

Approaching the Limit of Image Rescaling via Flow Guidance

Shang Li¹²

lishang2018@ia.ac.cn

Guixuan Zhang²

guixuan.zhang@ia.ac.cn

Zhengxiong Luo¹²

zhengxiong.luo@cripac.ia.ac.cn

Jie Liu^{✉2}

jie.liu@ia.ac.cn

Zhi Zeng²

zhi.zeng@ia.ac.cn

Shuwu Zhang²³

shuwu.zhang@ia.ac.cn

¹ School of Artificial Intelligence

University of Chinese Academy of
Sciences

Beijing, China

² Institute of Automation

Chinese Academy of Sciences
Beijing, China

³ Beijing Engineering Research Center

for Digital Content Technology

Beijing, China

Abstract

Image downscaling and upscaling are two basic rescaling operations. Once the image is downscaled, it is difficult to be reconstructed via upscaling due to the loss of information. To make these two processes more compatible and improve the reconstruction performance, some efforts model them as a joint encoding-decoding task, with the constraint that the downscaled (*i.e.* encoded) low-resolution (LR) image must preserve the original visual appearance. To implement this constraint, most methods guide the downscaling module by supervising it with the bicubically downscaled LR version of the original high-resolution (HR) image. However, this bicubic LR guidance may be sub-optimal for the subsequent upscaling (*i.e.* decoding) and restrict the final reconstruction performance. In this paper, instead of directly applying the LR guidance, we propose an additional invertible flow guidance module (FGM), which can transform the downscaled representation to the visually plausible image during downscaling and transform it back during upscaling. Benefiting from the invertibility of FGM, the downscaled representation could get rid of the LR guidance and would not disturb the downscaling-upscaling process. It allows us to remove the restrictions on the downscaling module and optimize the downscaling and upscaling modules in an end-to-end manner. In this way, these two modules could cooperate to maximize the HR reconstruction performance. Extensive experiments demonstrate that the proposed method can achieve state-of-the-art (SotA) performance on both downscaled and reconstructed images.

1 Introduction

With the tremendous advances of mobile devices and web applications, the demand has surged for image downscaling to downscale high-resolution (HR) images to smaller-sized

ones, which can save storage, accelerate transmission speed, or fit low-resolution (LR) screens while preserving the original visual appearance. Meanwhile, the inverse process, image upscaling, is also indispensable to reconstruct the underlying HR image of the down-scaled LR input. In previous researches, image downscaling and upscaling are studied as two independent tasks, and many methods have been proposed for downscaling [10, 16, 23, 25] and upscaling. One of the most well-known downscaling methods is bicubic interpolation [23], and upscaling is studied as super-resolution (SR) [8, 9, 12, 14, 21, 22]. However, due to the information loss during downscaling, the reconstruction performance of the subsequent upscaling is heavily restricted [21]. More and more efforts focus on the impact of downscaling for the final reconstruction performance and attempt to find the upscaling-optimal downscaling method.

To make downscaling and upscaling more compatible, some recent methods model these two processes as a joint encoding-decoding task [11, 18, 29]. Specifically, the downscaling module is regarded as the encoder to downscale the original HR image into a smaller-sized representation, while the upscaling module is the decoder to reconstruct a more detailed HR image from the downscaled representation. However, there is an additional constraint: the downscaled representation must preserve the visual appearance of the original HR image, that is, it should be a visually plausible downscaled image. To implement this constraint, most existing methods employ the bicubically downscaled version of the original HR image to supervise downscaling module, which is termed bicubic LR guidance [11]. This guidance helps the downscaling module produce a visually plausible LR image, but this may not be optimal for the subsequent upscaling and restrict the HR reconstruction performance.

To address this problem, we design a flow guidance rescaling network (FGRN), which can remove the aforementioned restriction on the downscaled representation. Specifically, FGRN is an encoding-decoding based network with an invertible flow guidance module (FGM). FGM can transform the downscaled representation into a visually pleasing LR image during downscaling and transform this image back to the original representation during upscaling, which could guarantee the visual quality of the downscaled image. Moreover, the lossless transformation of FGM could help relax restrictions on the downscaled representation. Thus, the downscaling and upscaling modules can be jointly optimized in an end-to-end manner without applying any LR guidance to the downscaling module. In this way, the downscaled representation could be upscaling-optimal and maximize the HR reconstruction performance.

In summary, our main contributions are as follows:

- To the best of our knowledge, our proposed flow guidance rescaling network (FGRN) is the first image rescaling method without any constraints on the downscaling process. With the help of the invertible flow guidance module, the downscaling and upscaling modules can be optimized in an end-to-end manner and cooperate to maximize the HR reconstruction performance.
- We design an invertible flow guidance module (FGM), which can perform lossless transformation between downscaled representations and visually favorable LR images. It enables us to remove the restriction on the downscaling process and helps improve the final reconstruction performance.
- Extensive experiments on different benchmark datasets demonstrate that the proposed method can achieve SotA performance on both downscaled and reconstructed images with fewer parameters.

2 Related Work

2.1 Image Rescaling

Image downscaling and upscaling are two basic image rescaling tasks. Image downscaling aims to reduce the resolution of the HR image while keeping its visual appearance. Image upscaling aims to reconstruct a visually pleasing HR image through the given LR image. In early researches, they are treated as independent tasks. One of the most well-known downscaling methods is Bicubic [23]. And upscaling is studied as super-resolution (SR) [9] and plays a significant role in low-level vision tasks [4, 9]. However, due to the information loss in downscaling, the performance of the subsequent upscaling is heavily restricted.

Toward this issue, some recent efforts attempt to model these two processes as a united task, thus the downscaling method could be “customized” for the upscaling method. Although these methods make downscaling and upscaling cooperate better, there are still some limitations. For instance, [10] and [18] jointly optimize downscaling and upscaling with an encoding-decoding based framework, and they apply the bicubic LR guidance to guarantee the visual quality of the downsampled image. However, a visually plausible downsampled image may not be the optimal representation for upscaling. Besides, [3] uses an invertible network to model downscaling and upscaling according to their reciprocal nature. But the invertibility of [3] restricts that these two processes must share the same network parameters in opposite directions, which is rigid and may largely limit the representation capability. In this paper, the structure of our rescaling module is encoding-decoding based and free to be adjusted. Moreover, we remove the restriction on the downscaling module by proposing an invertible flow guidance module. It can guarantee the visual quality of the downsampled image and help the rescaling module maximize the HR reconstruction performance.

Difference to SR. Note that image upscaling in this paper is different from SR. In our scenario, the ground-truth HR image is available at the beginning but somehow is discarded and we store/transmit the LR version instead. Our goal is to recover the HR image using this LR image. While for SR, the real HR is unavailable in applications and the task is to generate new HR images for LR ones.

2.2 Normalizing Flow

Normalizing flows aim to map the complex data distribution $p_X(x)$ to a tractable prior distribution $p_Z(z)$ using the bijective function F , that is $z = F(x)$. Thus the generative process over X is defined as $x = F^{-1}(z)$. Given such bijective function F , the *change of variable formula* defines the model distribution on X by:

$$p_X(x) = p_Z(F(x)) \left| \det \left(\frac{\partial F(x)}{\partial x^T} \right) \right| \quad (1)$$

where $\left| \det \left(\frac{\partial F(x)}{\partial x} \right) \right|$ is the absolute value of the determinant of the Jacobian matrix $\frac{\partial F(x)}{\partial x}$. A stack of such invertible transformations can form a complex final density. NICE [9] is a flow model that stacks non-linear additive coupling and other transformation layers. Based on NICE, RealNVP [6] further upgrades additive coupling to affine coupling without loss of invertibility and achieves better performance. Then Glow [15] replace the fixed permutation layer in [6] with the more flexible 1×1 convolutional layer. Nowadays, normalizing flows have been successfully applied in generating images [8, 20], audio data [13] and point cloud data [26, 34]. As for image rescaling, IRN [3] is an invertible rescaling network to model image downscaling and upscaling. However, this invertible model only allows these two

processes to share the same parameters in opposite directions, and the prior distribution in IRN could not compensate for the information lost in downscaling. In this paper, we use flow to perform lossless transformation just between downsampled representations and visually pleasing LR images, instead of the overall image rescaling network.

3 Method

3.1 Bicubic LR guidance

To alleviate the loss of information during downscaling, recent researches [10, 18, 33] adopts an encoding-decoding based framework. The downscaling module G_d is jointly optimized with the upscaling module G_u . They hope that G_d could learn to be upscaling-optimal. The overall downscaling-upscaling process can be expressed as:

$$\hat{x} = G_d(y), \quad \hat{y} = G_u(\hat{x}), \quad (2)$$

where y is the original HR image, \hat{x} is the downsampled image, and \hat{y} is the reconstructed image. To ensure the reconstruction performance, \hat{y} is directly supervised with y via L1 loss. To ensure that \hat{x} keeps the visual appearance of y , it is usually supervised with the bicubically downsampled version of y , which is what we call bicubic LR guidance. Although G_d is jointly optimized with G_u in these methods, the reconstruction performance may still be restricted, because the bicubic LR guidance may get conflict with the reconstruction quality. To maximize the reconstruction performance, it is better to minimize the influence of constraint on G_d .

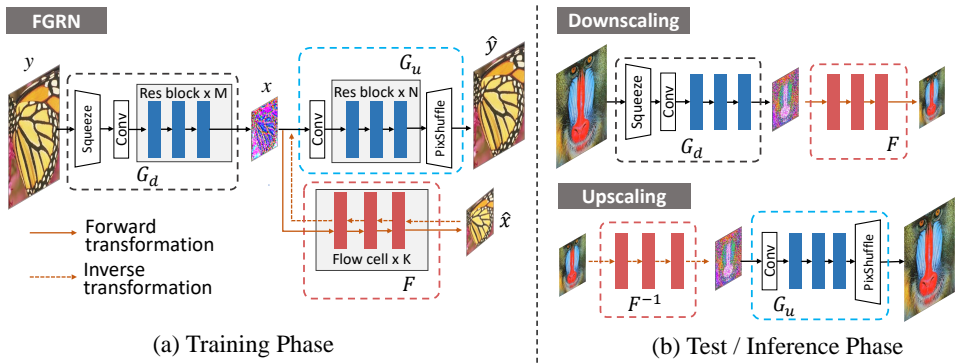


Figure 1: The network architecture of FGRN.

3.2 Invertible Flow Guidance

Different from the methods above, we do not directly apply bicubic LR guidance on G_d . In our method, G_d and G_u are optimized in an end-to-end manner. They are only responsible for improving the quality of the upscaled HR image \hat{y} without considering the visual quality of the downsampled representation, which can be formulated as:

$$x = G_d(y), \quad \hat{y} = G_u(x), \quad (3)$$

where x is downsampled representations. We call it “representations” instead of “image”, because x may not have any visual appearances, since we do not apply any constraint on it.

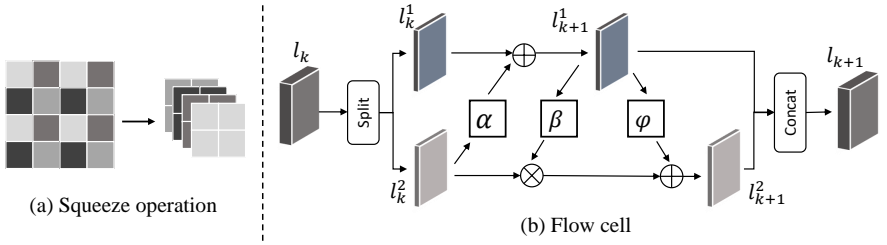


Figure 2: (a) The squeeze operation in G_d . (b) The structure of flow cell in F .

In this way, it is easier for G_d to become optimal for the upscaling task and help maximize the HR reconstruction performance. It will be experimentally proved in Sec 4.4.

As for achieving a visually plausible downscaled image, we further introduce an invertible flow guidance which can perform lossless transformation between the downscaled representation and the visually plausible LR image:

$$\hat{x} = F(x), \quad x = F^{-1}(\hat{x}). \quad (4)$$

where F is our flow guidance module (FGM). During downscaling, F can transform the downscaled representation x (*i.e.* $G_d(y)$) to a visually plausible LR image \hat{x} . While during upscaling, F^{-1} can inversely transform the given LR image \hat{x} to the original representation x which will be subsequently upscaled into an excellent reconstructed HR image.

3.3 Network Architecture

As can be seen in Figure 1 (a), the proposed FGRN includes three parts: 1) the downscaling module G_d , 2) the upscaling module G_u , and 3) the flow guidance module F .

Downscaling module. G_d is responsible for mapping the input HR image to a down-scaled representation. It consists of a squeeze layer [5], a shallow convolutional layer, and a stack of simple residual blocks (ResBlock) [19, 56]. As shown in Figure 2 (a), the squeeze layer can reshape the input HR image of size $sH \times sW \times C$ into the representation of size $H \times W \times s^2C$ (s denotes the scale factor), which compresses the spatial resolution but will not cause information loss.

Upscaling module. G_u aims to reconstruct the original HR image through the down-scaled LR image. The structure of G_u is similar to many SR models [19, 65] that contain a stack of residual blocks for different feature extraction and the pixel shuffle (PixShuffle) [28] operation for enlarging the spatial resolution.

Flow guidance module. F is composed of stacked invertible flow cells [6, 63], which can be defined as $F = f_1 \circ f_2 \circ \dots \circ f_K$. As shown in Figure 2 (b), for the k^{th} cell, the input l_k is split into l_k^1 and l_k^2 along the channel axis, and they undergo the additive affine transformations [9]. Here, to enhance the transformation, we employ an augmented version of the identity branch referring to [9]. The forward and inverse transformations can be expressed as Eq. 5 and Eq. 6 respectively:

$$\begin{aligned} l_{k+1}^1 &= l_k^1 + \alpha(l_k^2), \\ l_{k+1}^2 &= l_k^2 \otimes \exp(\beta(l_{k+1}^1)) + \phi(l_{k+1}^1). \end{aligned} \quad (5)$$

$$\begin{aligned} l_k^1 &= l_{k+1}^1 - \alpha(l_k^2), \\ l_k^2 &= (l_{k+1}^2 - \varphi(l_{k+1}^1)) \otimes \exp(-\beta(l_{k+1}^1)). \end{aligned} \quad (6)$$

where $\alpha(\cdot)$, $\beta(\cdot)$ and $\varphi(\cdot)$ can be arbitrary transformation functions. Here we use Dense Block [57] for all three functions. Due to the invertibility, F can perform lossless transformation between the downsampled representation and the visually pleasing LR image.

As shown in Figure 1 (a), during training, the original HR image is firstly downsampled by G_d to the 3-channel representation x . Then x will be processed by G_u and F respectively. G_u aims to reconstruct a visually pleasing HR image through \hat{x} . While F tries to transform it forward into the visually plausible LR image. During the test (as shown in Figure 1 (b)), the HR image is downsampled by G_d to the upscaling-optimal representation and then transformed by F to the downsampled LR image, which is the downscaling process. For upscaling, the given LR image is inversely transformed by F^{-1} to the corresponding upscaling-optimal representation, and then it is reconstructed by G_u into the final upscaled image.

Quantization. To convert the values of output LR images from floating-point numbers to 8-bit unsigned integers, we adopt the same quantization module in IRN [53]. Although the rounding operation is non-differentiable, the quantization module straightly uses the gradients before rounding to approximate the gradients after rounding. In this way, the quantization module can be inserted into any trainable network. To ensure the transformation of the flow guidance module is lossless and invertible, we also use the same module to convert the downsampled representations to quantized values.

3.4 Training Objectives

Since the joint downscaling-upscaling module and the flow guidance module serve different purposes, we optimize them with different optimization objectives.

HR Reconstruction. Since G_d and G_u are responsible for improving the HR reconstruction performance without considering the visual quality of the downsampled image, we jointly optimize these two modules as follows:

$$\mathcal{L}_{rec} = \mathbb{E}_{y \sim p_Y} \|y - G_u((G_d(y; \theta_{G_d})); \theta_{G_u})\|_1 \quad (7)$$

Here we employ the original HR image y as supervision and $L1$ loss as the difference metric.

Flow Guidance. Instead of the commonly used Bicubic LR guidance, we employ flow guidance to supervise the downscaling process. To optimize the invertible flow guidance module F , we need to specify each item in Eq 1. We assume the prior distribution $p_Z(z) \sim N(y_{bic}, \sigma^2)$. Here y_{bic} denotes the Gaussian mean of which the value is the bicubically downsampled version of the original HR image, and σ^2 is the variance. This assumption makes sense because y_{bic} can ensure the visual rationality of the downsampled image, while σ increases uncertainty which can alleviate limitations of the Bicubic downscaling method. Thus, the downscaling performance of our FGRN is more likely to outperform Bicubic, which will be experimentally demonstrated in Sec 4.4. As for the determinant of the Jacobian matrix, the transformation expressed in Eq 5 can make $\frac{\partial F(x)}{\partial x^T}$ triangular and efficient to calculate the determinant [2]. Accordingly, the specific objective of Eq 1 for maximizing the

log-likelihood can be reformulated as follows:

$$\begin{aligned} \log p_X(x) &= \log p_Z(F(x)) + \log \left| \det \left(\frac{\partial F(x)}{\partial x^T} \right) \right| \\ &= -\frac{1}{2\sigma^2} \|F(x) - y_{bic}\|_2 + \eta + \sum_{i=1}^K \log \left| \det \frac{\partial f_i(l_{i-1})}{\partial l_{i-1}} \right| \end{aligned} \quad (8)$$

where η is a constant, and we define $l_0 = x$, $l_i = f_i(l_{i-1})$ and $l_K = F(x)$. Thus, we optimize F by modifying the the negative log-likelihood of Eq 8:

$$\mathcal{L}_{guide} = \mathbb{E}_{x \sim p_X} \|F(x; \theta_F) - y_{bic}\|_1 - \gamma \sum_{i=1}^K \log \left| \det \frac{\partial f_i(l_{i-1}; \theta_{f_i})}{\partial l_{i-1}} \right| \quad (9)$$

where γ is a weight coefficient positively related to the variance σ^2 . If $\sigma^2 \rightarrow 0$, \mathcal{L}_{guide} will approximately degenerate to the LR guidance loss [14]. Here we set $\gamma = 1e - 3$.

3.5 Difference to IRN

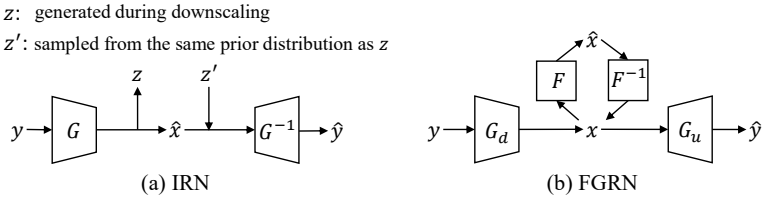


Figure 3: The framework of IRN [14] (a) and our FGRN (b).

There are several differences and advantages of FGRN compared to the popular rescaling method IRN [14]. 1) As shown in Figure 3, IRN models downscaling and upscaling with an invertible network, which is inspired by the reciprocal nature of this pair of rescaling tasks. While the motivation of our FGRN is to find the upscaling-optimal downscaled representation for the encoding-decoding based image rescaling model. 2) IRN relies on an extra latent variable z to help recover the lost information during downscaling. However, We experimentally find that the latent variable fails to formulate the lost information and contributes little to the reconstruction performance. The detailed experiments can be referred to Sec 4.6. While in our work, the lost information could be adaptively recovered by the upscaling module, since the downscaling and upscaling modules are jointly optimized. 3) To keep its invertibility, the downscaling and upscaling modules in IRN must share the same parameters in opposite directions, which may largely limit its representation capability. While our downscaling and upscaling modules are composited of ordinary convolutional layers, which allows us to freely adjust their structures. Besides, the two modules do not share any parameters and potentially have larger representation capability.

4 Experiments

4.1 Experimental Setup

Datasets and Metrics. We employ 3450 high-quality (2K resolution) images as the training data, including 800 images from the training set of DIV2K [14] and 2650 images from




Figure 4: PSNR/SSIM and visual comparison of upscaling the $\times 4$ downscaled image.

Flickr2K [35]. For the test, we use five standard benchmark datasets: Set5, set14, BSD100, Urban 100, and the validation set of DIV2K. The visual quality of upscaled and downsampled results is evaluated on the Y channel (*i.e.* luminance) of transformed YCbCr space. Since the downsampled results have no corresponding ground truth, they are evaluated with no-reference metrics: Natural Image Quality Evaluator (NIQE) [24] and Perceptual-based Quality Evaluator (PIQE) [30]. Lower NIQE and PIQE indicate higher perceptual quality. While the upscaled results are evaluated with full-reference metrics: PSNR and SSIM [32].

Implement Details. For the training data, the batch size is 36 and the size of input HR patches is 192×192 . During training, the input images are randomly rotated by 90° , 180° or 270° and randomly flipped. We use Adam [14] as the optimizer. The initial learning rate is 3×10^{-4} and decayed by half every 2×10^5 iterations.

4.2 Evaluation of Reconstructed HR Images

We compare the HR reconstruction performance of FGRN with two kinds of downscaling and upscaling methods: *Type1*: the downscaling module is fixed (*i.e.* Bicubic interpolation) and the upscaling module is the SotA SR method [3, 6, 19, 31, 35, 36]; *Type2*: the downscaling module is jointly optimized with the upscaling module [11, 18, 29, 33].

Quantitative Results We evaluate the average PSNR and SSIM on five benchmark datasets. Besides, we also provide a comparison between the number of parameters. As shown in Table 1, compared with methods in Type1, the jointly trained methods in Type2 can obtain downsampled images that are more conducive to improving the reconstruction performance. However, the direct LR guidance still restricts the reconstruction performance of the jointly trained methods. In Type2, IRN used to be a strong baseline for image rescaling models, but FGRN can outperform it by a large margin (PSNR: $+1.7 \sim 2.3dB$ for $\times 2$; $+0.8 \sim 1.1dB$ for $\times 4$). Meanwhile, FGRN consumes fewer parameters than other methods. It indicates that the invertible flow guidance can help relax restrictions on the downscaling module and efficiently boost the reconstruction performance.

Qualitative Results In Figure 4, we intuitively present visual comparisons of these methods on upscaling the $4\times$ downscaled image. Methods that the downscaling and upscaling modules are separately trained tend to produce undesirable artifacts and distortion. Even the jointly trained methods can not avoid this problem. In contrast, our FGRN reconstructs the texture clearer and more visually pleasing.

Table 1: Quantitative results of different downscaling and upscaling methods for HR image reconstruction. Bold text: best.

	Downscaling & Upscaling	Scale	Param	Set5	Set14	BSD100	Urban100	DIV2K
				PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑
Type1	Bicubic & Bicubic	×2	-	33.66 / 0.9299	30.24 / 0.8688	29.56 / 0.8431	26.88 / 0.8403	31.01 / 0.9393
	Bicubic & SRCNN [10]		57.3K	36.66 / 0.9542	32.45 / 0.9067	31.36 / 0.8879	29.50 / 0.8946	-
	Bicubic & EDSR [14]		40.7M	38.20 / 0.9606	34.02 / 0.9204	32.37 / 0.9018	33.10 / 0.9363	35.12 / 0.9699
	Bicubic & RDN [16]		22.1M	38.24 / 0.9614	34.01 / 0.9212	32.34 / 0.9017	32.89 / 0.9353	-
	Bicubic & RCAN [15]		15.4M	38.27 / 0.9614	34.12 / 0.9216	32.41 / 0.9027	33.34 / 0.9384	-
	Bicubic & SAN [9]		15.7M	38.31 / 0.9620	34.07 / 0.9213	32.42 / 0.9028	33.10 / 0.9370	-
Type2	TAD & TAU [13]	×2	-	38.46 / -	35.52 / -	36.68 / -	35.03 / -	39.01 / -
	CNN-CR & CNN-SR [12]		-	38.88 / -	35.40 / -	33.92 / -	33.68 / -	-
	CAR & EDSR [14]		51.1M	38.94 / 0.9658	35.61 / 0.9404	33.83 / 0.9262	35.24 / 0.9572	38.26 / 0.9599
	IRN [13]		1.66M	43.99 / 0.9871	40.79 / 0.9778	41.32 / 0.9876	39.92 / 0.9865	44.32 / 0.9908
	FGRN (Ours)		1.33M	44.15 / 0.9902	42.28 / 0.9840	41.87 / 0.9887	41.71 / 0.9904	45.08 / 0.9917
Type1	Bicubic & Bicubic	×4	-	28.42 / 0.8104	26.00 / 0.7027	25.96 / 0.6675	23.14 / 0.6577	26.66 / 0.8521
	Bicubic & SRCNN [10]		57.3K	30.48 / 0.8628	27.50 / 0.7513	26.90 / 0.7101	24.52 / 0.7221	-
	Bicubic & EDSR [14]		43.1M	32.62 / 0.8984	28.94 / 0.7901	27.79 / 0.7437	26.86 / 0.8080	29.38 / 0.9032
	Bicubic & RDN [16]		22.3M	32.47 / 0.8990	28.81 / 0.7871	27.72 / 0.7419	26.61 / 0.8028	-
	Bicubic & RCAN [15]		15.6M	32.63 / 0.9002	28.87 / 0.7889	27.77 / 0.7436	26.82 / 0.8087	30.77 / 0.8460
	Bicubic & RRDB [14]		16.3M	32.74 / 0.9012	29.00 / 0.7915	27.84 / 0.7455	27.03 / 0.8152	30.92 / 0.8486
	Bicubic & SAN [9]		15.7M	32.64 / 0.9003	28.92 / 0.7888	27.78 / 0.7436	26.79 / 0.8068	-
Type2	TAD & TAU [13]	×4	-	31.81 / -	28.63 / -	28.51 / -	26.63 / -	31.16 / -
	CAR & EDSR [14]		52.8M	33.88 / 0.9174	30.31 / 0.8382	29.15 / 0.8001	29.28 / 0.8711	32.82 / 0.8837
	IRN [13]		4.35M	36.19 / 0.9451	32.67 / 0.9015	31.64 / 0.8826	31.41 / 0.9157	35.07 / 0.9318
	FGRN (Ours)		3.35M	36.97 / 0.9505	33.77 / 0.9168	31.83 / 0.8907	31.91 / 0.9253	35.15 / 0.9322

4.3 Evaluation of Downscaled LR Images

We also compare the downscaling performance of FGRN with other downscaling methods. Since the downscaled results have no corresponding ground truth images, we employ the no-reference NIQE and PIQE as assessment metrics. As shown in Table 2, our downscaled images have comparable or even better quality than other methods. Besides, Figure 5 also demonstrates the pleasant visual effects of our method. It indicates that FGRN does not sacrifice the quality of downscaled LR images for the prominent reconstruction performance. When the downscaling factor is larger, the advantage of our method is more obvious.










Figure 5: NIQE/PIQE and visual comparison of the ×4 downscaled image.

4.4 Study on the Downscaling Guidance

In this section, we experimentally investigate the influence of different downscaling LR guidances on the HR reconstruction performance. The baseline model is the pure encoding-decoding model without any downscaling guidance, denoted as DU_{none} . DU_{bic} uses the Bicubically downscaled image as the downscaling guidance, *i.e.* the commonly used bicubic LR guidance, while DU_{flow} use our flow guidance as the downscaling guidance to supervise the downscaling process. The downscaling and reconstruction performance are shown in Figure 6 and Table 3 respectively.

Table 2: Quantitative results of different downscaling methods. Bold text: best.

Downscaling Method	Scale	Set5	Set14	BSD100	Urban100	DIV2K
		NIQE↓ / PIQE↓	NIQE↓ / PIQE↓	NIQE↓ / PIQE↓	NIQE↓ / PIQE↓	NIQE↓ / PIQE↓
Biubic	×2	16.027 / 38.561	8.519 / 39.802	5.282 / 32.494	5.708 / 42.788	2.944 / 34.618
TAD 		16.027 / 39.861	8.469 / 39.356	5.273 / 32.180	5.700 / 42.248	2.947 / 34.357
CAR 		16.100 / 48.496	8.943 / 41.016	5.859 / 30.597	6.401 / 44.446	3.578 / 35.780
IRN 		16.017 / 34.623	8.558 / 37.001	5.818 / 30.046	6.143 / 38.161	3.391 / 30.244
FGRN (Ours)		16.166 / 30.124	8.621 / 32.982	5.568 / 27.183	5.995 / 35.051	3.167 / 27.951
Biubic	×4	18.876 / 51.101	17.932 / 45.335	18.878 / 41.527	17.378 / 46.609	3.983 / 38.161
TAD 		18.876 / 49.691	17.936 / 45.016	18.878 / 41.190	17.320 / 46.360	3.964 / 37.784
CAR 		18.873 / 59.884	17.922 / 45.479	18.878 / 41.431	20.284 / 48.508	5.499 / 37.821
IRN 		18.876 / 45.112	17.909 / 40.213	18.879 / 35.713	17.755 / 39.769	4.292 / 30.233
FGRN (Ours)		18.876 / 40.627	17.952 / 35.464	18.878 / 34.282	17.132 / 36.991	4.358 / 27.314

DU_none only needs to focus on enhancing the HR image quality from the upscaling module, which can maximize the reconstruction performance. The reconstruction performance of DU_none in Table 3 can be seen as the upper limit. However, without the restriction of downscaling guidance, its downscaling performance is unsatisfactory as illustrated in Figure 6. The bicubic LR guidance in DU_bic can largely improve the quality of the downscaled LR image, but the reconstruction performance is compromised because the visually favorable downscaled image is not optimal for the upscaling process. As for the DU_flow , both excellent downscaling performance and HR reconstruction performance can be achieved. It indicates the superiority of our flow guidance that can decouple the mutual interference of the downscaling and upscaling, thus maximizing the capabilities of both processes.


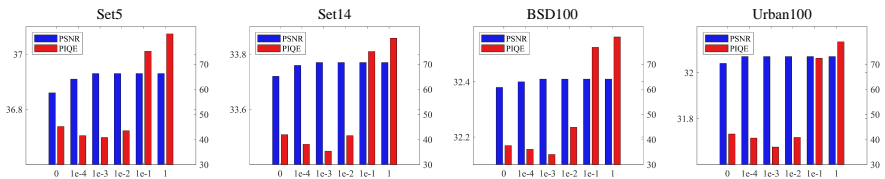


Figure 6: Visual results of methods with different downscaling guidance.

Figure 7: The $\times 4$ downscaling (PIQE↓) and reconstruction (PSNR↑) performance for different settings of γ .

4.5 Parameter Selection for γ

We assume the prior distribution of the downscaled image is $p_Z(z) \sim N(y_{bic}, \sigma^2)$, and the parameter γ in Eq 9 is positively related to the variance σ^2 . In this section, we further investigate the influence of γ on downscaling and reconstruction performance. As shown in Figure 7, $\gamma = 0.001$ is the best choice. When γ is too small or too large, it will adversely affect the downscaling or reconstruction performance. For instance, when $\gamma = 0$ (*i.e.* $\sigma^2 = 0$), the transformed LR image from FGM is supervised by the bicubically downscaled image y_{bic} without uncertainty, which will destroy the invertibility of the flow guidance module

and simultaneously reduce downscaling and reconstruction performance. On the other hand, when γ is too large (*i.e.* $\gamma = 0.1$ or $\gamma = 1$), the transformed LR image is more likely to deviate from the visually plausible image y_{bic} , resulting in downscaling performance decline.

Table 3: Reconstruction results of training the joint downscaling-upscaling model with different downscaling guidance.

	Scale	Set5	Set14	BSD100	Urban100	DIV2K
		PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow
DU_none	$\times 4$	36.91 / 0.9538	33.79 / 0.9215	32.41 / 0.9045	32.07 / 0.9307	35.35 / 0.9338
DU_bic		34.31 / 0.9258	30.95 / 0.8645	29.99 / 0.8365	29.38 / 0.8823	33.14 / 0.9020
DU_flow		36.97 / 0.9505	33.77 / 0.9168	31.83 / 0.8907	31.91 / 0.9253	35.15 / 0.9322

4.6 Study on the Information Loss of IRN

IRN claims that the high-frequency information lost during downscaling could be embedded into a latent variable z obeying a prior distribution, which could compensate for information during upscaling. In this section, we experimentally investigate the role of the latent variable for IRN. We fix the latent variable as 0 (denoted as z_0) and retrain IRN. In this case, z_0 becomes a part of the parameters of the upscaling module. As shown in Table 4, the retrained IRN with z_0 can still achieve the same results as the original one with z obeying Gaussian distribution. It indicates that the latent variable in IRN could not provide specific high-frequency information for different images to improve the reconstruction performance. The lost information in IRN is actually implicitly learned in the upscaling module.

Table 4: Reconstruction results of training IRN [53] with different latent variable.

	Scale	Set5	Set14	BSD100	Urban100	DIV2K
		PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow	PSNR \uparrow / SSIM \uparrow
z	$\times 4$	36.19 / 0.9451	32.67 / 0.9015	31.64 / 0.8826	31.41 / 0.9157	35.17 / 0.9318
z_0		36.21 / 0.9455	32.70 / 0.9013	31.62 / 0.8820	31.36 / 0.9147	35.16 / 0.9316

5 Conclusion

In this paper, we argue that in image rescaling, commonly used bicubic LR guidance will limit the reconstruction performance of the subsequent upscaling task. Towards this issue, we propose a flow guidance image rescaling network (FGRN) that can eliminate the restriction of the bicubic LR guidance. In our method, the invertible flow guidance module (FGM) can make a lossless transformation between the downsampled representation and the visually pleasing LR image, which can not only guarantee the quality of the downsampled image but also remove the restrictions on the downsampled representation. Thus, the downscaling and upscaling can be jointly optimized in an end-to-end manner without considering the visual quality of the downsampled image. In this way, the downsampled representation can become upscaling-optimal, which makes the reconstruction capability of the upscaling module approach the upper limit. Extensive experiments demonstrate that our model significantly improves both downscaling and HR reconstruction performance while being light-weighted.

Acknowledgments This work was supported by the National Key R&D Program of China (2019YFB1406200). It was also the research achievement of the Key Laboratory of Digital Rights Services.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 126–135, 2017.
- [2] Chaofeng Chen, Dihong Gong, Hao Wang, Zhifeng Li, and Kwan-Yee K Wong. Learning spatial attention for face super-resolution. *IEEE Transactions on Image Processing*, 30:1219–1231, 2020.
- [3] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11065–11074, 2019.
- [4] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [5] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [6] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199. Springer, 2014.
- [7] Daniel Glasner, Shai Bagon, and Michal Irani. Super-resolution from a single image. In *2009 IEEE 12th international conference on computer vision*, pages 349–356. IEEE, 2009.
- [8] Aditya Grover, Manik Dhar, and Stefano Ermon. Flow-gan: Combining maximum likelihood and adversarial learning in generative models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [9] Ke Han, Yan Huang, Chunfeng Song, Liang Wang, and Tieniu Tan. Adaptive super-resolution for person re-identification with low-resolution images. *Pattern Recognition*, 114:107682, 2021.
- [10] Xianxu Hou, Jiang Duan, and Guoping Qiu. Deep feature consistent deep image transformations: Downscaling, decolorization and hdr tone mapping. *arXiv preprint arXiv:1707.09482*, 2017.
- [11] Heewon Kim, Myungsub Choi, Bee Lim, and Kyoung Mu Lee. Task-aware image downscaling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 399–414, 2018.
- [12] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016.
- [13] Sungwon Kim, Sang-Gil Lee, Jongyoon Song, Jaehyeon Kim, and Sungroh Yoon. Flowavenet: A generative flow for raw audio. *arXiv preprint arXiv:1811.02155*, 2018.
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [15] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.
- [16] Johannes Kopf, Ariel Shamir, and Pieter Peers. Content-adaptive image downscaling. *ACM Transactions on Graphics (TOG)*, 32(6):1–8, 2013.
- [17] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [18] Yue Li, Dong Liu, Houqiang Li, Li Li, Zhu Li, and Feng Wu. Learning a convolutional neural network for image compact-resolution. *IEEE Transactions on Image Processing*, 28(3):1092–1107, 2018.
- [19] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017.
- [20] Rui Liu, Yu Liu, Xinyu Gong, Xiaogang Wang, and Hongsheng Li. Conditional adversarial generative flow for controllable image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7992–8001, 2019.
- [21] Zhengxiong Luo, Yan Huang, Shang Li, Liang Wang, and Tieniu Tan. Unfolding the alternating optimization for blind super resolution. *arXiv preprint arXiv:2010.02631*, 2020.
- [22] Tomer Michaeli and Michal Irani. Nonparametric blind super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 945–952, 2013.
- [23] Don P Mitchell and Arun N Netravali. Reconstruction filters in computer-graphics. *ACM Siggraph Computer Graphics*, 22(4):221–228, 1988.
- [24] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012.
- [25] A Cengiz Oeztireli and Markus Gross. Perceptually based downscaling of images. *ACM Transactions on Graphics (TOG)*, 34(4):1–10, 2015.
- [26] Albert Pumarola, Stefan Popov, Francesc Moreno-Noguer, and Vittorio Ferrari. C-flow: Conditional generative flow models for images and 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7949–7958, 2020.
- [27] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

- [28] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. 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*, pages 1874–1883, 2016.
- [29] Wanjie Sun and Zhenzhong Chen. Learned image downscaling for upscaling using content adaptive resampler. *IEEE Transactions on Image Processing*, 29:4027–4040, 2020.
- [30] N Venkatanath, D Praneeth, Maruthi Chandrasekhar Bh, Sumohana S Channappayya, and Swarup S Medasani. Blind image quality evaluation using perception based features. In *2015 Twenty First National Conference on Communications (NCC)*, pages 1–6. IEEE, 2015.
- [31] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018.
- [32] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [33] Mingqing Xiao, Shuxin Zheng, Chang Liu, Yaolong Wang, Di He, Guolin Ke, Jiang Bian, Zhouchen Lin, and Tie-Yan Liu. Invertible image rescaling. In *European Conference on Computer Vision*, pages 126–144. Springer, 2020.
- [34] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4541–4550, 2019.
- [35] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 286–301, 2018.
- [36] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018.