
Improving the Learning Capability of Small-size Image Restoration Network by Deep Fourier Shifting

Man Zhou

Aerospace Information Research Institute, Chinese Academy of Sciences
University of Science and Technology of China

Abstract

State-of-the-art image restoration methods currently face challenges in terms of computational requirements and performance, making them impractical for deployment on edge devices such as phones and resource-limited devices. As a result, there is a need to develop alternative solutions with efficient designs that can achieve comparable performance to transformer or large-kernel methods. This motivates our research to explore techniques for improving the capability of small-size image restoration standing on the success secret of large receptive field.

Targeting at expanding receptive field, spatial-shift operator tailored for efficient spatial communication and has achieved remarkable advances in high-level image classification tasks, like S^2 -MLP [1] and ShiftVit [2]. However, its potential has rarely been explored in low-level image restoration tasks. The underlying reason behind this obstacle is that image restoration is sensitive to the spatial shift that occurs due to severe region-aware information loss, which exhibits a different behavior from high-level tasks. To address this challenge and unleash the potential of spatial shift for image restoration, we propose an information-lossless shifting operator, i.e., Deep Fourier Shifting, that is customized for image restoration. To develop our proposed operator, we first revisit the principle of shift operator and apply it to the Fourier domain, where the shift operator can be modeled in an information-lossless Fourier cycling manner. Inspired by Fourier cycling, we design two variants of Deep Fourier Shifting, namely the amplitude-phase variant and the real-imaginary variant. These variants are generic operators that can be directly plugged into existing image restoration networks as a drop-in replacement for the standard convolution unit, consuming fewer parameters. Extensive experiments across multiple low-level tasks including image denoising, low-light image enhancement, guided image super-resolution, and image de-blurring demonstrate consistent performance gains obtained by our Deep Fourier Shifting while reducing the computation burden. Additionally, ablation studies verify the robustness of the shift displacement with stable performance improvement.

1 Introduction

The spatial shift operator [3] is a technique that shifts channels from a pixel to its adjacent pixels. It is celebrated for its efficient facilitation of spatial information exchange. Due to its parameter-free nature and computational efficiency, this operator has found widespread application as a substitute for standard convolution units, particularly in high-level image classification tasks. A representative work, S^2 -MLP [1], integrates the spatial shift operator into a channel-mixing MLP framework. This approach serves as an alternative to token-mixing MLPs, effectively mitigating intrinsic overfitting issues and significantly enhancing recognition accuracy. ShiftVit [2] explores the role of the self-attention mechanism in Vision Transformers (ViTs) for high-level tasks. It proposes a replacement using a modified spatial shift operation [4] that facilitates the exchange of a subset of channels




Figure 1: **Comparison between the spatial shift operator and the proposed deep Fourier shift operator.** (a) Traditional spatial shift operator involves a spatial shift mechanism that moves each channel of the input tensor in a distinct spatial direction, thus suffering from severe region-aware information loss and conflicting with the requirements of image restoration tasks. (b) Deep Fourier Shifting/Cycling operator is a more ingenious information-lossless operator, which is tailored for image restoration tasks. (c), (d) Deep Fourier shifting achieves a more stable performance gain than the spatial shifting mechanism with varying “ns” shift displacements and “n” basic units over image de-noising task where the cut-off is for compressing the vertical axis scale to better illustrate the contrast effect clearly.

between neighboring features. Despite these impressive achievements, the utility of spatial shift operators in low-level image restoration tasks remains unexplored.

The main difficulty in introducing the spatial shift mechanism to low-level vision tasks lies in its inherent conflict with the objectives of image restoration tasks, as depicted in Figure 1. In particular, this mechanism operates by moving each channel of its input tensor in distinct spatial directions, leading to a mixing of spatial information across channels. Notably, image restoration is fundamentally a standard regression problem where both features and channels play critical roles in determining the final output. This contrasts with the behavior observed in high-level vision tasks. A key limitation of the spatial shift operator is its inherent loss of region-aware information, where the regions subject to shifting are filled with zero values. Consequently, image restoration tasks, which are particularly sensitive to this kind of information loss, can experience a decline in performance due to spatial shifting.

In response to these challenges, we introduce a novel deep shifting operator specifically crafted for low-level image restoration tasks, named **Deep Fourier Shifting**. This operator revisits and enhances the fundamental principles of the traditional shift operator by extending its application into the Fourier domain. Here, the shift operation is reformulated as an *information-lossless* process executed through Fourier cycling. Deep Fourier Shifting is composed of two distinct variants, each featuring three key elements: a 2D discrete Fourier transform, Fourier cycling rules, and a 2D inverse Fourier transform. This operator is designed to be generic, and seamlessly integrable into existing image restoration architectures as a replacement for standard convolution units, offering the added advantage of reduced parameter usage. To assess the effectiveness of Deep Fourier Shifting, we undertake comprehensive experiments across a spectrum of low-level image restoration tasks. These include image denoising, low-light image enhancement, guided image super-resolution, and image de-blurring. Our experiments consistently reveal performance improvements achieved through the integration of Deep Fourier Shifting, while alleviating computational burden. Furthermore, we conduct ablation studies to evaluate the robustness of Deep Fourier Shifting against varying shift displacements. These studies demonstrate its ability to consistently enhance performance, underscoring its stability and effectiveness.

Our results suggest that Deep Fourier Shifting is a promising tool for image restoration, showing potential for varied real-world applications. We hope that Deep Fourier Shifting could contribute




Figure 2: **(a) The information-lossless cycling mechanism.** The discrete Fourier transform of a signal exhibits period-extended and cycling properties. Specifically, in the Fourier domain, the two pixels in sequence beginning and end may not appear adjacent, but due to the period property, they are actually considered adjacent, as indicated by the upper right corner. This inherent period-extended and cycling behavior of the Fourier transform enables us to model the shifting mechanism in a manner that is information-lossless, making it well-suited for image restoration tasks. Consider the Fourier transform of a discrete time-domain signal, represented as $\begin{pmatrix} 0 & 1 & 2 \\ 3 & 4 & 5 \end{pmatrix}$. It may appear that the values 3 and 5 are not adjacent within the main period. However, owing to the property of period extension, the 3 from the previous period and the 5 from the current period are theoretically considered adjacent. It is reasonable to move the removed area from the end to the front, meeting the cycling mechanism. **(b) Our deep Fourier shifting operator.** Our operator borrows the principle of the spatial shifting mechanism and models the shifting mechanism in information-lossless Fourier cycling rules. The cycling is coded as 2D queue rolling.

to advancements in neural network designs for image restoration, particularly in improving spatial communication interactions.

2 Deep Fourier Shifting

Building on the principle of the shift operator, we extend its application to the Fourier domain. In this domain, the shift operator is conceptualized as an information-lossless operation, which we term Fourier cycling. To substantiate this approach, we present a theorem along with its corresponding proof. Additionally, we introduce two distinct variants of Deep Fourier shifting: i) the magnitude and phase variant, and ii) the real and imaginary variant, which are derived from the transformation rules we have identified within the Fourier domain.

Definitions. $f(x, y) \in \mathbb{R}^{H \times W \times C}$ is the spatial signal and $F(u, v) \in \mathbb{R}^{H \times W \times C}$ denotes its Fourier transform where (x, y) and (u, v) represent the space coordinates and Fourier spectrum, respectively.

Theorem. The Fourier transform of a discrete signal is a period-extended and cycling: $F(u, v) = F(u + nH, v) = F(u, v + mW) = F(u + nH, v + mW)$ where $u = 0, 1, 2, \dots, H - 1$, $v = 0, 1, 2, \dots, W - 1$ and $n, m \in \mathbb{N}$. $\mathbb{N}$ is the set of positive integers starting from zero.

2.1 Proof: The Fourier Transform of a Discrete Signal is Period-Extended and Cycling

We show the periodicity and cycling properties of the Fourier transform of a discrete signal, as illustrated in Figure 2(a). Note that the Fourier transform $F(u, v)$ of $f(x, y)$ is expressed as

$$F(u, v) = \frac{1}{HW} \sum_{x=0}^{H-1} \sum_{y=0}^{W-1} f(x, y) e^{-j2\pi(\frac{ux}{H} + \frac{vy}{W})}. \quad (1)$$

Then, we show the periodicity of $F(u, v) \in \mathbb{R}^{H \times W}$ with H and W . It means $F(u, v) = F(u + nH, v) = F(u, v + mW) = F(u + nH, v + mW)$ where $u = 0, 1, 2, \dots, H - 1$,

$v = 0, 1, 2, \dots, W-1$ and $n, m \in \mathbb{N}$ that records the set of non-negative integers. We take the $F(u, v) = F(u + nH, v + mW)$ for example and recall Eq. (1) as

$$\begin{aligned} & F(u + nH, v + mW) \\ &= \frac{1}{HW} \sum_{x=0}^{H-1} \sum_{y=0}^{W-1} f(x, y) \cdot e^{-j2\pi(\frac{(u+nH)x}{H} + \frac{(v+mW)y}{W})} \\ &= \frac{1}{HW} \sum_{x=0}^{H-1} \sum_{y=0}^{W-1} f(x, y) e^{-j2\pi(\frac{ux}{H} + \frac{vy}{W})} e^{-2j\pi m} e^{-2j\pi n} \\ &= F(u, v) e^{-2j\pi m} e^{-2j\pi n}, \end{aligned} \quad (2)$$

where for any integer z , it has $e^{-2j\pi z} = 1$.

Further, $e^{-2j\pi n} = 1$ and $e^{-2j\pi m} = 1$ for $n, m \in \mathbb{N}$. Therefore,

$$F(u + nH, v + mW) = F(u, v) e^{-2j\pi m} e^{-2j\pi n} = F(u, v). \quad (3)$$

Similarly, we can prove the periodicity of $F(u, v)$ as well.

$$\begin{aligned} F(u, v) &= F(u + nH, v + mW) \\ &= F(u + nH, v) = F(u, v + mW). \end{aligned} \quad (4)$$

Furthermore, deep Fourier transform can be expressed in Cartesian and polar coordinates by an equivalent form as

$$F(u, v) = Ae^{jP} = a + bj. \quad (5)$$

The period-extended and cycling property holds over the amplitude-phase and real-imaginary format.

2.2 Architectural Design

Recall the **Theorem-1**, we propose two deep Fourier shifting variants: amplitude-phase variant and real-imaginary variant.

Amplitude-phase shifting variant. This variant is illustrated in Figure 2(b). The pseudo-code is shown in Figure 3 (left). Given an image $X \in \mathbb{R}^{H \times W \times C}$, we first adopt the Fourier transform $\text{FFT}(X)$ to obtain its amplitude component A and phase component P . We then evenly split the generated A and P in 4 folds A_g and P_g by the channel dimension, perform the Fourier cycling over A_g and P_g in both the H and W dimensions:

$$\begin{aligned} A_g[0] &= \text{torch.roll}(A_g[0], \text{dim}=1, s) \\ A_g[1] &= \text{torch.roll}(A_g[1], \text{dim}=1, -s) \\ A_g[2] &= \text{torch.roll}(A_g[0], \text{dim}=2, s) \\ A_g[3] &= \text{torch.roll}(A_g[0], \text{dim}=2, -s), \end{aligned} \quad (6)$$

where $\text{torch.roll}(\cdot)$ accounts for the cycling function by dim parameter for shifting dimension and s for shifting displacement. The transformed A_g and P_g are then fed into two independent convolution modules with 1×1 kernel and followed by the inverse Fourier transform $\text{iFFT}(\cdot)$ to project the shifting ones back to spatial domain.

Real-imaginary shifting variant. The pseudo-code for the real-imaginary shifting variant is presented on the right side of Figure 3. In this variant, we perform Fourier cycling separately on the real component a and the imaginary component b , while keeping the remaining processing steps the same as the amplitude-phase shifting variant.

Concerning two variants, the first one entails trigonometric function calculations, wherein minor numerical alterations could potentially result in computational instability in engineering applications. However, it offers more accurate physical interpretations over amplitude-phase operation from signal processing perspective.

<pre> def DFS_AP(X): # X: input with shape [N, C, H, W] # A and P are the amplitude and phase A.e{jp} = FFT(X) # Fourier shifting transform rules A_g = torch.spilt(A, 4, dim=1) P_g = torch.spilt(P, 4, dim=1) A_fc = Fourier-cycling(A_g) P_fc = Fourier-cycling(P_g) A_fc = Convs_1x1(A_fc) P_fc = Convs_1x1(P_fc) # Inverse Fourier transform Y = iFFT(A_fc, P_fc) Return Y #[N, C, H, W] </pre>	<pre> def DFS_ab(X): # X: input with shape [N, C, H, W] # a and b are the real and imaginary part a+bj = FFT(X) # Fourier shifting transform rules a_g = torch.spilt(a, 4, dim=1) b_g = torch.spilt(b, 4, dim=1) a_fc = Fourier-cycling(a_g) b_fc = Fourier-cycling(b_g) A_fc = Convs_1x1(a_fc) P_fc = Convs_1x1(b_fc) # Inverse Fourier transform Y = iFFT(a_fc, b_fc) Return Y #[N, C, H, W] </pre>
---	--

Figure 3: **Pseudo-code of the two variants of the proposed deep Fourier shifting.** The left is the *amplitude-phase variant* while the right is the *real-imaginary variant*.

Table 1: Quantitative comparisons on low-light image enhancement. The arrow $\rightarrow$ denotes the generalization setting by training on the data before the arrow and testing directly on the data after the arrow.

Model	Config	LOL $\rightarrow$		$\rightarrow$ Huawei		Huawei $\rightarrow$		$\rightarrow$ LOL		#Paras
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
DRBN	Original	19.7931	0.8361	17.7929	0.6247	20.1549	0.6851	18.0856	0.7543	0.55M
	Shift-sa	19.7072	0.8343	17.6221	0.6071	20.2165	0.6873	17.9112	0.7532	0.41M
	Fcycle-AP	22.4274	0.8448	19.3252	0.6472	<u>20.5855</u>	<u>0.6872</u>	<u>18.8666</u>	<u>0.7587</u>	0.41M
	Fcycle-ab	<u>22.2054</u>	<u>0.8429</u>	<u>19.3125</u>	<u>0.6431</u>	20.6651	0.6876	19.1535	0.7681	0.41M
SID	Original	20.1062	0.7895	16.5874	0.5925	20.1742	0.6659	18.5468	0.7441	7.76M
	Shift-sa	20.0148	0.7911	16.8214	0.5911	20.1517	0.6651	18.4998	0.7434	7.53M
	Fcycle-AP	22.8565	0.8019	<u>19.1707</u>	<u>0.6238</u>	<u>20.9068</u>	0.6708	18.8161	0.7494	7.53M
	Fcycle-ab	<u>22.6313</u>	<u>0.7995</u>	19.2471	0.6242	20.9271	<u>0.6691</u>	<u>18.5741</u>	<u>0.7443</u>	7.53M

Table 2: Comparisons on image denoising.						Table 3: Comparisons on image deblurring.					
Dataset	Metric	DNCNN				Dataset	Metric	DeepDeblur			
		Original	Shift-sa	Fcycle-AP	Fcycle-ab			Original	Shift-sa	Fcycle-AP	Fcycle-ab
SIDD	PSNR	37.1992	37.2247	<u>37.6891</u>	38.1837	GoPro	PSNR	28.9423	28.9037	29.2123	<u>29.1939</u>
	SSIM	0.8954	0.8980	0.9013	<u>0.9066</u>		SSIM	0.8716	0.8712	0.8777	<u>0.8765</u>
DND	PSNR	38.33	38.42	38.69	<u>38.65</u>	HIDE	PSNR	26.9770	26.9991	27.2860	<u>27.2547</u>
	SSIM	0.8974	0.8985	<u>0.942</u>	0.949		SSIM	0.8468	0.8476	0.8541	<u>0.8536</u>
	#Paras	1.51M	1.43M	1.43M	1.43M		#Paras	11.72M	10.61M	10.61M	10.61M

3 Experiments

3.1 Experimental Settings

Image enhancement. We evaluate our Fourier shifting operator on two popular image enhancement benchmarks: LOL [5] and Huawei [6]. The LOL dataset consists of 500 low-/normal-light image pairs. Following the original setting, we use 485 pairs for training and 15 pairs for testing. The Huawei dataset contains 2480 paired images, with 2200 pairs for training and 280 pairs for testing. We compare with the representative approaches SID [7] and DRBN [8].

Image deblurring and denoising. For the image deblurring task, we employ DeepDeblur [9] in our experiments. We use the GoPro dataset [9] for training. To demonstrate the generalizability of our operator, we also apply the model trained on the GoPro dataset directly to the test images of the HIDE dataset [10]. For image denoising, we use the SIDD dataset [11] as the training benchmark.

Table 4: Quantitative comparisons on guided image super-resolution.

Model	Config	WorldView-II				GaoFen2				WorldView-III			
		PSNR	SSIM	SAM	ERGAS	PSNR	SSIM	SAM	ERGAS	PSNR	SSIM	SAM	ERGAS
PANNET	Original	40.8172	0.9630	0.0257	1.0555	42.1699	0.9569	0.0192	0.9565	29.68	0.907	0.085	3.426
	Shift-sa	40.8791	0.9631	0.0255	1.0511	42.2107	0.9577	0.0185	0.9549	29.32	0.897	0.103	3.734
	Fcycle-AP	<u>41.2633</u>	<u>0.9650</u>	<u>0.0242</u>	<u>1.0080</u>	42.6361	0.9680	0.0173	0.9005	<u>30.46</u>	<u>0.915</u>	<u>0.078</u>	<u>3.253</u>
	Fcycle-ab	41.3184	0.9671	0.0238	1.0032	<u>42.5594</u>	<u>0.9648</u>	<u>0.0185</u>	<u>0.9118</u>	30.55	0.918	0.077	3.187
MutNet	Original	41.4967	0.9692	0.0232	0.9781	47.1069	0.9883	0.0106	0.5626	30.59	0.924	0.0741	3.0798
	Shift-sa	41.5227	0.9695	0.0230	0.9766	47.1371	0.9886	0.0098	0.5617	29.87	0.912	0.082	3.325
	Fcycle-AP	41.6686	0.9712	0.0228	0.9598	47.4225	0.9890	0.0101	0.5453	<u>30.71</u>	<u>0.926</u>	<u>0.0737</u>	<u>3.0767</u>
	Fcycle-ab	<u>41.6579</u>	<u>0.9701</u>	<u>0.0228</u>	<u>0.9611</u>	<u>47.3489</u>	<u>0.9887</u>	<u>0.0104</u>	<u>0.5479</u>	30.74	0.926	0.0737	3.0764

Performance evaluation is conducted on the remaining validation samples from the SIDD dataset and the DND benchmark dataset [12]. The selected baseline for comparison is DnCNN¹ [13].

Guided image super-resolution. We adopt the pan-sharpening task as a representative task of guided image super-resolution. The WorldView II and GaoFen2 datasets [14] are used in our experiments. The baselines are state-of-the-art INNformer [15] and SFNet [14].

In evaluating the performance of different approaches, we use image quality assessment metrics such as the relative dimensionless global error in synthesis (ERGAS) [16], the peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), the spectral angle mapper (SAM) [17].

3.2 Implementation Details

Based on the above competitive baselines, we perform the comparison over the following configurations by replacing the standard convolution with the spatial or Fourier shifting operator:

- 1) **Original**: the baseline without any changes;
- 2) **Fcycle-AP**: replacing the original model's standard convolution operator with the amplitude-phase variant of Deep Fourier shifting;
- 3) **Fcycle-ab**: replacing the original model's standard convolution operator with the real-imaginary variant of Deep Fourier shifting;
- 4) **Shift-sa**: replacing the variants of Deep Fourier shifting in the settings of 2)/3) with the spatial shifting operator;
- 5) **Shift-ns**: replacing "ns" standard convolution operator with the spatial shifting or Fourier shifting-ab/AP operator;
- 6) **Shift-n**: replacing the spatial shifting or Fourier shifting-ab/AP operator in the settings of 5) with varying shifting displacement "n".

3.3 Comparison and Analysis

Quantitative comparison. We evaluate model performance across various configurations as detailed in the implementation section (Section 3.2). The quantitative results of this analysis are systematically presented in Tables 1, 2, 3, and 4. In these tables, the best and second-best results are highlighted in bold and underlined, respectively. Values that fall below the established baseline are marked with a gray background for clear distinction. We observe a uniform trend of performance enhancement across all the tasks and datasets tested when our two novel deep Fourier shifting variants are integrated. This improvement is particularly noteworthy as it is achieved with reduced computational costs, suggesting the efficiency and effectiveness of our approach. For instance, in the low-light image enhancement baseline SID (Table 1), our "Fcycle-AP" and "Fcycle-ab" models surpass the "Original" model by achieving higher PSNR values of 2.75dB/2.53dB and 2.6dB/2.7dB on the LOL and Huawei datasets, respectively. In contrast, the integration of the spatial shifting operator into the SID baseline results in a slight reduction in PSNR by 0.1dB/0.02dB on the LOL and Huawei datasets, as evident in the 3rd and 7th rows. These findings validate the superiority of our proposed Fourier shifting approach, demonstrating its suitability and effectiveness in addressing the unique challenges of low-level image restoration tasks. Additionally, the integration of our proposed methods has shown to enhance the training performance, in Figure 4.

¹We followed the official code to reproduce the results <https://github.com/cszn/KAIR>.




Figure 5: Visual comparison over image enhancement.




Figure 6: The effectiveness of information preservation. **Left:** we compare mutual information levels before and after employing Fcycle-ab and Shift-sa operators on the LOL test set, respectively. Our operator exhibits significantly higher mutual information than Shift-sa, showcasing its efficacy in information preservation. **Right:** we visualize feature maps and their amplitude components before and after operations. This demonstrates that our Fcycle-AP promotes frequency information in the low domain.




Figure 4: The proposed operators improve the training performance. It shows the training PSNR on the image enhancement task on the LOL and Huawei datasets in the top and bottom.

on the LOL test set. The statistical analysis, illustrated in Figure 6 (left), demonstrates a signifi-

Qualitative comparison. Due to space constraints, we present only a subset of the visual results in Figure 5, which effectively illustrate the efficacy of our proposed deep Fourier shifting operator. Additional visual results are available in the supplementary materials. As demonstrated in these figures, the integration of the deep Fourier shifting operators (Fcycle-ab and Fcycle-AP) with the original baseline models yields results that are visually more appealing compared to the baselines and the spatial shift operator (Shift-sa). Models enhanced with our operators exhibit a superior ability to restore fine texture details and mitigate degradation effects. In contrast, models using the spatial shift operator tend to produce significant artifacts.

Information preservation. To offer a deeper understanding of the efficacy of our frequency cycling mechanism, we begin by comparing the mutual information of features before and after the application of our proposed Fcycle-AP operator

Model	Config	Fcycle-ab		Fcycle-AP		Shift-sa		#Paras
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
Original	#0	37.1992	0.8954	37.1992	0.8954	37.1992	0.8954	1.51M
Shift-ns	#2	37.2309	0.8958	37.3781	0.8983	37.2247	0.8980	1.43M
	#4	37.3981	0.8974	37.4832	0.8981	<u>37.4256</u>	<u>0.8994</u>	1.35M
	#6	37.6964	0.9011	37.5891	0.9013	37.5091	0.8993	1.28M
	#12	<u>38.0837</u>	<u>0.9064</u>	37.4676	0.9003	37.2996	0.8983	1.06M
	#18	38.2810	0.9081	<u>37.5251</u>	<u>0.9044</u>	23.6867	0.3340	0.84M
Shift-n	#1	37.2309	0.8958	<u>37.3781</u>	<u>0.8983</u>	37.2247	0.8980	1.43M
	#2	37.2663	0.8961	<u>37.3551</u>	<u>0.8978</u>	<u>37.3379</u>	<u>0.8995</u>	1.43M
	#3	37.4277	0.8985	37.3139	0.8974	37.4009	0.8998	1.43M
	#4	<u>37.3781</u>	<u>0.8977</u>	37.4754	0.9001	37.1247	0.8909	1.43M

Table 5: Ablation studies of image denoising network, DNCNN.

cantly higher level of mutual information as outlined in [18] for our method when compared to the conventional shifting operation in the spatial domain. This notable increase in mutual information underscores the strength of our Fourier cycling mechanism in preserving information and minimizing information loss.

Moreover, an analysis of the features and their corresponding amplitude components, as depicted in Figure 6 (right), reveals distinct grid effects in the feature maps generated by the Shift-sa operator. In contrast, our method substantially mitigates these grid effects. This is achieved by facilitating interactions among frequency information and improving spatial communication. Together, these observations compellingly show the efficiency of our proposed Fourier cycling mechanism in both preserving information and enhancing frequency interaction.

3.4 Ablation Studies




Figure 7: The effect of shifting displacement shift-n on SID.

evaluation are provided in Table 5. We highlighted the best and second-best results in bold and underline, respectively. A visual depiction of these results is provided in Figure 1 (c). A key observation from these results is that as “ns” increases, models incorporating our proposed operators not only

With varying “n”. To assess the robustness of our method, we conduct a comparative analysis with the image enhancement baseline SID and the image denoising baseline DNCNN. This comparison involved varying the shifting displacement “n”. The corresponding quantitative results are presented in Figure 7 and Table 5. The results clearly show that integrating our proposed Fourier shifting solutions into these baseline models consistently yields better performance than the original baselines alone. Conversely, when spatial shifting is employed within these baselines, there is a marked decline in performance compared to their original versions, particularly at larger shifting displacements (e.g., “n=4”). These findings not only validate the efficacy of our approach but also highlight its distinct robustness over traditional spatial shifting.

With varying “ns”. We examine the robustness of our deep Fourier shifting variants by comparing them with the image de-noising baseline DNCNN. For this comparison, we substitute the variable “ns” standard convolution units in the baseline with both the spatial shifting operator or our proposed deep Fourier shifting variants, as detailed in Section 3.2. The results of this

exhibit consistent performance improvements but also achieve these enhancements with fewer model parameters. In contrast, models using the spatial shifting operator display a notable decrease in performance, particularly when “ns” is increased to 18. This trend demonstrates the robustness and efficiency of our proposed Fourier shifting solutions. To further elucidate this point, values falling below the baseline are marked with a gray background, highlighting their relevance and impact.

4 Related Work

Spatial-shifting operator. The spatial-shift operator, initially proposed by Wu et al.[3], enhances efficient spatial information communication. It has gained traction in high-level image classification tasks [1], where they innovated upon MLP Mixer [19]. They replaced the spatial-wise token mixer with a spatial-shift operation, enabling inter-patch communication and effectively addressing overfitting challenges associated with spatial-specific token-mixing MLPs. Similarly, Wang et al. [2] explored the role of attention mechanisms in Vision Transformers (ViTs). They introduced a partial shift operation, exchanging a subset of channels among adjacent features. This simple yet effective operation led to the development of ShiftViT, a backbone network that replaces traditional attention layers in ViTs with shift operations. These approaches illustrate the adaptability and versatility of spatial-shift operators in enhancing deep learning models for image analysis.

Deep Fourier transform over image restoration. There has been a growing interest in integrating the Fourier transform to refine deep learning-based image restoration models. A notable example is the work of Zhou et al. [20], which re-examines the interplay between spatial and Fourier domains. This study uncovers transformation rules applicable to various resolution features within the Fourier domain, leading to the development of a theoretically sound Deep Fourier Up-Sampling method applicable across multiple restoration tasks. Similarly, the research in Zhou et al. [14] investigates the degradation dynamics of guided image super-resolution. The authors propose a novel spatial-frequency dual-domain integration network tailored to this specific task. The idea leads to a deep Fourier-based exposure correction network designed to tackle exposure-related issues. These pioneering studies harness the synergistic potential of deep learning and the intrinsic attributes of the Fourier transform.

5 Limitations

Our study acknowledges that there is room for more comprehensive experiments and exploration of representative baselines on broader computer vision tasks. It is important to note that this work represents the initial endeavor to delve into the utilization of shifting mechanisms and the development of tailored deep Fourier shifting operators specifically for low-level image restoration tasks. Moreover, our focus extends beyond designing a plug-and-play module for integration into existing networks to achieve performance improvements. We aim to provide a powerful and efficient spatial communication interaction choice by offering an alternative to the basic convolution operator pool when constructing new models from scratch.

Broader Impact

Our research has the potential to facilitate efficient image restoration on edge devices and resource-limited platforms, democratizing access to high-quality capabilities. This benefits critical domains like mobile photography, remote sensing, and medical imaging. Additionally, reduced computational requirements contribute to energy efficiency, extending battery life and reducing environmental impact. We emphasize responsible development, evaluation, and deployment to ensure equitable and ethical use of our methods.

6 Conclusion

We have presented Deep Fourier Shifting, a shifting operator grounded in solid theoretical principles, specifically tailored for image restoration tasks. This operator, characterized by its information-lossless nature, leverages the Fourier cycling approach to shift operations. A key feature of our Deep Fourier Shifting is its versatility; it can be seamlessly integrated as a replacement for standard convolution units in existing image restoration networks, offering the added benefit of reduced parameter usage. Our experimental evaluations have consistently demonstrated the enhancements in performance attributable to the incorporation of our Deep Fourier Shifting.

References

- [1] Tan Yu, Xu Li, Yunfeng Cai, Mingming Sun, and Ping Li. S2-mlp: Spatial-shift mlp architecture for vision. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 297–306, January 2022.
- [2] Guangting Wang, Yucheng Zhao, Chuanxin Tang, Chong Luo, and Wenjun Zeng. When shift operation meets vision transformer: An extremely simple alternative to attention mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2423–2430, 2022.
- [3] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, and Kurt Keutzer. Shift: A zero flop, zero parameter alternative to spatial convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9127–9135, 2018.
- [4] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7083–7093, 2019.
- [5] Wenhan Yang Jiaying Liu Chen Wei, Wenjing Wang. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference*. British Machine Vision Association, 2018.
- [6] Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Fengzhu Zou, Fang Lin, and Songchen Han. R2rnet: Low-light image enhancement via real-low to real-normal network. *arXiv preprint arXiv:2106.14501*, 2021.
- [7] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3291–3300, 2018.
- [8] Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, and Jiaying Liu. From fidelity to perceptual quality: A semi-supervised approach for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3063–3072, 2020.
- [9] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3883–3891, 2017.
- [10] Ziyi Shen, Wenguan Wang, Xiankai Lu, Jianbing Shen, Haibin Ling, Tingfa Xu, and Ling Shao. Human-aware motion deblurring. In *IEEE International Conference on Computer Vision*, 2019.
- [11] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1692–1700, 2018.
- [12] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1586–1595, 2017.
- [13] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.
- [14] Zhou Man, Huang Jie, Yan Keyu, Yu Hu, Fu Xueyang, Liu Aiping, Wei Xian, and Zhao Feng. Spatial frequency domain information integration for pan-sharpening. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–250, 2022.
- [15] Man Zhou, Xueyang Fu, Jie Huang, Feng Zhao, Aiping Liu, and Rujing Wang. Effective pan-sharpening with transformer and invertible neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–15, 2022.

- [16] L. Alparone, L. Wald, J. Chanussot, C. Thomas, P. Gamba, and L. M. Bruce. Comparison of pansharpening algorithms: Outcome of the 2006 grs-s data fusion contest. *IEEE Transactions on Geoscience and Remote Sensing*, 45(10):3012–3021, 2007.
- [17] A. F. Goetz J. R. H. Yuhas and J. M. Boardman. Discrimination among semi-arid landscape endmembers using the spectral angle mapper (sam) algorithm. *Proc. Summaries Annu. JPL Airborne Geosci. Workshop*, pages 147–149, 1992.
- [18] Tyrone E. Duncan. On the calculation of mutual information. *SIAM Journal on Applied Mathematics*, 19(1):215–220, 1970.
- [19] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems*, 34:24261–24272, 2021.
- [20] Hu Yu, Jie Huang, Feng Zhao, Jinwei Gu, Chen Change Loy, Deyu Meng, Chongyi Li, et al. Deep fourier up-sampling. *Advances in Neural Information Processing Systems*, 35:22995–23008, 2022.

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Abstract

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Section 4.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Section 2.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Section 3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: Codes will be released after acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: The detailed descriptions are provided in Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: Section 3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.

- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Section 3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have read and followed the NeurIPS Code of Ethics

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Broader Impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.