4820	7.3462
	LPIPS	-	0.0868	0.1003	0.0884	0.0841	0.1020	0.0776
Set14 [17]	PSNR	27.77	28.07	27.56	27.57	27.53	27.61	27.87
	SSIM	0.8325	0.8356	0.8206	0.8322	0.8228	0.8280	0.8282
	NRQM	7.5575	7.6827	7.5042	7.5301	7.5936	7.6074	7.7957
	LPIPS	-	0.1658	0.1757	0.1612	0.1539	0.1754	0.1423
BSD100 [27]	PSNR	26.46	26.82	26.59	26.33	26.44	26.46	26.84
	SSIM	0.7048	0.7085	0.6926	0.6980	0.7002	0.6991	0.7010
	NRQM	8.4452	8.5948	8.1790	8.3298	8.3034	8.3770	8.4406
	LPIPS	-	0.2140	0.2238	0.2045	0.1903	0.2129	0.187
Urban100 [12]	PSNR	25.96	26.26	25.15	25.88	26.08	25.38	26.28
	SSIM	0.9620	0.9649	0.9495	0.9610	0.9624	0.9546	0.9636
	NRQM	6.4317	6.4556	6.2190	6.3190	6.1762	6.5140	6.6220
	LPIPS	-	0.1604	0.1775	0.1402	0.1519	0.1506	0.1235
Manga109 [28]	PSNR	-	-	29.07	28.77	29.72	29.49	30.20
	SSIM	-	-	0.8815	0.9795	0.9772	0.9789	0.9804
	NRQM	-	-	6.4073	6.6071	6.3909	6.5344	6.6532
	LPIPS	-	-	0.0779	0.0634	0.0610	0.0719	0.0627

Quantitative Comparison with Single-focused Methods: We further compare our method with single-focused state-of-the-art methods. Quantitatively comparing our balanced method with methods that only consider one type of quality is non-trivial. Due to the PD trade-off, our method will be inferior to the single-focused method in its optimized direction while surpassing it in the other. We thus make the comparison on a PD trade-off plane [2]. We generate a curve on the plane by interpolating the outputs from the state-of-the-art objective-focused and perceptual-focused methods. The interpolated image Y is produced through the following function:

$$Y = \alpha \cdot Y^O + (1 - \alpha) \cdot Y^P, \tag{14}$$

where Y^O is an HR image super-resolved by objective-focused model, Y^P is the perceptual-focused counterpart and the value of α is in [0, 1]. Freirich *et al.* [9]

10 Y. Zhang et al.

found this interpolation curve to be an optimal PD trade-off boundary achievable using single-focused estimators. Methods that sit in the region below the curve are better than those on or above the curve in terms of PD trade-off. In Fig. 3, we plot the curve between the best estimators, NLSA [31] and RankSRGAN [49], and show that our method is the only one below the estimated trade-off boundary considering all single-focused state-of-the-art methods on the plane.



Fig. 3: Comparison with single-focused start-of-the-art SR methods on Urban100 [12] and Manga109 [28] datasets for $\times 4$ super-resolution. We plot the curve by interpolating results from the best objective estimator [31] and the best perceptual estimator [49] on the plane. Our model can reach a better balance than the trade-off boundary approximated by single-focused state-of-the-art methods, as it sits well below the boundary.

Qualitative Results: We compare image samples generated from our method with those from NLSA [31], RankSRGAN [49] and WDST [7] on BSD100 [27] and Urban100 [12] in Fig. 4. Our method generates clear structures with fewer artifacts compared to single-focused methods and more realistic high-frequency details than WDST does. We provide frequency subbands comparison in Fig. 5. While single-focused methods only achieve good results at one of the subbands, our method learns both low- and high-frequency information effectively.

Inference Complexity: A significant advantage of our method is that it is one-shot. We directly optimize a single model and bypass the complex fusing of two outputs, unlike competing work WDST [7] and its predecessor, Deng [6]. Central to their inference is an image style transfer procedure [10], which requires thousands of iterations to update the initial input signal towards a fused wavelet channel or image. In contrast, our method super-resolves the LR image through a single model in a one-shot manner. We used a 24GB NVIDIA RTX A5000 GPU for running the models. Table 2 provides a quantitative comparison of the computational complexity, including the FLOPs and runtime for $\times 4$ superresolution of a 128×128 LR input. Our model's FLOPs and runtime are three



Fig. 4: Image sample comparison for $\times 4$ super-resolution on BSD100 [27] (first two rows) and Urban100 [12] (thrid and forth rows) dataset. NLSA [31] is objective-focused, RankSRGAN [49] is perceptual-focused, while WDST [7] and our method are balanced approaches. Our method retains relatively sharp detailing without introducing unnatural artifacts.

orders of magnitude lower than WDST, assuming only 1000 iterations of updates for WDST (the actual number of updates is usually much more).

4.3 Ablation Studies

Verification of low-frequency constraint and ADMM optimization: In this experiment, we kept a constant architecture as described in Sec. 4.1. Then we implemented the equivalence constraint through the second-stage regularizer in Eq. (7) and trained the models with different weights λ_r for the regularizer. Fig. 6 shows the performance of regularizer-based methods on the PD trade-off plane, together with our PD-ADMM method; for each hyperparameter setting, we also illustrate the LF difference between the two stages' outputs, *i.e.* Y' and Y, with a circle. The more different the two stages, the larger the circle. Compared to the no-constraint version ($\lambda_r = 0$), both using PD-ADMM and adding the low-frequency constraint yield higher LF similarity. Fig. 7 also provides visual evidence of the efficacy of adding the low-frequency constraint. The increase of λ_r results in higher LF similarity, transitioning from better perceptual quality to better objective quality. However, the model optimized by PD-ADMM achieves a superior balance of perceptual- and objective-quality to this transition.

Stage Order and Constrained Subband: This experiment validates the reconstruction ordering of frequency subbands in our proposed model. Without



Fig. 5: Visual comparison of low- and high-frequency subbands for $\times 4$ superresolution by NLSA [31], RankSRGAN [49] and our method. We visualize the LL (first row) and HH (second row) wavelet channels for low- and high-frequency comparison. Within each image, we show the HR image (top left), wavelet channel (top right) and the zoom-in region of the wavelet channel under them. NLSA generates high-fidelity low-frequency information but lacks high-frequency detailing in the HH channel; RankSRGAN produces twisted structures in the LL channel but can reconstruct more high-frequency information. Our method predicts both high-fidelity LL channel and abundant high-frequency detailing.



Fig. 6: Comparison with regularizer-based method on BSD100 [27] and Urban100 [12]. λ_r is the weight of regularizer; numbers below λ_r are mean absolute error (MAE) of low-frequency subbands between outputs of two stages. A lower value means a higher similarity. Our method achieves appropriate LF similarity and better balance than the curve plot by regularizer-based models.

changing model architecture, we exchanged the optimization objectives of the two stages and converted the equivalence constraint to HF subbands extracted by DWT. Representing our original method as the O-P model, this swapped P- Table 2: Comparison of the inference complexity of our method and WDST [7] for $\times 4$ super-resolution of a 128×128 LR image. The SR model super-resolves LR inputs and gives HR images, while post-processing further process super-resolved images. Our model's FLOPs and runtime are three orders of magnitude less than WDST. We refer to the supplementary material for detailed analysis.

Method	Values	SR model	Post-processing	All
WDST	FLOPs	$50.6 \mathrm{M}$	pprox 37448.1M	$\approx 37500.3M$
	Run-time	0.918s	$\approx 1152s$	$\approx 1153s$
Ours	FLOPs	$26.8 \mathrm{M}$	-	26.8M
	Run-time	0.566s	-	0.566s



Fig. 7: Visualization of *difference* in low-frequency subband compared to the ground-truth for $\times 4$ super-resolution. Each image shows the HR result *(top left)*, grayscale map of difference of the LL channel between the variant and the ground truth *(top right)*, and the zoom-in region of the grayscale map. The darker the map, the better. Limiting the change of the LF subband keeps higher accuracy of the LF information while learning the HF information at the perceptual stage.

O model completes perceptual-focused reconstruction first, and then optimizes LF subbands while keeping HF equivalence to the first stage. We also optimized the P-O model through PD-ADMM. Borrowing the comparison with single-focused models in Sec. 4.2, we plotted the results of our original model and P-O model on a PD trade-off plane with the trade-off boundary. As shown in Fig. 8, the performance of the order-exchanged model is dominated by the objective optimization and fails to achieve a good PD trade-off.

Low-frequency Subbands Extraction: To verify that our results are not dependent on specific frequency decomposition method, *i.e.* the DWT, we applied Gaussian blur on the outputs from the two stages to extract the LF subbands that should satisfy the equivalence constraint in PD-ADMM. Experimentally, we applied convolution with a 21×21 Gaussian kernel with $\sigma = 3$ to extract the LF information. As shown in Fig. 8, models with DWT and with Gaussian Blur (GB) have extremely close positions on the PD trade-off plane, and both of them are under the PD trade-off boundary introduced in Sec. 4.2. It shows that using different extraction methods can reach similar PD balances. We chose DWT for our model because of its popularity in image frequency analysis.



Fig. 8: Comparison between our model and its two variants. The order-exchanged version (P-O) optimizes HF subband for perceptual quality first. Its performance is close to the top-left corner and above the estimated trade-off boundary, *i.e.* it fails to balance the perceptual and objective quality. The other variant (GB) implements the low-frequency constraint of LFc-SR by Gaussian Blur rather than DWT used in our original model. The close performance of two models shows that our constraint design does not rely on a specific extraction method.

5 Conclusion & Limitations

In this paper, we learned a SISR model for both objective and perceptual quality. We proposed LFc-SR, a two-stage model trained with a low-frequency constraint to implicitly optimize low- and high-frequency bands for objective and perceptual quality in a successive way. The training of LFc-SR is formulated as a constrained multi-objective optimization problem. We designed PD-ADMM, an alternating algorithm, to solve for a solution that balances the perception-distortion trade-off. Our method is effective compared with other PD trade-off related methods and single-focused models without the expense of heavy fusion procedure.

Despite being more efficient than trade-off related fusion methods, our method is still far from optimal compared with some efficiency-oriented SISR models. Due to the successive processing in stages, some parts of the two-stage model may extract similar information and cause redundancy. However, such a design also offers the potential to reuse information efficiently, which we will leave for future work.

Acknowledgement This research is supported by Singapore Ministry of Education (MOE) Academic Research Fund Tier 1 T1251RES1819.

15

References

- 1. Bevilacqua, M., Roumy, A., Guillemot, C., Alberi-Morel, M.L.: Low-complexity single-image super-resolution based on nonnegative neighbor embedding (2012)
- Blau, Y., Michaeli, T.: The perception-distortion tradeoff. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6228–6237 (2018)
- 3. Boyd, S., Parikh, N., Chu, E.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Now Publishers Inc (2011)
- 4. Chen, T.: A Fast Reduced-Space Algorithmic Framework for Sparse Optimization. Ph.D. thesis, Johns Hopkins University (2018)
- Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-order attention network for single image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11065–11074 (2019)
- Deng, X.: Enhancing image quality via style transfer for single image superresolution. IEEE Signal Processing Letters 25(4), 571–575 (2018)
- Deng, X., Yang, R., Xu, M., Dragotti, P.L.: Wavelet domain style transfer for an effective perception-distortion tradeoff in single image super-resolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3076–3085 (2019)
- Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. IEEE transactions on pattern analysis and machine intelligence 38(2), 295–307 (2015)
- 9. Freirich, D., Michaeli, T., Meir, R.: A theory of the distortion-perception tradeoff in wasserstein space. Advances in Neural Information Processing Systems **34** (2021)
- Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2414–2423 (2016)
- Guo, T., Seyed Mousavi, H., Huu Vu, T., Monga, V.: Deep wavelet prediction for image super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 104–113 (2017)
- 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)
- Kiaee, F., Gagné, C., Abbasi, M.: Alternating direction method of multipliers for sparse convolutional neural networks. arXiv preprint arXiv:1611.01590 (2016)
- Kim, J., Lee, J.K., Lee, K.M.: 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)
- Kim, J., Lee, J.K., Lee, K.M.: Deeply-recursive convolutional network for image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1637–1645 (2016)
- Kim, J., Lee, S.: Deep learning of human visual sensitivity in image quality assessment framework. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1676–1684 (2017)
- Kim, K.I., Kwon, Y.: Single-image super-resolution using sparse regression and natural image prior. IEEE transactions on pattern analysis and machine intelligence 32(6), 1127–1133 (2010)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)

- 16 Y. Zhang et al.
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image superresolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4681–4690 (2017)
- Lee, H., Kim, T., Son, H., Baek, S., Cheon, M., Lee, S.: Smoother network tuning and interpolation for continuous-level image processing. arXiv preprint arXiv:2010.02270 (2020)
- Leng, C., Dou, Z., Li, H., Zhu, S., Jin, R.: Extremely low bit neural network: Squeeze the last bit out with admm. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
- Li, C., Chen, C., Zhang, B., Ye, Q., Han, J., Ji, R.: Deep spatio-temporal manifold network for action recognition. arXiv preprint arXiv:1705.03148 (2017)
- Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 136–144 (2017)
- Liu, P., Zhang, H., Zhang, K., Lin, L., Zuo, W.: Multi-level wavelet-cnn for image restoration. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 773–782 (2018)
- Ma, C., Yang, C.Y., Yang, X., Yang, M.H.: Learning a no-reference quality metric for single-image super-resolution. Computer Vision and Image Understanding 158, 1–16 (2017)
- Ma, C., Rao, Y., Cheng, Y., Chen, C., Lu, J., Zhou, J.: Structure-preserving super resolution with gradient guidance. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7769–7778 (2020)
- 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: Proc. 8th Int'l Conf. Computer Vision. vol. 2, pp. 416–423 (July 2001)
- Matsui, Y., Ito, K., Aramaki, Y., Fujimoto, A., Ogawa, T., Yamasaki, T., Aizawa, K.: Sketch-based manga retrieval using manga109 dataset. Multimedia Tools and Applications 76(20), 21811–21838 (2017). https://doi.org/10.1007/s11042-016-4020-z
- Mechrez, R., Talmi, I., Shama, F., Zelnik-Manor, L.: Maintaining natural image statistics with the contextual loss. In: Asian Conference on Computer Vision. pp. 427–443. Springer (2018)
- Mechrez, R., Talmi, I., Zelnik-Manor, L.: The contextual loss for image transformation with non-aligned data. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 768–783 (2018)
- Mei, Y., Fan, Y., Zhou, Y.: Image super-resolution with non-local sparse attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3517–3526 (2021)
- Navarrete Michelini, P., Zhu, D., Liu, H.: Multi-scale recursive and perceptiondistortion controllable image super-resolution. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018)
- Niu, B., Wen, W., Ren, W., Zhang, X., Yang, L., Wang, S., Zhang, K., Cao, X., Shen, H.: Single image super-resolution via a holistic attention network. In: European conference on computer vision. pp. 191–207. Springer (2020)
- Park, S.J., Son, H., Cho, S., Hong, K.S., Lee, S.: Srfeat: Single image superresolution with feature discrimination. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 439–455 (2018)

- Ramamonjisoa, M., Firman, M., Watson, J., Lepetit, V., Turmukhambetov, D.: Single image depth prediction with wavelet decomposition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11089– 11098 (2021)
- Shoshan, A., Mechrez, R., Zelnik-Manor, L.: Dynamic-net: Tuning the objective without re-training for synthesis tasks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3215–3223 (2019)
- Singh, V., Mittal, A.: Wdn: A wide and deep network to divide-and-conquer image super-resolution. IEEE Journal of Selected Topics in Signal Processing 15(2), 264– 278 (2020)
- Takapoui, R., Moehle, N., Boyd, S., Bemporad, A.: A simple effective heuristic for embedded mixed-integer quadratic programming. International journal of control 93(1), 2–12 (2020)
- Timofte, R., Agustsson, E., Van Gool, L., Yang, M.H., Zhang, L.: Ntire 2017 challenge on single image super-resolution: Methods and results. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 114–125 (2017)
- Tong, T., Li, G., Liu, X., Gao, Q.: Image super-resolution using dense skip connections. In: Proceedings of the IEEE international conference on computer vision. pp. 4799–4807 (2017)
- Vasu, S., Thekke Madam, N., Rajagopalan, A.: Analyzing perception-distortion tradeoff using enhanced perceptual super-resolution network. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018)
- Vu, T., Luu, T.M., Yoo, C.D.: Perception-enhanced image super-resolution via relativistic generative adversarial networks. In: Proceedings of the European conference on computer vision (ECCV) workshops. pp. 0–0 (2018)
- Wang, W., Guo, R., Tian, Y., Yang, W.: Cfsnet: Toward a controllable feature space for image restoration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4140–4149 (2019)
- Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: Esrgan: Enhanced super-resolution generative adversarial networks. In: Proceedings of the European conference on computer vision (ECCV) workshops. pp. 0–0 (2018)
- 45. Yang, X., Xiang, W., Zeng, H., Zhang, L.: Real-world video super-resolution: A benchmark dataset and a decomposition based learning scheme. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4781–4790 (2021)
- Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparserepresentations. In: International conference on curves and surfaces. pp. 711–730. Springer (2010)
- 47. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
- Zhang, T., Ye, S., Zhang, K., Tang, J., Wen, W., Fardad, M., Wang, Y.: A systematic dnn weight pruning framework using alternating direction method of multipliers. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 184–199 (2018)
- Zhang, W., Liu, Y., Dong, C., Qiao, Y.: Ranksrgan: Generative adversarial networks with ranker for image super-resolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3096–3105 (2019)

- 18 Y. Zhang et al.
- Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: Proceedings of the European conference on computer vision (ECCV). pp. 286–301 (2018)
- Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2472–2481 (2018)
- Zhao, P., Liu, S., Wang, Y., Lin, X.: An admm-based universal framework for adversarial attacks on deep neural networks. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 1065–1073 (2018)
- Zhong, Z., Shen, T., Yang, Y., Lin, Z., Zhang, C.: Joint sub-bands learning with clique structures for wavelet domain super-resolution. Advances in neural information processing systems **31**, 165–175 (2018)
- Zhou, Y., Deng, W., Tong, T., Gao, Q.: Guided frequency separation network for real-world super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 428–429 (2020)