

# Image Segmentation by Latent Space Phase-Gating with Applications in High-Content Screening

Jiadong Yu<sup>1</sup> and Rahul Singh<sup>1,2</sup>

<sup>1</sup> Department of Computer Science, University of Iowa, Iowa City, Iowa 52242-1419

<sup>2</sup> Center for Discovery and Innovation in Parasitic Diseases, UC San Diego, CA 92093  
rahul-singh@uiowa.edu

**Abstract.** Schistosomiasis is a parasitic disease with significant global health and socio-economic implications. Drug discovery for schistosomiasis typically involves high-content whole-organism screening. In this approach, parasites are exposed to various chemical compounds and their systemic, whole-organism-level responses are captured via microscopy and analyzed to obtain a quantitative assessment of chemical effect. These effects are multidimensional and time-varying, impacting shape, appearance, and behavior. Accurate identification of object boundaries is essential for preparing images for subsequent analysis in high-content studies. Object segmentation is one of the most deeply studied problems in computer vision where recent efforts have incorporated deep learning. Emerging results indicate that acquiring robust features in spectral domain using Fast Fourier Transform (FFT) within Deep Neural Networks (DNNs) can enhance segmentation accuracy. In this paper, we explore this direction further and propose a latent space Phase-Gating (PG) method that builds upon FFT and leverages phase information to efficiently identify globally significant features. While the importance of *phase* in analyzing signals has long been known, technical difficulties in calculating phase in manners that are invariant to imaging parameters has limited its use. A key result of this paper is to show how phase information can be incorporated in neural architectures that are compact. Experiments conducted on complex HCS datasets demonstrate how this idea leads to improved segmentation accuracy, while maintaining robustness against commonly encountered noise (blurring) in HCS. The compactness of the proposed method also makes it well-suited for application specific architectures (ASIC) designed for high-content screening.

**Keywords:** Fast Fourier Transform, Deep Neural Networks, biological imaging, parasitic diseases, whole-organism screening, drug discovery.

## 1 Introduction

### 1.1 Background

Schistosomiasis is a parasitic disease with global health and socio-economic impacts. Current estimates suggest that more than 200 million individuals are presently infected, with over 700 million people at risk across 78 countries. The primary method of

treatment revolves around the administration of the drug praziquantel (PZQ). However, its efficacy is mainly targeted at the adult stage of the worm's life cycle [1], and instances of drug resistance have been documented [2]. Consequently, both the World Health Organization (WHO) and the National Institutes of Health (NIH) have classified schistosomiasis as a condition urgently necessitating new treatment modalities.

Drug discovery for schistosomiasis (and other helminthic diseases) typically involves conducting whole organism phenotypic screening. Unlike target-based screening, which uses a purified protein target to identify drug candidates, phenotypic screening involves the entire disease system. This method entails exposing live parasites to various compounds and then recording and analyzing the resulting multidimensional and systemic phenotypic changes to assess the efficacy of the compounds and identify potential new drugs. Phenotypic screenings are inherently holistic and have been shown to be more effective than target-based screening [3]. The development of phenotype-analysis methods in this field has garnered considerable attention, particularly since the pioneering work in [4].

Developing automated, high content phenotypic screening techniques for schistosomiasis require segmentation to separate individual parasites. Subsequently, their appearance, shape, and motion can be characterized at any desired time point while being tracked across the video. A significant technical obstacle lies in devising image segmentation methods that precisely delineate each parasite and accurately identify their boundaries. The lack of accurate segmentation can increase downstream measurement errors of the parasite phenotypes.

In recent years, deep learning based segmentation approaches [5] have garnered widespread attention and demonstrated exceptional performance. However, the application of such methods to real-time segmentation often necessitates a trade-off between speed and accuracy, posing challenges for use in HCS. Furthermore, HCS can operate at a broad frame rate ranging between 10 to 100 frames per second (FPS). Compared to the standard 24 FPS videos, HCS recordings can be particularly intricate. The development of segmentation methods that fulfill these multifaceted requirements is imperative within the domain of HCS.

## 1.2 Problem formulation

For making precise measurements of complex parasite phenotypes, the initial segmentation process must possess certain characteristics. Firstly, it should be insensitive to image artifacts introduced by the culture medium. Secondly, it must accurately identify boundaries, particularly those touching or overlapping parasites. Lastly, the segmentation process should demonstrate real-world robustness, particularly against blurry images. Additionally, consideration needs to be given to computation speed, as slower methods take significantly longer to produce results, thus potentially limiting the number of compounds that can be analyzed. Readers are directed to the figure in the Results section illustrating the aforementioned segmentation challenges.

## 2 Related work

The first novel algorithm for segmenting somules – the juvenile stage of the parasite causing schistosomiasis was proposed by Asarnow and Singh [6]. This method employed phase congruency (PC) [7] – the fact that at edge pixels, all Fourier components of the image signal tend to be *in-phase*, to accurately identify touching or overlapping edges, achieving excellent accuracy compared to extant segmentation methods. Hereafter, we refer to this method as the ASA method. Due to their reliance on capturing certain fundamental edge characteristics, methods such as ASA may struggle when the positioning of multiple objects in an image, render it difficult to determine the fundamental edge characteristics being determined. On the other hand, deep neural network (DNN) based segmentation methods take a different approach by taking advantage of the supervised step in learning the patterns that need to be found and thus circumventing the challenge more fundamentally oriented methods like ASA have to deal with. One such example is Mask R-CNN [5], which has demonstrated high segmentation accuracy in many computer vision tasks. These deep learning methods, however, require large training datasets to perform well and can suffer from underfitting and overfitting with limited training data.

A related problem of cell segmentation, has been an active research area for some time. Cellpose [8], which employed a convolutional neural network (CNN) and adopted the network architecture from U-net [9], incorporated a “style” representation in the smallest convolutional maps to obtain “style” vectors and aimed to learn intrinsic features. It demonstrated higher segmentation accuracy than Mask R-CNN using 70,000 objects for training. Nevertheless, CNNs tend to focus on local receptive fields and are thus sensitive to image noise. In domain generalization, Deep Frequency Filtering (DFF) was purposed to learn robust features through Fast Fourier Transforms (FFT). It extended the use of Fast Fourier Convolution (FFC) [10] and added a spatial attention module (SAM) [11], demonstrating the ability to learn domain-generalizable features. In this paper, we focus not on domain generalization but on utilizing the robust features learned through FFT to improve segmentation accuracy. While both FFC and DFF utilized real and imaginary components directly in FFT, we have found that utilizing phase directly results in improved segmentation. Our research shows how deep learning can be used to take advantage of an old idea in image and signal analysis – the importance of phase in signals [12].

## 3 Methods

### 3.1 Core Idea

The image signal in the spatial domain can be transformed into the spectral domain using the Fast Fourier Transform (FFT), and subsequently decomposed into its real and imaginary components. Studies [10, 13] have demonstrated that concatenating these components, followed by CNN, Batch Normalization (BN), and ReLU activation, can utilize global receptive fields in the spectral domain by learning the intercorrelation.

However, it is important to leverage relationships such as the amplitude and phase spectra derive directly from the real and imaginary components rather than relying solely on learned features. This motivates us, as the amplitude spectrum conveys image intensity information, while the phase spectrum reflects highly informative structural details [7], which correlates with human perception of visual information.

### 3.2 Latent Spectrum Representations

Given an image  $X \in \mathbf{R}^{C \times H \times W}$ , where  $C, H, W$  represents the channel, height, and width of the image, respectively. We apply a 2D FFT independently to each channel and get the corresponding spectrum representation  $F_x$ . We denote the FFT and the inverse FFT operations as  $\mathcal{F}, \mathcal{F}^{-1}$ , respectively, and formulate this transform  $F_x = \mathcal{F}(x)$  as follows:

$$F_x(u, v) = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} e^{-j2\pi(\frac{uh}{H} + \frac{vw}{W})} \cdot x(h, w) \quad (1)$$

where  $x(h, w), F_x(u, v)$  denotes the value in the spatial and spectral domain at the coordinates  $(h, w)$  and  $(u, v)$ , respectively. Additionally, we use  $\phi_x, A_x$  to represent the phase spectrum and amplitude spectrum, which are defined as:

$$\phi_x(u, v) = \arctan \left[ \frac{\text{Im}(F_x(u, v))}{\text{Re}(F_x(u, v))} \right] \quad (2)$$

$$A_x = \sqrt{\text{Re}(F_x(u, v))^2 + \text{Im}(F_x(u, v))^2} \quad (3)$$

In Eqs.(2)-(3), Re and Im represents the real and imaginary parts of  $F_x$ . The image or features  $x$  can be recovered directly from  $F_x$  via  $\mathcal{F}^{-1}$  or more intuitive through

$$F_x = \mathcal{F}(x) = A_x e^{j2\pi\phi_x} \quad (4)$$

$$x = \mathcal{F}^{-1}(F_x) = \mathcal{F}^{-1}(A_x e^{j2\pi\phi_x}) \quad (5)$$

For the phase and amplitude spectrum  $\phi_x, A_x$ , there are two notable properties: 1)  $\phi_x$  itself captures structural information, enabling consistent enhancement of global receptive fields. 2)  $A_x$  captures magnitude of informative features, which can be utilized with a simple attention module to regulate the flow.

### 3.3 Latent space Phase-Gating

Our goal is to adaptively modulate phase and amplitude during training. To achieve this, we propose a latent space Phase-Gating (PG) operation to enhance the global features. Given phase  $\phi_x$  and amplitude  $A_x$ , the proposed PG operation is formulated as follows:

$$A'_x = G(A_x) \quad (6)$$

$$\phi'_x = \text{Conv}_{1 \times 1}(\phi_x) \quad (7)$$

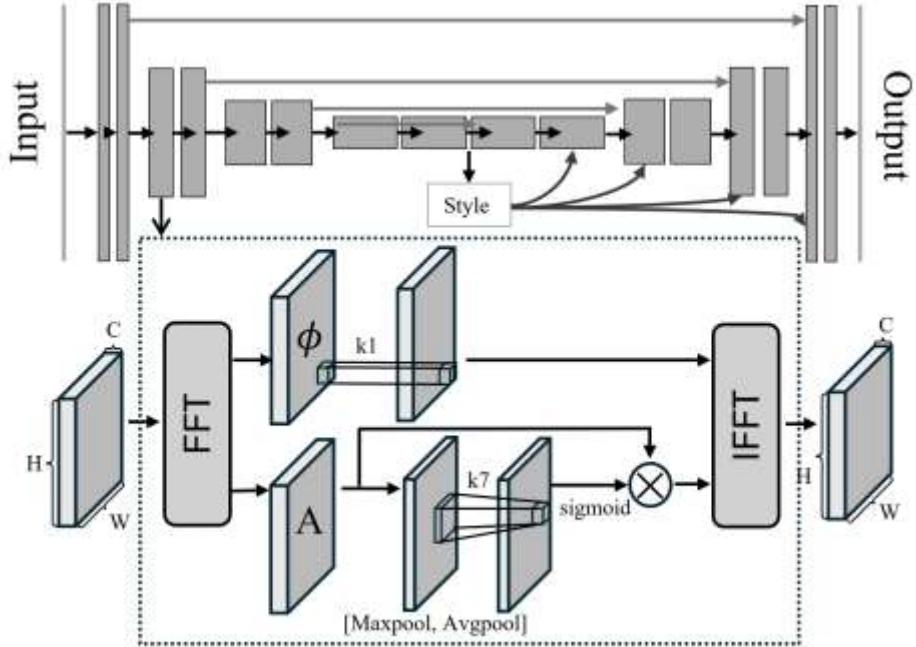
$$F'_x = A'_x e^{j2\pi\phi'_x} \quad (8)$$

where  $G$  refers to a component that regulate the flow of globally important features. In Eq. 7, a simple convolution layer with kernel size  $1 \times 1$  and BN is used to magnify the phase. For operation  $G$ , we adopt a lightweight attention module designed for features learning in the latent space [11].

To clarify the use of this module, we expand Eq. 3. An aggregation of  $A_x$  is performed to obtain  $A_x^{max}$  and  $A_x^{avg}$  using max-pooling and average-pooling operations over all channels. These serve as compressed representation of  $A_x$ . We then concatenate them and apply a convolution layer with kernel size  $7 \times 7$  to access its relative importance with respect to its neighbors. This convolution layer, despite its large kernel size, remains computationally efficient as it only applies to two channels in practice. The result is then passed to a sigmoid function and multiplied back with  $A_x$ . Thus, we can rewrite Eq. 6 as:

$$A'_x = A_x \otimes \sigma(\text{Conv}_{7 \times 7}([A_x^{max}, A_x^{avg}])) \quad (9)$$

where  $\sigma$  is the sigmoid function,  $[,]$  denotes concatenation, and  $\otimes$  is element-wise multiplication.



**Fig. 1.** Illustration of the proposed Phase-Gating (PG) component and the overall network architecture. The PG adaptively modulates phase and amplitude during training to capture globally important features. It involves four key operations: a 2D FFT to convert the input into phase  $\phi$  and amplitude  $A$ , a gating operation to regulate the flow, a simple convolution layer with kernel size  $1 \times 1$  and BN to enhance phases, and a 2D inverse FFT to transform back to the original latent space.

### 3.4 Network architecture with dual-branch features aggregation

The features captured using PG are generally global. Therefore, we use a dual-branch features aggregation method from [10] to ensemble both local and global features. This architecture has been proven effective in [10, 13]. It uses a split-and-merge concept to capture complementary features with different receptive fields and intentionally exchange information between them. Given features  $X \in \mathbf{R}^{C \times H \times W}$ , it is first split along the features channels into local and global part, denoted as  $X^l, X^g$ , respectively, where  $X = [X^l, X^g]$  with  $X^l \in \mathbf{R}^{(1-\alpha)C \times H \times W}$  and  $X^g \in \mathbf{R}^{\alpha C \times H \times W}$ . Similarly, output features  $Y \in \mathbf{R}^{C \times H \times W}$  can be split into  $Y^l, Y^g$ , where  $Y = [Y^l, Y^g]$  with  $Y^l \in \mathbf{R}^{(1-\beta)C \times H \times W}$  and  $Y^g \in \mathbf{R}^{\beta C \times H \times W}$ . Both  $\alpha, \beta$  are in range between 0 and 1, controlling the degree of complementarity. This dual-branch features aggregation can be formulated as:

$$Y^l = f_l(X^l) + f_{g \rightarrow l}(X^g) \quad (10)$$

$$Y^g = f_g(X^g) + f_{l \rightarrow g}(X^l) \quad (11)$$

where  $f$  denote four different transformation functions.  $f_g$  corresponds to spectral transform implemented in [10] while other three are regular convolutions. We evaluate the proposed PG by integrating it into the spectral transform branch.

For the overall network architecture, we build upon the Cellpose architecture [8], which extends from U-net [9]. U-net is compact and efficient for producing binary mask and separating foreground/background pixels. However, it was not initially designed to separate different instances of the same category. Cellpose adds two additional layers along the mask prediction process to compute the horizontal and vertical gradients. These three components are then combined to produce the final results. Additionally, Cellpose utilizes a “style” vector, defined as the average pooling at the smallest feature map with normalization. This vector is subsequently used until the last convolutional layer, except during the upsampling operation.

## 4 Experiments

### 4.1 Data

The proposed method was evaluated with a comprehensive chemical-phenotype dataset described in [14]. The data included 5779 parasites from the training and validation set, and the test set, referred to as the ASA set, was the same as used in [6], and contained 7015 parasites. A summary of the dataset is provided in Table 1. The datasets were obtained by hand-labeling on a high-resolution touch-screen device. On an average, it took about 20 minutes to label one image, although some images required significantly longer labeling time due to the number of parasites present in it.

Additionally, evaluating deep learning-based segmentation methods on images corrupted by different types of noise allows us to gain an understanding of real-world performance. A number of noise models have been considered in the literature [15], including: blur, noise, digital, and weather corruptions, each with five severity levels controlling the corruption intensities. In this paper, we focus on blurring, such as motion

blur, defocus blur, and Gaussian blur. This is due to the wells in which the parasites were located during screening imaged have depth, not all parasites were co-planar. Consequently, some parasites appear blurred and testing the segmentation accuracy on blurred images provides an indication of real-world robustness.

**Table 1. Summary of the data**

	Training and Validation	ASA
Total parasites	5779	7015
Image size	1040 x 1388	1040 x 1388
Parasites in each image	14 - 68	23 - 56
The size of parasites	1494 - 7383	942 - 6672
Perimeter of parasites	144 - 400	118 - 410
Proportion of the bbox filled	21% - 86%	20% - 83%
Grayscale mean intensity	94 - 205	79 - 211
Exposure time in days	1, 2, 3, 4	1, 3, 7
Concentrations in $\mu M$	0.01, 0.1, 1, 10	0, 0.01, 0.1, 1
Compound	Acepromazine, Alimemazine, Chlorophenothiazine, Hycanthone, Methiothepin Mesylate, Mevastatin, Pravastatin, Praziquantel (PZQ), Promazine, Rosuvastatin, Triflupromazine	Acepromazine, Alimemazine, Chlorophenothiazine, K777, Pravastatin Praziquantel (PZQ), Promazine, Rosuvastatin, Simvastatin, Control

## 4.2 Evaluation Metrics

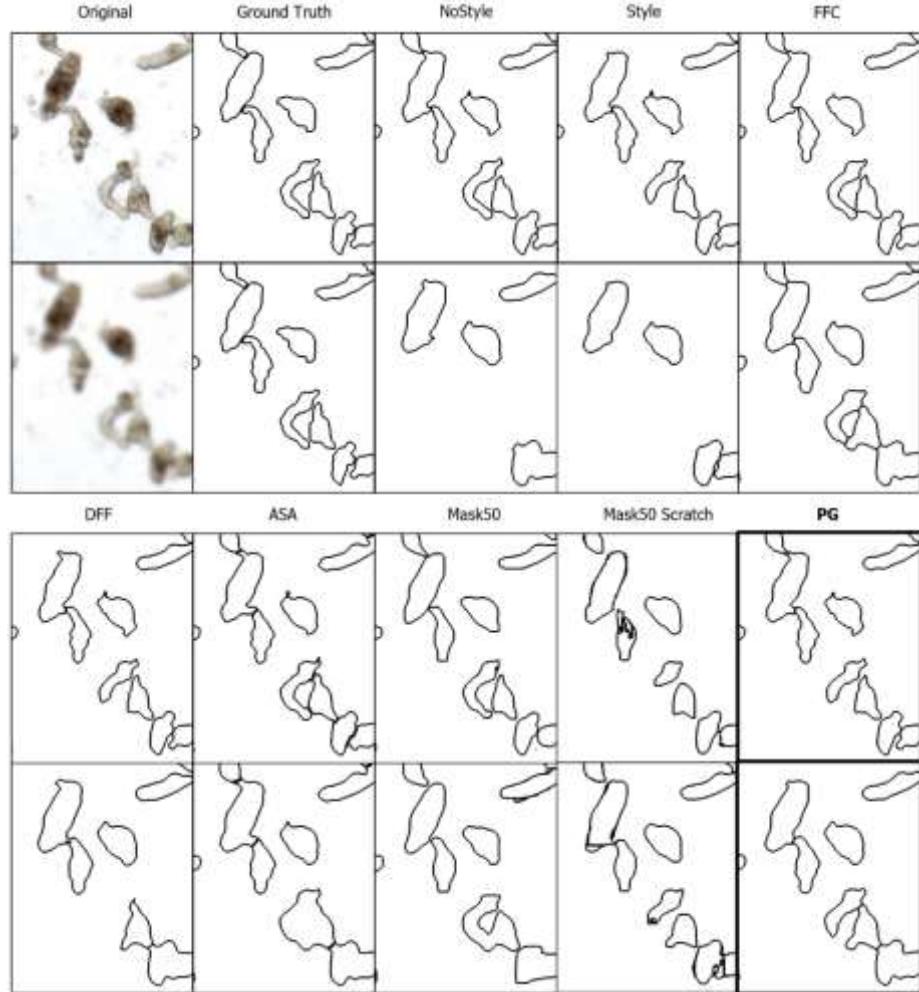
Following [6], we use the following measures to assess segmentation accuracy.

1. Precision, Recall, and F1 values for the foreground pixels.
2. Precision, Recall, and F1 values for the edge pixels.
3. Object-count agreement. Defined as the fractional object count discrepancy  $O = |N_s - N_o|/N_o$ , where  $O$  is the discrepancy,  $N_o$  is the true number of objects and  $N_s$  is the number of objects in each trial segmentation.
4. Mean deviation of object boundaries. Estimated by calculating the average absolute difference between the Euclidean distance transform of the ground truth object boundaries and the boundaries of the objects from the trial segmentation.
5. Degradation. Defined as the fraction between the performance of one method relative to a reference.

Each metric is computed as the average value across all images in the test set, and this value is weighted by the number of objects present in each ground truth image. This approach ensures a comprehensive and nuanced evaluation, taking into consideration the varying complexities of the ground truth images.

### 4.3 Implementation

We implemented the ASA methods using Matlab and the PG method using Python with PyTorch. We utilized the codebases provided by the authors of Cellpose [8] and FFC [10]. Additionally, we implemented the DFF [13] based on the description in the paper, as no codebase was provided. We also used the implementations of Mask R-CNN [5] from MMDetection frameworks [16].



**Fig. 2.** Segmentation results for parasites subjected to Praziquantel (PZQ) with 1 and 7 days of exposure. Each column displays results obtained from the original image, the ground truth labels, and the segmentation methods, respectively. The rows represent clean image and image with defocused blur at severity level 3.

#### 4.4 Training parameters

We trained two versions of Mask R-CNN: one using pretrained weights and the other trained from scratch. Both versions were trained using the AdamW optimizer with a learning rate of 0.001, a weight decay of 0.01, and a batch size of 2. The learning rate was reduced by a factor of 10 at 8, 15, and 20 epochs, for a total of 24 epochs. Both versions employed multi-scaling training methods based on the size of the objects in the image.

The Cellpose architecture, along with regular convolution, FFC, DFF and PG, were trained with SGD optimizer with a learning rate of 0.01, a momentum of 0.9, a weight decay of 0.0001, a batch size of 8, and a total of 200 epochs. The learning rate started at zero and was linearly increased to 0.01 over the first ten epochs. At epoch 160, the learning rate was reduced by a factor of 2 every ten epochs. For FFC, DFF, and PG, an  $\alpha, \beta$  value of 0.5 was used to evenly split the local and global channels.

#### 4.5 Results

We initially tested the effectiveness of our proposed latent space Phase-Gating (PG) on the ASA set. The results can be found in Table 2. It was observed that Mask R-CNN with pretraining yielded the best performance, while Mask R-CNN trained from scratch showed the worst performance. Our proposed PG method outperformed all other variations of the Cellpose architecture. However, it's worth noting that Mask R-CNN is almost 20 times larger than the Cellpose architecture, making it significantly slower. An illustration of the segmentation outcomes is depicted in Figure 2.

**Table 2. Evaluation measures on the ASA set**

Model Name	Precision	Recall	F1	Edge Precision	Edge Recall	Edge F1	OCA	MBD
Style*	0.969	0.932	0.950	0.304	0.296	0.300	0.018	2.282
NoStyle*	0.968	0.935	0.951	0.302	0.294	0.298	0.020	2.294
FFC*	0.961	0.939	0.950	0.291	0.289	0.290	0.021	2.147
DFF*	0.964	0.936	0.949	0.293	0.289	0.291	0.019	2.165
PG*	0.966	0.938	0.951	0.297	0.293	0.295	0.016	2.049
ASA	0.898	0.967	0.930	0.169	0.172	0.170	0.049	2.918
Mask	0.951	0.955	0.953	0.276	0.278	0.277	0.015	1.656
Mask Scratch	0.934	0.879	0.905	0.197	0.214	0.205	0.119	5.082

\* indicates Cellpose architecture, Mask means Mask R-CNN with ResNet 50

All the deep learning-based segmentation methods attained good F1 scores for edge pixels, with the Cellpose architectures being the most effective. On the other hand, the ASA methods achieved the lowest F1 scores for edge pixels due to their tendency to segment the outer boundaries of the parasites instead of the center. This could be an issue as it could alter the segmented shape, particularly when dealing with parasites that are in contact with image artifacts introduced by the culture medium.

We also evaluated these methods against the ASA set with blurring. The Cellpose architecture with regular convolution and “style” representation was used as a reference method. We calculated the average degradation of all methods across all five severity

levels in comparison to the reference method. The results can be found in Table 3. DFF appears to be excessively sensitive and therefore performed poorly compared to FFC and PG. Cellpose with no style representation performed slightly worse across all the blur types. Mask R-CNN with pretraining produced the best degradation results, while Mask R-CNN trained from scratch showed results similar to FFC and PG. On the other hand, ASA methods performed slightly worse than FFC and PG.

**Table 3. Evaluation degradation on the ASA set with image blur**

Model Name	Motion Blur			Defocus Blur			Gaussian Blur		
	MBD	F1	Edge F1	MBD	F1	Edge F1	MBD	F1	Edge F1
Style*	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
NoStyle*	1.010	1.008	0.997	1.045	1.065	1.009	1.071	1.067	1.008
FFC*	0.730	1.009	1.007	0.490	0.834	0.996	0.615	0.887	1.002
DFF*	0.851	1.036	1.005	0.735	0.973	1.009	0.739	0.980	1.012
PG*	0.752	0.996	0.998	0.562	0.928	1.027	0.662	0.932	1.017
ASA	0.817	1.072	1.009	0.671	1.080	1.072	0.819	1.179	1.129
Mask	0.551	0.725	0.942	0.481	0.855	0.992	0.574	0.904	1.023
Mask Scratch	0.734	1.033	1.007	0.519	0.861	0.965	0.842	1.137	1.008

## 5 Conclusion

This paper describes latent space Phase-Gating (PG), which leverages phase information and gates the amplitude flow in the spectral domain, thereby improving the enhancement of globally significant features. In the context of specific spectral domain component, PG, while being half the size, delivers higher segmentation accuracy compared to FFC and DFF applied to phenotypic drug screening data. Furthermore, robustness analysis against blurring, using average degradation as a metric, reveals that PG delivers results comparable to FFC and surpasses other Cellpose architectures. The findings presented in this paper underscore the potential of the proposed methods. Its compact design and fast inference capabilities make it particularly desirable for high-content screening applications and ASIC hardware implementations.

**Acknowledgements.** The authors thank Conor R. Caffrey for the screening data from which was reported in [14]. This work was funded in part by the grants IIS 1817239 (NSF) and AI146719 (NIH).

## References

1. Vale, N., Gouveia, M.J., Rinaldi, G., Brindley, P.J., Gartner, F., Correia Da Costa, J.M.: Praziquantel for Schistosomiasis: Single-Drug Metabolism Revisited, Mode of Action, and Resistance. *Antimicrob Agents Chemother.* 61, e02582-16 (2017). <https://doi.org/10.1128/AAC.02582-16>

2. Fallon, P.G.: Schistosome resistance to praziquantel. *Drug Resistance Updates*. 1, 236–241 (1998). [https://doi.org/10.1016/S1368-7646\(98\)80004-6](https://doi.org/10.1016/S1368-7646(98)80004-6)
3. Swinney, D.C., Anthony, J.: How were new medicines discovered? *Nat Rev Drug Discov.* 10, 507–519 (2011). <https://doi.org/10.1038/nrd3480>
4. Singh, R., Pittas, M., Heskia, I., Xu, F., McKerrow, J., Caffrey, C.R.: Automated image-based phenotypic screening for high-throughput drug discovery. In: 2009 22nd IEEE International Symposium on Computer-Based Medical Systems. pp. 1–8. IEEE, Albuquerque, NM, USA (2009)
5. He, K., Gkioxari, G., Dollar, P., Girshick, R.: Mask R-CNN. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 2980–2988. IEEE, Venice (2017)
6. Asarnow, D.E., Singh, R.: Segmenting the Etiological Agent of Schistosomiasis for High-Content Screening. *IEEE Trans. Med. Imaging.* 32, 1007–1018 (2013). <https://doi.org/10.1109/TMI.2013.2247412>
7. Kovesi, P.: Image Features from Phase Congruency. *Videre: Journal of Computer Vision Research*. 1, 1–26 (1999)
8. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. *Nat Methods.* 18, 100–106 (2021). <https://doi.org/10.1038/s41592-020-01018-x>
9. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W.M., and Frangi, A.F. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 234–241. Springer International Publishing, Cham (2015)
10. Chi, L., Jiang, B., Mu, Y.: Fast Fourier Convolution. In: *Advances in Neural Information Processing Systems*. pp. 4479–4488. Curran Associates, Inc. (2020)
11. Woo, S., Park, J., Lee, J.-Y., Kweon, I.S.: CBAM: Convolutional Block Attention Module. In: Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. pp. 3–19. Springer International Publishing, Cham (2018)
12. Oppenheim, A.V., Lim, J.S.: The importance of phase in signals. *Proc. IEEE.* 69, 529–541 (1981). <https://doi.org/10.1109/PROC.1981.12022>
13. Lin, S., Zhang, Z., Huang, Z., Lu, Y., Lan, C., Chu, P., You, Q., Wang, J., Liu, Z., Parulkar, A., Navkal, V., Chen, Z.: Deep Frequency Filtering for Domain Generalization. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11797–11807. IEEE, Vancouver, BC, Canada (2023)
14. Lee, H., Moody-Davis, A., Saha, U., Suzuki, B.M., Asarnow, D., Chen, S., Arkin, M., Caffrey, C.R., Singh, R.: Quantification and clustering of phenotypic screening data using time-series analysis for chemotherapy of schistosomiasis. *BMC Genomics.* 13, S4 (2012). <https://doi.org/10.1186/1471-2164-13-S1-S4>
15. Hendrycks, D., Dietterich, T.: Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. Presented at the International Conference on Learning Representations September 27 (2018)
16. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open MMLab Detection Toolbox and Benchmark, <http://arxiv.org/abs/1906.07155>, (2019)