

# BokehOrNot: Transforming Bokeh Effect with Image Transformer and Lens Metadata Embedding

Zhihao Yang Wenyi Lian Siyuan Lai  
Uppsala University, Sweden

{zhihao.yang.8094,wenyi.lian.7322,siyuan.lai.3205}@student.uu.se

## Abstract

Bokeh effect is an optical phenomenon that offers a pleasant visual experience, typically generated by high-end cameras with wide aperture lenses. The task of bokeh effect transformation aims to produce a desired effect in one set of lenses and apertures based on another combination. Current models are limited in their ability to render a specific set of bokeh effects, primarily transformations from sharp to blur. In this paper, we propose a novel universal method for embedding lens metadata into the model and introducing a loss calculation method using alpha masks from the newly released Bokeh Effect Transformation Dataset (BETD) [3]. Based on the above techniques, we propose the BokehOrNot model, which is capable of producing both blur-to-sharp and sharp-to-blur bokeh effect with various combinations of lenses and aperture sizes. Our proposed model outperforms current leading bokeh rendering and image restoration models and renders visually natural bokeh effects. Our code is available at: <https://github.com/indicator0/bokehornot>.

## 1. Introduction

The bokeh effect refers to the aesthetic quality of the out-of-focus areas in a photograph. This effect is achieved by using a shallow depth of field in photography, which creates a blurred background while keeping the subject in focus. The bokeh effect requires a wide aperture lens, which is typically only available on high-end cameras with sophisticated optics and usually large sensors. Fig. 1a and Fig. 1b illustrate the impact of aperture size on the bokeh effect, whereby Fig. 1a was captured using a narrow aperture (f/16) and Fig. 1b was taken with a wide aperture (f/1.8). As these images show, a wide aperture produces a shallow depth-of-field with a blurry background, while a narrow aperture creates a deeper depth-of-field with a clearer background. This comparison underscores the significance of the aperture size in achieving desired bokeh effects. The bokeh effect is not

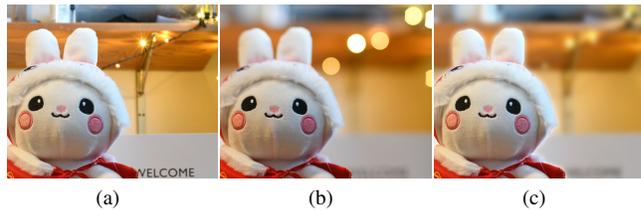


Figure 1. The effect of different aperture sizes on the bokeh effect and Gaussian blur. (a) With narrow aperture (f/16). (b) With wide aperture (f/1.8). (c) Artificial Gaussian blur.

identical to the blurring effect. For example, Fig. 1c demonstrates an artificial blurred image using Gaussian blur with the same blurring strength as Fig. 1b. However, the bokeh effect still can be simulated computationally using algorithms, making it possible to produce this effect on devices such as portable cameras or mobile phones without accessories.

To learn the bokeh style directly from high-quality DSLR camera photos, a large-scale Everything is Better with Bokeh! (EBB!) dataset was proposed in [9], where the authors also built a deep learning-based method (namely PyNET) capable of transforming wide to shallow depth-of-field images automatically. However, this method relies on an accurate depth map estimation of an image, which is time-consuming and computationally costly. To address it, several works [6, 18, 31] propose to render bokeh effects in an end-to-end manner without relying on depth maps. By leveraging fancy network structures (e.g., Hierarchical Network [6] and Transformer [18]), they can achieve fast and realistic results on the EBB! dataset. However, their approaches have not yet gained widespread adoption due to the dataset limitation, where the bokeh effect transformation is fixed (on both lens models and aperture sizes) and thus cannot fit the requirements of real-world applications.

More specifically, the EBB! dataset has collected 5,000 pairs of aligned images captured by a Canon 70D DSLR camera. For each training pair, the source image is captured using a narrow aperture (f/16) and the target is cap-

tured using a wide aperture (f/1.8). This means the only bokeh transformation is to render blurry backgrounds from sharp photos. Based on this dataset, many solutions were presented to address the bokeh effect rendering problem [6, 10, 20].

However, as we mentioned above, a major limitation of this dataset is the fixed lens and aperture parameters, meaning that there is little variation in the types of scenes and objects captured.

To address it, NTIRE 2023 Bokeh Effect Transformation challenge [3] proposes a new dataset, Bokeh Effect Transformation Dataset (BETD) [3], which contains image pairs with an artificially synthesized foreground portrait, a background captured by different lenses (Sony50mm and Canon50mm) and aperture sizes. Moreover, this dataset also provides foreground masks and metadata which contains lens parameters and the extra disparity values measuring how much two images of the same scene differ. One potential advantage of this dataset is that it provides a wider range of blur strengths than the EBB! Dataset, which is essential for real-world scenarios. Additionally, the inclusion of foreground masks can also be helpful for advanced training losses or evaluations.

Our work is built upon BETD, with full consideration of the metadata and foreground masks in a comprehensive manner. Specifically, we propose to embed the metadata in image inputs and assemble them into the network using affine transformations. Further, we leverage the alpha masks provided in the dataset and propose a novel loss-calculating method.

Our primary contributions are summarized as follows:

- We propose a novel model, named *BokehOrNot*, based on Restormer [29]. Our model is capable of transforming and rendering the natural bokeh effect among multiple aperture sizes and lenses without compromising the sharpness of the foreground and the global image quality.
- We build a Lens Embedding Module (LEM) and a Dual-Input Transformer Block (DITB) which integrate lens metadata to original image inputs with sinusoidal embedding and affine transforms, allowing the model to learn different bokeh styles and strengths efficiently.
- Alpha-masked loss that calculates loss excluding foreground areas. This helps the model to learn bokeh effect transformation more accurately without the influence of the foreground, especially with the cropping training technique.

## 2. Related Work

### 2.1. Bokeh Rendering

Bokeh rendering is growing to be an important technique in computational photography, which usually takes photos with out-of-focus areas to highlight regions of interest. Hach *et al.* [8] introduced a Point-Spread-Function

based approach to render different styles of bokeh. However, this method requires high-quality lens point-spread-function measurements and specific cameras. To render bokeh effect with minimal requirements, Ignatov *et al.* [9] proposed a large-scale *Everything is Better with Bokeh!* (EBB!) dataset which contains 5,000 aligned shallow / wide depth-of-field image pairs collected in the wild with the Canon 70D DSLR camera and 50mm f/1.8 fast lens. Subsequently, abundant works were proposed to tackle the bokeh rendering problem based on the EBB! dataset [6, 11, 18, 21, 31]. Specifically, PyNET [9] is the first method proposed along with the EBB! dataset, which utilizes a multi-scale CNN architecture and meanwhile adds the pre-computed depth map to the input. Luo *et al.* [16] proposed a method based on defocus map estimation with radiance and upsampling to learn shallow depth-of-field from a single bokeh-free image. BGGAN [21] combines a Glass-Net with Generative Adversarial Network (GAN) to improve the visual effect of synthetic bokeh rendering. To get rid of the requirement of depth estimation, Dutta *et al.* [6] created a Stacked DMSHN which does not depend on the depth map estimation and incorporates multiple encoder-decoder structures to gradually refine the output, resulting in an efficient bokeh rendering model. Moreover, Hariharan *et al.* [18] also proposed a depth map free transformer-based approach that is built on the PyNET but achieves better results in retaining the details of the foreground, namely BRViT. Moreover, BokehMe [20] and AMPN [7] were proposed to generate photo-realistic bokeh effects. In particular, BokehMe utilizes imperfect disparity maps and can adjust blur sizes, focal planes, and aperture shapes. AMPN employs a mask prediction module which enables users to control their bokeh rendering with a user-guided mask.

### 2.2. Vision Transformer and Image Restoration Network

Vision transformer (ViT) has shown surprising performance in both high-level computer vision tasks [1, 5, 15, 25] and low-level image restoration tasks [2, 13, 14, 17, 28, 29]. IPT [2] is the first method that employs ViT in image restoration, but it is not practical since the complexity of performing full self-attention is quadratic to the image size. To meet the requirement of limited GPU resources and a fast inference, SwinIR [14] successfully combines Swin Transformer [15] with residual convolution layers to perform self-attention on large feature maps and achieves a reliable performance on several image restoration tasks. Meanwhile, Uformer [28] and Restormer [29] were designed to address the problems that were usually solved by U-Net [22] structures, such as denoising and deblurring. The former also applies local window attention and utilizes U-structure to capture global dependencies. And the latter applies self-attention across channels rather than

the spatial dimension to make the complexity increase linearly. Meanwhile, lightweight networks like LPIENet [4] remarkably reduce computational workload than state-of-the-art networks on image restoration tasks while maintaining close image quality, making it possible to run certain tasks on mobile devices.

### 3. Dataset

We assess our method on Bokeh Effect Transformation Dataset (BETD) [3], which consists of 20,000 training image sets and 500 validation images. Each image has a size of  $1920 \times 1440$  pixels and each training pair includes a source image, a target image, and a metadata tuple of id, source lens, target lens and disparity value. The source and target images have the same foreground and background, but the latter is taken under a different set of lenses and apertures. Moreover, each training pair is accompanied by an alpha mask, which can be utilized for more advanced training losses or evaluations. Disparity value indicates the level of blur strength difference between the target image and the source image. To explore the dataset in more detail, we examined the dataset features and statistics. We create cross-tabulations (Fig. 2a) for each pair of transformation combinations and plot the distribution of disparity (Fig. 2b). Furthermore, we summarized the evaluation result by using the three following evaluation metrics, which are also used in our experiments to evaluate the performance of models:

- (i) Peak Signal-to-Noise Ratio (PSNR). By calculating the ratio of the maximum possible power of a pixel to the power of corrupting noise. A higher PSNR indicates better image quality.
- (ii) Structural Similarity (SSIM) [27]. This metric quantifies the similarity between two images. A higher SSIM ensures the image is similar to its reference. SSIM is significant especially for blur-to-sharp bokeh transformation.
- (iii) Learned Perceptual Image Patch Similarity (LPIPS) [30]. Unlike metrics such as PSNR or SSIM, LPIPS considers the human perception of image quality, making the evaluation system more comprehensive.

Dataset	PSNR	SSIM	LPIPS
BETD	34.691	0.903	0.162

Table 1. Overall statistics of BETD [3].

Table 1 shows the overall statistics of BETD [3]. All values are calculated between source images and the corre-

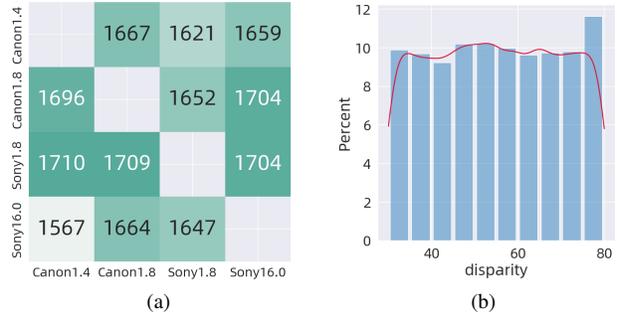


Figure 2. (a): The transformation distribution of lens pairs in BETD [3]. Each number in the grids indicates the number of transformation occurrences from the vertical lens labels to the horizontal lens labels. “Canon1.4” is the abbreviation for “Canon50mmf1.4BS”, and so forth. (b): The percentages of different disparity values.

Target \ Source	Canon1.4	Canon1.8	Sony1.8	Sony16
Canon1.4	-	46.013	42.062	24.293
Canon1.8	46.016	-	45.624	24.623
Sony1.8	42.030	45.693	-	25.052
Sony16	24.431	24.741	24.995	-

Table 2. PSNR of different transformation pairs

Target \ Source	Canon1.4	Canon1.8	Sony1.8	Sony16
Canon1.4	-	0.993	0.989	0.808
Canon1.8	0.993	-	0.992	0.810
Sony1.8	0.989	0.993	-	0.821
Sony16	0.809	0.814	0.817	-

Table 3. SSIM [27] of different transformation pairs

Target \ Source	Canon1.4	Canon1.8	Sony1.8	Sony16
Canon1.4	-	0.005	0.011	0.324
Canon1.8	0.005	-	0.007	0.319
Sony1.8	0.011	0.007	-	0.311
Sony16	0.327	0.318	0.315	-

Table 4. LPIPS [30] of different transformation pairs

sponding target images. Table 2,3 and 4 illustrate the statistics on different transformations. These statistics can be referenced as baselines to evaluate model performance.

### 4. Proposed Framework

The bokeh effect, which involves sophisticated factors, is an optical phenomenon generated by the light passing

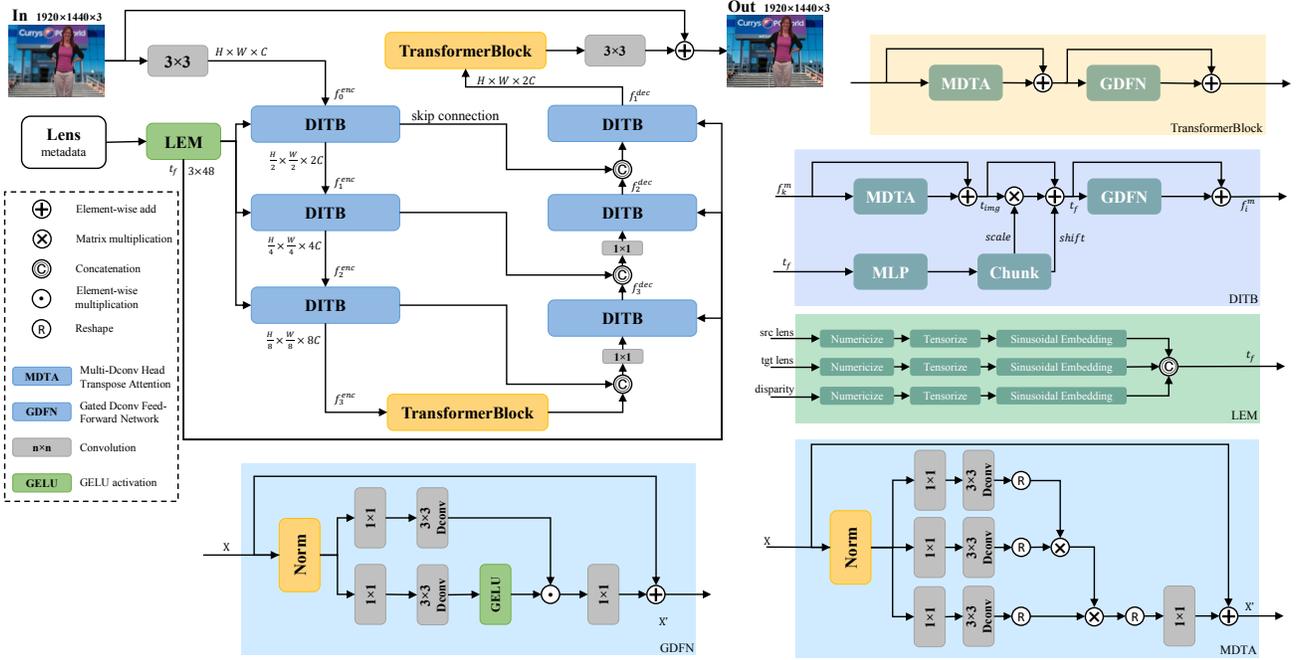


Figure 3. Overview of our BokehOrNot Framework. Our BokehOrNot consists of multi-scale U-net [22] design incorporating dual-input transformer block (DITB) and lens embedding module (LEM). GDFN and MDTA are blocks adopted from [29].

through the optical module of the lens. Bokeh effects gain explicit differences from multiple aspects, including the sharpness of the bokeh, the aberration around the bokeh, the distortion of the bokeh *etc.*

**Overview.** The proposed BokehOrNot framework is based on the Restormer [29] framework, a Transformer-based network, and combined with factors related to bokeh effect. BokehOrNot contains a lens information embedding module and a symmetric encoder-decoder structure with dual-input Transformer blocks. Our proposed model framework is shown in Fig. 3.

#### 4.1. Lens Embedding Module (LEM)

Optics and photography professionals have verified that different bokeh effects are caused by multiple factors within lenses [24], including aperture sizes, the material of lens elements, the number of diaphragm blades [19] *etc.* Among these factors, the aperture size has been identified as a crucial factor. This observation has led to an approach in which aperture sizes and lens models are extracted from the BETD [3] to leverage lens information.

To utilize the lens information, the aperture size, which significantly affects the intensity of the bokeh effect, is expressed as the value of the corresponding size, *e.g.* a f/1.8 aperture will be expressed as 1.8. Furthermore, to distinguish lenses, we use one-hot encoding on the lens metadata

by the following pattern:

$$meta = [k_1, k_2, \dots, k_{n-1}, aperture], \quad (1)$$

where  $n$  is the kinds of lenses,  $k_i \in \{-1, 1\}$  and  $k_i = 1$  for the  $i$ -th lens whereby the others are  $-1$ . Considering the new BETD [3], Sony lenses are assigned  $[1, aperture]$  and Canon lenses are assigned  $[-1, aperture]$ . For example, a Canon f/1.4 lens is transformed into a value of  $[-1, 1.4]$  (under a binary scenario,  $-1.4$  is identical to  $[-1, 1.4]$ ), allowing the model to distinguish lens models that significantly influence the bokeh style through an explicit value. By calculating the Hamming distance between the source lens tensor and the target lens tensor while a longer distance typically represents bigger differences, this encoding ensures that:

- (i) When the source lens model and the target lens model are the same, the magnitude of the bokeh transformation is expressed by the absolute value of the difference between aperture values. A larger absolute value signifies a greater transformation magnitude. For example, a transformation from Sony50mm f/1.4 to Sony50mm f/1.8 has a difference of 0.4 while a transformation from Sony50mm f/1.8 to Sony50mm f/16.0 has a difference of 14.2. The latter applies to a larger transformation magnitude.
- (ii) When aperture sizes are same, the transformation gains the same magnitude between different lenses.

- (iii) Transformations involving both lens type and aperture size obtain a greater Hamming distance than changing a single factor alone.

The disparity, which is an indicator of blur strength, receives identical processing as the aperture size. This approach facilitates the indication of differences in bokeh intensity at various aperture sizes through the difference in values.

To learn the transformation from the source lens to the target lens with a certain disparity, we use sinusoidal positional embedding [26] to process source lens values, target lens values and disparity values. Each of these values is “tensorized” to an  $1 \times 48$  float tensor and three tensors are then concatenated into a  $3 \times 48$  tensor. A two-layer perceptron helps to extract features from this tensor and returns an  $1 \times 48$  tensor.

## 4.2. Dual-Input Transformer Block (DITB)

Upon obtaining the embedding tensor from the LEM, we incorporate lens information into the network forwarding process. To this end, a dual-input transformer block based on vision transformer [5] is introduced to leverage both the image input and the lens input. The lens input is subject to processing by a two-layer perceptron, which produces an intermediate tensor  $t_p$  with dimensions  $2 \cdot dim$ . Subsequently,  $t_p$  is chunked into two distinct segments, namely the *scale* and *shift*. The image input  $t_{img}$  is only updated by multi-Dconv head transposed attention (MDTA) [29] in this phase.

$$t_{img} = t_{img} + MDTA(t_{img}) \quad (2)$$

In order to achieve a comprehensive utilization of lens information, we propose a novel approach that involves the fusion of both lens and image inputs, enhanced by scaling and shifting operations. The outcome of this process  $t_f$  is a fused feature that encapsulates a more holistic representation of the input data:

$$t_f = t_{img} \cdot (1 + scale) + shift \quad (3)$$

which makes key information propagate further, gated-Dconv feed-forward network (GDFN) [29] is applied to the  $t_f$ , with residual added after. The output of DITB is a feature map with a size corresponding to the depth of the block:

$$f_i^m = t_f + GDFN(t_f) \quad (4)$$

where  $i \in \{1, 2, 3\}$  and  $m \in \{enc, dec\}$ .

## 4.3. Alpha-masked Loss

Our proposed framework employs  $L_1$  loss function and  $L_1$  loss with foreground area removed in different training

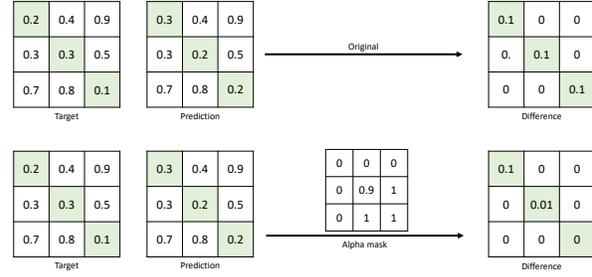


Figure 4. Demonstration of the alpha-masked loss. The loss of background area (upper-left grid in this example) obtains a larger proportion of the loss. The edge area (central grid) still contributes to the loss with less significance while the foreground area (lower-right grid) is no longer considered.

stages. Specifically, when performing bokeh effect transformation, the foreground, *e.g.* people, pets, and other salient objects, should not be altered to achieve a higher verisimilitude. Along with the BETD [3], it motivates us to exclude the foreground area which is covered by the alpha mask when calculating loss.

Figure 4 illustrates how the alpha-masked loss is generated. Grids in green represent pixels that failed to predict accurately. In the original  $L_1$  loss, only one of the three green pixels contributes to transformation in the background and all areas contribute equally, while in alpha-masked loss, green grids in the foreground area are no longer considered. The loss now involves only the background area and slightly considers the edge area.

The alpha-masked loss can be represented as:

$$L = \frac{1}{N} \sum_{i=1}^N |pred_i - gt_i| * (1 - alpha_i) \quad (5)$$

where  $pred_i$  is the  $i$ -th pixel of the transformation result from our model,  $alpha$  is the mask from BETD [3],  $gt$  is the ground truth and  $N$  is the number of pixels. The alpha-masked loss performs element-wise multiplication between the alpha mask and the prediction image or the target image, resulting in a zero-filled area.

With this, the learning process benefits from two perspectives: alleviating the transformation on the foreground area and eliminating the impact of loss increment due to the foreground area while learning bokeh effect within background areas.

## 5. Experimental Results

### 5.1. Implementation Details

BokehOrNot is implemented using PyTorch and trained with the Adam [12] optimizer. Two training stages are involved: the precise detecting stage and the global transformation stage.

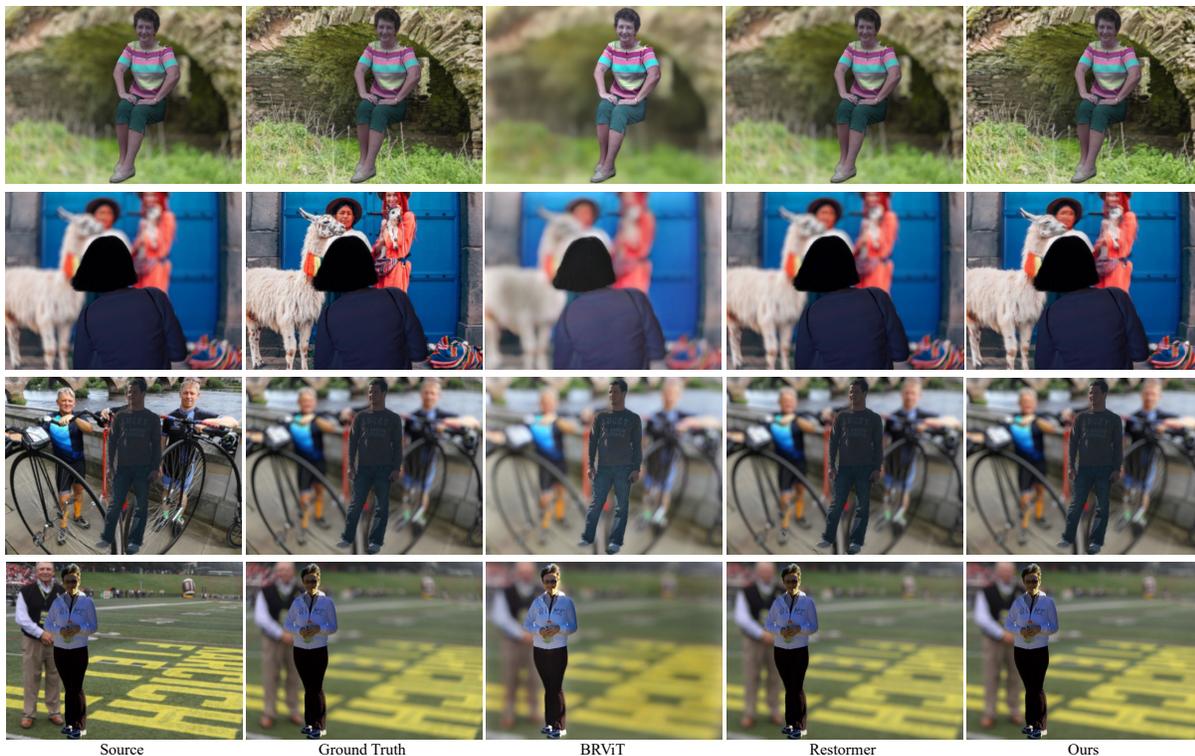


Figure 5. Qualitative comparison of bokeh effect transformation. The first image applies a transformation from Sony50mmf1.8BS to Sony50mmf16.0BS and the second image applies a transformation from Canon50mmf1.4BS to Sony50mmf16.0BS, *i.e.* blur-to-sharp. The following two images are transformed from Sony50mmf16.0BS to Sony50mmf1.8BS and Canon50mmf1.4BS, respectively. BRViT [18] failed to render realistic bokeh effect and Restormer [29] cannot render blur-to-sharp transformation ideally.

Model	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
BRViT(i) [18]	12.620	0.765	0.214
Restormer [29](i)	36.594	0.942	0.106
Ours(i)	38.090	0.984	0.041
Ours(ii)	39.320	0.967	0.049

Table 5. Quantitative results on different models. (i) a total of 500 “f/16.0 to f/1.8” images are tested in the experiments, (ii) a total of 500 images of any bokeh transformation are tested.

During the precise detecting stage, the batch size is set to 4 with a learning rate of  $1e-4$ . The loss function is the original  $L_1$  loss, with an input image size randomly cropped to  $256 \times 256$ . The global transformation stage takes advantage of our alpha-masked loss shown in Eq. 5 with  $384 \times 384$  input size, enabling the model to learn holistically and ignoring the foreground. The learning rate is set to  $5e-5$  and the batch size is reduced to 2. The total training time is 70 hours, with approximately 35 hours for each stage. The training is performed on 2 NVIDIA A100 40GB GPUs. An average inference time for a  $1920 \times 1440$  size input is 0.94s on a single NVIDIA A100 40GB GPU.

Model	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Restormer [29]	36.594	0.942	0.106
Ours(i)	38.534	0.965	0.058
Ours(ii)	38.558	0.966	0.051
Ours(iii)	39.137	0.964	0.055
Ours(iv)	39.320	0.967	0.049

Table 6. Ablation experiments results on proposed approaches. The Restormer [29] model applies the same settings in Table 5. (i) only with LEM+DITB, (ii) model (i) with alpha-masked loss, (iii) model (i) with enlarged inputs applied, (iv) proposed BokehOrNot model. A total of 500 images of any bokeh transformation are tested.

## 5.2. Results and Analysis

To the best of our knowledge, no publicly available deep network models now involve lens metadata. Therefore, we select two recently published network structures for bokeh effect rendering and image restoration with publicly available code as references. Considering current models are primarily focusing on rendering shallow depth-of-field from a

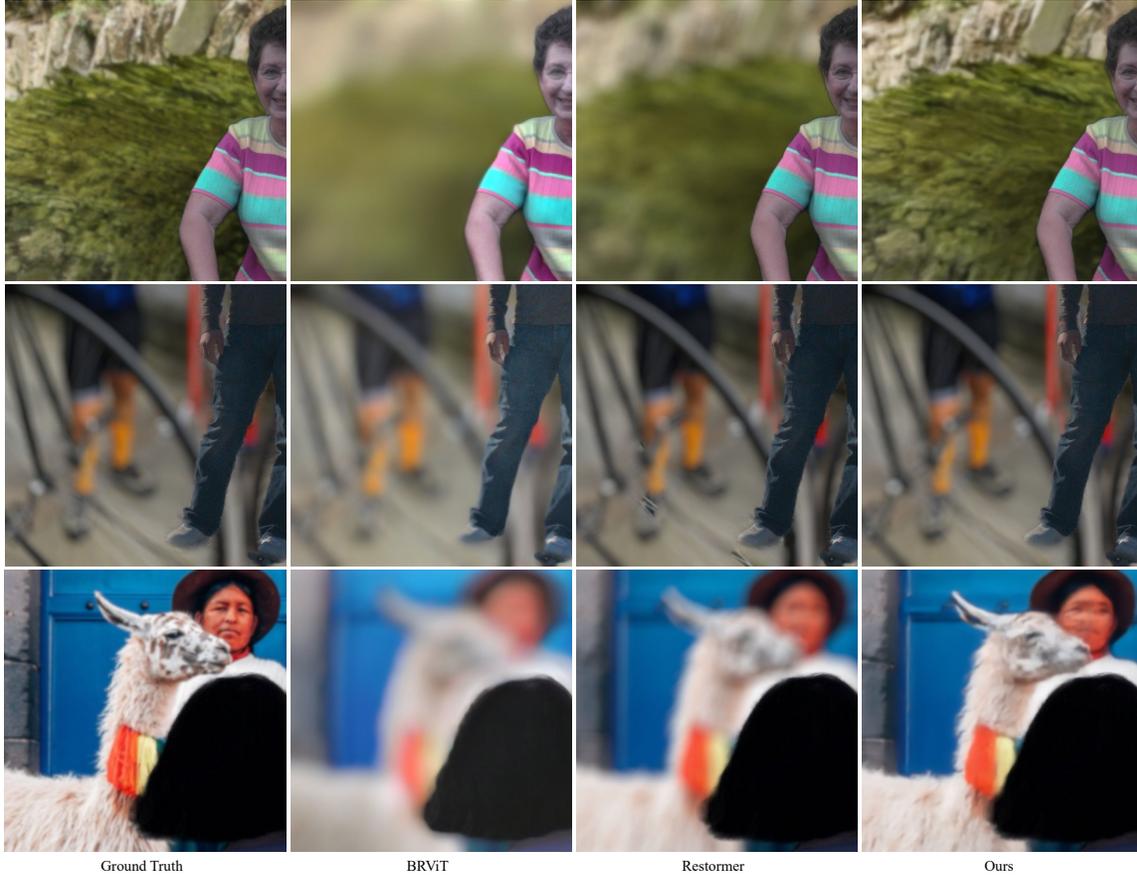


Figure 6. Magnified samples of typical transformations. BokehOrNot generates higher quality images than BRViT [18] and Restormer [29].

Team	All images			Real images	
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Ours (BokehOrNot)	32.288	0.9327	0.1130	0.8423	0.2199
EBokehNet [23]	34.543	0.9350	0.1039	0.8414	0.2206
IR-SDE	30.866	0.9297	0.1301	0.8427	0.2387
BIGbaodan	30.327	0.9281	0.1249	0.8415	0.2178
AIA-Smart	34.572	0.9361	0.0966	0.8435	0.2224
Samsung Research China	35.264	0.9362	0.0985	0.8416	0.2186
NUS-LV-Bokeh	32.326	0.9333	0.1076	0.8420	0.2230
IPAL-Bokeh	32.076	0.9324	0.1076	0.8419	0.2161
Baseline	28.599	0.9128	0.1878	0.8412	0.2181

Table 7. Quantitative results from all teams above baseline. This result is calculated on an undisclosed test set with both artificial synthetic images and real-captured images. A total of 155 images are tested, 95 are synthetic images and 60 are real-captured.

rather sharp image, we preclude blur-to-sharp transformations in the dataset. All models are trained with recommended settings or similar settings as ours.

Our proposed BokehOrNot model outperforms on all three metrics. The LEM and DITB which utilize lens meta-

data contribute greatly to the evaluation results shown in Table 5. As bokeh effect transformation quality is the subjective perception of people to some extent, we also conduct qualitative comparisons, shown in Fig. 5.

BRViT [18] model blurs images under any scenario and

engages brightness distortion. Restormer [29] can distinguish the direction of transformation to some extent, *i.e.* “to blur” or “to sharp”. However, it still cannot perform blur-to-sharp transformation with satisfactory quality. Our model renders both types of transformation desirably with natural bokeh effect.

Figure 6 shows magnified images with details. Comparing to BRViT [18] and Restormer [29], our model significantly restores the texture of rocks and mosses while distinguishing a clear edge between the background and the foreground in the first image. In the second image, the Restormer [29] model failed to render the bokeh effect on all background areas, especially the area adjacent to wheel spokes, whereby our model renders natural bokeh effect. In general, our model is capable of rendering bokeh effects under different scenarios while retaining a sharp foreground, creating a natural and pleasing visual effect.

### 5.3. Ablation Study

As we integrate multiple new approaches in our framework, ablation experiments are conducted to further verify and evaluate the performance and robustness of these methods, including: applying LEM and DITB, applying alpha-masked loss and applying enlarged inputs.

Table 6 indicates our proposed methods still improve the performance when applied separately. LEM and DITB notably elevate all metrics in comparison with Restormer [29]. The incorporation of alpha-masked loss contributes notably to a reduction in the LPIPS [30] metric and the enlarged inputs method results in a significant enhancement in the PSNR metric.

## 6. NTIRE 2023 Challenge

We enroll in NTIRE 2023 Bokeh Transformation Challenge [3]. The model proposed in this paper participates in the joint comparison among different models. Quantitative results are provided by challenge organizers, shown in Table 7. The result contains two parts: all images including synthesized images, and only real images. Real images are taken directly from a camera, and only apply transformations between  $f/1.4$  and  $f/1.8$ .

In this challenge, we use the direct output of the model, without any model ensemble the test-time augmentation. Our result surpasses the baseline for 3.689 dB on PSNR, ranking the fifth place in this competition.

## 7. Conclusion

Bokeh Effect Transformation Dataset [3] is the first dataset ever incorporating lens metadata and multi-style bokeh effect transformation. We introduce the BokehOrNot model with two new methods to transform and render beautiful bokeh effect, namely lens embedding module with

dual-input Transformer block and alpha-masked loss, along with this novel dataset. The former focuses on integrating lens metadata efficiently, allowing the model to learn different bokeh styles, while the latter aims to learn specific bokeh effect more accurately. Quantitative and qualitative experiments demonstrate that our proposed model can transform and render natural bokeh effect given metadata and outperforms cutting-edge models on three metrics. Considering images taken by either portable cameras or mobile phones are embedded with lens metadata, it is feasible to apply this model to real-world scenarios.

## Acknowledgements

The computations were enabled by the *Berzelius* resources provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre. And thanks to Ziwei Luo for insightful feedback.

## References

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 2
- [2] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12299–12310, 2021. 2
- [3] Marcos V Conde, Manuel Kolmet, Tim Seizinger, Tom E Bishop, and Radu Timofte. Lens-to-lens bokeh effect transformation. ntire 2023 challenge report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2023. 1, 2, 3, 4, 5, 8
- [4] Marcos V Conde, Florin Vasluianu, Javier Vazquez-Corral, and Radu Timofte. Perceptual image enhancement for smartphone real-time applications. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1848–1858, 2023. 3
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 5
- [6] 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
- [7] Konstantinos Georgiadis, Albert Saà-Garriga, Mehmet Kerim Yucel, Anastasios Drosou, and Bruno Manganelli. Adaptive mask-based pyramid network for realistic bokeh rendering. In *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 429–444. Springer, 2023. 2

- [8] 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
- [9] 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
- [10] 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. 2
- [11] Andrey Ignatov, Radu Timofte, Jin Zhang, Feng Zhang, Gaocheng Yu, Zhe Ma, Hongbin Wang, Minsu Kwon, Hao-tian Qian, Wentao Tong, et al. Realistic bokeh effect rendering on mobile gpus, mobile ai & aim 2022 challenge: report. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*, pages 153–173. Springer, 2023. 2
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [13] Jingyun Liang, Jiezhong Cao, Yuchen Fan, Kai Zhang, Rakesh Ranjan, Yawei Li, Radu Timofte, and Luc Van Gool. Vrt: A video restoration transformer. *arXiv preprint arXiv:2201.12288*, 2022. 2
- [14] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1833–1844, 2021. 2
- [15] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 2
- [16] 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 16*, pages 245–261. Springer, 2020. 2
- [17] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 998–1008, 2022. 2
- [18] Hariharan Nagasubramaniam and Rabih Younes. Bokeh effect rendering with vision transformers. *TechRxiv preprint techrxiv.17714849.v1*, 2022. 1, 2, 6, 7, 8
- [19] HH Nasse. Depth of field and bokeh. *Carl Zeiss camera lens division report*, 1, 2010. 4
- [20] Juewen Peng, Zhiguo Cao, Xianrui Luo, Hao Lu, Ke Xian, and Jianming Zhang. Bokehme: When neural rendering meets classical rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16283–16292, 2022. 2
- [21] 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 *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 229–244. Springer, 2020. 2
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 2, 4
- [23] Tim Seizinger, Marcos V Conde, Manuel Kolmet, Tom E Bishop, and Radu Timofte. Efficient multi-lens bokeh effect rendering and transformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2023. 7
- [24] Viktor P Sivokon and Michael D Thorpe. Theory of bokeh image structure in camera lenses with an aspheric surface. *Optical Engineering*, 53(6):065103–065103, 2014. 4
- [25] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7262–7272, 2021. 2
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [27] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 3
- [28] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17683–17693, 2022. 2
- [29] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5728–5739, 2022. 2, 4, 5, 6, 7, 8
- [30] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 3, 8
- [31] Bolun Zheng, Quan Chen, Shanxin Yuan, Xiaofei Zhou, Hua Zhang, Jiyong Zhang, Chenggang Yan, and Gregory Slabaugh. Constrained predictive filters for single image bokeh rendering. *IEEE Transactions on Computational Imaging*, 8:346–357, 2022. 1, 2