

High-Resolution Document Shadow Removal via A Large-Scale Real-World Dataset and A Frequency-Aware Shadow Erasing Net

Zinuo Li* Xuhang Chen* Chi-Man Pun† Xiaodong Cun†
University of Macau

Abstract

Shadows often occur when we capture the document with casual equipment, which influences the visual quality and readability of the digital copies. Different from the algorithms for natural shadow removal, the algorithms in document shadow removal need to preserve the details of fonts and figures in high-resolution input. Previous works ignore this problem and remove the shadows via approximate attention and small datasets, which might not work in real-world situations. We handle high-resolution document shadow removal directly via a larger-scale real-world dataset and a carefully-designed frequency-aware network. As for the dataset, we acquire over 7k couples of high-resolution (2462×3699) images of real-world documents pairs with various samples under different lighting circumstances, which is 10 times larger than existing datasets. As for the design of the network, we decouple the high-resolution images in the frequency domain, where the low-frequency details and high-frequency boundaries can be effectively learned via the carefully designed network structure. Powered by our network and dataset, the proposed method shows a clearly better performance than previous methods in terms of visual quality and numerical results. The code, models, and dataset are available at <https://github.com/CXH-Research/DocShadow-SD7K>.

1. Introduction

When occluders obstruct the light source, shadows are cast on the paper. These shadows reduce the readability and hinder subsequent text-related intelligent tasks such as document enhancement, optical character recognition (OCR), key information extraction, and semantic entity labeling [17, 36, 28], etc.

Removing the cast shadow of the digital copy has a long need in the real world and also there are a lot of meth-

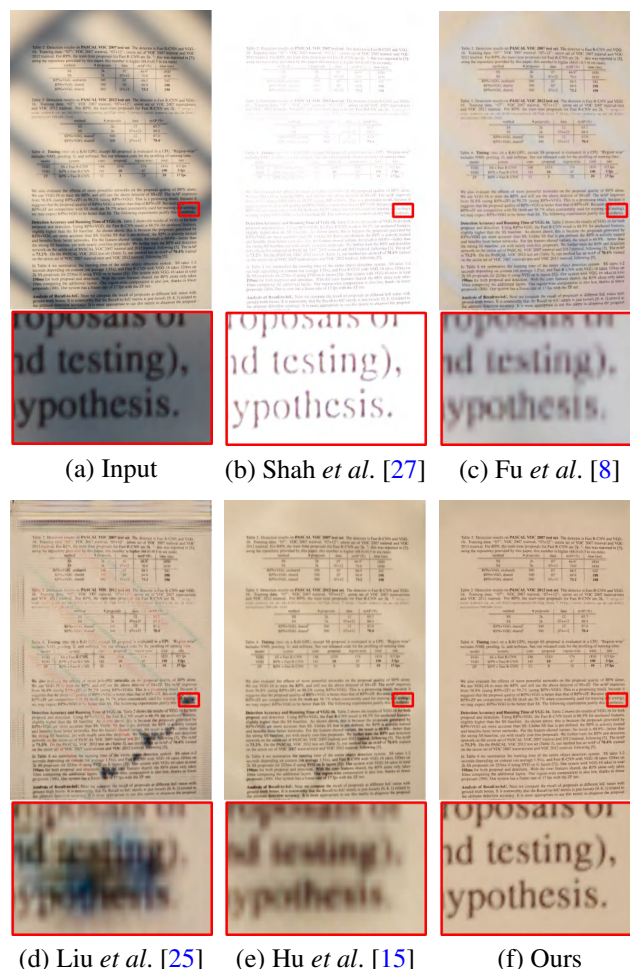


Figure 1. Visual results of input document shadow image (a), classic method (b), supervised method (c), weakly-supervised method (d), unsupervised method (e) and ours (f). Our model removes the shadow while preserving the original document’s content and aspect ratio.

ods, including the traditional methods that applied physics-based illumination models [16, 30, 32]. In reality, these approaches are significantly constrained due to the fact that the assumptions utilized in illumination models [27] are often too rigid to account for real-world shadow. Thus, they may

*Equal contributions

†Corresponding authors

incorrectly compute the exposure and reflectance as shown in Fig. 1.

Recent methods apply deep learning for this task since it has achieved remarkable progress in natural shadow removal [6] and other computer vision tasks [12, 13, 35]. However, several obstacles make it challenging to transfer previous natural shadow removal methods to the document domain. On the one hand, in contrast to natural shadow removal, documents are presented by high-resolution images in the real world, which needs to preserve fine-grained information such as fonts and figures. However, the majority of document shadow removal methods [24, 4] are only designed for relatively low-resolution images and is hard to handle high-resolution situations directly since they are relying on the low-resolution approximation of attention for guidance. On the other hand, there are no large-scale and high-resolution datasets designed specifically for document shadow removal at this time. Consequently, given limited image pairs for training, existing methods usually result in different artifacts (as shown in Fig. 1 (c)). Although weakly-supervised [18] and unsupervised [15, 25] shadow removal methods alleviate the requirements for data, they hold a strong assumption of the statistical similarity between shadowed and unshadowed domains. As illustrated in Fig. 1 (d) (e), once the training domain is very different from the test, these methods create hallucination contents and suffer from unstable results.

To solve the above problems, we first offer a large-scale real-world shadow dataset, Shadowed Document 7K (SD7K), which contains more than 7k real-world and high-resolution shadow and shadow-free document image pairs with also the manually annotated shadow masks for related research, *e.g.*, shadow detection [6]. Compared to existing document shadow removal datasets, SD7K includes various document types, *e.g.*, manga, papers, and figures. Also, different from previous work which only considers one specific light environment, we use three distinct light sources to increase the diversity of the training samples. Besides the dataset, we also propose a robust network for high-resolution image processing specifically, named Frequency-aware Shadow Erasing Net (FSENet). In detail, inspired by the laplacian pyramid [2, 23], we divide the document image into two frequency portions, each with its own domain-specific properties. The low-frequency component is primarily responsible for illuminations or colors while high-frequency components are more related to content details. As for the low-frequency deshading module, we involve the transformer [7] based network to model the global illuminations alternation and a high-frequency restoration module based on cascaded aggregations for adaptive pixel enhancement. For high-frequency modules, we involve several dilated convolution-based texture recovery modules to restore the refined details similar to [6]. With the aforementioned

techniques, our model achieves the state-of-the-art performance of previous methods on two common small benchmarks and the proposed SD7K.

We summarize the main contribution as follows:

- We provide SD7K, a large-scale real-world dataset consisting of high-resolution shadow and the associated shadow-free images under various illumination conditions.
- We propose FSENet, a frequency-aware network with a carefully designed network structure to handle high-resolution document shadows.
- Both qualitative and quantitative results on all available public datasets indicate that the proposed FSENet performs favorably against state-of-the-art methods.

2. Related Work

Document Shadow Dataset. The collection of genuine shadow images and their corresponding shadow-free counterparts is a challenging task due to the varying lighting conditions. As a result, existing datasets are limited in size. The Bako dataset [1] comprises outdated, inaccessible images in tiny quantities. Similarly, Jung’s dataset [16] is limited in size and favors light shadows. Kligler’s dataset [17] emphasizes colorful symbols, which do not convey a significant amount of information. Among these datasets, the RDSRD dataset [24], while still modest in size, has a better quality in comparison. Different from the existing datasets, we propose a high-resolution and large-scale dataset SD7K that addresses the aforementioned problem, which can facilitate the research in document shadow removal.

Natural Shadow Removal. Methods for natural shadow removal might be informative and applicable to document images. On the one hand, GAN-based methods such as ST-CGAN [31], use two stacked conditional GANs to jointly detect and eliminate shadows. Mask-ShadowGAN [15] extend the paired GAN to unpaired setting via re-formulated cycle-consistency constraints. To fully utilize the dataset, DHAN [6] synthesize pseudo shadow images and learned border artifacts-free images via a dual hierarchical aggregation network. SG-ShadowNet [29] studies the importance of normalization and learns the style representation of the non-shadow region to harmonize the shadow and shadow-free parts. On the other hand, CNN-based methods such as SP+M Net [19] and SP+M+I Net [20] utilize the physical linear transformation model to enhance the shadow region and image decomposition to reconstruct the shadow-free image. AEF-Net [8] predicts the exposure parameters of the shadowed regions and fused a series of overexposed shadow images with the original shadow image to restore the shadow region. ShadowFormer [10] is a transformer-

Dataset	Triples #	Entities #	Mean Resolution	Document Type	With Mask	Light Temperature Control
Bako [1]	81	N/A	N/A	Text	No	No
Kligler [17]	300	25	2424 × 2013	Text	No	No
Jung [16]	87	87	602 × 860	Text	No	No
RDSRD [24]	540	N/A	N/A	Text	Yes	N/A
SD7K	7620	350+	2462 × 3699	Text and picture	Yes	Different lamps & sunlight

Table 1. Comparison of SD7K with other real-world datasets. The proposed dataset covers a large number of document shadow images and it is the only dataset that satisfies all important data features about document shadow currently.

based network to exploit the context correlation between shadow and non-shadow regions.

Document Image Shadow Removal. Traditional document shadow removal methods usually demand prior knowledge regarding shadow and document context. This includes estimating the background color [30, 32] and modeling shading and lighting effects [16, 27]. Deep learning-based methods are also designed to remove the shadows of documents. Currently, BEDSR-Net [24] stands as the state-of-the-art model in document shadow removal. It predicts global background color and encodes shadow position as an attention map, allowing for the effective removal of shadows. However, as discussed in our introduction, it needs to predict the approximate attention maps from the feature map, which is not suitable for the high resolution inputs.

3. SD7K Dataset

As shown in Tab. 1 and Fig. 3 (a), there are only a limited number of datasets [16, 24, 1, 17] available for document shadow removal, which is typically small and comprises only hundreds of real-world samples. To train a network on such-scale datasets, synthesizing training samples via computer graphic engines is required [24]. However, the rendered samples still show different distributions from the real samples, making it difficult in real-world applications.

To overcome this limitation, we introduce a massive real-world dataset named SD7K for document shadow removal in terms of the occluder entities, the numbers, the resolution, the light source, *etc.* Below, we give the details of how to collect and process the dataset.

First, we use a Canon EOS M6 camera to capture printed document images from real-world scenarios. To capture paired shadow and shadow-free images, we make all the light close and leave one light source. All the settings of the camera are also fixed, including focal length, aperture, exposure, and ISO, to avoid the influence of the camera algorithm. Then, we use a camera stand to hold the camera and make it shoot on the paper vertically as shown in Fig. 2 (a). To capture the paired shadow and shadow-free image without involving the changes caused by the manual presses, as in Fig. 2 (b), we use a remote control to capture the photo between the shadowed and shadow-free paper by removing

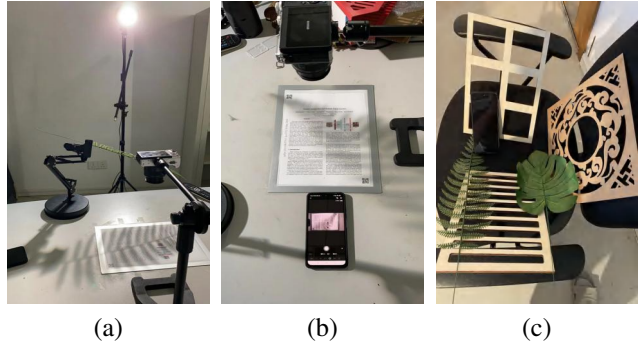


Figure 2. (a) our data acquisition setup for constructing the dataset. (b) a remote phone control shutter where the user only needs to click a button to capture a pair of shadow-free/shadow images. (c) some showcases of our occluders.




Figure 3. Data distribution of SD7K and quantitative comparison across all document shadow datasets.

the occluder. We also take the pairs as short as possible to avoid the light difference between captures.

To further enhance diversity, we alternated the positions of the document and occluder after capturing a few batches of documents. Overall, we use over 30 types of occluders along with more than 350 documents to contribute to the dataset, some of the occluders are shown in Fig. 2 (c). These occluders have the shape of both regular and irregular forms, which provides adequate coverage for various situations, we calculate the shadow area percentage of the whole dataset, which is $41.38\% \pm 13.58\%$, indicating that it is sufficiently divergent for our shadow removal task.

We also consider the fluctuations in sunlight intensity and random placement of shadows, which will be significantly influenced by the light variance. To simulate the results in different light sources, we introduce three types of light sources in our dataset for collection, including reg-



Cool light Warm light Sunlight

Figure 4. Example shadow and shadow-free images from SD7K.

ular sunlight, a stationary 80W LED yellow light having a color temperature of 3000K, and a stationary 80W LED white light having a color temperature of 6000K. The light distribution of the whole dataset is shown in Fig. 3 (b).

For data cleaning purposes, we first remove images with noticeable motion distortion or lack of focus. Since some photo still contains outliers and the context is duplicated, we then review the photographs in groups and remove them. Finally, we crop the areas outside of the margins of the document, resulting in a total of 7620 samples, where some of the samples are shown in Fig. 4. We also use a threshold-based method to extract the mask of the training pair since the shadow mask will also benefit related tasks.

4. Frequency-Aware Shadow Erasing Net

Since real-world document shadow removal needs huge GPU memory requirements, our network initially divides the input image into different frequency components using the Laplacian Pyramid (LP) similar to [23] to reduce the computational budgets. As shown in Fig. 5, we propose a powerful frequency-aware shadow removal network named FSENet by designing different modules to handle the high-frequency and low-frequency computation individually. After training, the multi-frequency components are mixed to produce the final results. In the rest of this section, we first give the details of the Laplacian pyramid in Sec. 4.1. Then, in Sec. 4.2 and Sec. 4.3, we give the details of the low-frequency deshading module and high-frequency restoration module to precisely modify the color details and the high-frequency details individually. Finally, we give the

loss functions in Sec. 4.4.

4.1. Laplacian Pyramid and Method Overview

Laplacian Pyramid (LP) [2] plays an important role in image processing. By decomposing and restoring images in a lossless manner, LP provides a rapid processing speed and is well-suited for dealing with high-resolution images. Inspired by [23] for image translation, we implement LP for document shadow removal and train the network on multi-frequency.

In detail, given an image $I \in \mathbb{R}^{H \times W \times 3}$, the D layers LP contains several high-frequency components $[L_0, \dots, L_{D-1}]$ and a low-frequency component $L_D \in \mathbb{R}^{\frac{H}{2^D} \times \frac{W}{2^D} \times 3}$, where L_D is D times downsample of I . Each high-frequency component can be described as $L_i = I_i - \text{PyrUp}(\text{PyrDown}(I_i))$, where I_i is the i -th down-sampled image and PyrUp and PyrDown stand for upsampling and downsampling, respectively. As illustrated in Fig. 5, the L_1 and L_2 include the high-frequency component edges and textures, while the low-frequency component $L_3 = I_3$ majorly contains the color information. For reversible reconstruction, The final \hat{I} can be obtained by reversed update the equation of $\hat{I}_{D-1} = \text{PyrUp}(\hat{I}_D) + \hat{L}_{D-1}$ from D to 0, where \hat{I}_D is the learned low-frequency features and \hat{L}_{D-1} is the learned high-frequency features from the proposed network. We give detailed instructions for each network below.

4.2. Low-Frequency Deshading Module

After getting the low-frequency components, we use several transformer-based blocks to reduce the color and illumination distortions of the shadow images. Compared to CNNs, transformers are superior at modeling non-local self-similarity and long-range dependencies. As shown in Fig. 5, we implement a transformer-based module called Dimension-Aware Transformer (DAT) inspired by [33] with a lightweight UNet-like sub-network [3] for low-frequency feature modeling.

Formally, given a low-frequency component $L_3 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 3}$ of input image $I \in \mathbb{R}^{H \times W \times 3}$, we first feed L_3 into three sequential DAT blocks to extract rich features. Each DAT block is built via a sequential height-attention and width-attention to learn the feature in different aspects individually and a convolution-based layer for local feature merging, where we denote the intermediate features of each DAT block as $F1, F2, F3 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$. Please refer to the original [33] for further details of the height-attention and width-attentions.

We then feed these features $F1, F2, F3$ into a Tri-layer Attention Alignment (TAA) block for feature remixing. Given concatenated features $F_{\text{in}} \in \mathbb{R}^{N \times \frac{H}{4} \times \frac{W}{4} \times C} = \text{Concat}([F1, F2, F3])$, where $N = 3$ in this case, F_{in} is first reshaped to $\hat{F}_{\text{in}} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 3C}$. Then, we send \hat{F}_{in} to a 1×1 convolution layer to aggregate pixel-level cross-

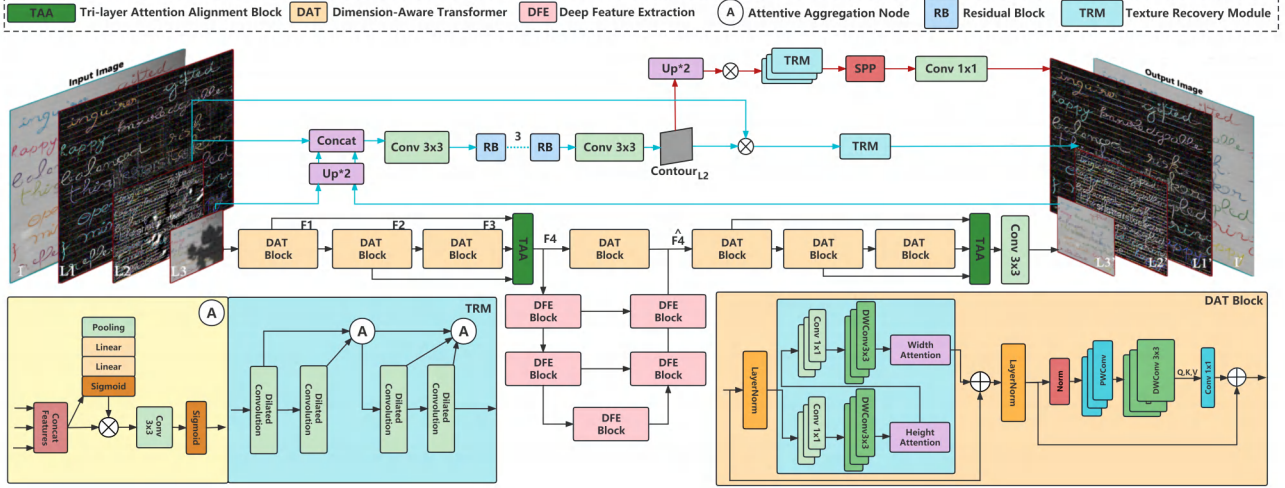


Figure 5. The network structure of our proposed FSENet. Following [23], given a high-resolution image $I \in \mathbb{R}^{H \times W \times 3}$, we first use Laplacian Pyramid (e.g. Depth $D = 2$ in this case) to decompose the images to multiple frequency components. The **Black** arrows: The low-frequency part $L_3 \in \mathbb{R}^{\frac{H}{2^D} \times \frac{W}{2^D} \times C}$ are refined to $L'_3 \in \mathbb{R}^{\frac{H}{2^D} \times \frac{W}{2^D} \times C}$ utilizing DAT, TAA and DFE block. **Cyan** arrows: For the high-frequency part $L_2 \in \mathbb{R}^{\frac{H}{2^{D-1}} \times \frac{W}{2^{D-1}} \times C}$, a contour $C_{L_2} \in \mathbb{R}^{\frac{H}{2^{D-1}} \times \frac{W}{2^{D-1}} \times 1}$ is learned to bridge the low-frequency and high-frequency features. **Red** arrows: For the remaining components with higher frequency, the learned contour is successively upsampled and refined using the proposed SPP and TRM.

channel context, followed by three 3×3 depth-wise convolution layers to generate query \mathbf{Q} , key \mathbf{K} , and value \mathbf{V} , individually. The \mathbf{Q} , \mathbf{K} are then reshaped to $\hat{\mathbf{Q}} \in \mathbb{R}^{N \times \frac{H}{4} \times \frac{W}{4} \times C}$ and $\hat{\mathbf{K}} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C \times N}$ respectively to calculate an attention matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. Next, we reshape the \mathbf{V} to $\hat{\mathbf{V}} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C \times N}$ to multiply $\hat{\mathbf{V}}$ and \mathbf{A} , and the input feature \hat{F}_{in} is added to generate the output feature $F4$, which can be described as:

$$F4 = Conv_{1 \times 1} LA(\hat{\mathbf{Q}}, \hat{\mathbf{K}}, \hat{\mathbf{V}}) + \hat{F}_{in}, \quad (1)$$

where $Conv_{1 \times 1}$ represents a convolution layer with the kernel of 1×1 . LA indicates Layer Attention [33], which can be written as:

$$LA(\hat{\mathbf{Q}}, \hat{\mathbf{K}}, \hat{\mathbf{V}}) = \hat{\mathbf{V}} \text{softmax}(\hat{\mathbf{Q}}\hat{\mathbf{K}}/\alpha), \quad (2)$$

where α is the temperature of the attention. Please refer to the supplementary material for more details of the TAA.

After getting $F4$, as shown in Fig. 5, We also use a UNet-like [3, 35] sub-network built by Deeper Feature extraction (DFE) blocks to learn multi-scale features. Each DFE block begins with a LayerNorm to stabilize the training procedure, followed by two convolutions and a deformable convolution. Then, it utilizes SimpleGate [3] to increase the diversity of the features by dividing the features along the channel dimension into two parts and multiplying them together. We also utilize Simplified Channel Attention (SCA) similar to [3] to capture global information in a computationally efficient manner, which can be denoted as:

$$SCA(\mathbf{X}) = \mathbf{X} * \mathbf{W} \text{pool}(\mathbf{X}), \quad (3)$$

where \mathbf{W} represents a fully-connected layer, pool represents the global average pooling procedure that combines spatial data into channels, and $*$ represents a channel-wise multiplication. The skip-connections [13] are connected before the LayerNorm to the output similar to [3]. We also give a detailed figure of the DFE block in the supplementary materials.

For the refined feature \hat{F}_4 , it then passes through three transformer blocks and a TAA Block to produce the enhanced features for image reconstruction using DAT and TAA blocks similar to the feature extractions of $F1, F2, F3$. Finally, a convolution layer with the kernel size 3×3 is performed on the enhanced features to produce the low-frequency component $L'_3 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$.

4.3. High-Frequency Restoration Module

After getting the low-frequency features, we also design novel components to recover the high-frequency details of the learned features. In detail, we first upsample L_3 and its corresponding result L'_3 to match the size of L_2 , then we learn a contour on the concatenation of $[up(L_3), up(L'_3), L_2]$ via several residual blocks [13], where $up(\cdot)$ means bilinear upsample operation, as illustrated in Fig. 5. This is a straightforward method for optimizing global correction compared to normal RGB images because the bands of the high-frequency feature map are very sparse, allowing us to reduce the volume of calculations as discussed in [23]. This contour will be expanded progressively into a collection of contours to match

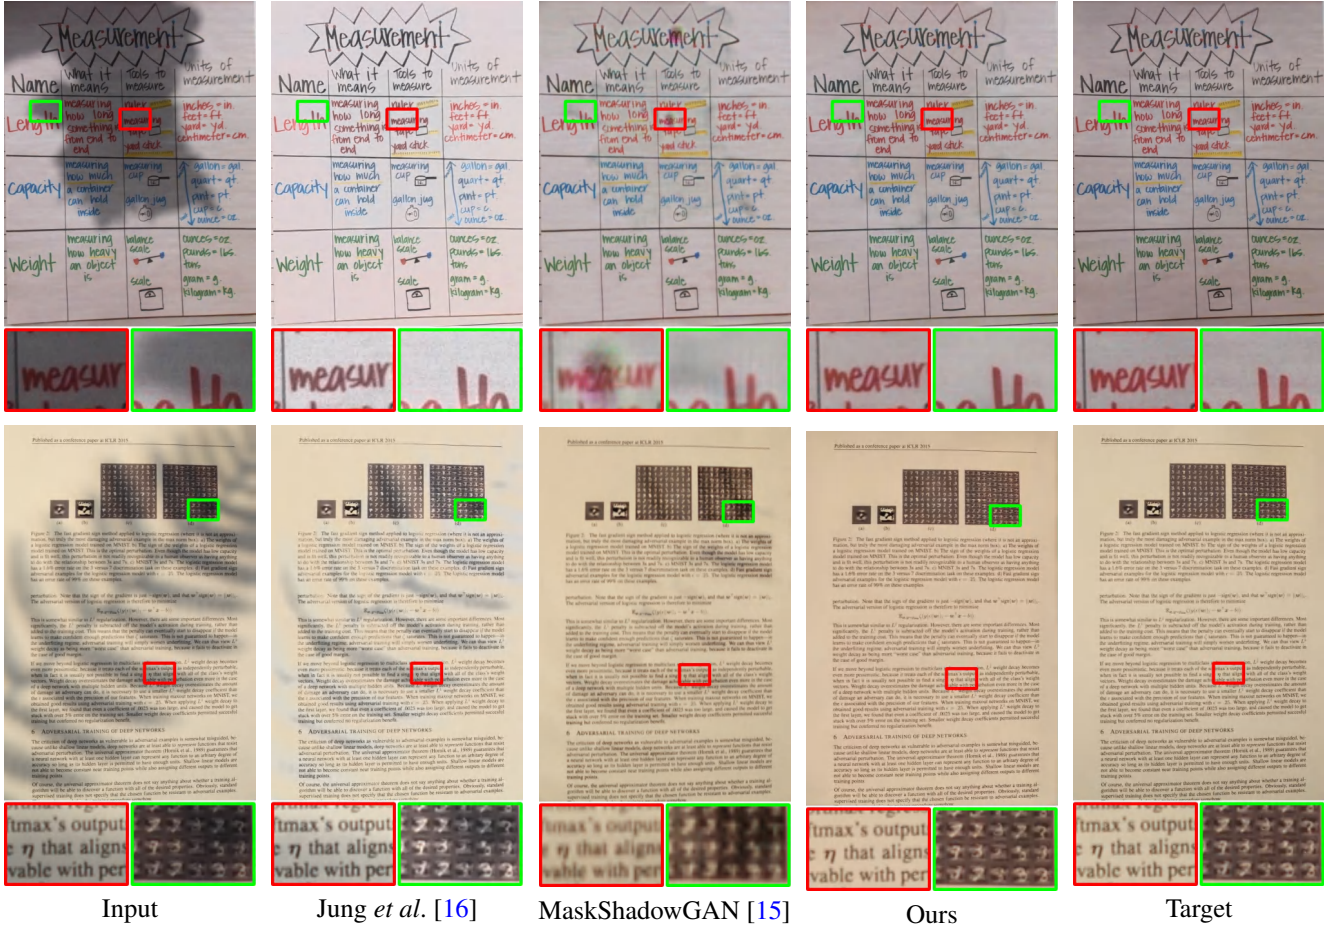


Figure 6. Qualitative results of the methods comparison in high resolution, the first row represents the results on the Kliger dataset and the second row depicts the results on the SD7K dataset.

the remaining high-frequency components. As illustrated in Fig. 5, given the high-frequency component L_2 and the contour C_{L_2} , the output L'_2 can be written as $L'_2 = L_2 \otimes C_{L_2}$, where \otimes represents the pixel-wise multiplication.

After multiplication, the Texture Recovery Module (TRM) is then designed to improve elaboration, inspired by the successful natural shadow removal methods [6, 37], which contain a series of the dilated convolutions and attentive aggregation nodes to recover the detailed features.

In detail, to increase the reception fields on the high-resolution map, we incorporate dilated convolution [38]. Dilated convolution inserts spaces between the kernel values, effectively increasing the receptive field of the filter without losing resolution or coverage over the input. The amount of dilation is controlled by a dilation rate parameter. For a dilation rate of 1, the layer behaves like a regular convolution. For rates larger than 1, the filter is effectively widened to have a larger field of view. This allows the network to incorporate a wider context and capture features at different scales in a computationally effi-

cient manner compared to using very large filters. As for Attentive Aggregation Node, to re-weight the significance of feature maps, each aggregation node employs a squeeze-and-excitation block [14] to reweight the feature. A 3×3 convolution layer is then employed to compress the features and match the original channels. Finally, a spatial pooling pyramid (SPP) [12] is introduced at the conclusion of the TRM to facilitate the remixing of multi-context features. After getting the learned high-frequency of each layer, we use the method introduced in Sec. 4.1 to recover the image and calculate the supervision.

4.4. Loss functions

For training, we utilize the loss functions of SmoothL₁ [9] $\mathcal{L}_{l1}^{smooth}$ and SSIM [34] \mathcal{L}_{SSIM} between the target image and the reconstructed image from LP, both two losses are helpful to preserve the details of the image, which can be represented as: $\mathcal{L}_{total} = \mathcal{L}_{l1}^{smooth} + \lambda * \mathcal{L}_{SSIM}$, where we set the weight of \mathcal{L}_{SSIM} λ to 0.4 empirically.

Method	SD7K												Param (M)
	512 × 512				1024 × 1024				2462 × 3699 (Full Size)				
	PSNR↑	SSIM↑	RMSE↓	Time(s) ↓	PSNR↑	SSIM↑	RMSE↓	Time(s) ↓	PSNR↑	SSIM↑	RMSE↓	Time(s) ↓	
Input	15.95	0.89	44.09	N/A	15.95	0.90	44.09	N/A	15.94	0.91	44.14	N/A	N/A
Wang <i>et al.</i> [32]	15.31	0.82	47.88	> 10	15.31	0.85	47.88	> 30	15.29	0.86	47.95	> 60	N/A
Wang <i>et al.</i> [30]	13.32	0.68	67.48	> 10	13.32	0.71	67.48	> 30	13.26	0.74	68.07	> 60	N/A
Shah <i>et al.</i> [27]	9.89	0.71	86.35	> 10	9.90	0.76	86.35	> 30	9.88	0.79	86.46	> 60	N/A
Jung <i>et al.</i> [16]	19.86	0.92	26.76	> 5	19.86	0.93	26.76	> 10	19.82	0.92	26.86	> 30	N/A
MaskGAN [15]	24.82	0.87	15.43	0.96	24.82	0.85	15.43	2.07	24.67	0.86	15.72	7.59	28.29
DHAN [6]	25.61	0.85	14.27	0.21	25.60*	0.83*	14.27*	N/A	25.42*	0.85*	14.60*	N/A	27.82
AEFNet [8]	24.18	0.95	16.83	0.11	23.00*	0.90*	19.14*	N/A	22.94*	0.90*	19.28*	N/A	391.10
BEDSR-Net [24]	21.50	0.90	30.52	0.85	19.81*	0.85*	33.03*	N/A	19.74*	0.86*	33.18*	N/A	29.44
ShadowFormer [10]	23.71	0.90	17.54	0.15	22.69*	0.84*	19.65*	N/A	22.64*	0.87*	19.79*	N/A	11.35
BMNet [39]	24.86	0.80	15.59	0.18	24.84*	0.79*	15.67*	N/A	24.70*	0.83*	15.96*	N/A	2.11
Ours	28.69	0.97	9.98	0.24	28.68	0.97	9.98	0.75	28.67	0.96	10.00	7.93	29.34

Table 2. Quantitative results of the ablation study on three datasets in different resolutions. The best result is highlighted in red and bold. The asterisk (*) indicates that the result is obtained from a low-resolution result (512×512) and then improved using a super-resolution technique [21] to address an out-of-memory error.

5. Experiments

5.1. Implementation Details

We evaluate the performance of different network structures with different resolutions on the proposed SD7K dataset, where the batch size is set to 4. We also give more experiments on the low-resolution of 512×512 of Jung [16], Kligler [17], and SD7K for fair comparison following previous methods, we give more details in the supplementary materials. To further utilize the dataset, we also incorporate a simple shadow synthesis pipeline to further increase the diversity of the training process. We investigate several potential options, including alpha composition [26] and Perlin noise generation [11] for comparison, and we empirically choose alpha composition to synthesize shadows through the shadow-free image and the mask on the fly via random alpha values in the range between [0.2, 0.7].

As for the evaluation metrics, we employ three widely-used metrics: Root Mean Square Error (RMSE), Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity (SSIM). While PSNR and SSIM are commonly used in image restoration evaluation [3, 5, 22], RMSE is the most important and frequently used metric specifically for shadow removal [6, 10].

5.2. Comparisons with State-of-the-Art Methods

As presented in Tab. 2, we compared FSENet with both traditional and learning-based models of various resolutions using publicly available implementations of these methods, where the proposed method outperforms all other methodologies in most metrics.

Furthermore, our analysis reveals that traditional methods, such as [16, 27], rely on physics-based assumptions that are not always correct, resulting in unstable performance across different datasets as shown in Fig. 6. Deep learning models, on the other hand, are highly dependent on

the dataset magnitude, and their performance is relatively consistent across different scenarios.

We also find that some deep learning models that are effective in low-resolution shadow removal do not scale well to high-resolution documents. They either compromise the aspect ratio of the document or lead to out-of-memory errors. While traditional methods can still remove shadows, they are time-consuming, requiring pixel-by-pixel calculations of illumination difference [27] and background estimation [30]. At present, only Mask-ShadowGAN [15] and the proposed FSENet can effectively remove document shadows in high-resolution scenarios. Despite having more parameters and requiring more time to produce results, FSENet outperforms other methods quantitatively.

5.3. Ablation Studies

As presented in Tab. 3 and Fig. 7, we conduct our ablation studies at the original high-resolution documents, examining the effectiveness of several components in our algorithm. These components include the depth of the Laplacian Pyramid, the efficiency of TRM, DFE, and DAT modules, and the effectiveness of our shadow synthesis pipeline.

We first evaluate the depth of the proposed Laplacian Pyramid. From Tab. 3, with the depth of the Laplacian Pyramid increases (Our w/ Depth=3), the magnitude of the low-frequency components decreases, affecting the low-frequency deshadowing and the reconstructed outcome. Consequently, the model’s performance progressively declines as the depth increases. We then evaluate the performance of the proposed DFE and DAT modules, which are also responsible for low-frequency deshadowing, primarily influencing the pixel-wise variation of the final result. Although each module individually contributes to enhancing shadow elimination performance, their combined network produces the best results. As for the high-frequency modules, the TRM module handles high-frequency components

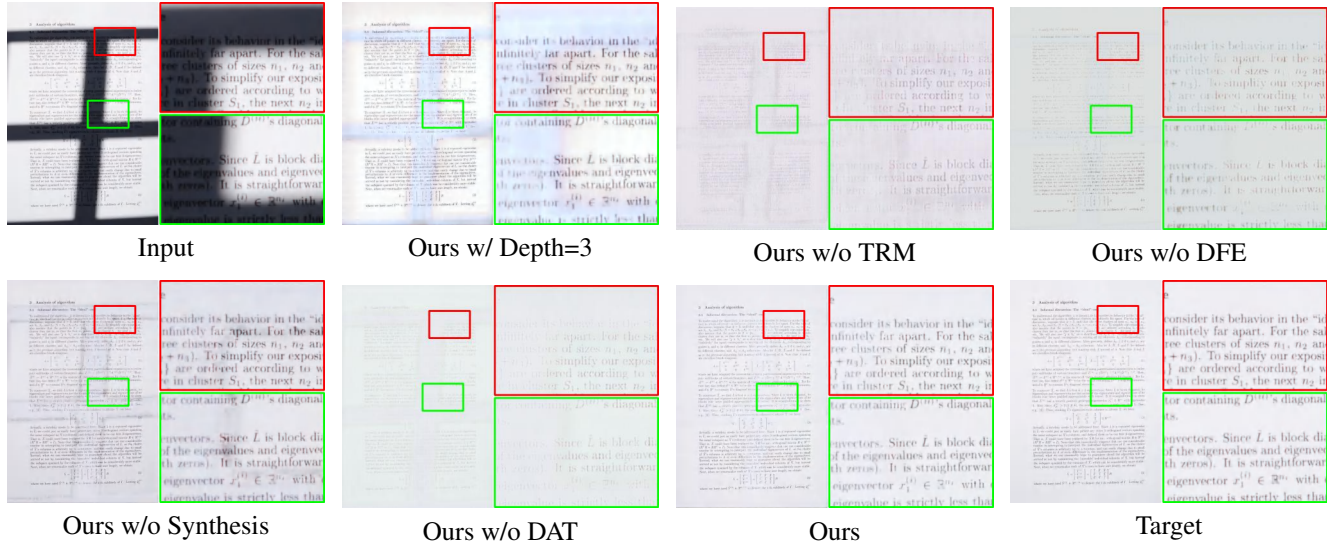


Figure 7. Qualitative results of the ablation studies on different modules of the proposed FSENet.

Ablation Study	Jung (Original)			Kliger (Original)			SD7K (Original)		
	PSNR \uparrow	SSIM \uparrow	RMSE \downarrow	PSNR \uparrow	SSIM \uparrow	RMSE \downarrow	PSNR \uparrow	SSIM \uparrow	RMSE \downarrow
Input	12.93	0.79	61.34	13.25	0.81	56.75	15.94	0.91	44.14
Ours w/ Depth=3	20.97	0.83	23.19	26.33	0.91	13.66	20.83	0.93	24.02
Ours w/o TRM	19.50	0.81	27.60	27.10	0.92	12.78	18.44	0.86	35.92
Ours w/o DFE	21.90	0.84	21.02	25.16	0.92	14.89	19.70	0.88	29.07
Ours w/o DAT	22.68	0.85	19.28	22.70	0.88	20.19	19.03	0.87	32.94
Ours w/o Synthesis	22.21	0.85	21.33	24.79	0.92	15.54	22.94	0.93	21.93
Ours Full	23.60	0.85	17.56	28.98	0.93	10.60	28.67	0.96	10.00

Table 3. Quantitative results of the ablation study on three datasets in high-resolution. The best result is highlighted in red and bold.

such as contours and edges, thus contributing significantly to the overall efficacy of our algorithm. Finally, we also evaluate the importance of the proposed shadow synthesis pipeline, where w/o shadow synthesis reduces the performance a lot, especially on the two small datasets.

5.4. Limitations

Although our model is capable of removing shadows in high-resolution images, the current network is not able to run on real-time or on edge devices, such as mobile phones, due to the larger parameters and heavy GPU computations. In the future, we will explore different approaches to further reduce the required parameters, making it usable on smart devices.

6. Conclusion

In this paper, we tackle the problem of removing shadows from high-resolution documents. To achieve this goal, we consider the contribution of the dataset and the network. As for the dataset, we gather a high-resolution dataset SD7K, which is comprised of over 7000 triplets of real-

world document images with distinct characteristics under various illumination conditions. This dataset is an order of magnitude greater than previous datasets and allows us to validate our approach in a more diverse and comprehensive manner. As for the network structure, we propose FSENet, a frequency-aware method to remove shadows in higher resolution via frequency decomposition. To separate the high-frequency contours from the low-resolution content, we employ a Laplacian Pyramid and a series of meticulously designed attention blocks, enabling each component to be learned separately. Powered by our dataset and network, our method achieves a new state-of-the-art for document shadow removal. We believe both the network structure and the dataset will benefit the community for research on this topic.

Acknowledgements. This work is supported in part by the University of Macau under Grant MYRG2022-00190-FST, in part by the Science and Technology Development Fund, Macau SAR, under Grant 0034/2019/AMJ, Grant 0087/2020/A2 and Grant 0049/2021/A.

References

- [1] Steve Bako, Soheil Darabi, Eli Shechtman, Jue Wang, Kalyan Sunkavalli, and Pradeep Sen. Removing shadows from images of documents. In *ACCV*, pages 173–183. Springer, 2016. [2](#), [3](#)
- [2] Peter J Burt and Edward H Adelson. The laplacian pyramid as a compact image code. In *Readings in computer vision*, pages 671–679. Elsevier, 1987. [2](#), [4](#)
- [3] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *ECCV*, pages 17–33. Springer, 2022. [4](#), [5](#), [7](#)
- [4] Xuhang Chen, Xiaodong Cun, Chi-Man Pun, and Shuqiang Wang. Shadocnet: Learning spatial-aware tokens in transformer for document shadow removal. In *IEEE ICASSP*, pages 1–5. IEEE, 2023. [2](#)
- [5] Xiaodong Cun and Chi-Man Pun. Improving the harmony of the composite image by spatial-separated attention module. *IEEE Transactions on Image Processing*, 29:4759–4771, 2020. [7](#)
- [6] Xiaodong Cun, Chi-Man Pun, and Cheng Shi. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting gan. In *AAAI*, pages 10680–10687, 2020. [2](#), [6](#), [7](#)
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. [2](#)
- [8] Lan Fu, Changqing Zhou, Qing Guo, Felix Juefei-Xu, Hongkai Yu, Wei Feng, Yang Liu, and Song Wang. Auto-exposure fusion for single-image shadow removal. In *CVPR*, pages 10571–10580, 2021. [1](#), [2](#), [7](#)
- [9] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. [6](#)
- [10] Lanqing Guo, Siyu Huang, Ding Liu, Hao Cheng, and Bihan Wen. Shadowformer: Global context helps image shadow removal. In *AAAI*, 2023. [2](#), [7](#)
- [11] John C Hart. Perlin noise pixel shaders. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*, pages 87–94, 2001. [7](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015. [2](#), [6](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [2](#), [5](#)
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141, 2018. [6](#)
- [15] Xiaowei Hu, Yitong Jiang, Chi-Wing Fu, and Pheng-Ann Heng. Mask-shadowgan: Learning to remove shadows from unpaired data. In *ICCV*, pages 2472–2481, 2019. [1](#), [2](#), [6](#), [7](#)
- [16] Seungjun Jung, Muhammad Abul Hasan, and Changick Kim. Water-filling: An efficient algorithm for digitized document shadow removal. In *Asian Conference on Computer Vision*, pages 398–414. Springer, 2018. [1](#), [2](#), [3](#), [6](#), [7](#)
- [17] Netanel Kliger, Sagi Katz, and Ayellet Tal. Document enhancement using visibility detection. In *CVPR*, pages 2374–2382, 2018. [1](#), [2](#), [3](#), [7](#)
- [18] Hieu Le and Dimitris Samaras. From shadow segmentation to shadow removal. In *ECCV*, August 2020. [2](#)
- [19] Hieu Le and Dimitris Samaras. From shadow segmentation to shadow removal. In *ECCV*, pages 264–281. Springer, 2020. [2](#)
- [20] Hieu Le and Dimitris Samaras. Physics-based shadow image decomposition for shadow removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9088–9101, 2021. [2](#)
- [21] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, pages 4681–4690, 2017. [7](#)
- [22] Jingtang Liang, Xiaodong Cun, Chi-Man Pun, and Jue Wang. Spatial-separated curve rendering network for efficient and high-resolution image harmonization. In *European Conference on Computer Vision*, pages 334–349. Springer, 2022. [7](#)
- [23] Jie Liang, Hui Zeng, and Lei Zhang. High-resolution photo-realistic image translation in real-time: A laplacian pyramid translation network. In *CVPR*, pages 9392–9400, 2021. [2](#), [4](#), [5](#)
- [24] Yun-Hsuan Lin, Wen-Chin Chen, and Yung-Yu Chuang. Bedsr-net: A deep shadow removal network from a single document image. In *CVPR*, pages 12905–12914, 2020. [2](#), [3](#), [7](#)
- [25] Zhihao Liu, Hui Yin, Yang Mi, Mengyang Pu, and Song Wang. Shadow removal by a lightness-guided network with training on unpaired data. *IEEE Transactions on Image Processing*, 30:1853–1865, 2021. [1](#), [2](#)
- [26] Jonathan Shade, Steven Gortler, Li-wei He, and Richard Szeliski. Layered depth images. In *SIGGRAPH*, pages 231–242, 1998. [7](#)
- [27] Vatsal Shah and Vineet Gandhi. An iterative approach for shadow removal in document images. *IEEE ICASSP*, pages 1892–1896, 2018. [1](#), [3](#), [7](#)
- [28] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(11):2298–2304, 2016. [1](#)
- [29] Jin Wan, Hui Yin, Zhenyao Wu, Xinyi Wu, Yanting Liu, and Song Wang. Style-guided shadow removal. In *ECCV*, pages 361–378. Springer, 2022. [2](#)
- [30] Bingshu Wang and CL Philip Chen. An effective background estimation method for shadows removal of document images. In *IEEE International Conference on Image Processing (ICIP)*, pages 3611–3615. IEEE, 2019. [1](#), [3](#), [7](#)
- [31] Jifeng Wang, Xiang Li, and Jian Yang. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *CVPR*, pages 1788–1797, 2018. [2](#)

- [32] Jian-Ren Wang and Yung-Yu Chuang. Shadow removal of text document images by estimating local and global background colors. *IEEE ICASSP*, pages 1534–1538, 2020. 1, 3, 7
- [33] Tao Wang, Kaihao Zhang, Tianrun Shen, Wenhan Luo, Bjorn Stenger, and Tong Lu. Ultra-high-definition low-light image enhancement: A benchmark and transformer-based method. *AAAI*, 2023. 4, 5
- [34] 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. 6
- [35] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *CVPR*, pages 17683–17693, 2022. 2, 5
- [36] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pre-training of text and layout for document image understanding. In *SIGKDD*, pages 1192–1200, 2020. 1
- [37] Yimin Xu, Mingbao Lin, Hong Yang, Ke Li, Yunhang Shen, Fei Chao, and Rongrong Ji. Shadow-aware dynamic convolution for shadow removal. *arXiv preprint arXiv:2205.04908*, 2022. 6
- [38] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *ICLR*, 2015. 6
- [39] Yurui Zhu, Jie Huang, Xueyang Fu, Feng Zhao, Qibin Sun, and Zheng-Jun Zha. Bijective mapping network for shadow removal. In *CVPR*, pages 5627–5636, June 2022. 7