

GBSD: Generative Bokeh with Stage Diffusion

Jieren Deng*
University of Connecticut

Xin Zhou
Baidu Research USA

Hao Tian
Baidu Research USA

Zhihong Pan
Baidu Research USA

Derek Aguiar
University of Connecticut

Abstract

The bokeh effect is an artistic technique that blurs out-of-focus areas in a photograph and has gained interest due to recent developments in text-to-image synthesis and the ubiquity of smartphone cameras and photo sharing apps. Prior work on rendering bokeh effects have focused on post hoc image manipulation to produce similar blurring effects in existing photographs using classical computer graphics or neural rendering techniques, but have either depth discontinuity artifacts or are restricted to reproducing bokeh effects that are present in the training data. More recent diffusion based models can synthesize images with an artistic style, but either require the generation of high-dimensional masks, expensive fine-tuning, or affect global image characteristics. In this paper, we present GBSD, the first generative text-to-image model that synthesizes photorealistic images with a bokeh style. Motivated by how image synthesis occurs progressively in diffusion models, our approach combines latent diffusion models with a 2-stage conditioning algorithm to render bokeh effects on semantically defined objects. Since we can focus the effect on objects, this semantic bokeh effect is more versatile than classical rendering techniques. We evaluate GBSD both quantitatively and qualitatively and demonstrate its ability to be applied in both text-to-image and image-to-image settings.

1. Introduction

The bokeh effect refers to an artistic styling in photography that creates an out-of-focus blurring in areas of a photograph. Photographers have traditionally achieved the bokeh blurring effect by widening the lens aperture or through lens aberrations. The ubiquity of portable smartphone cameras and photo sharing applications has gener-

ated increased interest in image manipulation in general and bokeh synthesis specifically [18, 19]. Prior work is primarily concerned with *post hoc* image manipulation to produce similar blurring effects in existing photographs using classical computer graphics [24, 56, 63] or neural rendering techniques [33, 31, 60, 17, 10, 55, 59]. Neural rendering resolves the depth discontinuity artifacts present in classical techniques, but are generally restricted to reproducing bokeh effects that are present in the training data [33] or typical of classical techniques (e.g., a bokeh ball effect). Recent work achieves arbitrary blur sizes and shapes, but requires high-dimensional maps that are difficult to generate [33]. Moreover, all prior methods assume there exists an input image to be manipulated, that is, they are not fully generative models.

Recently, diffusion models [8] have demonstrated an ability to generate photorealistic images given a text prompt [40, 46, 62]. The artistic characteristics of synthesized images can be controlled through *post hoc* image manipulation or within the generative process. Image editing tasks like image inpainting [26, 47, 34, 66] are typically cast as image-to-image translations or require a user specified mask (and thus presuppose an image) to define a location in the image to edit. Synthesizing images with a specified artistic styling can be achieved by conditioning on class labels or text descriptions [46, 62, 40]. However, these control signals typically affect global image characteristics like artistic style and bokeh effects have not been achieved in diffusion models.

In this work, we present GBSD, the first generative text-to-image model capable of synthesizing photorealistic images with a bokeh style. Motivated by how image synthesis occurs progressively in diffusion models, that is, image layout, shape, and color are generated before enhancing details [15], our approach combines latent diffusion models with a 2-stage conditioning algorithm to render bokeh effects on semantically defined objects (Fig. 1). The two

*The work was done when the 1st author was an intern at Baidu Research.

stages apply different text conditioning to the latent diffusion network; the first (global layout) stage generates the structure of the image (e.g., the shape and color of objects) and the second (focus) stage simultaneously focuses detail generation and bokeh on different objects. Since we can focus the effect on objects, this *semantic bokeh* effect is more versatile than classical rendering techniques. Due to the simplicity of our conditioning algorithm, GBSD does not require the specification of a high-dimensional mask or expensive retraining.

Our work makes the following contributions:

- We present a new problem, semantic bokeh, whose goal is to apply the bokeh blur effect to semantically distinct objects in an image.
- We propose GBSD, a generative photorealistic image synthesis method based on latent stage diffusion. GBSD is the first diffusion model capable of synthesizing photorealistic bokeh stylized images and can be applied in both text-to-image (Fig. 1, left) and image-to-image (Fig. 1, right) settings.
- We evaluate our bokeh stage diffusion model both quantitatively and qualitatively by varying stage time and diffusion prompts for text-to-image and image-to-image tasks.

2. Related Work

Bokeh Synthesis. Methods for rendering bokeh include classical techniques that are typically inefficient (e.g., ray tracing [24, 56]) or require highly structure prior information like 3D scenes or depth maps (e.g., image space blurring or defocusing [63, 61, 14, 5, 4]). Taking advantage of advances in computer vision, subsequent bokeh rendering methods improved on classical techniques through the integration with image segmentation [48, 49], depth perception [6, 27, 35, 57, 65], or both [54]. More recently, neural bokeh synthesis techniques were developed to address the depth discontinuity artifacts around boundaries and lack of scalability of classical methods. Neural translation of in-focus to bokeh images using depth maps are accurate [31, 59], but generating high-quality depth maps is unrealistic in most scenarios. Methods that performed depth prediction or employed encoder-decoder architectures followed [55, 33, 60, 17, 10, 18, 19], but are, by construction, limited to reproducing bokeh effects similar to the data used to train them or require the generation of high-dimensional maps. Additionally, while some prior work employ generative models [10, 18, 19, 38], all prior bokeh synthesis methods aim for image-to-image translation, assuming an input image exists to be manipulated.

Deep Generative Models. Early architectures for synthesizing images from text relied on generative adversarial networks (GANs) [11] and variational autoencoders (VAE) [22]. GANs are capable of generating high-quality and high-resolution images, but are difficult to optimize [13, 29] and can drop regions of the data distribution [30]. In contrast with GANs, VAEs [22] can efficiently synthesize high-resolution images but typically generate lower quality images [7]. More recently, probabilistic diffusion models (DM) [15, 51, 53], which are based on an iterative denoising process [8], have demonstrated state-of-the-art results across a variety of applications including text-to-image generation [40], natural language generation [25], time series prediction [2], medical image [37], audio generation [23], adversarial machine learning [58] and privacy-preserving machine learning [9].

Text Driven Image Generation and Editing Using DM.

A primary application of diffusion models is image synthesis and manipulation based on conditioning text, which includes text-to-image [28, 41] and image-to-image generation. The denoising task can be conditioned by text prompts [21] in image space (GLIDE [32] and Imagen [46]) or in latent space, which includes DALL-E 2 [40], latent diffusion models (LDM) [42] and vector quantized diffusion [12]. To improve computational efficiency, it is common practice to train a diffusion model using low-resolution images or latent variables, which are then processed by super-resolution diffusion models [16] or latent-to-image decoders [50].

Variable Text Prompt Conditioning. Adjusting the text conditioning during the denoising process has been considered in the image manipulation context. Imagic uses a 3-step process to linearly interpolate between a target and optimized textual embeddings based on a reference image and text prompt [20]. However, Imagic requires the optimization of text embeddings based on pre-trained diffusion models, followed by diffusion model fine-tuning using the optimized text embeddings for each text prompt input. The eDiff-I method changes the text prompt after a fixed percentage of denoising steps [3], though the focus is on evaluating the strength of text conditioning and denoising efficacy under different noise levels, rather than exploring how prompt switching affects the photographic properties of the generated image. In contrast with prior work, we investigate how to split a continuous denoising process into two stages and leverage the prompt in stage 2 to simultaneously sharpen a target object while introducing a bokeh effect in others for both text-to-image generation and image-to-image generation. Further, our model does not require the generation of a mask or expensive fine-tuning.

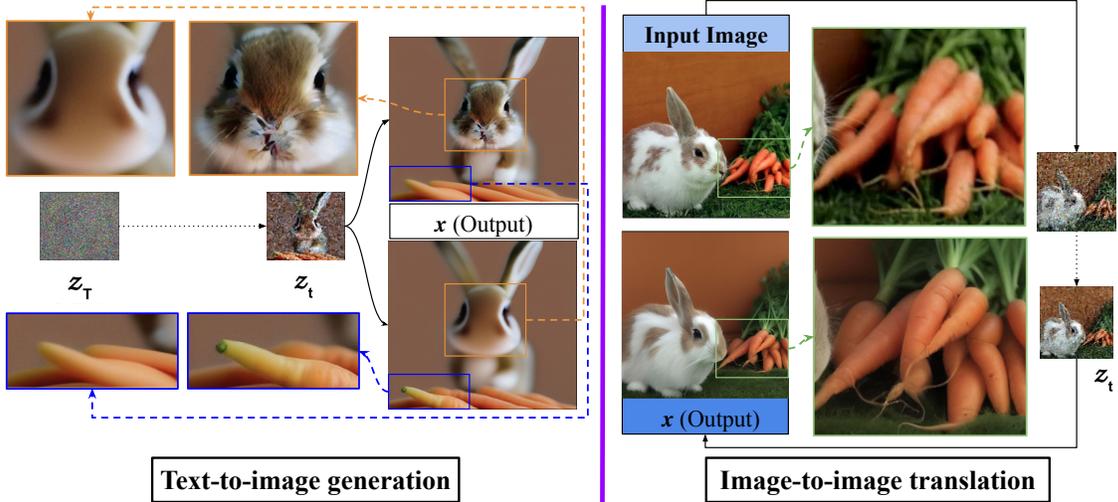


Figure 1. An illustration of stage diffusion for text-to-image and image-to-image generation with a bokeh style. A bokeh style image is generated by a two-stage semantic conditioning algorithm. The first stage (from z_T to z_t) generates the *global layout* of the image (e.g., shape and color) while stage two (from z_t to x) focuses detail and bokeh effects through semantic conditioning. In the text-to-image example (left), the stage 1 prompt was “A cute baby bunny standing on top of a pile of baby carrots under a spot light” and different prompts in stage 2 that focus either the carrots (bottom) or the bunny (top). In the image-to-image example (right), we use the prompt “A cute rabbit stands with carrots with green leaf” in stage 1 and “carrots with green leaf” in stage 2. The generated image demonstrates the previously blurry carrot coming into focus, revealing more clear and distinct textures, while creating a bokeh effect for the rabbit.

3. Methods

3.1. Diffusion Model Preliminaries

Probabilistic diffusion models estimate the data distribution $p(x)$ by denoising a normally distributed random variable with an input image $x \in \mathbb{R}^{H \times W \times 3}$ in RGB space. The denoising process is represented as the reverse of a length T Markovian diffusion process [42], with the best performing image generation models using a weighted variational lower bound on $p(x)$ [8, 15]. Let x_t be a noisy version of the input x and $\epsilon_\theta(x_t, t)$ be a denoising autoencoder with input x_t at step t . The diffusion process can be represented as denoising autoencoders $\{\epsilon_\theta(x_t, t)\}_{t=1}^T$, which are trained to predict a denoised variant of the input x_t . A simplified objective can be formulated as:

$$L_{DM} = \mathbb{E}_{x, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(x_t, t)\|_2^2]$$

Latent diffusion models (LDMs) leverage trained perceptual compression models ϵ and D , where $\epsilon(x)$ is an encoder for input x to latent space z and $D(z)$ decodes z from the latent space back to image space producing \bar{x} [42]. The model also uses a textual conditioning prompt y , which can be projected into an embedded representation through a parameterized domain-specific expert τ_θ . A new objective L using this latent space representation can be formulated as:

$$L_{LDM} = \mathbb{E}_{z, y, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(z_t, t, \tau_\theta(y))\|_2^2]$$

Most diffusion models denoise with consistent and continued conditioning. To achieve a bokeh effect in synthe-

sized images, we design a two-stage diffusion with a related, but distinct conditioning mechanism named *stage diffusion*.

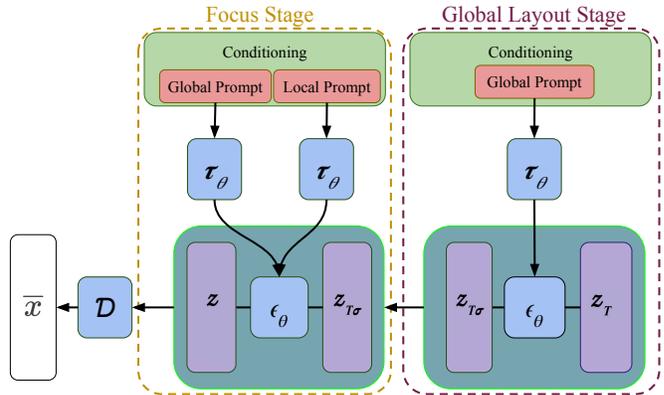


Figure 2. The generative bokeh with stage diffusion (GBSD) architecture. For brevity, we represent the diffusion process for each stage by a single representative denoising autoencoder.

3.2. Stage Diffusion

The goal of stage diffusion is to synthesize an image that simultaneously focuses on a target object while producing a bokeh effect on others in both text-to-image and image-to-image generative scenarios. Our stage diffusion method leverages LDMs [42], which provide a consistent text conditioning signal in each denoising autoencoder during image synthesis, and the progressive manner in which diffu-

sion models synthesize images (i.e., generating image layout, shape, and color before enhancing details [15]). We implement stage diffusion by decomposing the diffusion process into two distinct stages: a *global layout* stage and a *focus* stage (Fig. 2). The global layout stage generates the structure of the image (e.g., layout, shape, and color), whereas the focus stage outputs a final bokeh-styled image.

3.2.1 Global Layout Stage

In the global layout stage, we process a global prompt that completely describes the image to synthesize, y_{global} , through a domain-specific expert, τ_θ to obtain its corresponding textual embedding $\tau_\theta(y_{global})$ [39]. The model uses $\tau_\theta(y_{global})$ as a textual input for conditioning the diffusion process and employs a hyperparameter $\sigma \in (0, 1)$ to regulate the number of denoising steps during global layout stage. The denoising steps for a given global layout stage are set as the product of the total number of denoising steps T and σ . If we consider the global layout stage as a function of $g(\cdot)$ with inputs, $\tau_\theta(y_{global})$, z_T and σ , we represent the output of global layout stage $z_{\sigma \cdot T}$, as:

$$z_{\sigma \cdot T} = g(\tau_\theta(y_{global}), z_T, \sigma)$$

The output of the global layout stage is an intermediate image with a stable and consistent structure for synthesized objects.

3.2.2 Focus Stage

In order to simultaneously sharpen details on some objects while producing a bokeh effect in others, we pass a local prompt y_{local} to the text encoder [39], which produces the textual embedding, $\tau_\theta(y_{local})$; the local prompt y_{local} should be semantically related to the focused object. We linearly interpolate $\tau_\theta(y_{global})$ from the global layout stage and $\tau_\theta(y_{local})$ with a hyperparameter α as:

$$\bar{e} = (\tau_\theta(y_{local}) + \alpha \times \tau_\theta(y_{global})) / (1 + \alpha)$$

The resulting \bar{e} represents the textual conditioning in focus stage. The parameter α balances the global and local information in the focus stage, impacting the final generated image. The denoising diffusion objective L in focus stage is represented as:

$$L = \mathbb{E}_{z, y, \epsilon \sim \mathcal{N}(0, 1), t} [\|\epsilon - \epsilon_\theta(z_{\sigma \cdot T}, t, \bar{e})\|_2^2]$$

where $t = 1 \dots (\sigma \times T)$. Compared with other similar image synthesis or manipulation methods [20, 31, 59], stage diffusion does not require expensive fine-tuning or high-dimensional mask generation.

3.3. Image Generation

Stage diffusion is capable of both text-to-image and image-to-image synthesis. Text-to-image generation does not require the initial diffusion process step. Instead, an arbitrary text prompt y_{global} is input into the model and the initial random input is created by concatenating a sample from a Gaussian random variable with $\tau_\theta(y_{global})$ to create z_T . For image-to-image generation, we add diffusion noise based on the desired amount of image perturbation. In the global layout stage, a related or matched prompt is used as the global prompt to ensure the preservation of objects and prevent image distortion. This allows the model to generate outputs that are consistent with the given prompt, ensuring high-quality results.

4. Results

We implement our generative bokeh with stage diffusion (GBSD) method both qualitatively and quantitatively and in text-to-image and image-to-image scenarios. All experiments are conducted using NVIDIA V100 GPUs with an image output size of 512×512 , a batch size of 4 (i.e., the number of samples to generate for each prompt), and a fixed random seed of 42. The code repository leveraged the officially released version of LDM (Stable Diffusion v1-4) [42, 43]. For all experiments, the number of denoising diffusion implicit model (DDIM) sampling steps was set to 50, the number of timesteps was set to 1000, and the scale was set to 15 for all groups [52]. Additional results are provided in the appendix.

4.1. Evaluation Measures

Based on a recent evaluation of focus measure operators [36], we consider variance of Laplacian [44, 45] score and Brenner score [64] for quantitative evaluations.

4.1.1 Variance of Laplacian

The variance of Laplacian (VoL) is a measure that uses the variance of the image Laplacian for an evaluation of blur or focus [44, 45] and is calculated by

$$\sum_{(i, j) \in \Omega(x, y)} (\Delta I(i, j) - \overline{\Delta I})^2,$$

where $\Omega(x, y)$ is a neighborhood around pixel $I(i, j)$, $\Delta I(i, j)$ is the Laplacian at pixel $I(i, j)$, and $\overline{\Delta I}$ is the average value of the image Laplacian within the pixel neighborhood $\Omega(x, y)$. The Laplacian operator is effective in detecting blur due to its ability to measure regions with rapid intensity changes based on the second derivative of an image, similar to the Sobel and Scharr operators used for edge detection. This method assumes that an image with high variance contains both edge-like and non-edge-like features, indicating an in-focus image. Conversely, an image with low

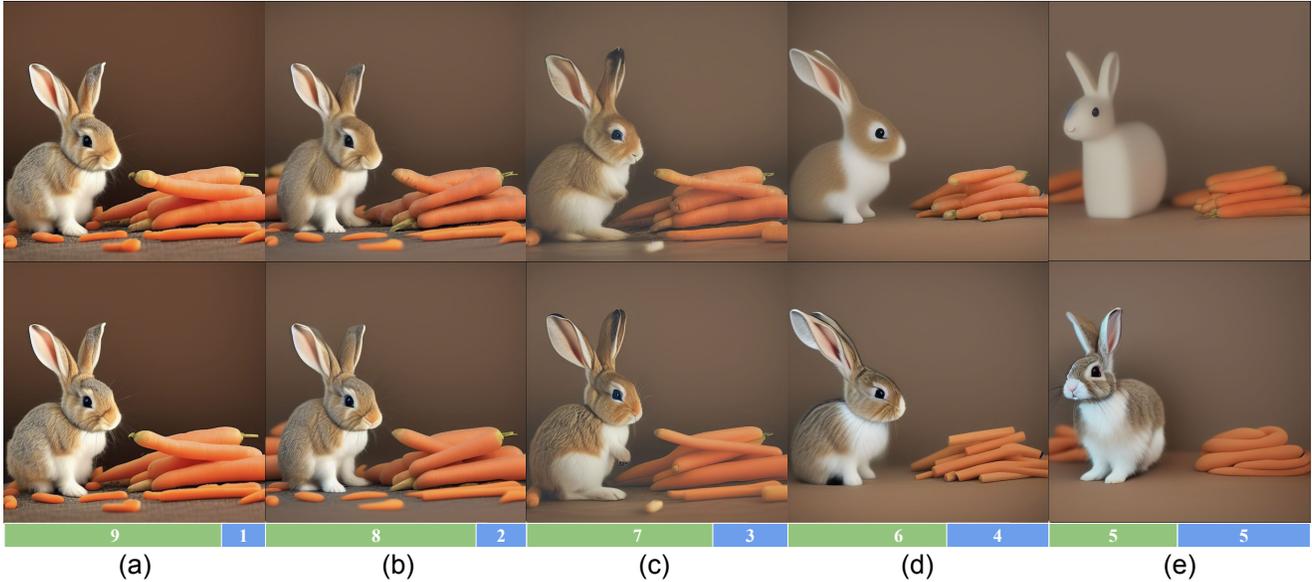


Figure 3. **Comparing image generation by adjusting σ .** The proportion of denoising steps allocated to the global layout stage was set to $\{0.9, 0.8, 0.7, 0.6, 0.5\}$ and the overall number of denoising steps and conditioning per stage was fixed. Both experiments (top and bottom rows) shared the same global prompt “a cute baby bunny standing on top of a pile of baby carrots under a spot light”. The top row was given a local prompt of “a pile of carrots” and the bottom row was given “a cute rabbit”. The proportion of denoising steps spent in the global layout (stage 1) and focus (stage 2) stages is given in the green and blue boxes, respectively. The *rabbit* features in the top experiment degenerated as stage 1 was shortened; similarly, the *carrots* features in the bottom experiment exhibited a similar trend.

variance has a small spread of responses, indicating a lack of edges and bokeh style. Therefore, as blur increases, the number of distinct edges in an image decreases, leading to lower variance and a higher degree of blurriness.

4.1.2 Brenner Score

The Brenner score [64] is a measure of the focus quality of a digital image. The Brenner score computes image textures at two different scales to characterize the amount of high and low resolution data contained within the image. For example, texture measurements may be calculated from the average of adjacent pixel pairs from the high resolution measurements compared with the average of adjacent pixels triplets from the low resolution measurements. A score that indicated the quality of focus is then generated as a function of the low- and high-resolution measurements.

4.2. The Effect of Adjusting σ

To investigate the impact of adjusting σ , which controls the proportion of the global layout (stage 1) and focus (stage 2) stages, we increased the length of the focus stage from 10% to 50% (Fig. 3). We performed two experiments with varied σ both containing the same global prompt: “a cute baby bunny standing on top of a pile of baby carrots under a spot light”. The first experiment (Fig. 3, top row) used a local prompt for the focus stage of “a pile of carrots”,

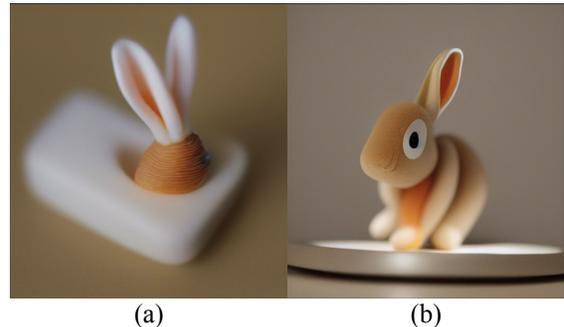


Figure 4. **Undesirable semantic mixing with a short global layout.** We used a global prompt that incorporated both the terms “carrots” and “rabbit”. In the subsequent focus stage, a local prompt that solely included the keyword “carrots” was added. When the global layout stage is too short (here, 20%), the “carrots” features merged with a “rabbit”-like object, producing an object that combined the features of both for two distinct random seeds (a) and (b).

whereas the second experiment (Fig. 3, bottom row) used the local prompt “a cute rabbit”.

Firstly, both “rabbit” and “carrots” objects are easily identifiable when the global layout stage encompasses at least 70% of the total denoising process (Fig. 3 (a-c)), indicating that a stage 1 length of 70%-90% of the denoising steps is sufficient to stabilize objects. When the global lay-

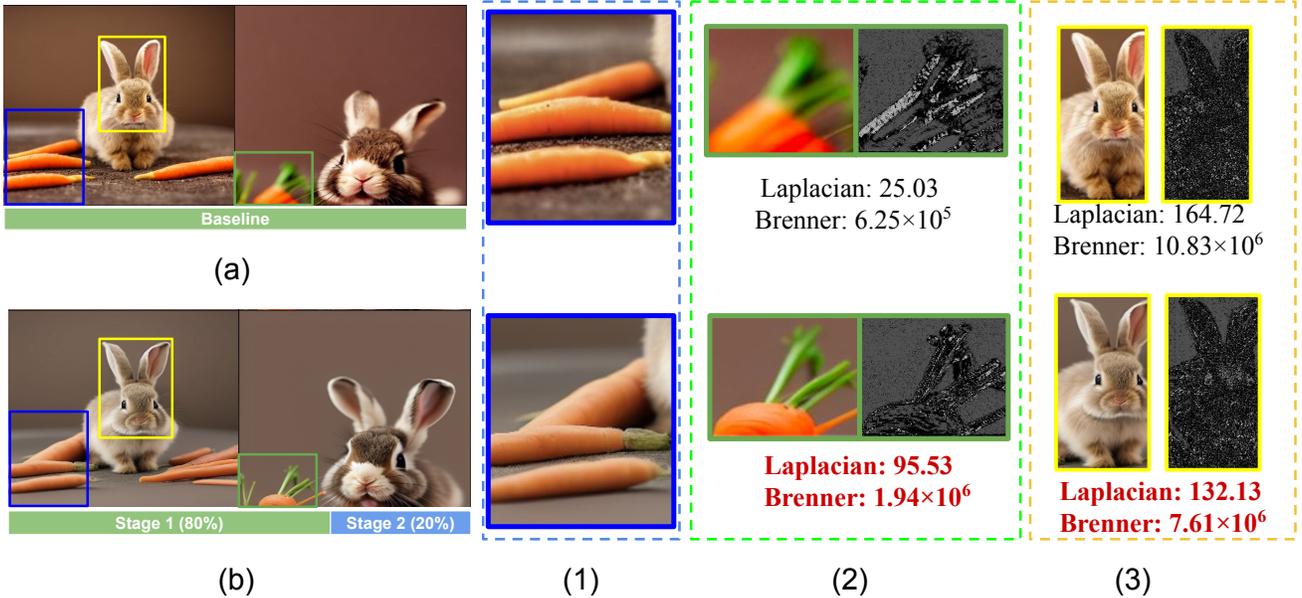


Figure 5. **Comparison between the baseline LDM and GBSD.** We used the text conditioning “a cute baby bunny standing on top of a pile of baby carrots under a spot light” both as the input prompt for LDM (a) and the global prompt for GBSD (b). In (b), a local prompt “a pile of baby carrots under a spot light” was used in stage 2 with 80% of denoising steps dedicated to stage 1 and 20% to stage 2. The highlighted image segments demonstrate that GBSD produces (1) stem details on the carrots, (2) a sharper image due to the focus on the carrots in stage 2, and (3) a bokeh effect on the bunny.

out stage was allocated 60% or less of the total denoising process, there existed insufficient evidence to confirm the presence of the feature that was the target of the bokeh effect, i.e., the “rabbit” in Figure 3, top row (d-e) and the “carrots” in Figure 3, bottom row (d-e). Furthermore, by comparing the different σ within each experiment, it can be observed that an increase in stage 2 leads to a degradation in the information that is present in the global prompt but not in the local prompt. When comparing the two experiments for each σ , we observed that the object position within each image are essentially identical since we use the same configuration to generate the images. However, the preservation of the object features are reversed with the first experiment retaining the features of “carrots” (Fig. 3, top row) and second experiment retaining the features of “rabbit” (Fig. 3, bottom row).

We also investigated the behavior of stage diffusion when the global layout stage was shortened to 20% of the denoising process (Fig. 4). Regardless of the local prompt, a short global layout stage produced objects that semantically intermixed the features of “carrots” and “rabbit” as a result of not being able to appropriately establish image layout before refining details.

4.3. Baseline Comparisons

We quantitatively compared the generated images between the baseline LDM [42] and GBSD (Fig. 5) using the VoL and Brenner measure. When a bokeh effect is desired,

the VoL and Brenner measure should be lower; conversely, when sharp details are desired, the VoL and Brenner measure should be larger. Using LDM and GBSD, we generated two images using two distinct random seeds (Fig. 5, (a) and (b), left and right). We used the prompt “a cute baby bunny standing on top of a pile of baby carrots under a spot light” as an input prompt for the baseline and global prompt for GBSD. We used a partial segment of the global prompt as the local prompt, “a pile of baby carrots under a spot light”, which encompassed 20% of denoising steps for GBSD in the focus stage.

Even when the random input seed is fixed across models, some segments of the image were too dissimilar to produce meaningful quantitative comparisons (Fig. 5 (1)). However, even these segments of carrots show enhanced detail in the GBSD synthesized image compared with LDM (e.g., the green stem). Next, we highlighted the highly similar carrot objects in the second experiment (Fig. 5, green boxes, right image) and computed their corresponding blur maps (Laplacian of the image normalized to grayscale). While the normalization of the blur map makes it difficult to compare the LDM and GBSD, the VoL, Brenner score, and qualitative comparisons demonstrate that the carrots are more in-focus for GBSD (Fig. 5 (2)) where our approach achieves 95.53 in Laplacian score and 1.94×10^6 in Brenner score which is $3.81 \times$ and $3.10 \times$ to baseline. When a bokeh effect is desired (here, in the bunny), the GBSD produces smaller (better) VoL and Brenner score values (Fig. 5 (3)).

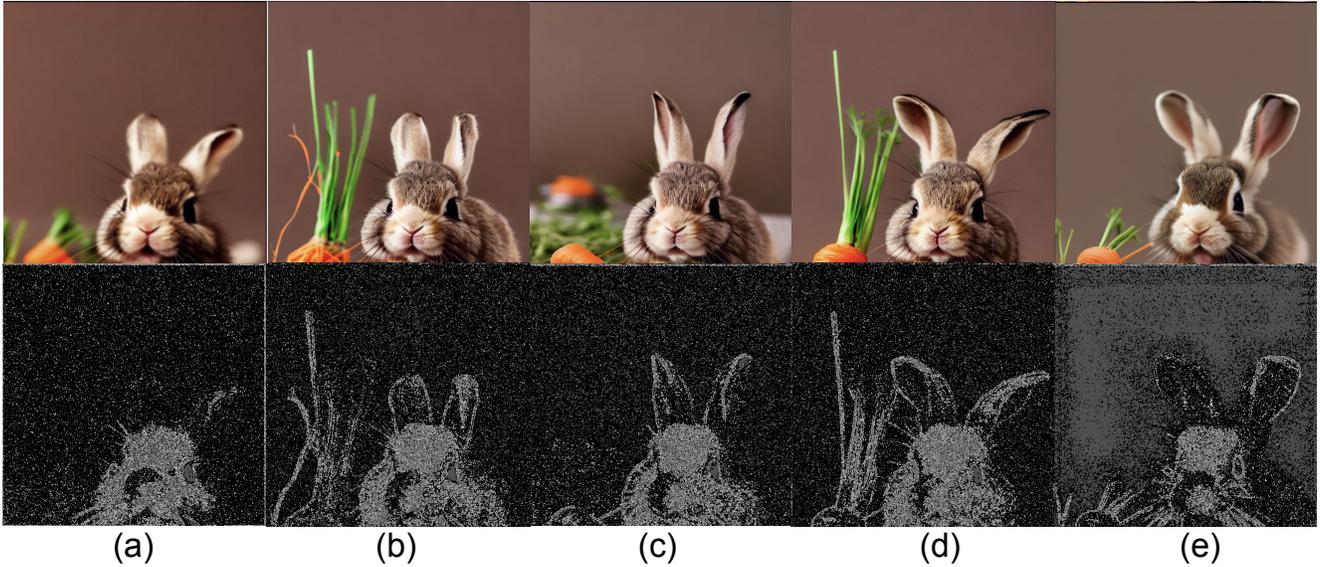


Figure 6. **Comparisons with prompt containing keyword “focus”.** We implemented variable focus using different prompts for LDM (a-d) and the same prompt as in Figure 5 for GBSD (e): (a) “a cute baby bunny standing on top of a pile of baby carrots under a spot light”, (b) “a cute baby bunny standing on top of a pile of baby carrots under a spot light, **rabbit is in focus, carrots out of focus**”, (c) “a cute baby bunny standing on top of a pile of baby carrots under a spot light, **carrots is in focus, rabbit out of focus**”, (d) “a cute baby bunny standing on top of a pile of baby carrots under a spot light, **rabbit and carrot in focus**” and (e) our output from Figure 5.

4.4. Comparison with “focus” Prompt

As discussed in the previous section, our proposed method demonstrated a significant score improvement in terms of Laplacian measures compared to the baseline approach when evaluated with the same prompt, that is, the LDM prompt and GBSD global prompts were identical. Furthermore, the focus prompt did not contain any additional information. We investigated whether adding text to the conditioning prompt could reproduce sharpening or blurring effects, we added “*in focus*” and “*out of focus*” phrases to the prompts (Fig. 6).

Using the same seed, we synthesized 4 images using LDM (Fig. 6 (a-d)). Image (a) was generated using the prompt “a cute baby bunny standing on top of a pile of baby carrots under a spot light”. To generate images (b-d), we added suffixes to change the image focus: (b) “**rabbit is in focus, carrots out of focus**”; (c) “**carrots is in focus, rabbit out of focus**”; (d) “**rabbit and carrot in focus**”. Finally, image (e) was generated using stage diffusion as described above (Fig. 5). Visual inspection of the images demonstrate undesirable results for images (a-c): (1) a bokeh effect in the carrots for image (a) when it was not desired; (2) an unnatural carrot structure being created in image (b) in response to the “*out of focus*” phrase; and (3) the opposite effect of what was desired in image (c), that is, the carrots are blurry then they should be in focus. Conversely, image (d) shows both rabbit and carrots in focus which, while desired based on the text prompt, does not achieve a bokeh effect. How-

ever, compared with the LDM baseline, the desired sharpening of the carrots and bokeh effect on the bunny is achieved in image (e) with GBSD (see also Fig. 5).

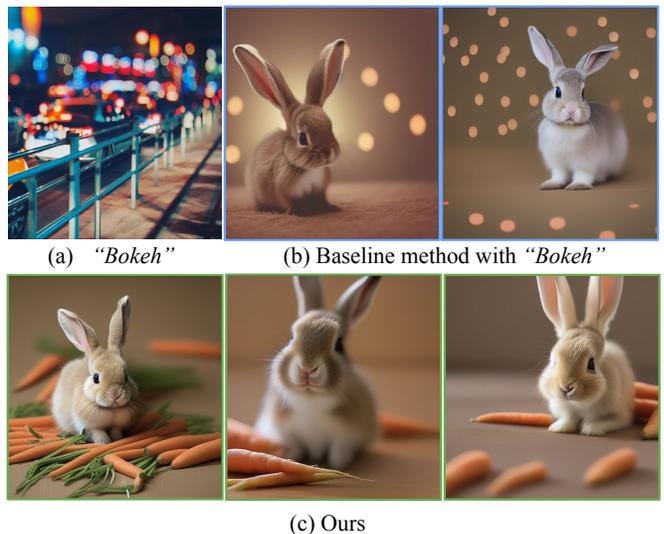


Figure 7. **Comparing the “bokeh effect” between LDM with “bokeh” in the prompt and GBSD.** (a) is an example of “bokeh” from a photography resource [1]. (b) is the output of LDM where the “bokeh” keyword is added to the text prompt. (c) is the output of GBSD with a focus stage prompt highlighting the bunny (left and right images) and carrots (middle image).

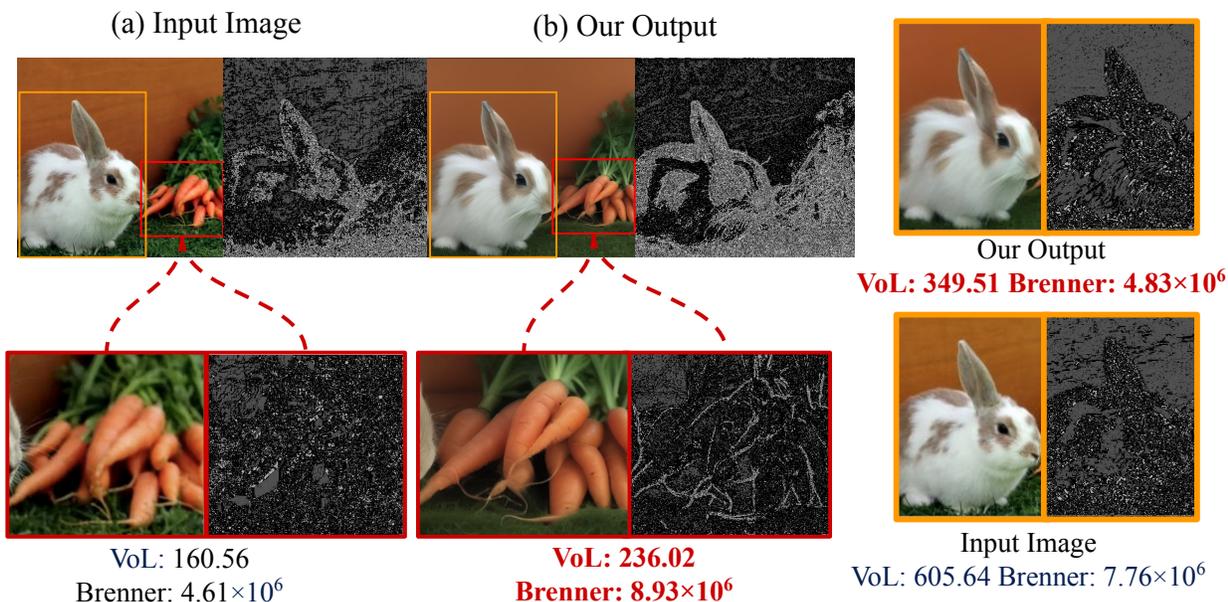


Figure 8. **Evaluating image-to-image generation.** The global prompt was “A cute rabbit stands with carrots with green leaf” and the local prompt is “carrots with green leaf” for (b). Our outputs achieve up to 73% improvement in VoL and 93% improvement in Brenner compared to the input image.

4.5. Comparisons with “bokeh” Prompt

Next, we investigated whether adding a “bokeh effect” prompt to the baseline would produce the desired blurring of semantically distinct objects (Fig. 7). Typical examples of the bokeh effect includes urban settings with a blurring ball out-of-focus effect on lighting in the background (Fig. 7 (a)). Including the keywords “bokeh effect” in the prompt of LDM interestingly removed carrot objects but generated a similar blurring effect with carrot-colored lighting (Fig. 7 (b)). The bokeh effect produces by LDM is also not realistic, as there are no physical light sources located in the background. Since the presence of points of light is a common feature in publicly available bokeh images, the bokeh artifact produced by LDM is likely due to its presence in the training data. In contrast, GBSD offers the advantage of producing a more realistic bokeh effect on either carrot or bunny objects instead of simply replicating a typical bokeh effect from the training data (Fig. 7 (c)).

4.6. Focus Shift in Image-to-Image Generation

Text-to-image generation involves generating an image from a text prompt, whereas image-to-image generation requires two types of conditioning: a text prompt and an input image. One challenge of image-to-image generation is identifying an appropriate level of noise to incorporate into the input image. The noise must be capable of modifying image features without causing significant deviations from the input image. An image was selected from an online source to use as the input (Fig. 8 (a)). To add focus to

the carrots, we set a global prompt “A cute rabbit stands with carrots with green leaf” in the global layout stage and local prompt of “carrots with green leaf” in focus stage. Overall, GBSD sharpens the detail of the carrots achieving a VoL of 236.02 and Brenner score of 8.93×10^6 , which are $1.73 \times$ and $1.93 \times$ larger than the input image, respectively (Fig. 8 (b), left). Further, the stage diffusion algorithm also achieves a bokeh effect on the rabbit, with a smaller VoL and Brenner score when compared with the input image (Fig. 8 (b), right).

5. Conclusions

In this paper, we presented GBSD, the first generative text-to-image model that synthesizes photorealistic images with a bokeh style. The approach combines latent diffusion models with a 2-stage conditioning algorithm to render blurring effects. Unlike prior bokeh methods, GBSD is able to produce a semantic bokeh effect, where semantically distinct objects are blurred based on the 2-stage text conditioning procedure. We evaluated GBSD both quantitatively and qualitatively and demonstrated its ability to be applied in both text-to-image and image-to-image settings. In sum, we believe that GBSD and other generative models of photorealistic images with artistic stylings can provide a valuable content generation resource to AI-assisted industries reliant on image synthesis.

References

- [1] Bokeh photography: The ultimate tutorial, Apr 2020. [7](#)
- [2] Juan Lopez Alcaraz and Nils Strodthoff. Diffusion-based time series imputation and forecasting with structured state space models. *Transactions on Machine Learning Research*, 2023. [2](#)
- [3] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. ediff-i: Text-to-image diffusion models with an ensemble of expert denoisers, 2022. [2](#)
- [4] Jonathan T Barron, Andrew Adams, YiChang Shih, and Carlos Hernández. Fast bilateral-space stereo for synthetic defocus. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4466–4474, 2015. [2](#)
- [5] Marcelo Bertalmio, Pere Fort, and Daniel Sanchez-Crespo. Real-time, accurate depth of field using anisotropic diffusion and programmable graphics cards. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004.*, pages 767–773. IEEE, 2004. [2](#)
- [6] Benjamin Busam, Matthieu Hog, Steven McDonagh, and Gregory Slabaugh. Sterefo: Efficient image refocusing with stereo vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#)
- [7] Rewon Child. Very deep vaes generalize autoregressive models and can outperform them on images, 2020. [2](#)
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc., 2021. [1, 2, 3](#)
- [9] Tim Dockhorn, Tianshi Cao, Arash Vahdat, and Karsten Kreis. Differentially Private Diffusion Models. *arXiv:2210.09929*, 2022. [2](#)
- [10] Saikat Dutta, Sourya Dipta Das, Nisarg A Shah, and Anil Kumar Tiwari. Stacked deep multi-scale hierarchical network for fast bokeh effect rendering from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2398–2407, 2021. [1, 2](#)
- [11] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. [2](#)
- [12] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis, 2021. [2](#)
- [13] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 5769–5779, Red Hook, NY, USA, 2017. Curran Associates Inc. [2](#)
- [14] Thomas Hach, Johannes Steurer, Arvind Amruth, and Artur Pappenheim. Cinematic bokeh rendering for real scenes. In *Proceedings of the 12th European Conference on Visual Media Production*, pages 1–10, 2015. [2](#)
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA, 2020. Curran Associates Inc. [1, 2, 3, 4](#)
- [16] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *arXiv preprint arXiv:2106.15282*, 2021. [2](#)
- [17] Andrey Ignatov, Jagruti Patel, and Radu Timofte. Rendering natural camera bokeh effect with deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 418–419, 2020. [1, 2](#)
- [18] Andrey Ignatov, Jagruti Patel, Radu Timofte, Bolun Zheng, Xin Ye, Li Huang, Xiang Tian, Saikat Dutta, Kuldeep Purohit, Praveen Kandula, et al. Aim 2019 challenge on bokeh effect synthesis: Methods and results. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3591–3598. IEEE, 2019. [1, 2](#)
- [19] Andrey Ignatov, Radu Timofte, Ming Qian, Congyu Qiao, Jiamin Lin, Zhenyu Guo, Chenghua Li, Cong Leng, Jian Cheng, Juewen Peng, et al. Aim 2020 challenge on rendering realistic bokeh. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 213–228. Springer, 2020. [1, 2](#)
- [20] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276*, 2022. [2, 4](#)
- [21] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2426–2435, June 2022. [2](#)
- [22] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*, 2014. [2](#)
- [23] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2021. [2](#)
- [24] Sungkil Lee, Elmar Eisemann, and Hans-Peter Seidel. Real-time lens blur effects and focus control. *ACM Transactions on Graphics (TOG)*, 29(4):1–7, 2010. [1, 2](#)
- [25] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori Hashimoto. Diffusion-LM improves controllable text generation. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. [2](#)
- [26] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceed-*

- ings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. 1
- [27] Xianrui Luo, Juewen Peng, Ke Xian, Zijin Wu, and Zhiguo Cao. Bokeh Rendering from Defocus Estimation. In *Computer Vision – ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III*, pages 245–261, Berlin, Heidelberg, Aug. 2020. Springer-Verlag. 2
- [28] Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Generating images from captions with attention. *arXiv preprint arXiv:1511.02793*, 2015. 2
- [29] Lars M. Mescheder. On the convergence properties of gan training. *ArXiv*, abs/1801.04406, 2018. 2
- [30] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. In *International Conference on Learning Representations*, 2017. 2
- [31] Oliver Nalbach, Elena Arabadzhyska, Dushyant Mehta, H-P Seidel, and Tobias Ritschel. Deep shading: convolutional neural networks for screen space shading. In *Computer graphics forum*, volume 36, pages 65–78. Wiley Online Library, 2017. 1, 2, 4
- [32] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 16784–16804. PMLR, 17–23 Jul 2022. 2
- [33] Juewen Peng, Zhiguo Cao, Xianrui Luo, Hao Lu, Ke Xian, and Jianming Zhang. BokehMe: When Neural Rendering Meets Classical Rendering. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16262–16271, New Orleans, LA, USA, June 2022. IEEE. 1, 2
- [34] Jialun Peng, Dong Liu, Songcen Xu, and Houqiang Li. Generating diverse structure for image inpainting with hierarchical vq-vae. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10775–10784, 2021. 1
- [35] Juewen Peng, Xianrui Luo, Ke Xian, and Zhiguo Cao. Interactive portrait bokeh rendering system. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 2923–2927. IEEE, 2021. 2
- [36] Said Pertuz, Domenec Puig, and Miguel Ángel García. Analysis of focus measure operators for shape-from-focus. *Pattern Recognit.*, 46:1415–1432, 2013. 4
- [37] Walter H. L. Pinaya, Petru-Daniel Tudosiu, Jessica Dafflon, Pedro F. Da Costa, Virginia Fernandez, Parashkev Nachev, Sebastien Ourselin, and M. Jorge Cardoso. Brain imaging generation with latent diffusion models. In Anirban Mukhopadhyay, Ilkay Oksuz, Sandy Engelhardt, Dajiang Zhu, and Yixuan Yuan, editors, *Deep Generative Models*, pages 117–126, Cham, 2022. Springer Nature Switzerland. 2
- [38] Ming Qian, Congyu Qiao, Jiamin Lin, Zhenyu Guo, Chenghua Li, Cong Leng, and Jian Cheng. BGGAN: Bokeh-Glass Generative Adversarial Network for Rendering Realistic Bokeh. In Adrien Bartoli and Andrea Fusiello, editors, *Computer Vision – ECCV 2020 Workshops*, volume 12537, pages 229–244. Springer International Publishing, Cham, 2020. 2
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 18–24 Jul 2021. 4
- [40] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv:2204.06125*, 2022. 1, 2
- [41] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International conference on machine learning*, pages 1060–1069. PMLR, 2016. 2
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022. 2, 3, 4, 6
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. Stable Diffusion v1-4 Model Card. <https://huggingface.co/CompVis/stable-diffusion-v1-4>, 2022. [Online; accessed 1-January-2023]. 4
- [44] Adrian Rosebrock. Blur detection with opencv, 2015. 4
- [45] Sagar. Laplacian and its use in blur detection, 2020. 4
- [46] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 1, 2
- [47] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1
- [48] Xiaoyong Shen, Aaron Hertzmann, Jiaya Jia, Sylvain Paris, Brian Price, Eli Shechtman, and Ian Sachs. Automatic portrait segmentation for image stylization. In *Computer Graphics Forum*, volume 35, pages 93–102. Wiley Online Library, 2016. 2
- [49] Xiaoyong Shen, Xin Tao, Hongyun Gao, Chao Zhou, and Jiaya Jia. Deep automatic portrait matting. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 92–107. Springer, 2016. 2

- [50] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-decoding models for few-shot conditional generation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 12533–12548. Curran Associates, Inc., 2021. [2](#)
- [51] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. [2](#)
- [52] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. [4](#)
- [53] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. [2](#)
- [54] Neal Wadhwa, Rahul Garg, David E Jacobs, Bryan E Feldman, Nori Kanazawa, Robert Carroll, Yair Movshovitz-Attias, Jonathan T Barron, Yael Pritch, and Marc Levoy. Synthetic depth-of-field with a single-camera mobile phone. *ACM Transactions on Graphics (ToG)*, 37(4):1–13, 2018. [2](#)
- [55] Lijun Wang, Xiaohui Shen, Jianming Zhang, Oliver Wang, Zhe Lin, Chih-Yao Hsieh, Sarah Kong, and Huchuan Lu. Deeplens: shallow depth of field from a single image. *arXiv preprint arXiv:1810.08100*, 2018. [1](#), [2](#)
- [56] Jiaye Wu, Changwen Zheng, Xiaohui Hu, and Fanjiang Xu. Rendering realistic spectral bokeh due to lens stops and aberrations. *The Visual Computer*, 29:41–52, 2013. [1](#), [2](#)
- [57] Ke Xian, Juewen Peng, Chao Zhang, Hao Lu, and Zhiguo Cao. Ranking-based salient object detection and depth prediction for shallow depth-of-field. *Sensors*, 21(5):1815, 2021. [2](#)
- [58] Chaowei Xiao, Zhongzhu Chen, Kun Jin, Jiong Xiao Wang, Weili Nie, Mingyan Liu, Anima Anandkumar, Bo Li, and Dawn Song. Densepure: Understanding diffusion models for adversarial robustness. In *The Eleventh International Conference on Learning Representations*, 2023. [2](#)
- [59] Lei Xiao, Anton Kaplanyan, Alexander Fix, Matt Chapman, and Douglas Lanman. Deepfocus: Learned image synthesis for computational display. In *ACM SIGGRAPH 2018 Talks*, pages 1–2. 2018. [1](#), [2](#), [4](#)
- [60] Xiangyu Xu, Deqing Sun, Sifei Liu, Wenqi Ren, Yu-Jin Zhang, Ming-Hsuan Yang, and Jian Sun. Rendering portraits from monocular camera and beyond. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–50, 2018. [1](#), [2](#)
- [61] Yang Yang, Haiting Lin, Zhan Yu, Sylvain Paris, and Jingyi Yu. Virtual dslr: High quality dynamic depth-of-field synthesis on mobile platforms. *Electronic Imaging*, 28:1–9, 2016. [2](#)
- [62] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022. [1](#)
- [63] Xuan Yu, Rui Wang, and Jingyi Yu. Real-time depth of field rendering via dynamic light field generation and filtering. In *Computer Graphics Forum*, volume 29, pages 2099–2107. Wiley Online Library, 2010. [1](#), [2](#)
- [64] Michael Zahniser. Method for assessing image focus quality, 2010. US Patent 8014583B2. [4](#), [5](#)
- [65] Xuaner Zhang, Kevin Matzen, Vivien Nguyen, Dillon Yao, You Zhang, and Ren Ng. Synthetic defocus and look-ahead autofocus for casual videography. *arXiv preprint arXiv:1905.06326*, 2019. [2](#)
- [66] Lei Zhao, Qihang Mo, Sihuan Lin, Zhizhong Wang, Zhiwen Zuo, Haibo Chen, Wei Xing, and Dongming Lu. Uctgan: Diverse image inpainting based on unsupervised cross-space translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5741–5750, 2020. [1](#)