This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Effective Invertible Arbitrary Image Rescaling

Zhihong Pan¹ Baopu Li^{2*} Dongliang He³ Wenhao Wu³ Errui Ding³ ¹Baidu Research (USA)

²Oracle Health and AI ³Baidu VIS

zhihongpan@baidu.com, baopu.li@oracle.com, {hedongliang01,wuwenhao01,dingerrui}@baidu.com

Abstract

Great successes have been achieved using deep learning techniques for image super-resolution (SR) with fixed scales. To increase its real world applicability, numerous models have also been proposed to restore SR images with arbitrary scale factors, including asymmetric ones where images are resized to different scales along horizontal and vertical directions. Though most models are only optimized for the unidirectional upscaling task while assuming a predefined downscaling kernel for low-resolution (LR) inputs, recent models based on Invertible Neural Networks (INN) are able to increase upscaling accuracy significantly by optimizing the downscaling and upscaling cycle jointly. However, limited by the INN architecture, it is constrained to fixed integer scale factors and requires one model for each scale. Without increasing model complexity, a simple and effective invertible arbitrary rescaling network (IARN) is proposed to achieve arbitrary image rescaling by training only one model in this work. Using innovative components like position-aware scale encoding and preemptive channel splitting, the network is optimized to convert the noninvertible rescaling cycle to an effectively invertible process. It is shown to achieve a state-of-the-art (SOTA) performance in bidirectional arbitrary rescaling without compromising perceptual quality in LR outputs. It is also demonstrated to perform well on tests with asymmetric scales using the same network architecture.

1. Introduction

Recent deep learning based image super-resolution (SR) methods have advanced the performance of image upscaling significantly but they are often limited to fixed integer scale factors and pre-determined downscaling degradation kernels. To work in real world applications where an image is commonly rescaled to arbitrary sizes, additional image resizing is often needed, which leads to degradation in both performance and efficiency. Lately there are growing interests in SR models that support arbitrary scale factors

and great successes have been achieved in recent works [5, 7, 21]. However, they are only optimized for the unidirectional upscaling process with the LR inputs either synthesized from a predefined downscaling kernel or in its native resolution. Considering the potential mutual benefits between downscaling and the inverse upscaling, some image rescaling models [9, 19, 24] are developed to optimize these two processes jointly and significant improvements in upscaling accuracy are achieved comparing to unidirectional SR models of the same scale factors. The state-ofthe-art (SOTA) performance for such bidirectional image rescaling is set by the invertible rescaling net (IRN) as proposed by Xiao et al. [24]. As shown in Fig. 1, it is able to achieve the best performance so far since both the Haar transformation and the invertible neural network (INN) [2] backbone are invertible processes, and its forward and invertible backward operations can model the downscaling and inverse upscaling cycle naturally. Denoting the forward downscaling process as $(\mathbf{y}_L, \mathbf{z}_L) = f(\mathbf{x}_H)$, the HR image can be fully restored as $\mathbf{x}_H = f^{-1}(\mathbf{y}_L, \mathbf{z}_L)$ if the latent variable \mathbf{z}_L is preserved. When the network is optimized to store as much information as allowed in y_L and convert \mathbf{z}_L as input-independent random variables, the optimal HR output $\hat{\mathbf{x}}_H$ can be restored as $f^{-1}(\mathbf{y}_L, \hat{\mathbf{z}}_L)$ where $\hat{\mathbf{z}}_L$ is randomly sampled with minimized loss in restoration accuracy. However, limited by the nature of INN architecture that the number of pixels must be equivalent when differences in resolution are accounted for (LR features have more channels), the applicable scale factors are limited to integers.

To overcome the above issues, we propose a new invertible arbitrary rescaling network (IARN) that is able to implement bidirectional arbitrary rescaling using similar INN backbone and it is able to train one model for a range of arbitrary scales, including asymmetric ones. As shown in Fig. 1, it replaces the Haar wavelet transformation blocks in IRN with a preemptive channel-splitting step denoted as

$$\mathbf{x}_{H}^{lf} = s(\mathbf{x}_{H})$$

$$\mathbf{x}_{H}^{hf} = \mathbf{x}_{H} - \mathbf{x}_{H}^{lf}$$
 (1)

where subscript H is used to specify that all images are preserved in original high-resolution. Note that this step is in-

^{*}Work completed at Baidu Research (USA).



Figure 1: Comparison between IRN [24] and the proposed IARN. \mathbf{x}_H , \mathbf{y}_L and $\mathbf{z}_{H,L}$ denote input HR image, output LR image and the latent variables. Note that subscripts H, L refer to high- and low-resolution respectively, and superscripts hf, lf refer to high and low frequency channels respectively. The main differences in IARN include channel-splitting in high-resolution in place of Haar transform, the additional near-invertible rescaling between \mathbf{y}_H and \mathbf{y}_L , and enhancement in transformation blocks inside the INN backbone.

trinsically invertible as $\mathbf{x}_H = \mathbf{x}_H^{lf} + \mathbf{x}_H^{hf}$ for the inverse direction. The feature transformation step using similar INN backbone, which is also invertible, is denoted as

$$(\mathbf{y}_H, \mathbf{z}_H) = f(\mathbf{x}_H^{lf}, \mathbf{x}_H^{hf})$$
(2)

The downscaling and inverse upscaling are included as the last step and it consists of two separate processes as

$$\mathbf{y}_L = d(\mathbf{y}_H)$$

$$\hat{\mathbf{y}}_H = u(\mathbf{y}_L)$$
(3)

To optimize the network for invertible arbitrary rescaling, in addition to the similar objectives to maximize information saved in y_H and make z_H signal independent, one key challenge is to make the last step invertible. If it satisfies that $\hat{\mathbf{y}}_H = \mathbf{y}_H$, the whole network is invertible. While this bidirectional process in Equation 3 is not invertible in general, for a given scale factor and specific downscaling function $d(\cdot)$ and upscaling $u(\cdot)$, like nearest neighbour (NN) or bicubic interpolation, we can find a set of images which are invertible to this rescaling process and our goal is to transform \mathbf{y}_H as one of these rescaling invertible images so the full process is effectively invertible. To help the transformation of y_H and make it approaching invertibility faster, a new preemptive channel-splitting step $s(\cdot)$ that converts \mathbf{x}_{H}^{lf} as rescaling invertible in advance is proposed and it is shown to be very effective. Lastly, as the aforementioned rescaling-invertible feature is scale-dependent, an innovative position-aware scale encoding is proposed as additional inputs of the network to make the network capable of handling large variations in arbitrary scale factors, including asymmetric ones. This component is not illustrated in Fig. 1 for simplicity reasons but will be discussed in details later. In summary, the main contributions of our work include:

- The first to use invertible neural network for bidirectional arbitrary image rescaling and set new SOTA performances in both generated LR and restored HR.
- A preemptive channel-splitting step, which separates the rescaling-invertible component from the input HR

image, is advanced to make the learning of invertible arbitrary rescaling more effective.

• A position-aware scale encoding, which is independent of input image size and compatible with asymmetric rescaling, is proposed to further boost model performance in a large range of arbitrary scales.

2. Related Works

Arbitrary Scale Super-Resolution. Single image superresolution, as a form of image rescaling with a fixed integer scale factor like $\times 2$ and $\times 4$, has been studied extensively. For the last few years, deep learning based methods like [6, 10, 15, 28, 27] have brought great successes in this field, but these methods commonly train one model for each scale factor. More recent models proposed by Lim et al. [15] and Li et al. [12] are capable of training one model for multiple scaling factors but only limited to integer ones. Inspired by the weight prediction techniques in meta-learning [11], a single Meta-SR model was proposed by Hu et al. [7] to solve image rescaling of arbitrary scale factors, by predicting weights of convolutional layers for arbitrary scale factors within certain range. Alternatively, Behjati et al. [3] proposed the OverNet to generate overupscaled maps from which HR images of arbitrary scales can be recovered using optimal downsampling. In the latest ArbSR, Wang et al. [21] proposed a plug-in module that can optimize existing SR models for arbitrary asymmetric rescaling, where scale factors along horizontal and vertical directions could be different. While these methods are often limited to a maximum scale factor like ×4 to maintain high performance, Chen et al. [5] proposed recently to use learned pixel representation features to replace pixel value features in previous methods. Using the innovative local implicit image function (LIIF), this model can extrapolate well to out-of-distribution large scales that are not seen in training. Different from the above models which are optimized for upscaling reconstruction only, we consider the bidirectional arbitrary rescaling of downscaling and upscaling as one process in this work.

Bidirectional Image Rescaling. To take advantage of the potential mutual benefits between downscaling and the inverse upscaling, Kim et al. [9] proposed an auto-encoder framework to jointly train image downscaling and upscaling together. Similarly, Sun et al. [19] suggested a new image downscaling method using a content adaptive-resampler, which can be jointly trained with any existing differentiable upscaling (SR) models. More recently, Xiao et al. [24] advanced an invertible rescaling net (IRN) that has achieved SOTA performance for learning based bidirectional image rescaling. Based on the invertible neural network (INN) [2], IRN learns to convert HR input to LR output and an auxiliary latent variable z. By mapping z to a case-agnostic normal distribution during training, inverse image upscaling is implemented by randomly sampling \hat{z} from the normal distribution without need of the case specific z. While the above bidirectional image rescaling methods are limited to a fixed integer scale factor like $\times 4$, Pan *et al.* [18] proposed the BAIRNet as the first to solve bidirectional arbitrary rescaling by utilizing local implicit image functions for both downscaling and upscaling with better overall cross-scale performance over IRN. Most recently, Xing et al. [25] proposed an encoder-decoder network (AIDN) to tackle the same challenge with consistent improvements over both IRN and BAIRNet. Instead of using separate encoder and decoder to model image downscaling and upscaling respectively, we propose here to model the downscaling and upscaling processes as forward and backward operations of one INN backbone. While very similar, their performances in the upscaling task are consistently lower than ours with very few exceptions in certain scales. Additionally, LR outputs from our IARN have better perceptual quality comparing to AIDN in both blind or non-blind image quality assessments. As it is shown in [13] that there is a performance trade-off between the generated LR and restored HR, our IARN is obviously the superior one comparing to AIDN when both downscaling and upscaling tasks are considered.

3. Proposed Method

3.1. Network Architecture

Elaborated from the general pipeline shown in Fig. 1, the detailed architecture of the proposed IARN is illustrated in Fig. 2. For the forward downscaling process to generate LR output y_L from HR input x_H , the overall process is summarized in the following four steps:

$$\mathbf{x}_{H}^{lf} = s(\mathbf{x}_{H})$$
$$\mathbf{x}_{H}^{hf} = \mathbf{x}_{H} - \mathbf{x}_{H}^{lf}$$
$$\mathbf{y}_{H}, \mathbf{z}_{H}, \tilde{\mathbf{p}}_{H}^{lf}, \tilde{\mathbf{p}}_{H}^{hf} = f(\mathbf{x}_{H}^{lf}, \mathbf{x}_{H}^{hf}, \mathbf{p}_{H}^{lf}, \mathbf{p}_{H}^{hf})$$
$$\mathbf{y}_{L} = d(\mathbf{y}_{H})$$
(4)

where $s(\cdot)$ is the preemptive channel splitting function detailed below, $f(\cdot)$ is the forward function of the INN backbone and $d(\cdot)$ is a downscaling function like nearest neighbour or bicubic interpolation. For variables, **x** is the input in forward direction, **y** is the output, and **p** and **p** are the position-aware scale encoding at the input and output side respectively. Details about this position-aware scale encoding will be introduced in later subsection. For subscripts and superscripts, *H* and *L* means high- and low-resolutions, while hf and lf represents high and low frequency components respectively. Similarly the inverse upscaling process consists of the following three steps:

$$\mathbf{y}_{H} = u(\mathbf{y}_{L})$$
$$\mathbf{x}_{H}^{lf}, \mathbf{x}_{H}^{hf}, \mathbf{p}_{H}^{lf}, \mathbf{p}_{H}^{hf} = f^{-1}(\mathbf{y}_{H}, \mathbf{z}_{H}, \tilde{\mathbf{p}}_{H}^{lf}, \tilde{\mathbf{p}}_{H}^{hf}) \qquad (5)$$
$$\mathbf{x}_{H} = \mathbf{x}_{H}^{lf} + \mathbf{x}_{H}^{hf}$$

where $u(\cdot)$ is an upscaling function and $f^{-1}(\cdot)$ is the inverse function of $f(\cdot)$.

The process presented in Equation 5 is an ideal inverse process, where \mathbf{y}_H is perfectly restored after the downscaling and upscaling cycle, and \mathbf{z}_H , $\mathbf{\tilde{p}}_H^{lf}$, $\mathbf{\tilde{p}}_H^{hf}$ are all preserved. This ideal situation does not exist in real application where only \mathbf{y}_L is saved and $\mathbf{\hat{y}}_H \neq \mathbf{y}_H$. Similar to the previous studies, a generated $\mathbf{\hat{z}}_H$ is used in place of \mathbf{z}_H . For scale encoding, as they are utilized to carry position and scale factor information for each pixel for both forward and backward directions, and they are independent of low and highfrequency branches, one general \mathbf{p} is used for \mathbf{p}_H^{lf} and \mathbf{p}_H^{hf} during forward process and the same for $\mathbf{\tilde{p}}_H^{lf}$ and $\mathbf{\tilde{p}}_H^{hf}$ during backward upscaling. As a result, the restored HR output $\mathbf{\hat{x}}_H$ is calculated as

$$\hat{\mathbf{y}}_{H} = u(\mathbf{y}_{L})$$

$$\hat{\mathbf{x}}_{H}^{lf}, \hat{\mathbf{x}}_{H}^{hf} = f^{-1}(\hat{\mathbf{y}}_{H}, \hat{\mathbf{z}}_{H}, \mathbf{p}_{H})$$

$$\hat{\mathbf{x}}_{H} = \hat{\mathbf{x}}_{H}^{lf} + \hat{\mathbf{x}}_{H}^{hf}$$
(6)

While the primary goal is to minimize the restoration error between $\hat{\mathbf{x}}_H$ and \mathbf{x}_H , the overall loss used to train the network includes multiple losses for different objectives as shown below

$$L = \lambda_1 L_r + \lambda_2 L_g + \lambda_3 L_d + \lambda_4 L_i.$$
(7)

Here L_r is the L1 reconstruction loss for upscaled HR output $\hat{\mathbf{x}}_H$ and L_g is the L2 guidance loss for downscaled LR output \mathbf{y}_L in reference to a downsampled LR reference $\overline{\mathbf{y}}_L$ using bicubic interpolation. For L_d , it is similar to the distribution regulation loss of latent variable \mathbf{z} as in IRN [24]. However, with the newly introduced position-aware scale encoding, in place of the L2 loss between \mathbf{z}_H and \mathbf{z}_H^0 , it is calculated between $(\mathbf{z}_H, \tilde{\mathbf{p}}_H^{lf}, \tilde{\mathbf{p}}_H^{hf})$ and reference $(\mathbf{z}_H^0, \mathbf{p}_H, \mathbf{p}_H)$. Here \mathbf{z}_H^0 is a zero tensor, following



Figure 2: Network architecture of the proposed IARN, including the full process of preemptive channeling splitting $s(\cdot)$, the feature transformation INN backbone and the rescaling-invertible module consists of downscaling $d(\cdot)$ and upscaling $u(\cdot)$, and the enhanced dense atrous block used for transformation blocks ϕ , ρ , η .

the practice in IRN for the surrogate distribution matching loss L'_{distr} . The last term L_i is a novel rescaling-invertible loss that aims to make the full process effectively invertible. As discussed earlier, the proposed process is only partially invertible since the downscaling and upscaling cycle is not invertible. In general cases, $\hat{\mathbf{y}}_H = u(d(\mathbf{y}_H)) \neq \mathbf{y}_H$. However, given specific pair of $d(\cdot)$ and $u(\cdot)$, there are certain subset of rescaling invertible images \mathbf{y}_H^i satisfies $\mathbf{y}_H^i = u(d(\mathbf{y}_H^i))$. For example, as explained later in Section 3.2, for NN resampling, we can generate a rescalinginvertible loss L_i is introduced here as the L2 loss between $\hat{\mathbf{y}}_H$ and \mathbf{y}_H . When the loss is zero, then \mathbf{y}_H is rescaling invertible.

Note that while the layers of Invertible Blocks (InvBlock) in our INN backbone is similar to the ones in IRN [22], there are some key differences. First, positionaware scale encoding is introduced to both the lower lowfrequency (LF) and upper high-frequency (HF) branches to enable the network adaptive to arbitrary scales. Secondly, for transformation functions between the upper and lower branches, including $\phi(\cdot)$, $\rho(\cdot)$ and $\eta(\cdot)$, the original densely connected convolutional block in IRN is enhanced with dilated convolutions where the dilation varies from 1 to l, which is the number of layers in each block. l is simplified as 2 here in Fig. 2. This enhanced Dense Atrous Block is introduced to increase effective receptive fields of the network with mixed dilation larger or equal to one. As the process in the INN backbone is applied to original high-resolution images, the receptive fields of our IARN is the same as IRN in pixel units, but would be smaller in terms of image regions given the resolution difference if the network structure and depth remains the same. Using a mixture of dilated convolutions in place of the original ones, it can increase the receptive field without changes in the number of model parameters and model complexity.

3.2. Preemptive Channel Splitting

As seen in the Fig. 2, the HR output y_H is located in the lower branch of INN backbone and matches the lowfrequency HR input \mathbf{x}_{H}^{lf} in channel orders. While the transformation functions in the INN backbone can transfer features between lower and upper branches quite efficiently, it may be beneficial to pre-process the inputs so that \mathbf{x}_{H}^{lf} is close to y_H even before feeding into the INN backbone, making the learning task of transformation functions even more efficient. As discussed above, it is ideal for y_H to be rescaling invertible for a given pair of resampling operation $d(\cdot)$ and $u(\cdot)$. Here a preemptive channel splitting function is proposed as $s(\cdot) = u(d(\cdot))$. As pointed out in BAIR-Net [18], in the case of integer scale factors, $s(\cdot)$ would be an idempotent operation when using bilinear downscaling and nearest-neighbour upscaling. That is, for any image x, s(s(x)) = s(x). In other words, s(x) is rescalinginvertible. For arbitrary scales here, $s(\mathbf{x})$ is only rescaling invertible when using nearest-neighbour for both downscaling and upscaling. Thus the preemptive channel splitting, which splits the input image to two branches, is proposed here as

$$\mathbf{x}_{H}^{lf} = u_{N}(d_{N}(\mathbf{x}_{H}))$$

$$\mathbf{x}_{H}^{hf} = \mathbf{x}_{H} - \mathbf{x}_{H}^{lf}$$
(8)

where N refers to nearest-neighbour resampling. Note that the pair of $d(\cdot)$ and $u(\cdot)$ must be the same for both channel splitting in the front and rescaling at the end. Experiments are also conducted to demonstrate the advantage of nearestneighbour over other choices of resampling choices.

3.3. Position-Aware Scale Encoding

Both the preemptive channel splitting and rescaling invertible loss help the network transform y_H to induce minimum losses during downscaling and rescaling steps and lead to optimal restoration of $\hat{\mathbf{x}}_H$ consequently. As this transformation is sensitive to scale factors, to enable the model robust to a large range of arbitrary scales, scale information are necessary during the transformation between the lower and upper branches. Additionally, depending on the scale factor and rescaling method, the position of the pixel inside the image is also needed to determine operations applied to the pixel itself. Using nearest-neighbour as an example, some pixels are resampled from neighbours located to the top-left direction while others resampled from different directions. The naive option of using absolution pixel locations would result in training in a limited image size while testing in larger out-of-distribution image size. To account for these factors, a position-aware and image size independent scale encoding **p** is defined as $(\mathbf{s}_h, \mathbf{s}_v, \mathbf{d}_h, \mathbf{d}_v)$. Here s_h and s_v are the scale factors along horizontal and vertical directions respectively, to accommodate asymmetric scales. For \mathbf{d}_h and \mathbf{d}_v , they are the relative horizontal and vertical distances from the input pixel to the closest resampled pixel to its bottom-right direction and can be calculated as below

$$\mathbf{d}_{h}(i,j) = \min_{\substack{i',i's'_{h} - is_{h} \ge 0}} i's'_{h} - is_{h} \\
 \mathbf{d}_{v}(i,j) = \min_{\substack{j',j's'_{v} - js_{v} \ge 0}} j's'_{v} - js_{v}$$
(9)

where (i, j) and (i', j') are indices of the input pixel and resampled pixel respectively, s_h and s_v are pixel sizes along horizontal and vertical axes for input image, and s'_h and s'_v are pixel sizes of resampled image. Experiments conducted later also demonstrate that the position-aware scale encoding is greatly beneficial to network performance and it is the best to include it in both lower and upper branches.

4. Experiments

4.1. Data and Settings

For fair comparisons with recent relevant works like IRN and BAIRNet, the DIV2K [1] training set is used for baseline training. Another dataset, Flickr2K [20], is also included for training the final model. For quantitative evaluation, we use HR images from six commonly used datasets for comprehensive comparison, including Set5 [4], Set14 [26], BSD100 [16], Urban100 [8], Manga109 [8], and the DIV2K validation set. Following previous practices, we take the peak noise-signal ratio (PSNR) and SSIM [23] on the luminance channel for all test sets, with the exception of DIV2K which uses average of RGB channels.

To accommodate the large range of arbitrary scale factors, a total of 20 layer of invertible blocks are included in the INN backbone, less than the total of 24 used in IRN $\times 2$ and IRN \times 4. For the Dense Atrous Block, 4 layers of dilated convolution are included with dilation setting as 1 to 4 consecutively. Mini-batches of 16 144×144 patches are initially used, each with a random scale sampled from an uniform distribution of $\mathcal{U}(1,4)$. It is upgraded to $24\ 192 \times 192$ patches for the final model when trained using both DIV2K and Flickr2K. There are two stages of training for the final models, each with 250k iterations where the learning rate is reduced by half after each 50k iterations. Settings of the two stages are the same except for the starting learning rate, which is 2×10^{-4} for the first and changed to 1×10^{-4} when resuming the second stage training. The weights of losses are set at empirically at 1, 16 and 2 for L_r , L_q and L_i respectively. As we have found out in experiments, and similar to the findings by Li *et al.* [13], setting \hat{z} as zero for training and inference achieves equivalent or better performance comparing to randomly sampled \hat{z} . As a result, L_d is set as zero for our experiments as it has little impact on learning.

4.2. Arbitrary Rescaling Performance

To assess the performance of our proposed method for arbitrary rescaling in symmetric scales, we compare restoration quality of rescaled HR images in a set of arbitrary scales and the results are included in Table 1. Similar to BAIRNet, models trained for fixed integer scales are evaluated for arbitrary scales using additional bicubic interpolations to keep the resolution of LR images fed into respective models identical across different methods. For models capable of arbitrary scales on their own, most of them are optimized for upscaling only so they are only assessed here for HR outputs. Other than our IARN, BAIRNet is the only one that is trained for bidirectional arbitrary image rescaling.

As shown in Table 1, results from our IARN are the best for all test cases except for a couple of cases where AIDN is slightly better. Using the challenging Manga109 test set as an example, it leads over BAIRNet in PSNR by +3.04, +1.16 and +0.89 for $\times 1.5, \times 2.5$ and $\times 3.5$ respectively. As shown in visual examples in Fig. 3, bidirectional methods like IRN, BAIRNet, AIDN and our IARN outperform others significantly. Among the four, IARN is able to restore details more accurately, making it easier to recognize details like the number 00. The IARN model used for comparisons here uses an asymmetric pre-training for the first stage. For intra-model comparison, two other variants, $IARN^{\dagger}$ and $IARN^{\ddagger}$, are also included in Table 1 with $\times 3.5$ as an example. IARN[†] takes the same pre-training strategy but only uses DIV2K with baseline batch and patch sizes, resulting a slightly worse performance behind IARN. IARN^{\ddagger} is the same as IARN^{\dagger} in data settings but it is exclusively trained using symmetric scales for both stages. It is consistently worse than IARN, proving that pre-training

Method	Scale	Set5	Set14	BSD100	Urban100	Manga109	DIV2K
RCAN [27]	1.5	40.97/0.9767	37.05/0.9578	35.59/0.9516	35.93/0.9660	42.33/0.9889	38.47/0.9701
Meta-SR [7]	1.5	41.47/0.9785	37.52/0.9601	35.86/0.9543	36.91/0.9696	43.17/0.9904	38.88/0.9718
LIIF [5]	1.5	41.23/0.9774	37.37/0.9591	35.76/0.9536	36.70/0.9684	42.84/0.9894	38.82/0.9717
ArbSR [21]	1.5	41.47/0.9786	37.51/0.9603	35.86/0.9547	36.92/0.9697	43.12/0.9904	38.84/0.9719
CAR [19]	1.5	40.50/0.9763	37.08/0.9596	35.72/0.9535	34.70/0.9635	40.90/0.9881	37.93/0.9683
IRN [24]	1.5	43.55/0.9891	39.52/0.9795	39.28/0.9833	36.52/0.9811	42.64/0.9936	40.18/0.9838
BAIRNet [18]	1.5	47.13/0.9849	43.12/0.9760	46.63/0.9959	44.01/0.9946	45.49/0.9948	44.99/0.9920
AIDN [25]	1.5	50.61/0.9961	46.70/0.9920	49.82/0.9983	46.26/0.9967	_1	47.01/0.9953
IARN	1.5	51.02/0.9968	47.25/0.9938	50.91/0.9986	47.58/0.9975	48.58/0.9975	46.74/0.9949
RCAN [27]	2.5	36.05/0.9436	31.69/0.8815	30.47/0.8508	30.42/0.8990	36.59/0.9634	32.72/0.9079
Meta-SR [7]	2.5	36.18/0.9441	31.90/0.8814	30.47/0.8508	30.57/0.9003	36.55/0.9639	32.77/0.9086
LIIF [5]	2.5	35.98/0.9434	31.64/0.8813	30.45/0.8510	30.42/0.8992	36.39/0.9630	32.78/0.9091
ArbSR [21]	2.5	36.21/0.9448	31.99/0.8830	30.51/0.8536	30.68/0.9027	36.67/0.9646	32.77/0.9093
CAR [19]	2.5	37.33/0.9548	33.78/0.9169	32.53/0.9020	32.19/0.9301	37.63/0.9717	34.32/0.9310
IRN [24]	2.5	39.78/0.9742	36.39/0.9553	35.56/0.9540	33.99/0.9589	39.33/0.9836	36.60/0.9607
BAIRNet [18]	2.5	40.11/0.9664	36.62/0.9469	36.29/0.9563	36.62/0.9679	40.26/0.9830	37.46/0.9627
AIDN [25]	2.5	40.77/0.9750	37.62/0.9588	36.65/0.9593	37.10/0.9710	_1	37.88/0.9659
IARN	2.5	40.93/0.9756	37.78/0.9598	36.81/0.9607	36.95/0.9703	41.38/0.9862	37.92/0.9662
RCAN [27]	3.5	33.47/0.9138	29.24/0.8141	28.42/0.7731	27.61/0.8348	32.74/0.9328	30.13/0.8511
Meta-SR [7]	3.5	33.59/0.9146	29.60/0.8140	28.42/0.7728	27.71/0.8356	32.75/0.9337	30.18/0.8524
LIIF [5]	3.5	33.41/0.9133	29.20/0.8131	28.39/0.7714	27.60/0.8334	32.60/0.9324	30.16/0.8517
ArbSR [21]	3.5	33.63/0.9149	29.58/0.8147	28.41/0.7744	27.69/0.8360	32.84/0.9339	30.14/0.8518
CAR [19]	3.5	34.98/0.9303	31.38/0.8643	30.14/0.8326	29.97/0.8871	35.00/0.9507	31.88/0.8865
IRN [24]	3.5	37.12/ <mark>0.9546</mark>	33.65/0.9196	32.54/ <mark>0.9047</mark>	31.84/0.9277	36.86/0.9690	33.84/0.9281
BAIRNet [18]	3.5	36.85/0.9472	32.97/0.9074	32.36/0.8986	32.71/0.9338	36.98/0.9671	33.87/0.9266
AIDN [25]	3.5	37.25/0.9538	33.87/0.9197	32.73/0.9032	33.20/0.9372	_1	34.19/0.9292
IARN	3.5	37.44/0.9547	34.04/0.9218	32.90/0.9058	33.27/0.9371	37.91/0.9705	34.33/0.9308
IARN [†]	3.5	36.82/0.9538	33.66/0.9194	32.69/0.9032	33.02/0.9350	37.72/0.9697	34.19/0.9292
IARN [‡]	35	37 27/0 9537	33 74/0 9175	32 68/0 9019	32 81/0 9326	37 59/0 9690	34 09/0 9282

Table 1: Quantitative comparisons of SOTA SR and rescaling methods with the best two results highlighted in red and blue respectively (IARN^{\dagger} and IARN^{\ddagger} are used for intra-model comparisons only and not ranked, and methods in **bold** means requirement of multiple models and additional interpolations to conduct tests in arbitrary scales).



Figure 3: Visual examples of arbitrary rescaling at two scales: $\times 2.5$ and $\times 3.5$ (Zoom in to see improved sharpness in IARN results comparing to AIDN, including the area marked by red arrows).

with asymmetric scales can help the final model performs better in tests of symmetric scales.

Results of continuous scales between $\times 1.1$ and $\times 4$ (increasing by 0.1 at each sampling point) are also illustrated in Fig. 5 for a more thorough assessment. As explained in BAIRNet, there is a significant performance boost when

the models are trained jointly for both downscaling and upscaling. Among this group, BAIRNet is obviously better than CAR and IRN across the range overall, except trailing slightly behind IRN for a limited choices of scales. Both AIDN and our IARN model are able to outperform other models across the full range of arbitrary scales, surpassing IRN even at fixed scales like $\times 2$ and $\times 4$. Between the two,

¹Incomplete AIDN results due to lack of model or data.



Figure 4: Visual examples of generated LR images with magnified view of worst-case visual artifacts.



Figure 5: Cross-scale $(\times 1.1 - \times 4)$ performance for arbitrary rescaling testing of DIV2K validation set.

our IARN has better performance for most of the scales, especially for small scales close to $\times 1.1$ and large scales around $\times 3$.

Table 2: Quantitative quality assessment (NIQE \downarrow /SSIM \uparrow) of generated LR outputs.

Method	S	Urban100	Manga109	DIV2K
CAR [19]	2	6.360/0.9658	4.365/0.9772	3.660/0.9748
IRN [24]	2	5.956/0.9941	4.214/0.9959	3.347/0.9945
BAIRNet [18]	2	6.475/0.9797	4.431/0.9892	3.507/0.9864
AIDN [25]	2	_1	_1	3.516/0.9920
IARN	2	5.783/0.9963	4.120/0.9971	3.359/0.9963
CAR [19]	4	22.731/0.9196	6.886/0.9529	5.549/0.9460
IRN [24]	4	18.035/0.9916	5.884/0.9932	4.094/0.9933
HCFlow [14]	4	18.475/0.9651	6.457/0.9738	4.744/0.9785
BAIRNet [18]	4	19.401/0.9716	7.632/0.9824	4.896/0.9841
AIDN [25]	4	_1	_1	4.165/0.9909
IARN	4	18.020/0.9928	5.961/0.9932	4.088/0.9944

4.3. LR Image Quality

For the group of six bidirectional rescaling methods, LR outputs are also generated from the learning based models respectively. Without a known ground-truth reference, the goal is to generate visually pleasant images which look similar to the LR reference that is downsampled from the HR input using conventional bicubic interpolation. As shown in Table 2, the SSIM calculated in related to the LR references and the blind image quality metric NIQE [17] are used for comparison. To accommodate the native integer scales in CAR and IRN, $\times 2$ and $\times 4$ are used for assessments. It is clearly shown that our IARN is the best overall, slightly ahead of IRN as the lose second best, while AIDN is trailing behind as the third best. An example of $\times 4$ LR images with the most visible artifacts are included in Fig. 4 to compare different models. While they are not much different from the bicubic reference in overall view, the CAR one is noticeably brighter. From zoomed-in views on the first row, HCFlow and BAIRNet are obviously the worst two with obvious false color artifacts. While the remaining three are similar visually, the artifacts are slightly more noticeable in AIDN, as marked by red arrows in the left zoomed-in window in particular.

Table 3: Ablation study for different modules (Colorhighlighted group in one column means they share the same setting in all other columns).

Channel	Scale	Atrous	Rescaling	#	PSNR
Splitting	Encoding	Convolution	Method	InvBlock	Urban100 $\times 4$
×	Dual	X	Bilinear	16	26.79
1	Dual	×	Bilinear	16	27.22
1	HF	×	Bilinear	16	26.93
1	LF	×	Bilinear	16	27.04
1	None	×	Bilinear	16	26.70
✓	Dual	1	Bilinear	16	27.41
1	Dual	1	NN	16	27.50
1	Dual	1	Bicubic	16	27.11
1	Dual	1	NN	20	27.76
✓	Dual	1	NN	24	27.79

4.4. Ablation Study

To show the effects of different modules proposed here, we conduct a comprehensive ablation study for different combinations using Urban100 \times 4 as the testing benchmark. For illustrative purpose, all models were trained for 100k iterations and the differences in performance is clearly shown in Table 3. For easy cross-reference, colored highlight are used to associate mini-groups for comparison. Two of the



Figure 6: Visual examples of arbitrary asymmetric scales.

most critical modules to boost performance are preemptive channel splitting and position-aware scale encoding. As highlighted in red, when channel-splitting is activated while keeping other configurations identical, the average PSNR increases by +0.43. Note that when no channelsplitting is applied, original RGB channels are fed to the LF branch while inputs to HR are zeros. For scale encoding options as highlighted in blue, it is the best to include it for both branches, boosting PSNR by +0.52. When atrous convolution is applied, another +0.19 is added in PSNR. For rescaling, NN is the best choice while bicubic is far behind the other two. This corroborates the methodology discussed earlier that channel-split function $s(\cdot)$ makes \mathbf{x}_{H}^{ij} rescaling-invertible when using NN resampling. While the final PSNR is always better when the number of InvBlock layers increases, it is set at 20 in other experiments for better efficiency with minimum performance loss.

Table 4: PSNR for symmetric and asymmetric scale factors.

	BSD100				Urban100			
	$\frac{\times 2}{\times 3}$	$\frac{\times 1.6}{\times 3.2}$	$\frac{\times 3.6}{\times 1.2}$	$\times 2.5$	$\frac{\times 2}{\times 3}$	$\frac{\times 1.6}{\times 3.2}$	$\frac{\times 3.6}{\times 1.2}$	$\times 2.5$
ArbSR [21]	30.58	30.87	30.24	30.51	30.59	30.60	29.74	30.68
IRN [24]	34.87	34.69	34.19	35.56	32.75	32.44	32.09	33.99
BAIRNet [18]	36.28	36.96	36.64	36.17	36.32	36.74	35.82	36.43
IARN*	36.96	38.04	39.17	36.65	37.06	37.92	38.56	36.83

4.5. Asymmetric Arbitrary Rescaling

As discussed earlier, our IARN model can be optimized for asymmetric scales using the exact same network architecture. For fair comparisons with ArbSR and BAIRNet, another model denoted as IARN* is trained using asymmetric scale in both stages. As shown in Table 4, two challenging large benchmark test sets are used for assessment in three asymmetric and one symmetric scales. Results from ArbSR, IRN and BAIRNet are used for comparisons, using the same evaluation protocol as in BAIRNet. It is demonstrated that, both in quantitative metrics and visual exampled as included in Fig. 6, our IARN* holds a clear advantage over other peers in restoring more accurate details.

4.6. Model Complexity and Efficiency

To assess the efficiency of our proposed IARN, it is compared with others in terms of model size and average inference time per image for the BSD100 test set. As illustrated in Table 5, three most relevant models which are optimized jointly for downscaling and upscaling are included so the inference time includes both downscaling and upscaling and listed separately. Note that for CAR and IRN, they have models of different sizes depending on the scale factor, while for BAIRNet and our IARN, only one model is needed for both tests. It is shown that our IARN is comparable to IRN ×4 in terms of number of parameters, much more smaller than CAR and BAIRNet. For inference speed, it is slower than CAR and IRN. But for the only two models capable of bidirectional arbitrary rescaling (AIDN not yet available for assessment), our model is able to outperform the previous SOTA BAIRNet consistently while using less than 20% parameters and reducing inference time by 66% and 60% for $\times 2$ and $\times 4$ respectively.

Table 5: Model size and inference time comparison.

Method	Scale	Param	Downscaling	Upscaling
CAR [10]	$\times 2$	51.1M	0.004s	0.005s
CAR [19]	$\times 4$	52.8M	0.004s	0.005s
IDN [24]	$\times 2$	1.66M	0.018s	0.021s
IKN [24]	$\times 4$	4.35M	0.025s	0.026s
BAIRNet [18]	$\times 2$	22.4M	0.506s	0.129s
DAIKINCI [10]	$\times 4$	22.4111	0.459s	0.061s
LADN	$\times 2$	4.32M	0.058s	0.154s
IAKIN	$\times 4$		0.058s	0.154s

5. Conclusions

Image arbitrary rescaling using deep learning is a relatively new and under explored topic in the field of low-level image processing due to its complexity. In this paper, we have presented the first invertible arbitrary image rescaling work. Based on an INN backbone enhanced with a novel preemptive channel splitting module and a new positionaware scale encoding method, the newly proposed IARN network is capable of handling bidirectional image arbitrary rescaling over a large range of scales using just one trained model. Extensive experiments on a comprehensive set of benchmark datasets validate a much better performance of arbitrary image rescaling over the current related SOTA methods in both HR and LR outputs, with reduced model size and faster inference comparing to BAIRNet. In addition, better performance is also achieved for asymmetric arbitrary rescaling tests.

References

- 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] Lynton Ardizzone, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W Pellegrini, Ralf S Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. arXiv preprint arXiv:1808.04730, 2018.
- [3] Parichehr Behjati, Pau Rodriguez, Armin Mehri, Isabelle Hupont, Carles Fernandez Tena, and Jordi Gonzalez. Overnet: Lightweight multi-scale super-resolution with overscaling network. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2694– 2703, 2021.
- [4] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie line Alberi Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *Proceedings of the British Machine Vision Conference*, pages 135.1–135.10. BMVA Press, 2012.
- [5] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8628–8638, 2021.
- [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 vi*sion, pages 184–199. Springer, 2014.
- [7] Xuecai Hu, Haoyuan Mu, Xiangyu Zhang, Zilei Wang, Tieniu Tan, and Jian Sun. Meta-SR: A magnificationarbitrary network for super-resolution. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1575–1584, 2019.
- [8] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5197–5206, 2015.
- [9] 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.
- [10] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeplyrecursive convolutional network for image super-resolution. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1637–1645, 2016.
- [11] Christiane Lemke, Marcin Budka, and Bogdan Gabrys. Metalearning: a survey of trends and technologies. *Artificial intelligence review*, 44(1):117–130, 2015.
- [12] Juncheng Li, Faming Fang, Kangfu Mei, and Guixu Zhang. Multi-scale residual network for image super-resolution. In Proceedings of the European Conference on Computer Vision (ECCV), pages 517–532, 2018.
- [13] Shang Li, Guixuan Zhang, Zhengxiong Luo, Jie Liu, Zhi Zeng, and Shuwu Zhang. Approaching the limit

of image rescaling via flow guidance. *arXiv preprint arXiv:2111.05133*, 2021.

- [14] Jingyun Liang, Andreas Lugmayr, Kai Zhang, Martin Danelljan, Luc Van Gool, and Radu Timofte. Hierarchical conditional flow: A unified framework for image superresolution and image rescaling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4076–4085, 2021.
- [15] 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.
- [16] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001.
- [17] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012.
- [18] Zhihong Pan, Baopu Li, Dongliang He, Mingde Yao, Wenhao Wu, Tianwei Lin, Xin Li, and Errui Ding. Towards bidirectional arbitrary image rescaling: Joint optimization and cycle idempotence. pages 17389–17398, 2022.
- [19] Wanjie Sun and Zhenzhong Chen. Learned image downscaling for upscaling using content adaptive resampler. *IEEE Transactions on Image Processing*, 29:4027–4040, 2020.
- [20] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. NTIRE 2017 challenge on single image super-resolution: Methods and results. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops, pages 114–125, 2017.
- [21] Longguang Wang, Yingqian Wang, Zaiping Lin, Jungang Yang, Wei An, and Yulan Guo. Learning a single network for scale-arbitrary super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4801–4810, 2021.
- [22] 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 Vi*sion (ECCV), pages 63–79, 2018.
- [23] 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.
- [24] 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.
- [25] Jinbo Xing, Wenbo Hu, and Tien-Tsin Wong. Scalearbitrary invertible image downscaling. arXiv preprint arXiv:2201.12576, 2022.
- [26] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Interna-*

tional conference on curves and surfaces, pages 711–730. Springer, 2010.

- [27] 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.
- [28] 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.