

Burst Super-Resolution with Diffusion Models for Improving Perceptual Quality

Kyotaro Tokoro

Toyota Technological Institute, Japan
sd24439@toyota-ti.ac.jp

Kazutoshi Akita

Toyota Technological Institute, Japan
sd21501@toyota-ti.ac.jp

Norimichi Ukita

Toyota Technological Institute, Japan
ukita@toyota-ti.ac.jp

Abstract—While burst LR images are useful for improving the SR image quality compared with a single LR image, prior SR networks accepting the burst LR images are trained in a deterministic manner, which is known to produce a blurry SR image. In addition, it is difficult to perfectly align the burst LR images, making the SR image more blurry. Since such blurry images are perceptually degraded, we aim to reconstruct the sharp high-fidelity boundaries. Such high-fidelity images can be reconstructed by diffusion models. However, prior SR methods using the diffusion model are not properly optimized for the burst SR task. Specifically, the reverse process starting from a random sample is not optimized for image enhancement and restoration methods, including burst SR. In our proposed method, on the other hand, burst LR features are used to reconstruct the initial burst SR image that is fed into an intermediate step in the diffusion model. This reverse process from the intermediate step 1) skips diffusion steps for reconstructing the global structure of the image and 2) focuses on steps for refining detailed textures. Our experimental results demonstrate that our method can improve the scores of the perceptual quality metrics. Code: <https://github.com/placerkyo/BSRD>.

Index Terms—Burst SR, Diffusion models, Perceptual quality

I. INTRODUCTION

Super-Resolution (SR) is a low-level vision task for super-resolving a Low-Resolution (LR) image to its High-Resolution (HR) image. SR is widely applicable to many real-world tasks, such as image enhancement on consumer cameras and remote sensing. SR for super-resolving an LR image is called Single-Image Super-Resolution (SISR) [1]–[5]. However, SISR is not an easy task due to its ill-posed nature. In particular, SR for a degraded LR image is much more difficult in SISR. Such degraded images are often captured by, for example, commercially available smartphones.

To improve the SR quality, SR can accept multiple LR images. Assume that a quasi-static scene is captured many times. Such a burst shot allows us to capture many burst images with subtle differences, such as sub-pixel displacements and small noise. By integrating all the burst images into one [6], the SR quality can be improved compared with SISR. This type of SR is called burst SR [7], [8]. Since the burst shot is allowed in recent smartphones, burst SR is a promising technique.

As with many computer vision tasks, burst SR can be improved by deep learning (e.g., [9]). In previous burst SR methods using deep learning, an SR image is reconstructed in a deterministic manner with a loss function based on the Mean Squared Error (Fig. 1 (a)). However, such an MSE-based loss function leads to blurred SR images because the MSE is minimized by the average of training HR images [10].

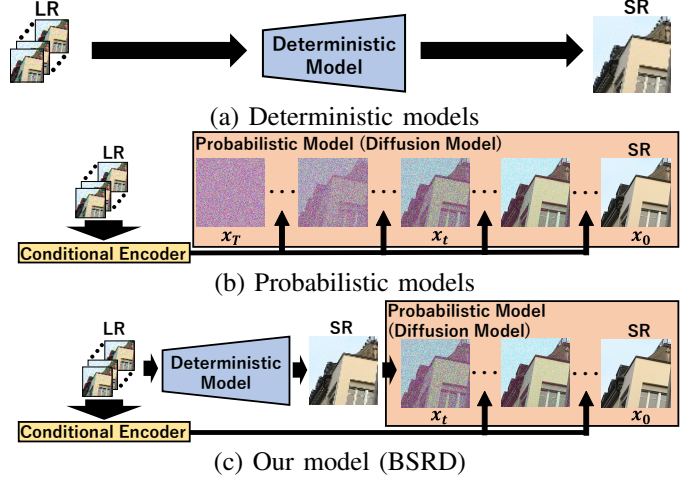


Fig. 1. Comparison between prior models and our model (i.e., BSRD).

To avoid this problem, this paper aims to improve the perceptual quality of the SR image by probabilistic modeling. Probabilistic modeling allows us to represent the probabilistic distribution of sharp SR images corresponding to each LR image. Among such probabilistic models, our method employs diffusion models [11], [12] (Fig. 1 (b)) because of their performance validated in various computer vision tasks.

While the diffusion model begins with a random noise for image synthesis, the diffusion model should be conditioned for an output image that fits with input images for image enhancement and restoration, including SR. In SISR using the diffusion model such as [13], [14], an input LR image is used for conditioning the reverse process. Unlike SISR, however, burst SR accepts multiple displaced LR images. If such displaced images are directly fed into the diffusion model, the reverse process may reconstruct the blurry SR image.

This paper proposes a burst SR method in which the aforementioned burst images are appropriately employed for improving the perceptual quality of the SR image. The core contributions of this work are as follows (Fig. 1 (c)):

- To use the diffusion model for burst SR, features useful for SR are extracted from input burst images and employed for conditioning the diffusion model.
- Since a diffusion model with between random noise and output images is not optimized for the SR task, the initial burst SR image reconstructed by simple burst SR is fed into the intermediate step of the diffusion model to focus on refining detailed textures on the burst SR image.

II. RELATED WORK

A. Burst SR

While a RAW image consisting of RGGB channels with high bits per pixel resolution (e.g., 14 bits, 16 bits) is captured by a standard digital camera, the RAW image is converted by an Image Signal Processing (ISP) to its 8-bit RGB image. General burst SR tasks accept a set of unprocessed RAW images instead of their processed RGB images for the following two reasons. 1) Since the function of ISP depends on the camera, there is a domain gap between the RGB images processed by different cameras [15], [16]. This domain gap makes it difficult to apply a universal SR model to various camera images. We expect that the domain gap between RAW images is smaller. 2) Image conversion from a raw image with high bits per pixel to its 8-bit RGB image drops useful cues for burst SR.

Many burst SR methods generally consist of four processes: feature extraction, alignment, fusion, and reconstruction. The alignment process rectifies image features extracted by the feature extraction process so that the features are spatially consistent. This spatial rectification is achieved to align all images to the reference frame. The aligned features are then merged by the fusion process. The fused features are fed into the reconstruction process to acquire the SR image.

Among the four processes, the alignment process is peculiar and essential to burst SR. If non-aligned features are directly used for SR, the SR image may be blurred. While the alignment process is required also for video SR [17]–[20], the properties of displacements in video frames and burst images are different. Therefore, the alignment process optimized to burst SR is required. For example, since simple optical flow estimation is employed for the alignment process in DBSR [6], its SR performance is limited because simple optical flow estimation is not optimized for sub-pixel level flows. In [21], [22], on the other hand, deformable convolution (DC) [23] is used for implicit spatial alignment. Adaptively optimized receptive fields in DC improve the alignment process for burst SR. The two-step alignment with optical flow estimation and DC is also proposed in [24].

However, all burst SR methods introduced in Sec. II-A reconstruct SR images in a deterministic manner, which prones to blurry SR images.

B. Diffusion Models

A diffusion model consists of the diffusion and reverse processes. The diffusion process changes an original data (e.g., a noiseless image denoted by x_0) towards its corresponding random noise (denoted by x_T) by gradually giving noise through T steps (e.g., $T = 1000$ steps). The reverse process, on the other hand, gradually recovers the original data from a random noise. Given x_{t-1} at $(t-1)$ -th step, t -th step in the diffusion process predicts x_t according to the conditional probability $q(x_t|x_{t-1})$ as follows:

$$q(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{\alpha_t}x_{t-1}, \beta_t), \quad (1)$$

where $\mathcal{N}(x; \mu, \sigma)$ denotes a Gaussian distribution whose mean and covariance matrix are μ and σ , respectively. $0 < \beta_1 <$

$\beta_2 < \dots < \beta_T < 1$ is a hyper-parameter that adjusts the variance so that $\alpha_t := 1 - \beta_t$. On the other hand, the reverse process in t -th step predicts x_{t-1} according to $p(x_{t-1}|x_t)$.

Given x_{t-1} , small noise, and x_t that is x_{t-1} added by small noise, the reverse process described above is trained so that t -th step can predict this small noise from x_t . The reverse process is generally implemented with U-Net [25].

As mentioned in Sec. I, while the reverse process begins from random noise for image synthesis, the reverse process should be controlled to reconstruct an output image that fits with an input image for SR, as described in what follows.

C. Image Conditioning for Diffusion Models

1) *Image Enhancement and Restoration with Diffusion Models*: Diffusion models are applied to various image enhancement and restoration tasks, such as image deblurring [26], [27], image inpainting [28], [29], and image colorization [30], [31]. In all of these methods, input degraded images are used to condition the diffusion model for each task.

Unlike these methods that begin with a random noise, SDEdit [32] begins from an intermediate step. SDEdit is designed for realistic image synthesis from sketches. While the diffusion model is not trained with sketches, its reverse process can convert a sketch to its realistic image by denoising the appropriately-noised sketch from the intermediate step. While SDEdit is originally proposed for sketch-to-image style transfer, our proposed method applies the SDEdit-style reverse process starting from the intermediate step to burst SR.

2) *SISR with Diffusion Models*: As with other image enhancement and restoration tasks mentioned above [26]–[31], an input degraded image (i.e., LR image in the SR task) is used for conditioning the diffusion model. In SR3 [13], the LR image is upsampled by Bicubic interpolation and concatenated with x_t for conditioning. In LDMs [33], features extracted from the LR image are fed into the middle layers of U-Net in the reverse process through the cross attention mechanism. As well as the features extracted from the LR image, the number of the diffusion step is also used for step-aware conditioning in Stable SR [14].

Unlike these SISR methods [13], [14], [33], this paper proposes the diffusion model conditioned with burst LR images.

III. PROPOSED METHOD

The feature extraction and alignment modules borrowed from the deterministic burst SR method (Burstormer [34]) are explained in Sec. III-A. Our probabilistic burst SR method (Burst Super-Resolution with Diffusion Model: BSRD) is proposed in Sec. III-B1. BSRD is augmented by the reverse process starting from an intermediate step (Sec. III-B2).

A. Feature Extraction and Alignment Modules

Burstormer [34] is one of the state-of-the-art deterministic burst SR methods. The feature extraction and alignment networks in Burstormer are illustrated in Fig. 2. Given B input burst raw LR images, each of which has four color channels with W and H pixels, features with $B \times f \times W \times H$ channels,

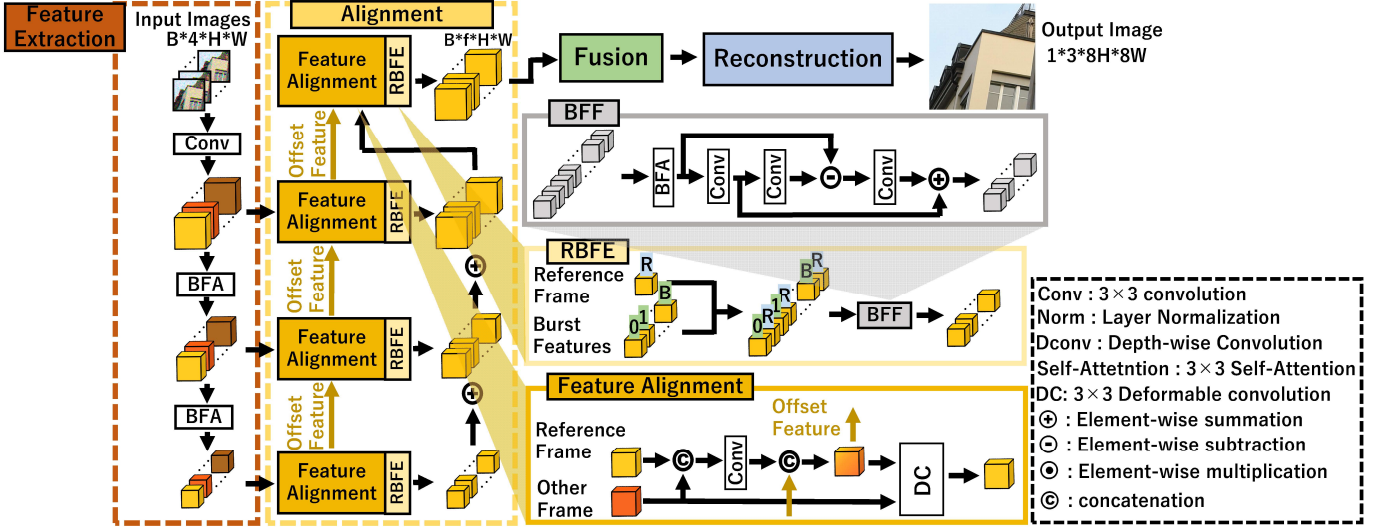


Fig. 2. Feature extraction and alignment in Burstormer [34], which are colored by red and yellow, respectively. Different colors in the feature extraction process mean spatial displacements.

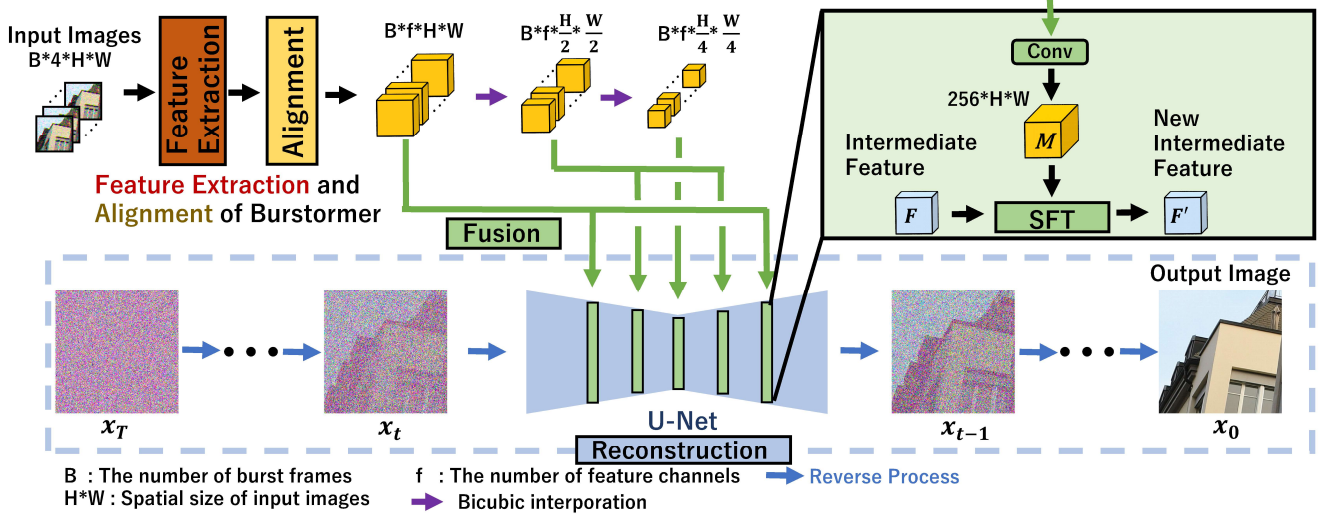


Fig. 3. Overview of the feature extraction, alignment, fusion, and reconstruction processes in BSRD. The feature extraction and alignment modules are borrowed from those of Burstormer, which are shown in Fig. 2. In our proposed fusion module, SFT [35] is included in U-Net and used for conditioning with the LR features. The reconstruction process is achieved by the reverse process of the diffusion model.

where f denotes the feature dimension, are extracted. Feature channels colored by different colors mean that these channels are spatially displaced. These features are fed into the Burst Feature Attention (BFA) modules for noise reduction and downsampling. Since BFA is implemented by Transformer [36], noise is reduced based on the global structure of the images. These noise reduction and downsampling can simplify feature alignment. The alignment process consists of the feature alignment and Reference-Based Feature Enrichment (RBFE) modules. The hierarchically-aligned features are fed into the RBFE modules in each of which the Burst Feature Fusion (BFF) module merges the features of the reference frame and other frames for feature alignment according to the

reference frame.

The output features of the four hierarchical feature alignment modules are fed into the fusion module in which the features are merged by average pooling. The merged features are fed into the reconstruction module for the final SR image.

Since the hierarchical feature representation is powerful for image alignment in burst SR as well as for various computer vision tasks, our method employs the feature extraction and alignment processes used in Burstormer. On the other hand, the fusion and reconstruction processes are achieved by the diffusion model in our method described in what follows.




Fig. 4. Reverse process from the intermediate step. Instead of the reverse process starting from T -th step, the initial burst SR image is appropriately noised and fed into the diffusion model from t -th step.

B. BSRD: Burst Super-Resolution with Diffusion Model

1) *Burst Feature Conditioning for Diffusion Models*: As mentioned in Sec. III-A, the feature extraction and alignment modules are borrowed from those of Burstormer, which are shown in Fig. 2. While the fusion and reconstruction processes in Burstormer are deterministic, BSRD employs the diffusion model for probabilistic burst SR, as illustrated in Fig. 3.

a) *Fusion*: In BSRD, features provided from burst LR images by the feature extraction and alignment modules are used for conditioning the diffusion model to reconstruct an SR image that fits with input burst LR image. This conditioning should be designed so that the conditioning features are properly formatted according to the network architecture of the reverse process in the diffusion model; inappropriate conditioning decreases the effect of this conditioning. As well as the reverse process of general diffusion models, that of BSRD is implemented with U-Net. We assume that hierarchical layers in U-Net should be conditioned for better SR reconstruction. For conditioning hierarchical layers in U-Net, the output of the alignment module (which is depicted as the orange cuboids with $B \times f \times W \times H$ channels in Fig. 3) is rescaled to the xy dimensions of these hierarchical layers by Bicubic interpolation as depicted by the purple arrows in Fig. 3. The B channels in each rescaled feature map are merged as done in Burstormer. For example, the feature map with $B \times f \times W \times H$ channels is converted to the one with $256 \times W \times H$. The 256 channels are merged by a convolutional layer. The merged feature map (denoted by M) is fed into U-Net for conditioning. This conditioning is achieved through Spatial Feature Transformation (SFT) [35]. In SFT, the conditioning feature map with $256 \times W \times H$ (i.e., M in BSRD) is transformed to two feature maps with $f_U \times W \times H$ channels, where f_U denotes the feature dimension of each U-Net layer, by two independent convolutional layers. Given these two feature maps (denoted by α and β) and the original feature map in U-Net (denoted by F with $f_U \times W \times H$ channels), SFT acquires the conditioned feature map denoted

by F' as follows:

$$F' = (1 + \alpha) \odot F + \beta \quad (2)$$

The aforementioned fusion process is performed in all steps in the diffusion model.

b) *Reconstruction*: While the reverse process of the diffusion model in BSRD is conditioned as described above, the continuous reverse steps are achieved in BSRD as with in the original diffusion model. This reverse process is regarded as the reconstruction process in BSRD.

2) *Efficient and High-quality SR Reconstruction by the Reverse Process from Intermediate Steps*: In the reverse process of the diffusion model for image synthesis, it is known that different steps take on different roles [37]. During steps handling noisy images (e.g., at and near T -th step), overall color and structure are synthesized. The overall image appearance is determined during these steps. In the middle steps, contextual information is reconstructed. Finally, more detailed textures and boundaries are synthesized during steps handling less-noisy images. Such an iterative process with many steps requires a huge computational cost, which is one of the major drawbacks of the diffusion model [38]. Furthermore, in the reverse process starting from a random sample, the conditioning scheme is required to significantly adjust the reverse process for SR (and other image enhancement and restoration methods) compared with free image synthesis.

To suppress the difficulty in computational cost and SR using the diffusion model, BSRD accepts an initial burst SR image instead of a random noise. To start the reverse process from the initial burst SR image, it is fed into an intermediate step, instead of T -th step, following SDEdit [32]. With this reverse process from the intermediate step, the number of execution steps is reduced, resulting in reducing the computational cost. Furthermore, since the difficult steps near T -th steps are skipped, the difficulty in reconstructing the SR image close to the input LR image can be reduced, and the reverse process can be trained to focus on reconstructing fine details that are important for the SR task.

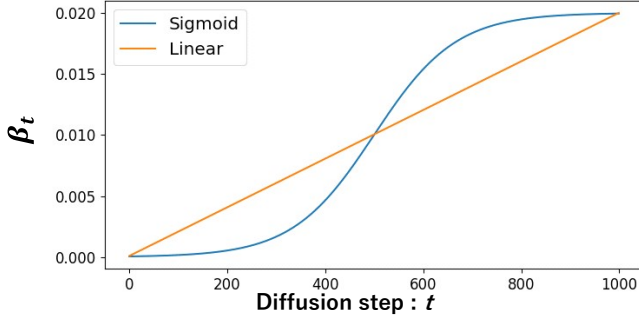


Fig. 5. Comparison between the linear and sigmoid schedulers.

The aforementioned reverse process in BSRD is illustrated in Fig. 4. Remember x_t denotes the image at t -th step in the diffusion model. In BSRD, x_τ is replaced by the initial burst SR image with noise (denoted by x'_τ). This time step τ is called the reverse start step. Any SoTA burst SR method can be used for obtaining x'_τ . Let x'_0 denote the initial burst SR image. While x'_0 corresponds to x_0 with no diffusion noise, x'_0 may be more blurred than x_0 because x'_0 is reconstructed in a deterministic manner. However, x'_τ computed from x'_0 must contain the diffusion noise that is included in x_τ . We assume that x_τ can be approximated by x'_τ by giving the diffusion noise at τ -th step to x'_0 . That is, small differences caused by the blur between x_0 and x'_0 can be drowned out by the diffusion noise if τ -th step is sufficiently apart from 0-th step. x'_τ is computed from x'_0 as follows:

$$x'_\tau = \sqrt{\bar{\alpha}_\tau} x'_0 + \sqrt{\bar{\beta}_\tau} \epsilon, \quad (3)$$

$$\bar{\alpha}_\tau := \prod_{s=1}^{\tau} \alpha_s, \quad (4)$$

$$\bar{\beta}_\tau := 1 - \bar{\alpha}_\tau, \quad (5)$$

where ϵ denotes the zero-mean Gaussian noise with $\sigma = 1$.

While the computational cost for computing the initial burst SR image is required, its cost is much smaller than the iterative steps in the diffusion model.

Since BSRD uses the diffusion steps between τ -th and 1-th steps, $(\tau + 1)$ -th and greater steps are not required. Therefore, the diffusion model can be trained only between τ -th and 1-th steps in the training stage. This training strategy allows the diffusion model to focus only on reconstructing detailed appearances represented near the 1-th step, resulting in better SR.

IV. EXPERIMENTAL RESULTS

A. Details

a) Evaluation Metrics: As standard evaluation metrics, LPIPS [39], FID [40], PSNR, and SSIM [41] are used. While PSNR and SSIM are for image distortion-based evaluation, LPIPS and FID are proposed for perceptual quality evaluation. Lower scores are better in LPIPS and FID, while higher scores are better in PSNR and SSIM.

b) Training Details: The number of diffusion steps is $T = 1000$. As noise schedulers for the diffusion model, the linear scheduler between $\beta_1 = 10^{-4}$ and $\beta_T = 0.02$ and the Sigmoid scheduler between $\beta_1 = 10^{-5}$ and $\beta_T = 0.02$ are used (Fig. 5). The mini-batch size is 50. The optimizer is AdamW [42], and its learning rate is 1×10^{-4} .

Since our BSRD requires only early steps (i.e., between 1-th and τ -th steps) in the diffusion model, we can train it only between 1-th and τ -th steps, as described at the end of Sec. III-B2. For all experiments shown in Sec. IV-B and Sec. IV-C, however, the diffusion model was trained between 1-th and 100-th steps for simplifying experiments (i.e., for conducting experiments using different τ with one trained model). Note that several models evaluated in Sec. IV-D are trained between different steps to validate the effectiveness of our proposed reverse process from intermediate steps; see Table III for details.

c) Datasets: The SyntheticBurst and BurstSR datasets are used [6]. On the SyntheticBurst dataset, each sRGB HR image is degraded to its LR image. This degradation process, including ISP, follows the one proposed in [6] as follows. Burst LR images, except for the reference frame, are randomly (1) translated up to 24 pixels along x and y axes and (2) rotated up to one degree. The SyntheticBurst dataset has 46,839 training images and 300 test images. In the BurstSR dataset, both LR and HR images are real images captured. Each set of LR images was captured by a burst shot mode of Samsung Galaxy S8. These LR images are subtly different from each other due to hand shake effects. Each HR image was captured by CANON 5D Mark IV. The BurstSR dataset has 5,405 training images and 882 test images.

d) Burst SR Condition: In all experiments, the dimensions of LR and HR images are $32 \times 32 \times 4$ RAW images and $256 \times 256 \times 3$ RGB images, respectively. The number of burst LR images is eight. The dimension of the conditioning feature is $f = 48$. For comparison, DBSR [6], BIPNet [43], and Burstormer [34] are evaluated. While the authors' pre-trained weights are used for DBSR and BIPNet, the pre-trained weight of Burstormer is further finetuned for better performance.

Initial burst SR images given to our method (BSRD) are given by BIPNet and Burstormer. For BSRD, an intermediate step (τ in Sec. III-B2) from which the reverse process starts is an important hyperparameter. In our experiments, τ is determined empirically. Since our goal is to improve the perceptual quality such as LPIPS and FID, τ giving the best perceptual score in LPIPS and FID is selected. LPIPS and FID in different steps on the SyntheticBurst dataset are shown in Fig. 6. In the tables in Fig. 6, the score at $\tau = 0$ equals that of each initial burst SR image. Based on the results shown in Fig. 6, τ for “the linear scheduler for LPIPS,” “the linear scheduler for FID,” “the sigmoid scheduler for LPIPS,” and “the sigmoid scheduler for FID” are 1, 3, 3, and 100, respectively, for BIPNet with BSRD. The results with these parameters on the SyntheticBurst dataset are shown in Table I. For Burstormer with BSRD, 1, 5, 30, and 100. For “BIPNet” and “Burstormer” with BSRD in the BurstSR dataset, “6, 30,

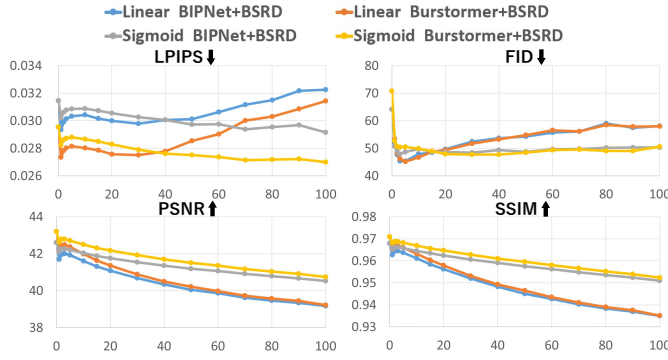


Fig. 6. Results on the SyntheticBurst dataset. The vertical and horizontal axes indicate the reverse start step, τ , and the quality measures, respectively. While the perceptual scores (i.e., LPIPS and FID) are important for our work, PSNR and SSIM are also shown just for reference.

TABLE I
QUANTITATIVE COMPARISON RESULTS ON THE SYNTHETICBURST DATASET. THE BEST AND SECOND BEST SCORES ARE COLORED RED AND BLUE, RESPECTIVELY.

Methods	τ	Noise Schedule	LPIPS↓	FID↓	PSNR↑	SSIM↑
DBSR			0.07212	99.66	40.05	0.9482
BIPNet			0.03145	64.17	42.61	0.9682
BIPNet +BSRD	1	linear	0.02934	51.92	41.69	0.9627
	3	linear	0.03013	45.32	42.00	0.9645
	3	sigmoid	0.03077	47.88	42.26	0.9662
	100	sigmoid	0.02915	50.26	40.53	0.9509
Burstormer			0.02954	70.76	43.21	0.9709
Burstormer +BSRD	1	linear	0.02734	53.42	42.17	0.9653
	5	linear	0.02815	45.16	42.37	0.9660
	30	sigmoid	0.02791	47.74	41.92	0.9628
	100	sigmoid	0.02699	50.61	40.74	0.9524

60, and 70” and “6, 60, 70, and 80,” respectively, as shown in Table II.

B. Results: SyntheticBurst Dataset

Quantitative comparison results on the SyntheticBurst dataset are shown in Table I. Independently of the reverse start step, the noise scheduler, and the base method (i.e., BIPNet or Burstormer), both LPIPS and FID are improved by BSRD from the base method. BSRD improves LPIPS up to about 8.6% with the sigmoid scheduler and $\tau = 100$ compared with the original Burstormer (i.e., 0.02954 in Burstormer vs. 0.02699 in Burstormer + BSRD). Furthermore, BSRD improves FID up to about 36% with the linear scheduler and $\tau = 5$ compared with the original Burstormer (i.e., 70.76 in Burstormer vs. 45.16 in Burstormer + BSRD).

However, the performance change depending on τ differs between the noise schedulers. With the linear scheduler, both LPIPS and FID degrade from earlier steps (i.e., around 1 and 5 steps in LPIPS and FID, respectively). With the sigmoid scheduler, on the other hand, these scores improve until later steps. In particular, LPIPS improves until around 100 steps.

TABLE II
QUANTITATIVE COMPARISON RESULTS ON THE BURSTSR DATASET.

Methods	τ	Noise Schedule	LPIPS↓	FID↓	PSNR↑	SSIM↑
DBSR			0.05467	90.34	50.72	0.9846
BIPNet			0.04975	75.43	51.55	0.9857
BIPNet +BSRD	6	linear	0.04813	43.34	50.91	0.9850
	30	linear	0.04682	46.65	50.32	0.9836
	60	sigmoid	0.04688	43.34	50.54	0.9835
	70	sigmoid	0.04719	43.11	50.48	0.9839
Burstormer			0.05593	74.21	50.49	0.9845
Burstormer +BSRD	6	linear	0.05445	46.85	50.11	0.9837
	60	linear	0.05017	48.51	49.45	0.9815
	70	sigmoid	0.05238	44.54	49.94	0.9826
	80	sigmoid	0.05178	45.20	49.91	0.9823

This is because the sigmoid scheduler finely changes the noise at early steps, as shown in Fig. 5. This scheduling allows the reverse process to focus on detailed textures and boundaries optimized in early steps, as described in Sec. III-B2.

Different from the perceptual scores (i.e., LPIPS and FID), the image distortion scores (i.e., PSNR and SSIM) are degraded by our method. This result is unavoidable due to the tradeoff between the perceptual and image distortion scores, as demonstrated in [44].

Visual results are shown in Fig. 7. The result of BSRD is obtained with the sigmoid scheduler with $\tau = 100$. In the upper results, while the boundaries between white and red pixels are blurred in the result of Burstormer, our method can reconstruct the sharp boundaries. In the lower results, the boundary between the different colors on the road is blurred in the result of Burstormer, while our method can reconstruct the clear boundary. In addition, the road surface is flat and textureless in the result of Burstormer, while the texture of the road surface is reconstructed by our method. These results demonstrate that our method can improve the perceptual quality with clear boundary lines.

C. Results: BurstSR dataset

Quantitative comparison results on the BurstSR dataset are shown in Table II. As with the SyntheticBurst dataset, the BurstSR dataset demonstrates that our method can improve the perceptual quality. For example, our method improves LPIPS and FID by 5.8% and 43%, respectively (i.e., 0.04975 with BIPNet vs. 0.04642 with our method in LPIPS, and 75.43 with BIPNet vs. 43.11 with our method in FID). With Burstormer, our method can also improve the perceptual quality scores, while its improvement is lower than that of BIPNet.

Visual results on the BurstSR dataset are shown in Fig. 8. Our method (i.e., “BIPNet+BSRD” in Fig. 8) is combined with BIPNet because this combination produces the best results on the BurstSR dataset, as shown in Table II. Our method outperforms all other methods, including its base method (i.e., BIPNet), regarding the sharpness of boundary lines.

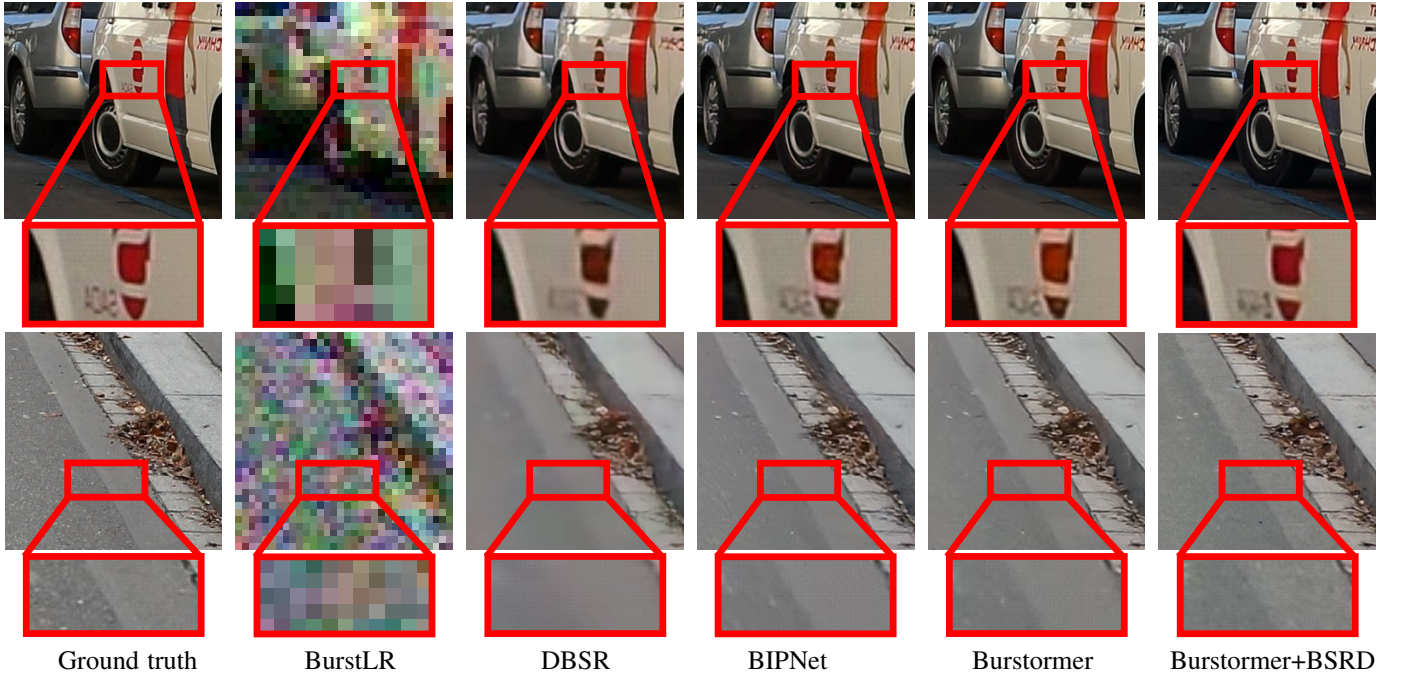


Fig. 7. Visual results on the SyntheticBurst dataset. In both the upper and lower results, a part of each image is zoomed in for detail.




Fig. 8. Visual results on the BurstSR dataset.

TABLE III
DETAILED ANALYSIS ON THE SYNTHETICBURST DATASET FOR
VALIDATING THE EFFECTIVENESS OF EACH COMPONENT IN OUR BSRD.
RESULTS IN EACH ROW ARE ACQUIRED WITH THE DIFFUSION MODEL
TRAINED BETWEEN 1-TH AND τ_L -TH STEPS.

	Methods	τ_L	τ	LPIPS↓	FID↓	PSNR↑	SSIM↑
(a)	From random noise	1000	1000	0.06679	85.02	33.86	0.8730
(b)	From bicubic	1000	400	0.07164	85.79	35.60	0.8334
(c)	upscaled image	100	5	0.31882	206.37	31.03	0.8178
(d)	From Burstormer	1000	5	0.02851	49.19	42.19	0.9649
(e)	SR image	100	5	0.02815	45.16	42.37	0.9660

D. Detailed Analysis

To validate the effectiveness of each implementation in our proposed BSRD, the following experiments are conducted:

- (a) begins from random noise (i.e., $\tau = 1000$) instead of from the intermediate step.
- (b) and (c) use the initial image upscaled by bicubic interpolation instead of the burst SR image reconstructed by Burstormer.
- (d) and (e) use the burst SR image reconstructed by Burstormer. (e) is our proposed implementation of BSRD.

We can see the following observations in Table III:

- 1) In comparison between (a) and (d,e), the reverse process from the intermediate step is better.
- 2) In comparison between (b,c) and (d,e), the reverse process with the initial burst SR image is better.
- 3) In comparison between (d) and (e), the model trained only from τ -th step is better.

V. CONCLUDING REMARKS

This paper proposed a combination of burst SR and diffusion models for improving the perceptual quality of burst SR. Unlike prior diffusion-based burst SR methods, which start from random noise, our BSRD starts the reverse process of the diffusion model from an intermediate step. This allows the diffusion model to efficiently focus on reconstructing detailed textures and boundaries in earlier diffusion steps. Furthermore, hierarchical burst SR features are employed to condition the reverse process to optimize it for the burst SR task.

Future work includes more efficiency improvement. It is validated that latent diffusion is useful for SISR as well as for other computer vision tasks [33], [45]. For burst SR also, the effectiveness of the latent diffusion should be explored.

This work is partly supported by JSPS KAKENHI 22H03618.

REFERENCES

- [1] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(2):295–307, 2016.
- [2] Radu Timofte et al. NTIRE 2018 challenge on single image super-resolution: Methods and results. In *ICCVW*, 2018.
- [3] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Deep back-projection networks for super-resolution. In *CVPR*, 2018.
- [4] Shuhang Gu et al. AIM 2019 challenge on image extreme super-resolution: Methods and results. In *ICCVW*, 2019.
- [5] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Deep back-project networks for single image super-resolution. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(12):4323–4337, 2021.
- [6] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In *CVPR*, 2021.
- [7] Goutam Bhat et al. NTIRE 2021 challenge on burst super-resolution: Methods and results. In *CVPRW*, 2021.
- [8] Goutam Bhat et al. NTIRE 2022 burst super-resolution challenge. In *CVPRW*, 2022.
- [9] Andrea Bordone Molini, Diego Valsesia, Giulia Fracastoro, and Enrico Magli. Deepsum: Deep neural network for super-resolution of unregistered multitemporal images. *IEEE Trans. Geosci. Remote. Sens.*, 58(5):3644–3656, 2020.
- [10] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.
- [12] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015.
- [13] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(4):4713–4726, 2023.
- [14] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin C. K. Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *CoRR*, abs/2305.07015, 2023.
- [15] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T. Barron. Unprocessing images for learned raw denoising. In *CVPR*, 2019.
- [16] Andrey Ignatov, Luc Van Gool, and Radu Timofte. Replacing mobile camera ISP with a single deep learning model. In *CVPR Workshops*, 2020.
- [17] Seungjun Nah et al. NTIRE 2019 challenge on video super-resolution: Methods and results. In *CVPR Workshop*, 2019.
- [18] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Recurrent back-projection network for video super-resolution. In *CVPR*, 2019.
- [19] Dario Fuoli et al. AIM 2020 challenge on video extreme super-resolution: Methods and results. In *ECCV Workshop*, 2020.
- [20] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Space-time-aware multi-resolution video enhancement. In *CVPR*, 2020.
- [21] Ziwei Luo, Lei Yu, Xuan Mo, Youwei Li, Lanpeng Jia, Haoqiang Fan, Jian Sun, and Shuaicheng Liu. EBSR: feature enhanced burst super-resolution with deformable alignment. In *CVPR Workshops*, 2021.
- [22] Shi Guo, Xi Yang, Jianqi Ma, Gaofeng Ren, and Lei Zhang. A differentiable two-stage alignment scheme for burst image reconstruction with large shift. In *CVPR*, 2022.
- [23] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017.
- [24] Ziwei Luo, Youwei Li, Shen Cheng, Lei Yu, Qi Wu, Zhihong Wen, Haoqiang Fan, Jian Sun, and Shuaicheng Liu. BSRT: improving burst super-resolution with swin transformer and flow-guided deformable alignment. In *CVPR Workshops*, 2022.
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [26] Hyungjin Chung, Jeongsol Kim, Sehui Kim, and Jong Chul Ye. Parallel diffusion models of operator and image for blind inverse problems. In *CVPR*, 2023.
- [27] Naoki Murata, Koichi Saito, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Yuki Mitsufuji, and Stefano Ermon. Gibbsddrm: A partially collapsed gibbs sampler for solving blind inverse problems with denoising diffusion restoration. In *ICML*, 2023.
- [28] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *ICLR*, 2023.
- [29] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In *NeurIPS*, 2022.
- [30] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *NeurIPS*, 2022.
- [31] Xiangming Meng and Yoshiyuki Kabashima. Diffusion model based posterior sampling for noisy linear inverse problems. *CoRR*, abs/2211.12343, 2022.
- [32] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2022.
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [34] Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Burstformer: Burst image restoration and enhancement transformer. In *CVPR*, 2023.
- [35] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *CVPR*, 2018.
- [36] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022.
- [37] Jooyoung Choi, Jungbeom Lee, Chaehun Shin, Sungwon Kim, Hyunwoo Kim, and Sungroh Yoon. Perception prioritized training of diffusion models. In *CVPR*, 2022.
- [38] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021.
- [39] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [40] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017.
- [41] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004.
- [42] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- [43] Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Burst image restoration and enhancement. In *CVPR*, 2022.
- [44] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *CVPR*, 2018.
- [45] Takeru Oba, Matthew Walter, and Norimichi Ukita. Read: Retrieval-enhanced asymmetric diffusion for motion planning. In *CVPR*, 2024.