Letter

ADVANCED IMAGING

3D Gaussian adaptive reconstruction for Fourier light-field microscopy

Chenyu Xu,[a,b,†] Zhouyu Jin,[a,c,†] Chengkang Shen,[b] Hao Zhu,[b] Zhan Ma,[b] Bo Xiong,[d,*] You Zhou,[a,c,*] Xun Cao,[b] and Ning Gu[a,c]

[a]Medical School, Nanjing University, Nanjing, China

[b]School of Electronic Science and Engineering, Nanjing University, Nanjing, China

[c]Jiangsu Key Laboratory for Cardiovascular Information and Health Engineering Medicine. Institute of Clinical Medicine, Nanjing Drum Tower Hospital, Medical School, Nanjing University, Nanjing, China

[d]State Key Laboratory for Multimedia Information Processing, Peking University, Beijing, China

Abstract. Compared to light-field microscopy (LFM), which enables high-speed volumetric imaging but suffers from non-uniform spatial sampling, Fourier light-field microscopy (FLFM) introduces sub-aperture division at the pupil plane, thereby ensuring spatially invariant sampling and enhancing spatial resolution. Conventional FLFM reconstruction methods, such as Richardson–Lucy (RL) deconvolution, may face challenges in achieving optimal axial resolution and preserving signal quality due to the inherently ill-posed nature of the inverse problem. While data-driven approaches enhance spatial resolution by leveraging high-quality paired datasets or imposing structural priors, physics-informed self-supervised learning has emerged as a compelling precedent for overcoming these limitations. In this work, we propose 3D Gaussian adaptive tomography (3DGAT) for FLFM, a 3D Gaussian splatting-based self-supervised learning framework that significantly improves the volumetric reconstruction quality of FLFM while maintaining computational efficiency. Experimental results indicate that our approach achieves higher resolution and improved reconstruction accuracy, highlighting its potential to advance FLFM imaging and broaden its applications in 3D optical microscopy.

Keywords: Fourier light-field microscopy; Gaussian adaptive tomography; self-supervised learning; 3D Gaussian splatting; 3D optical microscopy.

Received Apr. 1, 2025; revised manuscript received Jun. 16, 2025; accepted Aug. 21, 2025; published online Sep. 23, 2025.

[DOI: 10.3788/AI.2025.50000]

1. Introduction

Light-field microscopy (LFM)[1–4] has emerged in recent decades as an advanced optical imaging technique capable of capturing three-dimensional (3D) information in a single shot, enabling high-speed volumetric imaging while minimizing photobleaching and phototoxicity. However, its limited 3D spatial resolution and non-uniform spatial sampling scheme[5] have hindered its broader applicability. To overcome these limitations, Fourier light-field microscopy (FLFM)[6] has been developed, offering enhanced spatial resolution and spatially invariant sampling compared to conventional LFM. FLFM achieves multi-view imaging by dividing the sub-aperture at the pupil plane, thereby generating multiple parallax views on the camera sensor.

For 3D reconstruction in FLFM, the Richardson–Lucy (RL) deconvolution algorithm[7,8,9] and its variants are among the most widely used methods. However, due to the inherently ill-posed nature of the inverse problem, these approaches may face challenges in achieving optimal axial resolution and structural fidelity of the sample, which can ultimately compromise the volumetric imaging quality of FLFM. Recently, data-driven reconstruction techniques for FLFM[10] have demonstrated notable improvements in spatial resolution. However, their performance relies on access to high-quality paired datasets or specific structural assumptions about the sample, which restricts

*Address all correspondence to Bo Xiong, xiongbo@pku.edu.cn; You Zhou, zhouyou@nju.edu.cn

†These authors contributed equally to this work.

their generalizability to diverse sample types or those with significant structural differences.

In contrast, neural radiance field (NeRF)-based methods, which employ self-supervised learning by integrating physical imaging models, have been successfully applied in optical microscopy imaging, such as the artefact-free refractive-index field (DeCAF)[11], Fourier ptychographic microscopy with implicit neural representation (FPM-INR)[12], computational adaptive optics (CoCoA)[13], and volumetric wide-field microscopy with physics-informed ellipsoidal coordinate encoding implicit neural representation (PIECE-INR)[14]. Moreover, INR-based approaches are currently being extended to the FLFM modality[15,16] and scanning LFM[17]. While these methods achieve high-quality reconstructions without requiring large-scale paired training datasets, improving their computational efficiency remains an active topic of research[18]. Recent methods such as Instant-NGP[19 >], TensorRF[20], and FrugalNeRF[21] have reduced the training time and memory demands of early NeRFs by combining implicit and explicit representations. However, the need for high fidelity and fast training-rendering cycles in volumetric microscopy continues to drive the search for alternative approaches.

Recently, 3D Gaussian splatting (3DGS), an emerging multi-view 3D rendering technique, has shown exceptional performance in the field of computer vision, offering high computational efficiency while maintaining comparable rendering quality[22]. Furthermore, prior works have extended its use to volumetric reconstruction tasks, including CT reconstruction[23] and vessel reconstruction using digital subtraction angiography (DSA) images[24]. 3DGS achieves significant speed advantages by directly optimizing a compact set of Gaussian parameters and enabling neural-network-free rendering. Its projection and blending operations are highly parallelizable and well-suited for GPU acceleration. Besides, by effectively leveraging physical priors and benefiting from the compactness of 3D Gaussian point cloud representations[22 >], 3DGS is capable of accurately modeling complex 3D structures, making it a promising candidate for extending applications beyond computer vision to volumetric reconstruction in optical microscopy.

In this work, we propose 3D Gaussian adaptive tomography (3DGAT) as a 3D Gaussian splatting-based self-supervised learning method for 3D reconstruction of FLFM. By integrating efficient 3D Gaussian representation with the physical imaging model of FLFM, we demonstrate significant resolution improvement over conventional deconvolution-based methods. The main contribution of our 3DGAT relies on: 1) introducing the Gaussian-based representation into light microscopy and further applying it to FLFM modality; 2) implementing a robust initialization and loss constraint for the specific FLFM task; and 3) validating its effectiveness on both simulated data and real experimental datasets.

2. Methods

2.1. Forward Imaging Model of FLFM

By positioning a microlens array (MLA) at the pupil plane of the objective lens in an inverted fluorescence microscope, FLFM enables multi-view imaging of the observed object, generating multiple parallax views on the camera sensor, as illustrated in Fig. 1(a). Accordingly, FLFM can be approximated as a linear system, where the 3D spatial distribution of the fluorescence signal within the object space $O(\boldsymbol{r}_o)$ is sampled, propagates through the optical system, and is ultimately projected onto the camera sensor.

Therefore, the physical imaging model of FLFM can be described as the convolution of the object's distribution $O(\boldsymbol{r}_o)$ and the point spread function (PSF) of the FLFM system $h(\boldsymbol{r}_o, \boldsymbol{r}_i)$:

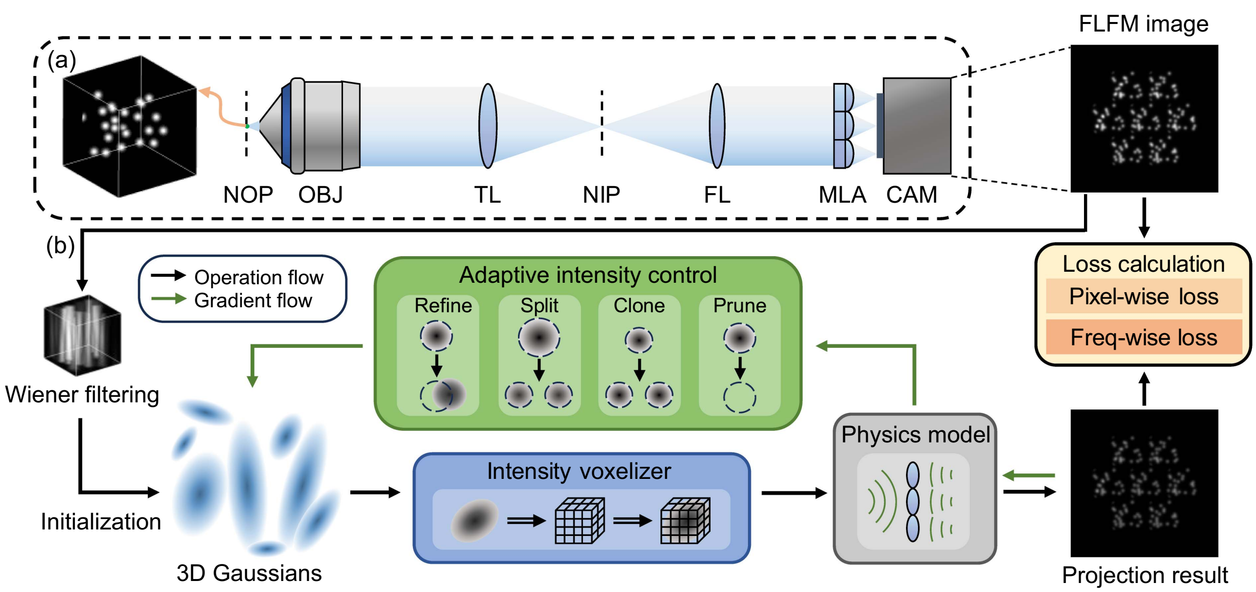

Fig. 1 Principle of 3D Gaussian adaptive tomography (3DGAT). (a) The optical setup and forward imaging process of the FLFM system. (b) The schematic of the training pipeline of 3DGAT. NOP, native object plane; OBJ, objective lens; TL, tube lens; NIP, native image plane; FL, Fourier lens; MLA, microlens array; CAM, camera.

$$I(\boldsymbol{r}_i) = \sum_z \int O(\boldsymbol{r}_o; z) h(\boldsymbol{r}_o; z, \boldsymbol{r}_i) \mathrm{d}\boldsymbol{r}_o, \tag{1}$$

where $\boldsymbol{r}_o = (x_o, y_o) \in \mathbb{R}^2$, $z \in \mathbb{R}$ represents the object space coordinates, and $\boldsymbol{r}_i = (x_i, y_i) \in \mathbb{R}^2$ represents the sensor plane coordinates.

The detailed computation of $h(\boldsymbol{r}_o; z, \boldsymbol{r}_i)$ can be referred to a previous work by Liu *et al.*[25]. For numerical computations, $O(\boldsymbol{r}_o; z)$ and $h(\boldsymbol{r}_o; z, \boldsymbol{r}_i)$ are discretized, enabling the relationship between the object volume $\boldsymbol{O}$ and the captured image $\boldsymbol{I}$ to be formulated as a discrete convolution:

$$\boldsymbol{I} = \sum_j \boldsymbol{H}_j * \boldsymbol{O}_j, \tag{2}$$

where $\boldsymbol{H}_j$ is the discretized format of $h(\boldsymbol{r}_o; \boldsymbol{r}_i)|_{z=j}$ and $*$ denotes the discrete 2D convolution operation. Therefore, the process of FLFM reconstruction involves solving an inverse problem for $\boldsymbol{O}$, using the measured image $\boldsymbol{I}$ along with the PSF $\boldsymbol{H}$, which is either simulated based on system parameters or experimentally obtained as the measured PSF. Considering that FLFM typically involves large imaging volumes and consequently large PSF sizes, we adopt a highly parallelized "Patch" strategy for forward imaging computation to mitigate edge aliasing and enhance computational efficiency. For more details, see Fig. S1 and Table S1 in the Supplement 1.

2.2 Realization of the 3DGAT Method

To solve this problem, we represent the target object $\hat{\boldsymbol{O}}$ with a set of 3D Gaussian kernels $\mathbb{G}^3 = \{G_i\}_{i=1}^m$; each kernel defines a Gaussian-shaped fluorescence intensity field in 3D space:

$$G_i(\mathbf{x}|\rho_i, \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) = \rho_i \cdot \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^{\mathrm{T}} \boldsymbol{\Sigma}_i^{-1} (\mathbf{x} - \boldsymbol{\mu}_i)\right], \tag{3}$$

where ρ_i, $\boldsymbol{\mu}_i$, and $\boldsymbol{\Sigma}_i$ are the learnable parameters, representing the central density, mean, and covariance of a 3D Gaussian ellipsoid. Considering the physical meaning of these parameters, $\boldsymbol{\Sigma}_i$ can be decomposed into scaling matrix $\boldsymbol{S}_i$ and rotation matrix $\boldsymbol{R}_i$, which are further represented as scaling vector $\boldsymbol{s}_i$ and rotation quaternion $\boldsymbol{r}_i$. These parameters control the ellipsoid's fluorescence intensity, central position, and scaling/rotation relative to a standard sphere in 3D space. In other words, ρ_i, $\boldsymbol{\mu}_i$, and $\boldsymbol{\Sigma}_i = \boldsymbol{R}_i \boldsymbol{S}_i \boldsymbol{S}_i^T \boldsymbol{R}_i^T$ can be used to generate arbitrary ellipsoids in 3D space.

In our proposed 3DGAT method, we first generate a coarse reconstruction by applying Wiener filtering to the raw FLFM measurement[26], which is formulated as

$$\boldsymbol{O}_{\mathrm{wiener\_filter}}(\boldsymbol{r}_o; z) = \mathcal{F}^{-1}\left\{\frac{\tilde{\boldsymbol{I}}(\boldsymbol{r}_i)\tilde{\boldsymbol{H}}^*(\boldsymbol{r}_i; z, \boldsymbol{r}_o)}{|\tilde{\boldsymbol{H}}(\boldsymbol{r}_i; z, \boldsymbol{r}_o)|^2 + w^2}\right\}, \tag{4}$$

where $\tilde{\cdot}$ denotes the 2D Fourier transform, $\mathcal{F}^{-1}\{\cdot\}$ denotes the 2D inverse Fourier transform, $\boldsymbol{O}_{\mathrm{wiener\_filter}}$ denotes the filtered result, and w is the Wiener parameter. Subsequently, we sample from this preliminary reconstruction result to generate a set of 3D Gaussian kernels, as a robust initialization approach of our Gaussian-based reconstruction framework. The effectiveness of the Wiener initialization strategy is validated through ablation experiments presented in the Supplement 1, Fig. S2, demonstrating that it substantially improves convergence speed and reconstruction accuracy compared to commonly used random or uniform initialization.

Inspired by the tile-based rasterizer on novel view synthesis tasks[22], we customize our intensity voxelizer based on the R^2-Gaussian method[23] to enable an efficient transformation from Gaussian-based to voxel-based representation. The voxelized data are then projected through the FLFM system's physical model to produce the final imaging results, as depicted in Eq. (2).

After this step, we introduce and adopt an objective loss function to mitigate blurring and artifacts during reconstruction, which follows the below formulation:

$$\mathcal{L} = \mathcal{L}_{\mathrm{MSE}}(I_{\mathrm{proj}}, I) + \alpha \mathcal{L}_{\mathrm{FDL}}(I_{\mathrm{proj}}, I), \tag{5}$$

where I_{proj} denotes the projected results and I denotes the raw FLFM measurement. The loss function in use is constructed with two terms, the mean square error (MSE) in the space domain $\mathcal{L}_{\mathrm{MSE}}$ to keep basic data similarity and the Fourier domain loss (FDL) $\mathcal{L}_{\mathrm{FDL}}$ to ensure further detail alignment in the frequency domain. The FDL is defined as $\mathcal{L}_{\mathrm{FDL}}(x, y) = \sum |\tilde{x} - \tilde{y}|$, where the $\tilde{\cdot}$ denotes the 2D Fourier transform and x, y denote the two input variables of FDL. α is a weight parameter, which empirically takes around 10^{-3}.

On the optimization step, the gradient of each Gaussian kernels is automatically back-propagated and accumulated owing to the differentiability of the intensity voxelizer and the physical model. In each iteration, all data points are sampled simultaneously, meaning the entire volume is propagated through the physical forward model in a single pass. The properties of the 3D Gaussians are optimized via gradient descent using the Adam optimizer, referred to as "refine" in Fig. 1(b). In addition, adaptive intensity control strategies named "split," "clone," and "prune" are applied every 100–200 iterations[22,23]. During this stage, accumulated gradients are compared against a predefined threshold. Larger gradients suggest that the corresponding Gaussians contribute more to the error and require adaptive adjustment. Then the density and scaling properties are checked together. Gaussians with densities below a certain threshold are considered nearly transparent and should be "pruned." Otherwise, the scaling values are used to determine whether Gaussians are under- or over-reconstructed, triggering "clone" or "split," respectively. The initial learning rates are set to 0.2 for position, density, and scaling, and 0.1 for rotation. A delayed exponential learning rate scheduler, as used in Plenoxels[27], is adopted.

2.3. Computation Details and Times

Most of the RL deconvolution and 3DGAT reconstruction experiments in this work are conducted on a workstation equipped with dual AMD EPYC 9654 CPUs and an NVIDIA A6000 GPU. The software environment consists of Python 3.9.21 and PyTorch 2.1.1 with CUDA 12.1 support. To leverage GPU acceleration, the tile-based intensity voxelizer is developed by customizing a CUDA kernel, building upon previous frameworks[22,23]. Experiments on the beads, reticulation, and zebrafish datasets are conducted on NVIDIA RTX A6000 GPUs due to high video random access memory (VRAM) demands. The line-pair and dandelion experiments are run on RTX 4090 GPUs for both methods. We provide a comparative analysis of the computational efficiency and runtime between 3DGAT and RL deconvolution, presented in the Supplement 1, Table S2. While 3DGAT achieves comparable runtimes to RL deconvolution in some cases, it becomes relatively slower as the dataset size increases, indicating opportunities for further optimization.

3. Results

To evaluate the performance and resolution of 3DGAT, we first use synthetically generated isotropic 3D fluorescent beads as the ground truth. Using the wave optics model of FLFM[25], we project the 3D images into 2D Fourier light-field images, with the FLFM parameters set to seven perspective views and a $20\times/0.45$ NA objective lens. The 3D reconstruction performance of 3DGAT is compared with RL deconvolution applying 100 iterations. As shown in Fig. 2(a), the results solved by 3DGAT are in good agreement with the ground truth. However, the RL deconvolution resolves the beads too finely in the lateral direction, while elongating them in the axial direction, as the line profiles shown in Figs. 2(b) and 2(c).

To further evaluate 3DGAT's reconstruction capability on complex samples, we simulate a 3D reticular structure (mesh-like formations exhibiting random curvature) using the same FLFM imaging parameters. As shown in Fig. 2(d), compared to RL deconvolution, 3DGAT more effectively recovers fine structural details and preserves contrast, producing results closer to the ground truth. This is further supported by the line profiles in Fig. 2(e) and the quantitative metrics in Fig. 2(f). We also assess the lateral resolution of 3DGAT using the Rayleigh criterion by simulating fluorescent lines with spacings from 2.40 to 1.20 μm and a line width of 0.48 μm. Reconstruction results are shown as x–y MIPs in Fig. 2(g) and side-view MIPs in Fig. 2(h). Intensity profiles along the yellow dashed lines are plotted in Figs. 2(i) and 2(j). Notably, 3DGAT clearly resolves three parallel lines spaced 1.44 μm apart, while RL deconvolution struggles to resolve even lines with 1.68 μm spacing.

To demonstrate that the strength of 3DGAT arises not only from its gradient descent framework but also from the 3D Gaussian representation, we implement a baseline that optimizes a voxel grid using the same loss and optimization strategy as 3DGAT. While this baseline improves axial contrast over RL deconvolution, it introduces more background artifacts. In contrast, 3DGAT better suppresses artifacts, achieving superior overall quality (see the Supplement 1, Fig. S5).

The simulation results clearly demonstrate that, compared to traditional RL deconvolution, the proposed 3DGAT method achieves higher-fidelity reconstruction with improved resolution. It effectively mitigates overfitting in the lateral direction while enhancing resolution in the axial direction.

To validate the effectiveness of our method on complex biological samples, we conduct a comparative simulation using a dandelion sample, employing FLFM with seven perspective views and a $20\times/0.45$ NA objective lens. A dandelion villi slice captured by a confocal microscope serves as the ground truth, and the corresponding data are synthesized by simulating the FLFM process. We compare the results of RL deconvolution [20 iterations for the best performance and highest peak signal-to-noise ratio (PSNR) value] and the results obtained using 3DGAT. We further evaluate the performance of 3DGAT with different loss functions, such as the commonly used MSE loss for reconstruction tasks, the mean absolute error (MAE) + structural similarity index measure (SSIM) loss from the original 3DGS paper[22], and our proposed MSE + FDL loss.

The MIP images from the x–y, x–z, and y–z planes of the reconstruction results, along with the ground truth, are shown in Fig. 3(a). Compared to RL deconvolution, 3DGAT demonstrates superior detail resolution and enhanced axial imaging capability. These improvements are more evident in the enlarged regions highlighted by the blue and white arrows in Fig. 3(b). Furthermore, the zoomed-in results in Figs. 3(b) and 3(c) indicate that our proposed MSE + FDL loss yields better visual fidelity compared to other loss functions. It closely matches the ground truth, particularly in the fine structural details indicated by the arrows in Fig. 3(b). The normalized intensity profiles along the white dashed line in Fig. 3(b) are plotted in Fig. 3(d), and quantitative evaluation metrics, including PSNR, SSIM, and learned perceptual image patch similarity (LPIPS), are summarized in Fig. 3(e). These results consistently support the conclusion that the proposed MSE + FDL loss achieves optimal quantitative performance. Notably, the intensity profile in Fig. 3(d) shows that our method reproduces all three distinct peaks at the correct positions, closely aligning with the ground truth.

This simulation demonstrates that the proposed 3DGAT method outperforms traditional RL deconvolution when applied to complex biological samples, offering both enhanced resolution and improved retention of fine details. We also validate that the proposed spatial-frequency domain simultaneous constraint loss (MSE + FDL loss) more accurately reconstructs the 3D structural information of the sample compared to existing loss functions.

Finally, we validate our method using real experimental zebrafish data. The data[28] are acquired using an FLFM system with 29 views and a $16\times/0.8$ NA water-immersion objective lens, as shown in Fig. 4(a). We further incorporate the effective rank regularization[29] (erank) to eliminate the needle-like artifacts caused by noise in the experimentally captured images. MIP images obtained using 3DGAT with (3DGAT-erank) and without (3DGAT) erank, along with RL deconvolution, are displayed together in Fig. 4(b). The enlarged views in Fig. 4(c) demonstrate that the 3DGAT-erank preserves more structural details than RL deconvolution while exhibiting fewer needle-like artifacts than the raw 3DGAT. This observation is further supported by the intensity profiles shown in Fig. 4(e). The Fourier domain visualization of the x–y and $x-z$ MIP images reconstructed by the respective methods in Fig. 4(d) highlights the effectiveness and high resolution of our proposed method. By better leveraging physical priors and the highly redundant views in the data, 3DGAT appears to better mitigate the missing cone problem in the Fourier domain compared to conventional methods. Additionally, we compute and present the Fourier ring correlation quality estimate (FRC-QE) scores in Fig. 4(f), where a higher score indicates superior recovery of frequency details.

To validate the noise robustness introduced by erank regularization, we evaluate 3DGAT under varying noise levels (30 dB to 15 dB), as shown in the Supplement 1, Fig. S3. Compared to the baseline, 3DGAT with erank more effectively preserves fine structural details and suppresses noise-induced needle-like artifacts, achieving superior performance both visually and across PSNR, SSIM, and LPIPS metrics. Given FLFM's well-known capability for high-speed volumetric imaging, we further assess the performance of our method on the time-series zebrafish data. Specifically, we reconstruct 120 consecutive volumes and analyze the temporal calcium activity traces of manually

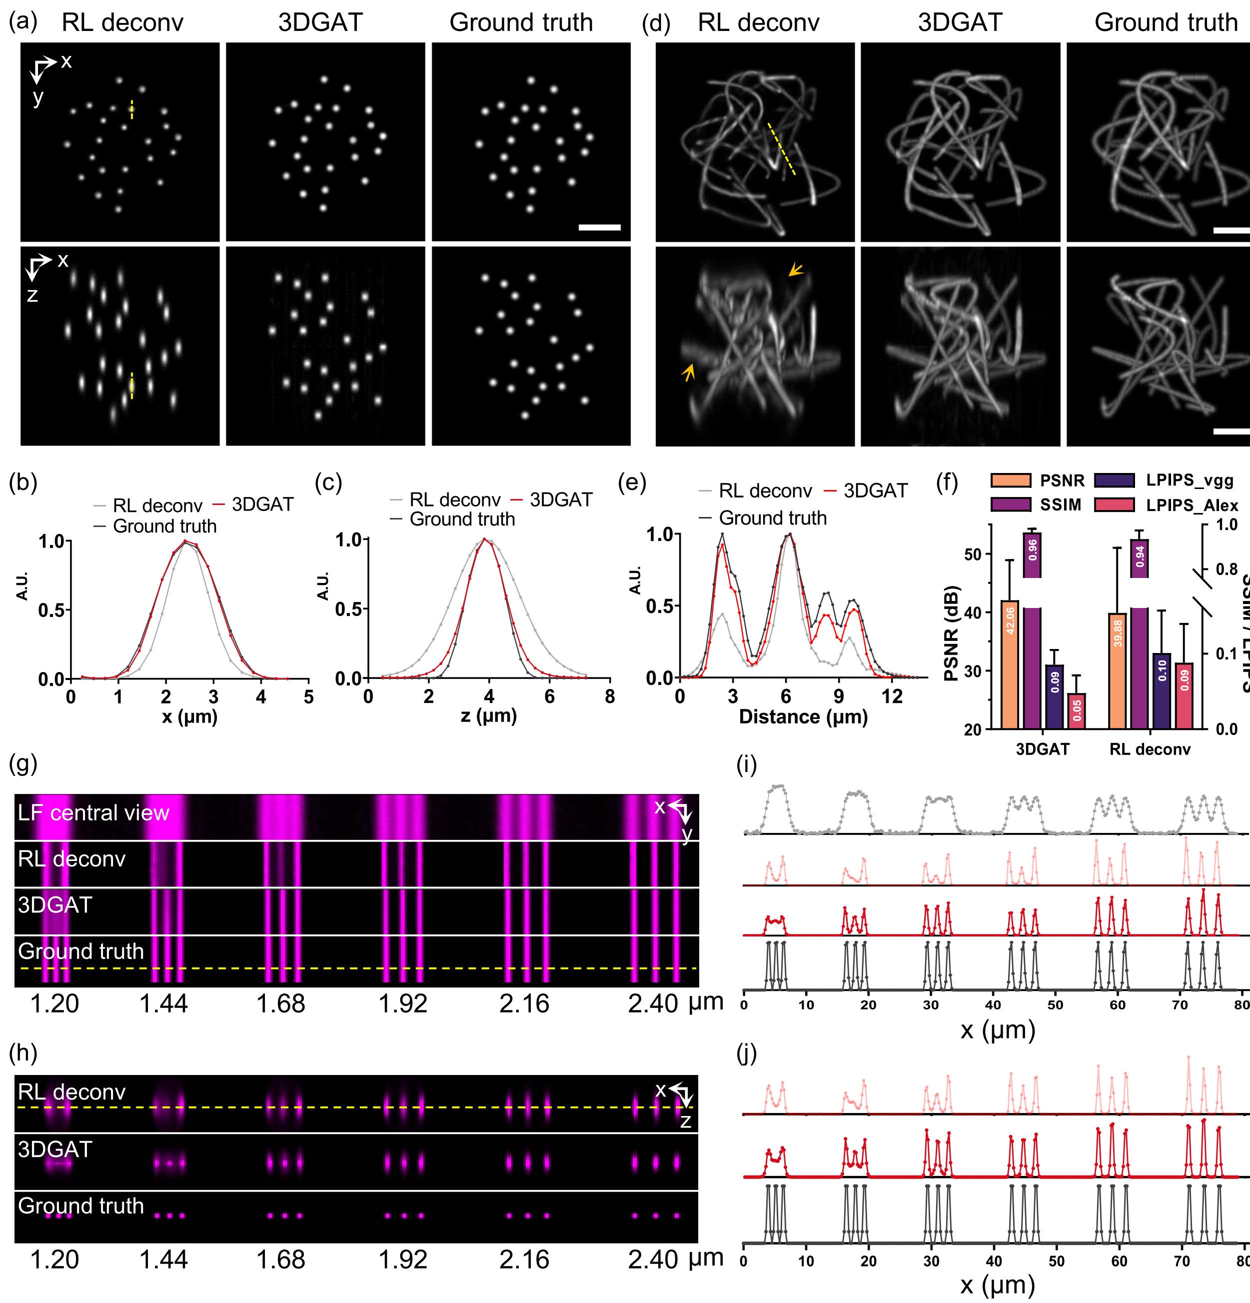

Fig. 2 Performance and resolution evaluation of 3DGAT on synthetic data. (a) Maximum intensity projections (MIPs) of synthetic fluorescent beads restored by RL deconvolution and 3DGAT, compared to the ground truth. (b), (c) Intensity profile comparisons of RL deconvolution (gray) and 3DGAT (red) with the ground truth (black) along the yellow dashed lines in (a). (d) MIPs of synthetic reticular structures reconstructed by RL deconvolution and 3DGAT, alongside the ground truth. (e) Intensity profile comparisons of RL deconvolution (gray) and 3DGAT (red) with the ground truth (black) along the yellow dashed line in (d). (f) Quantitative evaluation using PSNR, SSIM, and LPIPS metrics for RL deconvolution and 3DGAT across the depth range. (g), (h) x–y and x–z MIP images of synthetic fluorescent lines with varying intervals in the LF central view, as well as reconstructions by RL deconvolution, 3DGAT, and the ground truth. (i), (j) Intensity profile comparisons of the LF central view (gray), RL deconvolution (pink), and 3DGAT (red) with the ground truth (black) along the yellow dashed line in (g), (h). A.U., arbitrary units; Scale bar, (a), (d) 10 μm.

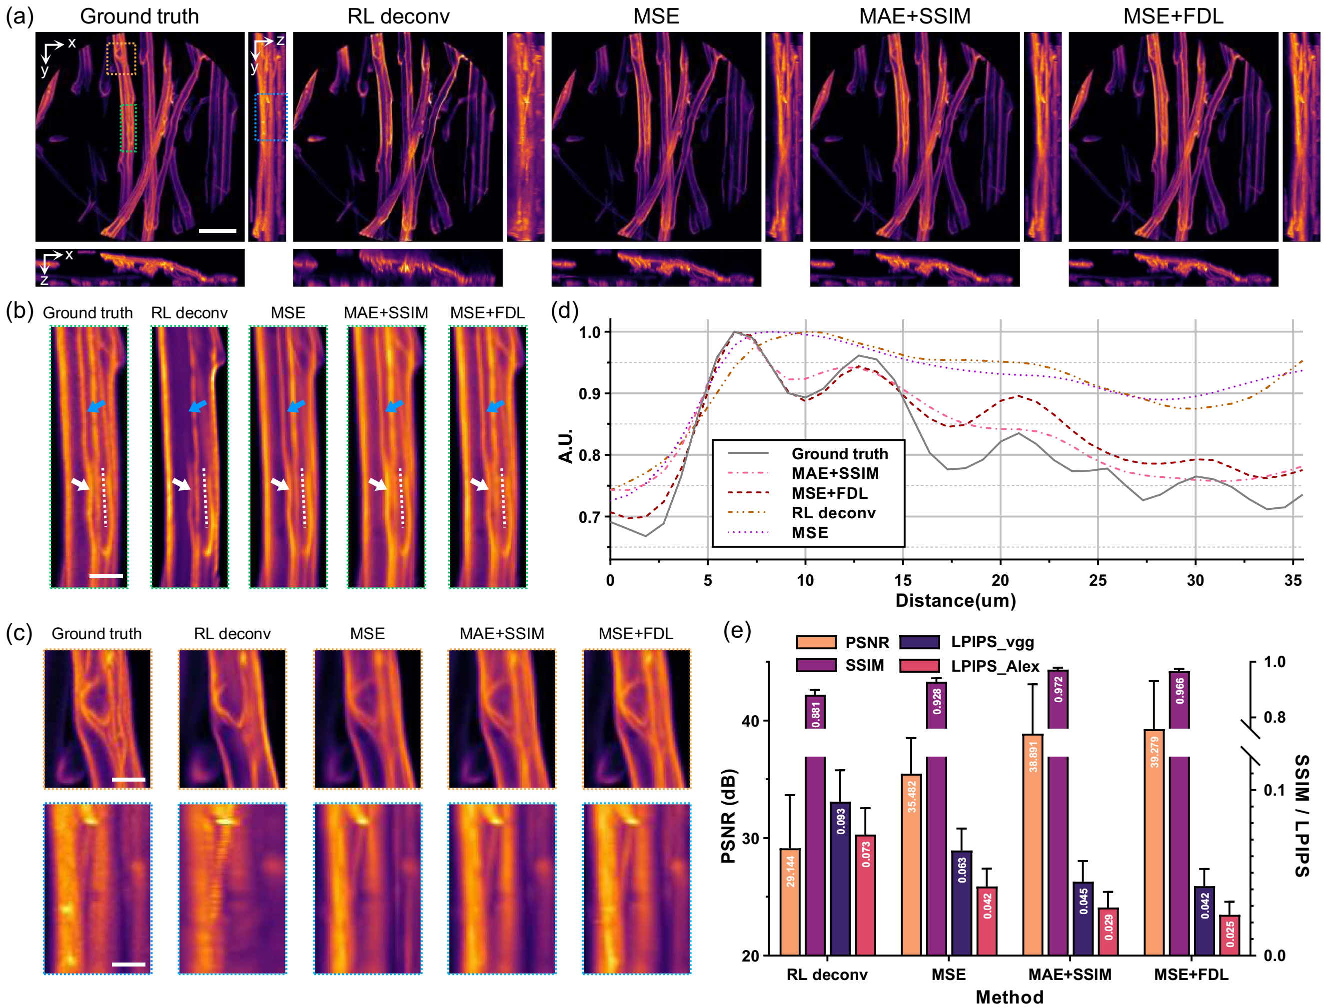

Fig. 3 Comparison between RL deconvolution and 3DGAT with different loss functions. (a) MIPs of the simulated dandelion sample restored by RL deconvolution and 3DGAT with different losses, such as MSE, MAE + SSIM, and MSE + FDL, compared with the ground truth. (b), (c) Enlarged views of the regions outlined by the dashed boxes in corresponding colors in (a). (d) Normalized intensity profiles of the ground truth (gray), RL deconvolution (brown), and 3DGAT with MSE loss (purple), MAE + SSIM loss (pink), and MSE + FDL loss (red) along the white dashed line in (b). (e) Quantitative metrics for reconstruction evaluation across the depth range. Scale bar: (a) 100 μm; (b), (c) 20 μm.

labeled neurons to validate the effectiveness of 3DGAT-erank. Detailed results and explanations are provided in the Supplement 1, Fig. S4.

4. Discussion and Conclusion

We introduce 3DGS into optical microscopy imaging and apply it to FLFM, proposing a novel 3DGAT framework. By leveraging robust initialization and loss constraints, while integrating the physical imaging model with efficient 3D Gaussian representations, our method enables high-quality 3D fluorescence reconstruction with adaptive intensity control. 3DGAT leverages 3D Gaussian primitives to embed implicit geometric priors that promote compactness and structural continuity. This regularization constrains the solution space in ill-posed problems, enabling finer detail reconstruction and reducing axial artifacts commonly observed in conventional approaches.

In our current implementation of 3DGAT, we first compute the 3D Gaussian distribution, then apply voxelization, followed by numerical simulation of the physical imaging process to estimate measurements for loss computation. The computational cost is largely dominated by the voxelization step and the forward imaging model. A promising research direction is to improve the voxelization heuristic by introducing explicit control over the number of Gaussian points during scene optimization. Additionally, we aim to develop analytical formulations of the imaging process across various microscopy modalities[30], mapping the 3D distribution to 2D acquisition, thereby improving both the efficiency and accuracy of 3DGS-based methods in this field.

We believe this work establishes a valuable foundation for extending 3DGS to a broader range of microscopic imaging applications. In the future, we plan to integrate the physical models of various multi-view acquisition schemes in optical microscopy

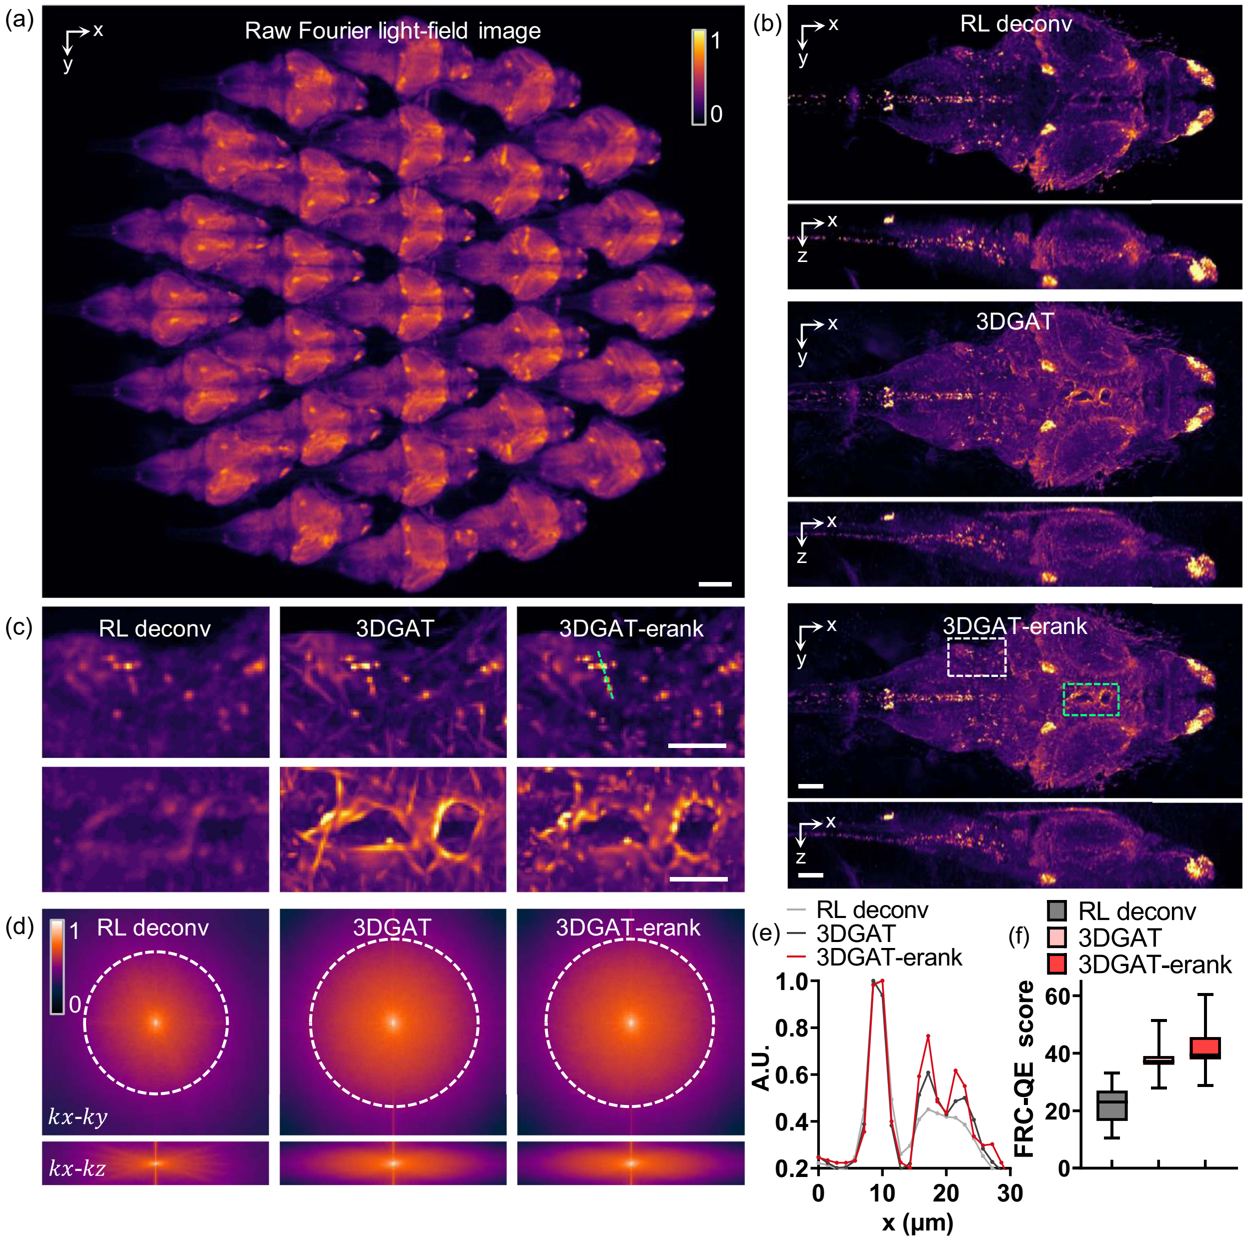

Fig. 4 Reconstruction of experimentally captured zebrafish data. (a) Raw Fourier light-field image of zebrafish data acquired from Ref. [28]. (b) x–y MIP images of results obtained by RL deconvolution, raw 3DGAT, and effective-rank-regularized 3DGAT (3DGAT-erank). (c) Enlarged views of the white and green dashed boxes in (b). (d) Fourier domain visualization of x–y and x–z MIP images recovered by the corresponding methods. (e) Normalized intensity profiles along the green dashed line in (c). (f) Fourier ring correlation quality estimate (FRC-QE) scores of three methods, respectively. Scale bar: (a), (b) 50 μm; (c) 30 μm.

into our proposed 3DGS-based framework, such as multi-view light-sheet microscopy[31,32] and other advanced techniques[33,34]. Moreover, given its inherent suitability for modeling complex 3D spatial structures, 3DGS offers capabilities beyond simple 3D reconstruction. For example, by incorporating temporal priors to leverage continuity between time-lapse frames, it enables robust 4D imaging of dynamic biological processes in live cells. It also supports downstream tasks such as motion artifact correction, dynamic surface reconstruction, and 3D cell segmentation and tracking, as well as neural calcium signal analysis. In parallel, we aim to incorporate efficient regularization informed by physical priors and modality-specific constraints, enabling an optimal trade-off among reconstruction quality, data efficiency, and computational cost.

Disclosures

The authors declare no conflicts of interest.

Data and Code Availability

The code and dataset for a representative example presented in this paper are available on GitHub at: https://github.com/Chaos1025/3DGAT.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (No. 2024YFF0508604), the Natural Science Foundation of Jiangsu Province (No. BK20222002), and the National Natural Science Foundation of China (Nos. 62071219, 62025108, and 62371006).

References

1. M. Levoy *et al.*, "Light field microscopy," *ACM Trans. Graph.* **25**, 924 (2006).
2. R. Prevedel *et al.*, "Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy," *Nat. Methods* **11**, 727 (2014).
3. Y. Zhou *et al.*, "Aberration modeling in deep learning for volumetric reconstruction of light-field microscopy," *Laser Photonics Rev.* **17**, 2300154 (2023).
4. Y. Zhou *et al.*, "Light-field micro-endoscopy using a fiber bundle: a snapshot 3D epi-fluorescence endoscope," *Photonics Res.* **10**, 2247 (2022).
5. M. Broxton *et al.*, "Wave optics theory and 3-D deconvolution for the light field microscope," *Opt. Express* **21**, 25418 (2013).
6. C. Guo *et al.*, "Fourier light-field microscopy," *Opt. Express* **27**, 25573 (2019).
7. A. Stefanoiu *et al.*, "What about computational super-resolution in fluorescence Fourier light field microscopy?," *Opt. Express* **28**, 16554 (2020).
8. L. B. Lucy, "An iterative technique for the rectification of observed distributions," *Astron. J.* **79**, 745 (1974).
9. W. H. Richardson, "Bayesian-based iterative method of image restoration," *J. Opt. Soc. Am.* **62**, 55 (1972).
10. C. Yi *et al.*, "Video-rate 3D imaging of living cells using Fourier view-channel-depth light field microscopy," *Commun. Biol.* **6**, 1259 (2023).
11. R. Liu *et al.*, "Recovery of continuous 3D refractive index maps from discrete intensity-only measurements using neural fields," *Nat. Mach. Intell.* **4**, 781 (2022).
12. H. Zhou *et al.*, "Fourier ptychographic microscopy image stack reconstruction using implicit neural representations," *Optica* **10**, 1679 (2023).
13. I. Kang *et al.*, "Coordinate-based neural representations for computational adaptive optics in widefield microscopy," *Nat. Mach. Intell.* **6**, 714 (2024).
14. Y. Zhou *et al.*, "Physics-informed ellipsoidal coordinate encoding implicit neural representation for high-resolution volumetric widefield microscopy," bioRxiv, bioRxiv:2024.2010.2017.618813 (2024).
15. C. Yi *et al.*, "High-fidelity generalizable light-field reconstruction of biological dynamics with physics-informed meta neural representation," bioRxiv, bioRxiv:2023.2011.2025.568636 (2023).
16. F. Zhong *et al.*, "Fast in vivo deep-tissue 3D imaging with selective-illumination NIR-II light-field microscopy and aberration-corrected implicit neural representation," bioRxiv, bioRxiv: 2025.2003.2016.643569 (2025).
17. J. Zhao *et al.*, "PNR: Physics-informed neural representation for high-resolution LFM reconstruction," arXiv preprint, arXiv:2409.18223 (2024).
18. B. Xiong *et al.*, "INeAT: an artifact-suppressed and resolution-enhanced computed tomography through iterative neural adaptive tomography," *Opt. Express* **32**, 32336 (2024).
19. T. Müller *et al.*, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Trans. Graph. (TOG)* **41**, 1 (2022).
20. A. Chen *et al.*, "TensoRF: tensorial radiance fields," in *European Conference on Computer Vision* (Springer, 2022), p. 333.
21. C.-Y. Lin *et al.*, "FrugalNeRF: Fast convergence for extreme few-shot novel view synthesis without learned priors," in *Proceedings of the Computer Vision and Pattern Recognition Conference* (2025), p. 11227.
22. B. Kerbl *et al.*, "3D Gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.* **42**, 139 (2023).
23. R. Zha *et al.*, "R^2-Gaussian: rectifying radiative gaussian splatting for tomographic reconstruction," *Adv. Neural Inform. Process. Syst.* **37**, 44907 (2025).
24. Z. Liu *et al.*, "4DRGS: 4D radiative gaussian splatting for efficient 3D vessel reconstruction from sparse-view dynamic DSA images," arXiv preprint, arXiv:2412.12919 (2024).
25. W. Liu *et al.*, "Fourier light-field imaging of human organoids with a hybrid point-spread function," *Biosens. Bioelectron.* **208**, 114201 (2022).
26. A. Llavador *et al.*, "Free-depths reconstruction with synthetic impulse response in integral imaging," *Opt. Express* **23**, 30127 (2015).
27. S. Fridovich-Keil *et al.*, "Plenoxels: radiance fields without neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), p. 5501.
28. J. Page Vizcaíno *et al.*, "Fast light-field 3D microscopy with out-of-distribution detection and adaptation through conditional normalizing flows," *Biomed. Opt. Express* **15**, 1219 (2024).
29. J. Hyung *et al.*, "Effective rank analysis and regularization for enhanced 3D Gaussian splatting," in *Conference on Neural Information Processing Systems (NeurIPS)* (2024).
30. J. Li, F. Xue, and T. Blu, "Fast and accurate three-dimensional point spread function computation for fluorescence microscopy," *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **34**, 1029 (2017).
31. U. Krzic *et al.*, "Multiview light-sheet microscope for rapid in toto imaging," *Nat. Methods* **9**, 730 (2012).
32. B. Yang *et al.*, "DaXi—high-resolution, large imaging volume and multi-view single-objective light-sheet microscopy," *Nat. Methods* **19**, 461 (2022).
33. Y. Wu *et al.*, "Multiview confocal super-resolution microscopy," *Nature* **600**, 279 (2021).
34. J. Wu *et al.*, "Iterative tomography with digital adaptive optics permits hour-long intravital observation of 3D subcellular dynamics at millisecond scale," *Cell* **184**, 3318 (2021).