Guided Augmentation for Monocular Depth Estimation in Cell Microscopy

Abhishek Viswanathan[1], Rajagopalan A.N.[1], Nikhil Yelamarthy[1], Ankit Rai[2], and Pradeep Ramachandran[2]

[1] Indian Institute of Technology Madras, Chennai, India
v.abhishek6678@gmail.com
[2] KLA Corporation

Abstract. Monocular Depth Estimation (MDE) in cell microscopy provides critical insights into cellular structures, with applications spanning cancer diagnostics, hematological analysis, and tumor margin assessment. However, it presents unique challenges such as sparse z-stacks with limited focal planes, optical aberrations degrading depth precision, and the inherently ill-posed nature of inferring depth from single 2D images. Existing MDE methods often rely on semantic priors, geometric modeling, or self-supervised learning. While effective in macroscopic applications, these approaches struggle with microscopy-specific challenges involving domain-specific feature distributions.
To address these limitations, we propose a novel deep learning-based physics-guided augmentation strategy leveraging Extended Depth of Field (EDOF) images to enhance MDE performance. To demonstrate the effectiveness of our approach, we employ a regression model trained to predict z-stack levels from individual cell images and a UNet-based model to synthesize blurred cell images at intermediate z-levels by modeling the point spread function (PSF) of the imaging process. Experiments on Giemsa-stained peripheral blood smear data demonstrate significant improvements in MDE over training without augmentation and simple augmentation strategies. Ablation studies validate the robustness of our approach, providing a promising framework for advancing medical microscopy-related applications.

Keywords: Cell Microscopy · Deep Learning · Augmentation · EDOF · Monocular Depth Estimation

1 Introduction

Monocular Depth Estimation (MDE) is a fundamental problem in computer vision that involves predicting depth information from a single 2D image. In cell microscopy, MDE has applications with respect to medical diagnostics, 3D cell analysis, and microscopic metrology. Accurate depth estimation is critical for tasks such as tumor margin assessment [1] during surgery, morphological analysis for hematological disorders, and therapy monitoring, detecting subtle changes linked to diseases caused through cell swelling or abnormal granules [2].

Unlike stereo depth estimation, which requires paired views, MDE in microscopy offers a simpler and more versatile approach by leveraging the relationship between image blur and depth [3, 4]. However, it poses unique challenges due to the complex physics of imaging systems and inherent data limitations. Image formation in microscopy is governed by PSFs, which describe how a point source of light is imaged through an optical system [5]. Optical aberrations such as chromatic distortions and diffraction effects degrade depth precision near focal planes [6, 7]. While z-stacks provide depth cues for MDE by capturing variations in focus across slices, limited focal planes exacerbate the inherently ill-posed nature of inferring depth from single 2D images [8].

Existing MDE methods often rely on geometric modeling, semantic priors, or self-supervised learning (SSL) approaches to improve generalization. Geometric models incorporate physical relationships between focal planes [9] but often fail to adapt to domain-specific feature distributions in microscopy [10]. Semantic priors leverage scene structure to constrain depth predictions but struggle with unstructured or textureless regions common in cell microscopy [11]. SSL approaches learn from unlabeled data by exploiting geometric constraints or temporal consistency but are limited by scale ambiguity and sparse supervision signals [12–14]. The lack of large-scale annotated datasets tailored for cell microscopy further complicates model training and generalization [15].

To address these limitations, we propose a novel physics-guided data augmentation strategy leveraging Extended Depth of Field images (which combine multiple focal planes into a single in-focus image), to enhance MDE performance in cell microscopy. Our approach focuses on generating realistic blurred cell images at intermediate z-stack levels by modeling the PSF of the imaging process, which bridges gaps in z-stack data and improves training sample diversity while maintaining physical consistency with the imaging system.

The main contributions of this work include: (i) A physics-guided data augmentation strategy using EDOF images to generate realistic blurred cell images at intermediate z-stack levels; (ii) A UNet-based model to learn the point spread function (PSF) of microscopy imaging systems; and (iii) Experiments to demonstrate how this augmentation strategy improves MDE performance when combined with a regression model trained to predict z-stack levels from individual cell images. Our approach is validated through experiments on Giemsa-stained peripheral blood smear data [16], showing significant improvements over training without augmentation as well as conventional augmentation techniques.

2 Related Works

Monocular Depth Estimation (MDE) has been extensively studied, with deep learning significantly advancing the field. Traditional methods relied on handcrafted features and geometric priors, but the integration of Convolutional Neural Networks (CNNs) has enhanced performance. Liu et al. [17] combined CNNs with Conditional Random Fields for depth estimation, while Chang and Wetzstein [18] introduced a deep optics approach integrating optical coding with

end-to-end optimization. Surveys by Rajapaksha et al. [19] and Zhao et al. [20] provide detailed overviews of advancements in MDE.

Depth estimation in microscopy presents unique challenges due to imaging physics and artifacts. Traditional focus-based methods lack robustness, while Shajkofci and Liebling [21] used spatially-variant CNNs for PSF estimation in optical microscopy, improving resolution and depth localization. Imtiaz et al. [22] proposed an Attention UNet for light field microscopy, and Imanishi et al. [23] explored deep learning for depth estimation in light field systems. Ghosh et al. [24] demonstrated DL based depth estimation in turbid media using epi-illuminated microscopy images, showcasing its applicability to specialized imaging scenarios.

Physics-guided approaches incorporate domain-specific knowledge into neural networks. Sun et al. [25] developed a hybrid deep-learning and physics-based model for computational microscopy, while Li et al. [26] proposed a physics-informed denoising diffusion model for image reconstruction.

Data augmentation has improved performance in medical imaging tasks as well. Garcea et al. [27] reviewed augmentation techniques, Welsman et al. [28] demonstrated its impact on electron microscopy classification, and Hussain et al. [15] explored differential augmentation tailored to medical imaging.

Building upon these advancements, our work introduces a physics-guided data augmentation strategy specifically tailored for microscopy, addressing the limitations of sparse z-stacks and domain-specific imaging artifacts. Unlike prior approaches that focus on general depth estimation or light field microscopy, we leverage Extended Depth of Field (EDOF) images to generate physically consistent blurred cell images at intermediate z-stack levels for improving the accuracy of monocular depth estimation in cell microscopy.

3 Methodology

Our methodology addresses the challenges of monocular depth estimation (MDE) in microscopy by employing a physics-guided deep learning approach to bridge gaps in discrete z-stack data and enhance training diversity, tackling the ill-posed nature of MDE and improving generalization. It consists of three key steps: (i) a z-level prediction model, (ii) an auxiliary blurring model, and (iii) an augmented training strategy.

3.1 z-Level Prediction Model

The first component of our methodology is a z-level prediction model, which we refer to as z-Net, that estimates the depth level of an input cell image based on its degree of blur (Figure 1a). Z-Net uses AlexNet architecture, a lightweight and efficient CNN suitable for small-scale datasets. The model takes 84×84 grayscale cell/blob images as input, a size chosen to balance computational efficiency and detail preservation, and outputs a continuous regression value representing the predicted z-level. Ground truth z-levels range from 1 to 10 in steps of 1, corresponding to discrete focal planes in the z-stack.

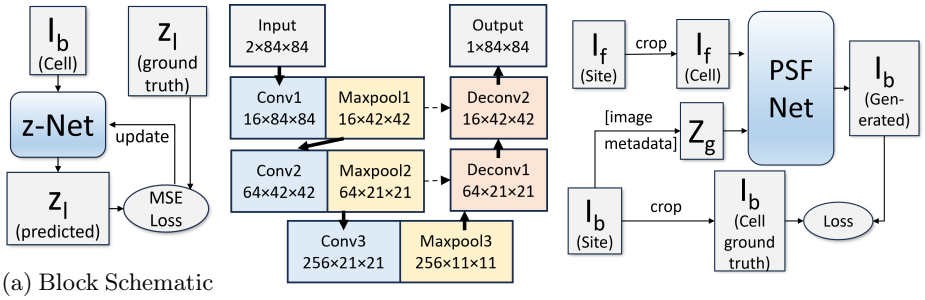

(a) Block Schematic of z-Net. (b) PSF-Net Architecture. (c) Block Schematic of PSF-Net.

Fig. 1: Pipeline of our methodology.

Z-stack imaging uses blur as a cue for depth estimation, offering robustness to lighting and texture variations, and computational efficiency. By regressing z-stack values using z-Net, we aim to approximate MDE in microscopy. The model is trained with a Mean Squared Error (MSE) loss function to minimize the difference between predicted and ground truth z-levels.

3.2 Auxiliary Blurring Model

To overcome the limitations of discrete z-stack levels, we develop an auxiliary blurring model to generate realistic images at intermediate focus levels. Figure 1 shows the training pipeline. Our model uses a U-Net architecture, well-suited for image-to-image translation tasks due to its ability to preserve spatial information via skip connections. The model takes as input: (i) The Extended Depth of Field (EDOF) image, $I_f \in \mathbb{R}^{H\times W}$, where H and W are the image height and width (set to 84 in our case). (ii) A single-channel image, $Z_g \in \mathbb{R}^{H\times W}$, where the desired focal z-level z_l is repeated across the grid.

The U-Net model, $f_{UNet}(\cdot)$, which we refer to as PSF-Net, takes a two-channel input formed by concatenating the EDOF image, I_f, and the z-level grid, Z_g, along the channel dimension, and outputs a synthesized blurred image, I_b, for the specified z-level:

$$I_b = f_{UNet}([I_f, Z_g]) \in \mathbb{R}^{H\times W}.$$

PSF-Net implicitly learns the Point Spread Function (PSF), which models how light spreads through an optical system in 3D space by describing intensity variations along depth. The PSF is mathematically defined as:

$$PSF(x, y, z) = |\mathcal{F}\{\text{ApertureFunction}(x, y, z)\}|^2,$$

where $\mathcal{F}$ denotes the Fourier Transform, and (x, y, z) represent spatial coordinates in 3D space. Aperture function, which is dependent on factors like numerical aperture (NA), wavelength (λ), and refractive index mismatches, describes

how diffraction and interference patterns determine the intensity distribution of light at different depths [5].

To generate intermediate z-level images, the object intensity distribution (approximated by the EDOF image) is convolved with a depth-specific PSF:

$$I_b(x, y, z') = I_f(x, y) \otimes PSF(x, y, z'),$$

where $\otimes$ denotes the convolution operation. Here, z' represents an intermediate depth level corresponding to the focal plane's z-level. While this convolution is theoretically computed in the Fourier domain using:

$$\mathcal{F}\{I_b(x, y, z')\} = \mathcal{F}\{I_f(x, y)\} \cdot OTF(x, y, z'),$$

where $OTF(x, y, z')$ is the Optical Transfer Function derived from the PSF. We do not explicitly compute OTF in practice. Instead, the proposed PSF-Net learns this mapping implicitly during training.

PSF-Net is trained by minimizing a pixel-wise Mean Squared Error (MSE) loss between the generated blurred image, I_b, and its ground truth counterpart, I_{GT}:

$$L = \frac{1}{HW} \sum_{x=1}^{H} \sum_{y=1}^{W} (I_b(x, y) - I_{GT}(x, y))^2.$$

3.3 Augmented Training Strategy

Using PSF-Net, we generate an augmented dataset by synthesizing cell images at intermediate z-levels (e.g., 1.5, 2.5, etc). These augmented images significantly increase training diversity and reduce gaps between discrete focal planes, enabling better generalization across unseen z-levels during testing.

The augmented images are combined with the original data and incorporated into the training pipeline of the z-Net. This augmentation strategy enhances MDE performance by providing more granular depth information during training while maintaining physical consistency with microscopy imaging systems.

4 Experiments and Results

4.1 Setup

Dataset Description The dataset used in this work consists of Giemsa-stained Peripheral Blood Smears (PBS) prepared for malaria parasite detection [16]. The Giemsa-stained cells retain their natural thickness, making them pseudo-3D objects when smeared onto a flat glass substrate. At each site, the 3D cells lie above the substrate and share a similar depth range. These characteristics, combined with z-stack images that provide depth estimation cues through varying degrees of blur, make this dataset well-suited for this task.

Each z-stack simulates different focus levels (blur) and includes 10 focal planes per site, along with an Extended Depth of Field (EDOF) image synthesized from the stack. High-resolution images (2560×2160 pixels) were acquired

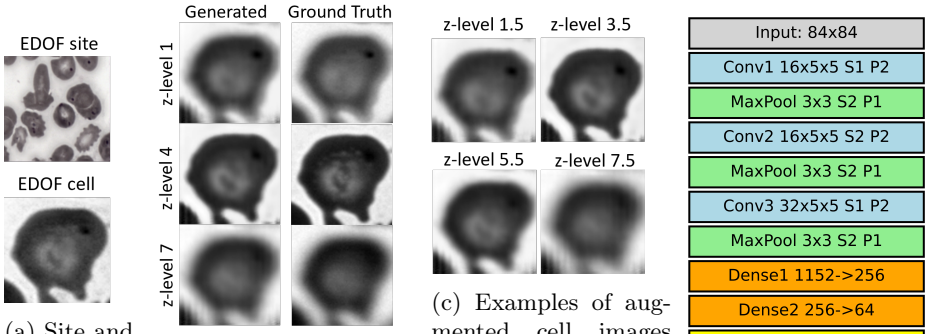

(a) Site and cell EDOF images.

(b) PSF-Net Generated vs Ground Truth Images.

(c) Examples of augmented cell images generated by PSF-Net.

(d) z-Net.

Fig. 2: Example cell images and the architecture of z-Net.

using a brightfield microscope with a focal offset of 0.5 μm between consecutive z-levels. The dataset includes approximately 10 sites, with each site yielding 150–200 extractable cells, resulting in thousands of individual cell samples.

Data Preparation and Organization The dataset is prepared for cell-level analysis using bounding box annotations provided in a JSON file. The preprocessing pipeline involves the following:

- Cells are extracted based on bounding box coordinates, resized to 84×84 grayscale images, and saved in both PNG and TIFF formats
- They are organized as 10 z-stack images per cell to represent varying degrees of blur, along with an EDOF image for a focused view, as shown in Fig 2a
- Approximately 1500 images per site (10 z-stack images $\times$ 150 cells) are extracted across all sites, with 20% of the dataset reserved for testing to ensure unbiased model evaluation.

4.2 Augmented Training

All training was conducted using the PyTorch framework on an NVIDIA Tesla V100 32 GB GPU. The first step involved training PSF-Net on 1690 cell images from Site 1, split into a 3:1:1 train-validation-test ratio. The model was trained for 300 epochs using the Adam optimizer with a learning rate of 0.001 until the loss curve saturated, achieving an R^2 score of 0.991 and a pixel wise mean squared error of 0.0001 on the test set. A qualitative comparison between PSF-Net generated images and the corresponding ground truth is made in Figure 2b. Additionally, simulations are included in the supplementary material, demonstrating PSF-Net's utility in generating depth cues for MDE through blur progression across focal planes.

Table 1: Performance metrics for different augmentation strategies on test data.

Site	Metric	No Augmentation	Traditional Augmentation	Linear Interpolation	Bicubic Spline Interpolation	Proposed Augmentation
1	MSE	0.769	0.981	0.443	0.359	0.126
	Sigma	0.871	0.986	0.655	0.592	0.354
2	MSE	0.777	1.031	0.343	0.447	0.181
	Sigma	0.879	0.951	0.586	0.668	0.423
3	MSE	0.601	3.135	0.372	0.337	0.123
	Sigma	0.772	1.765	0.598	0.578	0.351

Augmented cells were generated at intermediate z-levels (e.g., 1.5, 2.5, etc.) using PSF-Net, resulting in an additional 1521 training samples. Figure 2c illustrates examples of augmented cell images at intermediate z-levels generated by PSF-Net. These augmented images were incorporated into the training pipeline of z-Net (Fig. 2d), which was trained for 1000 epochs with a batch size of 33 using the same optimizer and learning rate as PSF-Net. The model's performance with and without augmentation was evaluated, showing significant improvements in test MSE. Table 1 shows the result for different sites, demonstrating the robustness and consistency of our augmentation strategy.

4.3 Comparisons with Other Augmentation Strategies

The efficacy of the proposed augmentation strategy was compared with two other approaches under similar training conditions. The first approach, traditional augmentation, involved applying standard techniques such as flips, rotations, and color jitter using `torchvision.transforms.Compose`. However, this resulted in performance degradation as these operations introduced irrelevant variations and failed to capture the depth-specific blur characteristics critical for microscopy-based depth estimation. The second approach involved generating intermediate z-level images through interpolation techniques, including linear and bicubic spline interpolation. While this approach gave improvements over no augmentation, it did not match the performance of the proposed physics-guided strategy. The test MSE and Sigma error were used as evaluation metrics to compare the training strategies. Table 1 summarizes the results.

4.4 Ablation Studies

Study 1: Effect of Augmentation Step Size To evaluate the impact of augmentation step size on model performance, we varied the step size used to generate intermediate z-level images. The original dataset contained z-levels ranging from 1 to 10 in steps of 1 (base with no augmented images). Augmented datasets were created with step sizes of 0.25, 0.33, 0.5 (used for augmented training previously), 2, and 3 (for a step size of 3, less blurry z-levels 1.5, 4.5, and 7.5 were chosen). Table 2 summarizes the performance metrics for different step sizes.

Table 2: Effect of augmentation step size on test performance (Site 1).

Step Size	Test MSE	Test Sigma
base	0.769	0.871
0.25	0.198	0.441
0.33	0.217	0.466
0.5	0.126	0.354
2	0.135	0.367
3	0.201	0.448

Table 3: Test MSE for different architectures and SOTA methods (Site 1).

Model	Without Augmentation	With Augmentation
AlexNet	0.769	0.126
ResNet	0.305	0.051
DenseNet	0.297	0.119
MobileNet	0.531	0.176
EfficientNet	0.659	0.320
Inception	0.560	0.226
ViT	1.263	0.507
MiDaS	0.371	0.114

Results showed that a step size of 0.5 achieved the best performance by balancing data diversity and model sensitivity. Smaller step sizes (e.g., 0.25, 0.33) provided diminishing returns due to the limited sensitivity of PSF-Net, while larger step sizes (e.g., 2 or 3) failed to generate sufficient training diversity.

Study 2: Robustness Across Model Architectures To assess the robustness of our proposed augmentation strategy across different architectures for z-Net, we replaced AlexNet with well known architectures such as ResNet, DenseNet, MobileNet, EfficientNet, and Inception networks. Additionally, partial state-of-the-art (SOTA) comparisons using Vision Transformer and MiDaS small MDE frameworks are included, though the higher computational cost of these SOTA models may limit their clinical utility compared to lightweight networks. Each architecture was trained with and without augmentation under identical conditions. Table 3 summarizes the results.

The results consistently showed that the proposed augmentation strategy improved performance across all architectures, as evidenced by lower MSE values in augmented training compared to non-augmented training. This highlights the robustness of our approach and its versatility in enhancing depth estimation performance regardless of the network design.

5 Conclusion

This work introduces a physics-guided augmentation strategy to improve monocular depth estimation (MDE) in microscopy. By generating realistic intermediate z-level images using Extended Depth of Field (EDOF) inputs, our approach bridges the gaps in discrete z-stack data, significantly reducing test mean squared error (MSE) and enhancing model robustness. The proposed framework is independent of the underlying deep network architecture, making it broadly applicable to various MDE methods in microscopy.

While our method was validated on Giemsa-stained peripheral blood smear data, the underlying physics-guided augmentation is not tied to a specific cell

type or imaging modality. Preliminary experiments on semiconductor scanning electron microscopy (SEM) images also indicated adaptability. From a practical perspective, the main computational cost lies in training PSF-Net; once trained, image generation and z-Net inference are efficient and suitable for real-time clinical workflows, balancing physical realism with computational efficiency.

Beyond depth estimation, the augmented images generated by this method offer potential for downstream tasks such as segmentation, classification, and cell tracking. The physics-guided design ensures realistic spatial representations, which could also benefit applications in industrial metrology and medical imaging workflows like super-resolution and deblurring.

Future work will focus on eliminating the reliance on precomputed EDOF images by leveraging shape-from-focus (SFF) models [29] or neural networks to generate all-focus images directly from z-stacks. Comprehensive validation on additional medical datasets, stains, and microscopy systems remains important to fully establish generalizability. Additionally, extending this approach to other domains beyond cell microscopy presents an exciting direction for further research and application.

Acknowledgments. This work was carried out while Abhishek Viswanathan and Nikhil Yelamarthy were students at IIT Madras.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Najafiaghdam, H., Rabbani, R., Gharia, A., Papageorgiou, E.P., Anwar, M.: 3d reconstruction of cellular images from microfabricated imagers using fully-adaptive deep neural networks. Scientific Reports **12**(1), 7229 (2022)
2. Lv, Z., Cao, X., Jin, X., Xu, S., Deng, H.: High-accuracy morphological identification of bone marrow cells using deep learning-based morphogo system. Scientific Reports **13**(1), 13364 (2023)
3. Palmieri, L., Scrofani, G., Incardona, N., Saavedra, G., Martínez-Corral, M., Koch, R.: Robust depth estimation for light field microscopy. Sensors **19**(3), 500 (2019)
4. Ban, Y., Liu, M., Wu, P., Yang, B., Liu, S., Yin, L., Zheng, W.: Depth estimation method for monocular camera defocus images in microscopic scenes. Electronics **11**(13), 2012 (2022)
5. Campus, Z.: Microscopy basics | the point spread function. Zeiss Microscopy (2025)
6. MicroscopyU, N.: The diffraction barrier in optical microscopy. Nikon MicroscopyU (2025)
7. Furieri, T., Ancora, D., Calisesi, G., Morara, S., Bassi, A., Bonora, S.: Aberration measurement and correction on a large field of view in fluorescence microscopy. Biomedical optics express **13**(1), 262–273 (2021)
8. Jiang, C., Lin, M., Zhang, C., Wang, Z., Yu, L.: Learning depth from focus with event focal stack. IEEE Sensors Journal (2024)
9. Wei, R., Li, B., Chen, K., Ma, Y., Liu, Y., Dou, Q.: Enhanced scale-aware depth estimation for monocular endoscopic scenes with geometric modeling. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 263–273. Springer (2024)

10. Zhang, J., et al.: Survey on monocular metric depth estimation. arXiv preprint arXiv:2501.11841 (2025)
11. Wang, L., Zhang, J., Wang, O., Lin, Z., Lu, H.: Sdc-depth: Semantic divide-and-conquer network for monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 541–550 (2020)
12. Hu, H., Feng, Y., Li, D., Zhang, S., Zhao, H.: Monocular depth estimation via self-supervised self-distillation. Sensors **24**(13), 4090 (2024)
13. Liu, X., Sinha, A., Ishii, M., Hager, G.D., Reiter, A., Taylor, R.H., Unberath, M.: Dense depth estimation in monocular endoscopy with self-supervised learning methods. IEEE transactions on medical imaging **39**(5), 1438–1447 (2019)
14. Sun, G., Liu, J., Liu, M., Liu, M., Zhang, Y.: Multiple prior representation learning for self-supervised monocular depth estimation via hybrid transformer. Engineering Applications of Artificial Intelligence **135**, 108790 (2024)
15. Hussain, Z., Gimenez, F., Yi, D., Rubin, D.: Differential data augmentation techniques for medical imaging classification tasks. In: AMIA annual symposium proceedings. vol. 2017, p. 979. American Medical Informatics Association (2017)
16. Manescu, P., Elmi, M., Bendkowski, C., Zajiczek, L., Shaw, M., Claveau, R., Brown, B.J., Fernandez-Reyes, D.: High magnification z-stacks from blood films (1 2022). https://doi.org/10.5522/04/13402301.v1, https://rdr.ucl.ac.uk/articles/dataset/High_Magnification_Z-Stacks_from_Blood_Films/13402301
17. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. IEEE transactions on pattern analysis and machine intelligence **38**(10), 2024–2039 (2015)
18. Chang, J., Wetzstein, G.: Deep optics for monocular depth estimation and 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10193–10202 (2019)
19. Rajapaksha, U., Sohel, F., Laga, H., Diepeveen, D., Bennamoun, M.: Deep learning-based depth estimation methods from monocular image and videos: A comprehensive survey. ACM computing surveys **56**(12), 1–51 (2024)
20. Zhao, C., Sun, Q., Zhang, C., Tang, Y., Qian, F.: Monocular depth estimation based on deep learning: An overview. Science China Technological Sciences **63**(9), 1612–1627 (2020)
21. Shajkofci, A., Liebling, M.: Spatially-variant cnn-based point spread function estimation for blind deconvolution and depth estimation in optical microscopy. IEEE Transactions on Image Processing **29**, 5848–5861 (2020)
22. Imtiaz, S.M., Hossain, F.F., Darkhanbaatar, N., Dashdavaa, E., Kwon, K.C., Jeon, S.H., Kim, N.: Estimating depth map from light field microscopic images using attention unet. In: Conference on Lasers and Electro-Optics/Pacific Rim. p. P2_142. Optica Publishing Group (2024)
23. Imanishi, H., Nishimura, T., Shimojo, Y., Awazu, K.: Deep learning based depth map estimation of protoporphyrin ix in turbid media using dual wavelength excitation fluorescence. Biomedical Optics Express **14**(10), 5254–5266 (2023)
24. Ghosh, A., Hohmann, M., Klämpfl, F., Schmidt, M.: Depth estimation in turbid media from stack of epi-illuminated microscopy images, using deep learning. In: Tissue Optics and Photonics III. vol. 13010, pp. 178–186. SPIE (2024)
25. Sun, R., Yang, D., Zhang, S., Hao, Q.: Hybrid deep-learning and physics-based neural network for programmable illumination computational microscopy. Advanced Photonics Nexus **3**(5), 056003–056003 (2024)

26. Li, R., della Maggiora, G., Andriasyan, V., Petkidis, A., Yushkevich, A., Kudryashev, M., Yakimovich, A.: Microscopy image reconstruction with physics-informed denoising diffusion probabilistic model. arXiv preprint arXiv:2306.02929 (2023)
27. Garcea, F., Serra, A., Lamberti, F., Morra, L.: Data augmentation for medical imaging: A systematic literature review. Computers in Biology and Medicine **152**, 106391 (2023)
28. Welsman, J.A., Weber, G.H., Amusat, O.O., Giannakou, A., Ramakrishnan, L.: Enhancing electron microscopy image classification using data augmentation. In: SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis. pp. 64–71. IEEE (2024)
29. Hariharan, R., Rajagopalan, A.: Shape-from-focus by tensor voting. IEEE Transactions on image Processing **21**(7), 3323–3328 (2012)