Splice, Focus and Relife: High-Resolution Periodic Pattern Generation

Xicheng Lan, Wenshuo Gao, Luyao Zhang, Jiaying Liu and Shuai Yang*
Wangxuan Institute of Computer Technology, State Key Laboratory of Multimedia Information Processing,
Peking University, Beijing, China

2025 IEEE International Symposium on Circuits and Systems (ISCAS) | 979-8-3503-5683-0/25/$31.00 ©2025 IEEE | DOI: 10.1109/ISCAS56072.2025.11043540

***Abstract*—The printing and dyeing industry requires periodic and high-resolution patterns to ensure seamless designs on large fabric sections and high-quality final products. However, current manual approaches to pattern creation are time-consuming and labor-intensive. Leveraging powerful image generative models, such as Latent Diffusion Models (LDMs), offers a promising alternative, but challenges persist in generating strictly periodic and high-resolution patterns due to the inherent randomness and high computational demands of LDMs. In this paper, we propose a novel text-driven framework for generating periodic and high-resolution patterns. We introduce a new training-free Splice-and-Focus Mechanism, which enhances the model by constraining latent features and modifying the attention mechanism to produce natural and strictly periodic patterns. Additionally, we present a ReLife Pipeline, which integrates super-resolution and guided image synthesis to enhance pattern resolution while eliminating artifacts and distortions. Experimental results demonstrate that our framework produces patterns of superior quality.**

I. Introduction

Printing and dyeing refer to the process of applying desired patterns onto fabric, such as cloth or silk, among which pattern design plays an esential role. Unlike ordinary designs, the patterns used in the printing and dyeing industry have more stringent requirements: **periodicity** and **high resolution**. First, since printing and dyeing are typically done on large sections of fabric, the pattern is expected to be composed of repeated splicing of periodic units, ensuring the overall design remains seamless and consistent across the entire surface. Second, the pattern needs to be of high definition, as clear and sharp designs are essential for producing high-quality final products.

Currently, printing and dyeing patterns are mainly created manually, whether through traditional hand-painting or digital drawing softwares. This results in significant time and labor costs for producing diverse patterns. Leveraging generative models to design patterns could dramatically reduce the costs. Moreover, generative models offer customization and versatility, allowing users to meet a wide range of design needs. Latent Diffusion Models (LDMs) [1] have made significant strides in the field of text-driven image generation, enabling the creation of high-quality visuals from prompts [1]–[3]. Meanwhile, existing customization technologies [4]–[7] such as Low-Rank Adaptation (LoRA) [7] enable the fine-tuning of LDMs tailored to specific styles based on curated datasets.

However, challenges remain in applying existing technologies to periodic and high-resolution pattern generation. While fine-tuning a LDM on a dataset of periodic images and incorporating "periodic" into the prompt as classifier-free guidance [8] could drive the model to produce roughly periodic patterns, the inherent randomness of LDM generation make it hard to achieve strict periodicity.

Meanwhile, the high computing power and graphics memory requirements of LDMs make it unrealistic to directly generate ultra-high-resolution patterns (*e.g.,* about 3960-pixel height for a US-Letter-sized pattern under 360 dpi). One solution is to alternatively apply super-resolution technologies [9]–[11] as post-processing. However, these methods mainly focus on faithfulness, assuming the correctness of the low-resolution inputs, which is rarely true in our case. Artifacts and pattern distortions are often produced during the diffusion process, which cannot be repaired and would even be magnified under high resolution by the aforementioned methods.

In this paper, we propose a novel text-driven periodic and high-resolution pattern generation framework. We first design a Splice-and-Focus Mechanism for LDMs that constrains latent features and modifies the attention mechanism, enabling the improved LDM to generate natural and periodic patterns. Notably, this mechanism requires no training, offering an elegant yet highly effective solution for generating patterns strictly adhere to scalable periodicity. In addition, we propose a new ReLife Pipeline. It combines super-resolution and guided image synthesis technologies, to simultaneously enhance image resolution and repair image elements to eliminate artifacts and distortions while maintaining the required periodicity, bringing the pattern details to life. Patterns processed through our framework demonstrate superior quality to meet the standards of printing and dyeing applications.

The rest of the paper is organized as follows. Sec. II and Sec. III introduce the proposed Splice-and-Focus Mechanism and ReLife Pipeline, respectively. Sec. IV presents experimental results, and Sec. V provides concluding remarks.

II. Splice-and-Focus Mechanism for Periodic Pattern Generation

To meet the key requirement for scalable and periodic pattern in design, printing and dyeing industries, this section

* Corresponding author. This work was supported in part by the National Natural Science Foundation of China under Grant 62471009, in part by CCF-Tencent Rhino-Bird Open Research Fund, and in part by the Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Technology).

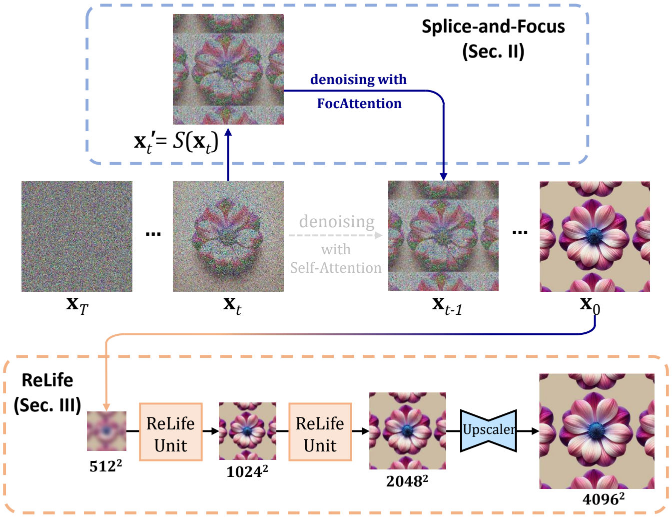

Fig. 1. Illustration of the proposed framework for periodic high-resolution image generation.

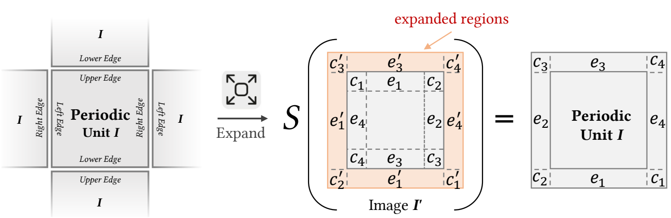

Fig. 2. The extensible periodic image maintains continuity after extension. The central valid area I remains unchanged, while the outer expanded regions are constrained, transforming the challenge from realizing continuity to keeping consistency.

focus on generating a primary periodic pattern. We would like the generated primary pattern to be repeated seamlessly over a large piece of canvas or fabric without any visible splicing. Although diffusion models have shown powerful generative abilities, they fail to synthesize strictly periodic even after fine-tuning on real periodic images and using specific key words in the prompt such as "periodic".

To address this challenge, we introduce a novel plug-and-play splice-and-focus mechanism of periodic constraints into the LDMs to ensure precise repetition. As shown in Fig. 2, the key idea is to expand the periodic image unit I outwards to a large image I'. Then the seamless alignment of the left and right edges, as well as the top and bottom edges of I can be naturally equivalent to the consistency between the expanded regions and their corresponding regions within I. In other words, if these corresponding regions of I' are the same (*e.g.*, $e_1 = e_1'$ and $c_1 = c_1'$), then I is strictly periodic. To realize this, we propose novel spatial constraints of **Latent Splice** and **FocAttention**, from the aspects of latent features and attentions respectively.

Latent Splice. LDMs project images to latent features [12] and generate in the latent space. To generate periodic images is to generate periodic latent features, which is equivalent to ensuring uniformity across different sections in the latent space

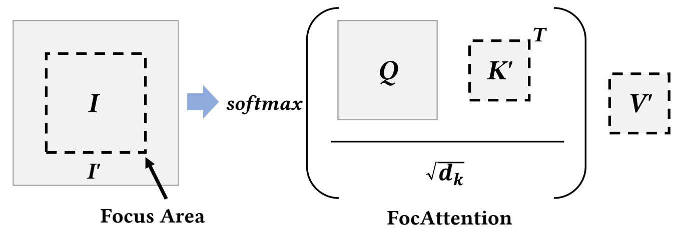

Fig. 3. Illustration of our FocAttention. K and V are clipped to focus attention on the valid region. For ease of understanding, we omit the process of reshaping tensors into vectors.

after our region expansion. Specifically, we begin by sampling a random latent feature $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$, whose resolution corresponds to I' after projecting back to the image space. Then, we define a substitution function $\mathcal{S}$, which operates in the latent space of the LDMs:

$$\mathbf{x}_t' = \mathcal{S}(\mathbf{x}_t), \tag{1}$$

where $\mathbf{x}_t$ is the noisy latent feature of LDMs at the time step t. As illustrated in Fig. 2, $\mathcal{S}$ splices $\mathbf{x}_t$ and directly copies the valid central part of $\mathbf{x}_t$ to the corresponding overlapped outer regions to ensure hard consistency. The resulting $\mathbf{x}_t'$ is then used in the next denoising step to obtain $\mathbf{x}_{t-1}$, until $t = 0$, when the fully denoised latent feature is generated [13]. As we will demonstrate in Sec. IV, this operation requires no training and does not compromise the efficiency of image generation.

FocAttention. FocAttention enforces a soft consistency by constraining the valid attention regions to ensure a smoother stitching effect in the generation process. Our experiments find that hard latent substitutions may result in unnatural seams. Since the valid portion of the final image is the central region, we aim to improve LDMs' standard self-attention mechanism:

$$\text{Self-Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V. \tag{2}$$

To focus on the valid region, we modify the K and V vectors to retain only the center valid area. These modified vectors are denoted as K' and V', respectively. This adjusted mechanism, we call it FocAttention, does not alter the attention calculation process, which is now expressed as:

$$\text{FocAttention}(Q, K, V) = \text{softmax}\left(\frac{QK'^T}{\sqrt{d_k}}\right)V'. \tag{3}$$

FocAttention modifies attention mechanism efficiently yet enables the model to focus on the valid region in the center, as illustrated in Fig. 3. In the experimental section of Sec. IV, we demonstrate that by incorporating Eq. (3) into the denoising process, we can generate natural periodic images.

Notably, our method requires no additional training, offering an elegant yet highly effective solution for constraining the generative process, thus enabling the creation of scalable periodic images.

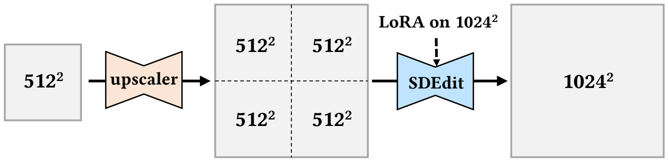

Fig. 4. One single unit of our ReLife Pipeline.

III. ReLife For Super-resolution and Refinement

This section super-resolves and refines the primary periodic pattern generated in Sec. II, to meet the resolution standards for printing and dyeing. Current super-resolution methods excel at restoring images [9]–[11] or enriching texture details [14]–[16], but they mainly focus on fidelity rather than generation. They tends to keep the artifacts and defects of the input's structures or textures. Thus, these methods cannot automatically repair artifacts or regenerate clear structures, such as replacing broken flower branches with realistic ones.

To address this unique challenge, we design a specialized ReLife Pipeline, as shown in Fig. 4, to simultaneously enhance image resolution and reconstruct image elements to meet the stringent quality requirements of printing and dyeing. Our pipeline is composed of two key stages: one dedicated to super-resolution, and the other to generation for defect correction. These stages are alternated to ensure optimal results.

Super-resolution Stage. We apply the pre-trained latent diffusion-based upscaler[1], which reliably increases the resolution of the original image by upsampling the latent code. Then, we segment the high-resolution image into overlapped sub-images to ensure each sub-image remains within processing capacity of LDMs in the next stage.

Generation Stage. We utilize SDEdit [17] for image-guided image synthesis. Specifically, SDEdit introduces noise into the super-resolved sub-image and then denoises and refines its content under the guidance of text prompts. In addition, we employ LoRA [7] to fine-tune the model with resolution-specific datasets, to ensure that the model can generate structures and details at the target resolution.

Finally, we stitch the segmented image. To ensure seamless and consistent stitching, we constrain the corresponding areas in subsequent sub-images based on the generated result of the previous sub-image. As shown in Fig. 5, at each step of the generation process, the latent features are replaced by a combination of the constrained part and the original part, ensuring consistency in the corresponding positions.

By alternating these steps, we are able to generate periodic patterns while preserving the artistic style and improving image resolution.

IV. Experimental Results

Experiment Setup. We incorporate our Splice-and-Focus Mechanism and ReLife Pipeline into the Stable Diffusion 1.5 model, and customize the model with LoRA on a curated dataset containing 100 printing and dyeing patterns. We

[1]https://huggingface.co/stabilityai/sd-x2-latent-upscaler

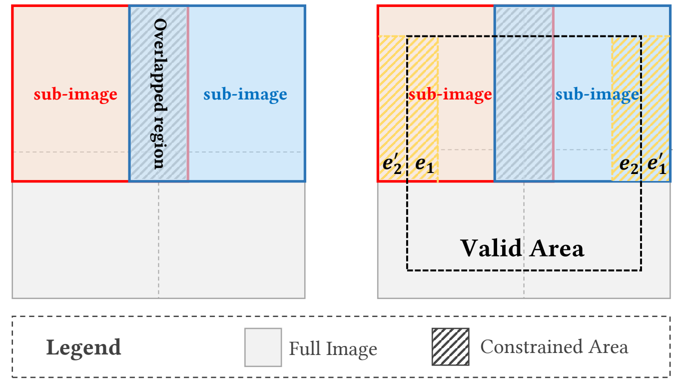

Fig. 5. Constraining the corresponding consistent parts of the sub-images for seamless stitching.

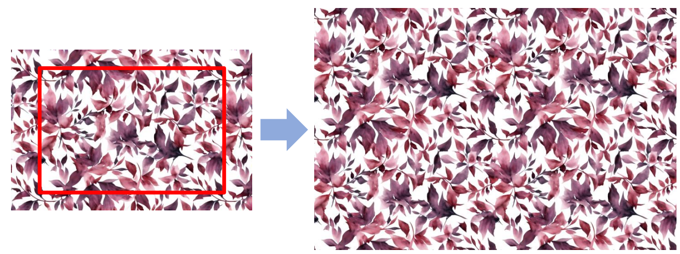

Fig. 6. In Splice-and-Focus mechanism, the periodic regions are spliced and extended. The central image is a seamless periodic pattern.

employed DDIM [18] for image generation and evaluation. Full experimental results are provided in our project page https://dongfengzy.github.io/SFaR/.

A. High-Resolution Periodic Pattern Generation

Figure 6 illustrates the stitching effect of the periodic pattern generated by our method. The results qualitatively demonstrate that our approach is capable of producing seamless periodic patterns while maintaining the intended style.

Figure 7 presents the result of our ReLife pipeline applied to an image, where we also compare it with one single upscaler. This process effectively corrects blurry defects and enhances the pattern resolution while preserving the original style.

B. Comparison with State-of-the-Art Methods

To demonstrate the advantages of our approach in generating high-quality periodic patterns, we compare our model with

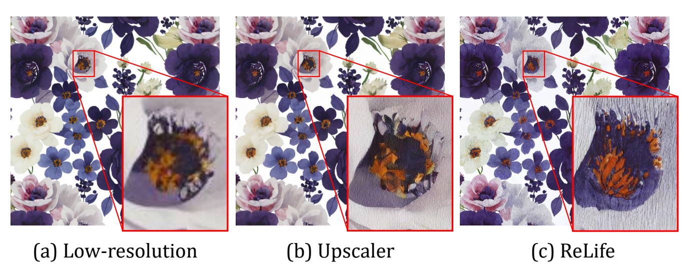

Fig. 7. Comparison direct super-resolution with our ReLife pipeline. Local regions are enlarged for better visual comparison. The red arrows indicate the confusing structures are corrected by our method.

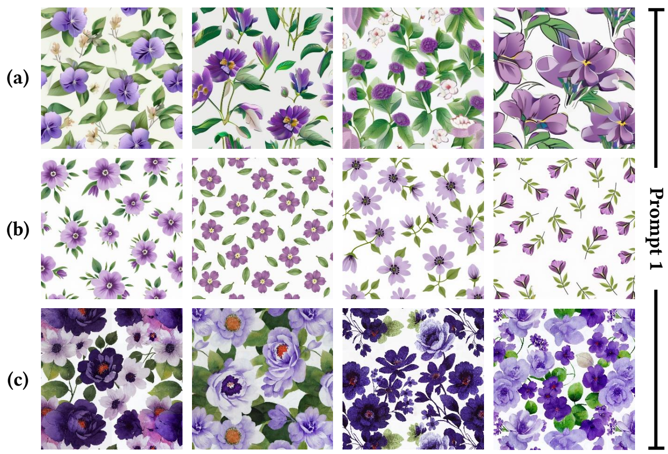

Prompt 1: *floral, white background, purple flowers, green leaves*

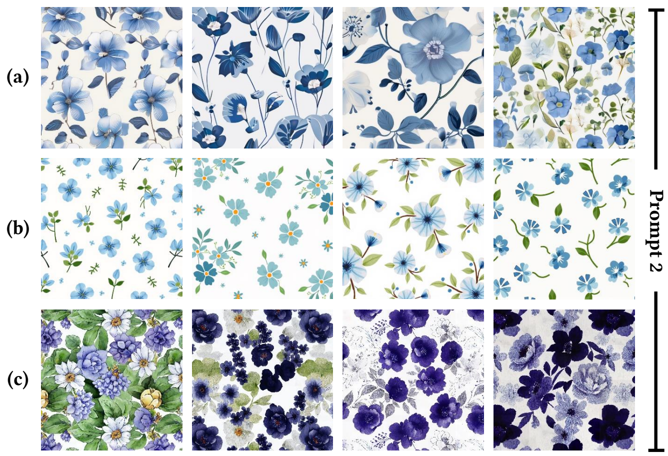

Prompt 2: *floral, white background, blue and white flowers, leaves*

Fig. 8. Qualitative comparisons with (a) SDXL-Turbo [19], (b) FLUX.1 and (c) the proposed method.

two state-of-the-art models, SDXL-Turbo [19] and FLUX.1[2]. To assess the generation of periodicity, we introduce the word "periodic" to the prompts for all models. For each set of prompts, we generate 100 images using our model, SDXL-Turbo and FLUX.1 for comparison.

Qualitative Analysis. Figure 8 presents qualitative results that illustrate the performance of each model under two prompt conditions. While all models generate patterns aligned with the text prompts, only our model consistently produces patterns that strictly adhere to periodicity, exhibiting superior quality and artistic value. Additionally, our model excels in the finer details of pattern generation.

Quantitative Analysis. For quantitative evaluation, we utilize the Fréchet Inception Distance (FID) [20] to measure the quality and naturalness of the generated images. We also conducted user study to evaluate the quality and periodicity of the results. Specifically, we asked users to select from the results from three methods the one with the highest visual quality and the one with the best periodicity. We report the

[2]https://github.com/black-forest-labs/flux

TABLE I
QUANTITATIVE ANALYSIS

Prompt Group	Method	FID↓	Quality↑	Periodicity↑
Prompt 1	SDXL-Turbo	266.00	0.03	0.02
	FLUX.1	259.87	0.20	0.01
	Ours	**164.13**	**0.77**	**0.97**
Prompt 2	SDXL-Turbo	229.39	0.02	0.01
	FLUX.1	255.78	0.17	0.01
	Ours	**142.65**	**0.81**	**0.98**

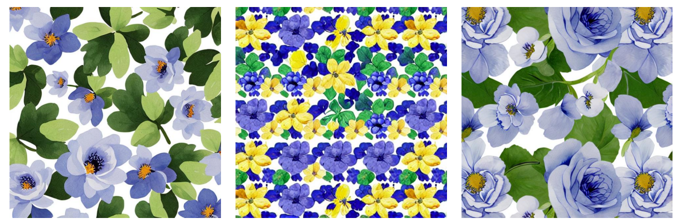

(a) w/o Latent Splice (b) w/o FocAttention (c) Full

Fig. 9. Ablation study of Latent Splice and FocAttention.

percentage of user selections in Table I over 20 users on 10 cases. It can be clearly seen that our model consistently designs higher-quality and strictly periodic images.

C. Ablation Study

We conducted ablation studies on different components of our method to evaluate their impact on the generated results. Figure 9 presents the images generated using identical prompt words and noise seeds, with the omission of Latent Splice and FocAttention, respectively.

- **Latent Splice**. Latent variable splicing is fundamental to ensuring the periodicity of the generated images. Without this component, the resulting images fail to maintain the desired periodicity.
- **FocAttention**. Removing FocAttention leads to a significant drop in image quality. Without this mechanism, the model performs rigid stitching without adjusting attention, causing over-repetitiveness and, in some cases, visible seams.
- **ReLife**. The impact of omitting ReLife is shown in Fig. 7. The ReLife pipeline enhances image quality. Its removal clearly results in lower image quality.

V. CONCLUSION

This paper presents a novel framework for generating periodic and high-resolution patterns tailored to the printing and dyeing industry. We propose a novel training-free Splice-and-Focus mechanism to ensure strict periodicity and seamless pattern generation when integrated with LDMs. Additionally, a ReLife Pipeline is proposed to enhance image resolution while effectively correcting artifacts. Experimental results demonstrate that our method outperforms existing techniques. By eliminating the need for manual design and training, our framework provides an efficient solution to the challenges of pattern generation.

REFERENCES

[1] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2022.

[2] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," *arXiv preprint arXiv:2204.06125*, 2022.

[3] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans *et al.*, "Photorealistic text-to-image diffusion models with deep language understanding," in *Advances in Neural Information Processing Systems*, 2022.

[4] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, "Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation," in *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2023.

[5] N. Kumari, B. Zhang, R. Zhang, E. Shechtman, and J.-Y. Zhu, "Multi-concept customization of text-to-image diffusion," in *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2023.

[6] L. Han, Y. Li, H. Zhang, P. Milanfar, D. Metaxas, and F. Yang, "SVDiff: Compact parameter space for diffusion fine-tuning," in *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2023.

[7] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *Proc. Int'l Conf. Learning Representations*, 2022.

[8] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *arXiv preprint arXiv:2207.12598*, 2022.

[9] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 295–307, 2015.

[10] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2017.

[11] B. Kawar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," in *Advances in Neural Information Processing Systems*, 2022.

[12] P. Esser, R. Rombach, and B. Ommer, "Taming transformers for high-resolution image synthesis," in *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2021.

[13] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, 2020.

[14] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 4, pp. 4713–4726, 2022.

[15] B. Xia, Y. Zhang, S. Wang, Y. Wang, X. Wu, Y. Tian, W. Yang, and L. Van Gool, "DiffIR: Efficient diffusion model for image restoration," in *Proc. Int'l Conf. Computer Vision*, 2023.

[16] J. Wang, Z. Yue, S. Zhou, K. C. Chan, and C. C. Loy, "Exploiting diffusion prior for real-world image super-resolution," *Int'l Journal of Computer Vision*, pp. 1–21, 2024.

[17] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon, "SDEdit: Guided image synthesis and editing with stochastic differential equations," in *Proc. Int'l Conf. Learning Representations*, 2022.

[18] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *Proc. Int'l Conf. Learning Representations*, 2021.

[19] A. Sauer, D. Lorenz, A. Blattmann, and R. Rombach, "Adversarial diffusion distillation," *arXiv preprint arXiv:2311.17042*, 2023.

[20] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local nash equilibrium," in *Advances in Neural Information Processing Systems*, 2017.