Shift-Invariant-Subspace Discretization and Volume Reconstruction for Light Field Microscopy

Herman Verinaz-Jadan, *Student Member, IEEE,* Pingfan Song, *Member, IEEE,* Carmel L. Howe, *Member, IEEE,* Amanda J. Foust, *Member, IEEE,* and Pier Luigi Dragotti, *Fellow, IEEE*

***Abstract*—Light Field Microscopy (LFM) is an imaging technique that captures 3D spatial information with a single 2D image. LFM is attractive because of its relatively simple implementation and fast volume acquisition rate. Capturing volume time series at a camera frame rate can enable the study of the behaviour of many biological systems. For instance, it could provide insights into the communication dynamics of living 3D neural networks. However, conventional 3D reconstruction algorithms for LFM typically suffer from high computational cost, low lateral resolution, and reconstruction artifacts. In this work, we study the origin of these issues and propose novel techniques to improve the performance of the reconstruction process. First, we propose a discretization approach that uses shift-invariant subspaces to generalize the typical discretization framework used in LFM. Then, we study the shift-invariant-subspace assumption as a prior for volume reconstruction under ideal conditions. Furthermore, we present a method to reduce the computational time of the forward model by using singular value decomposition (SVD). Finally, we propose to use iterative approaches that incorporate additional priors to perform artifact-free 3D reconstruction from real light field images. We experimentally show that our approach performs better than Richardson-Lucy-based strategies in computational time, image quality, and artifact reduction.**

***Index Terms*—Light Field Microscopy, shift-invariant subspaces, system discretization, volume reconstruction, artifact-free deconvolution.**

I. Introduction

LIGHT Field Microscopy (LFM) is a fast imaging technique used in microscopy to retrieve 3D spatial information from a single 2D image. LFM captures a 3D scene by placing a microlens array in front of the camera sensor of a standard microscope to simultaneously record both angular and spatial information from the light field.

This work was supported in part by the Biotechnology and Biological Sciences Research Council under Grant BB/R009007/1 and in part by the Royal Academy of Engineering under the RAEng Research Fellowships scheme (RF1415/14/26).

Herman Verinaz-Jadan and Pier Luigi Dragotti are with the Department of Electronic & Electrical Engineering, Imperial College London, London SW7 2AZ, U.K. (e-mail: herman.verinazjadan17@imperial.ac.uk; p.dragotti@imperial.ac.uk).

Pingfan Song was with Imperial College London and is now with University of Cambridge (e-mail: ps898@cam.ac.uk).

Carmel L. Howe and Amanda J. Foust are with the Department of Bioengineering, and Center for Neurotechnology, Imperial College London, London SW7 2AZ, U.K. (e-mail: carmel.howe@imperial.ac.uk; a.foust@imperial.ac.uk).

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by Home Office Project and Personal Licenses (project license 70/9095), and performed in line with the UK Animals (Scientific Procedures) Act 1986.

This paper has supplementary downloadable material available at http://ieeexplore.ieee.org., provided by the authors.

Light field imaging has been extensively investigated since the concept of light field was first proposed in [1]. Light field cameras allow separate rays coming from different directions to be distinguished, opening up a range of applications for light field photography. For instance, a light field image can be used to refocus, reconstruct super-resolved images, or change the aperture in post-processing [2], [3].

Even though the wave nature of light needs to be considered in LFM, concepts and optical designs proposed for photography are also suitable for microscopy. The application of light field imaging to microscopy was first proposed by Levoy et al. in [4]. This technique is particularly attractive due to its non-scanning nature, allowing a much faster volumetric imaging rate than conventional scanning microscopy. Ideally, the LFM acquisition rate is only limited by the camera frame rate. Such a fast rate of capturing volumetric time series implies a wide range of possible applications. For instance, it could generate new data and insights into the communication dynamics of living 3D neural networks [5].

LFM aims to recover a high-quality 3D volume from a single 2D measurement. As in many other fields, high spatial resolution in lateral and axial dimensions is a desirable property. However, using a single image to encode 3D spatial information limits the number of samples that can be recovered. Therefore, reconstruction approaches typically suffer from low resolution and reconstruction artifacts. Furthermore, conventional reconstruction methods require long computational times, which clashes with the goal of LFM to analyze volume time series or to study fast biological dynamics of specimens in-vivo. Time constraint limits the complexity of the reconstruction algorithm for many time-sensitive applications.

In this work, we present a generalized method to discretize the light field model using shift-invariant subspaces. Then, we propose a model simplification method that accelerates the computation of the forward model. Finally, we propose a novel optimization method based on the Alternating Direction Method of Multipliers (ADMM) to perform artifact-free 3D reconstruction. Note that the light field microscope used in our work is based on the original optical design proposed in [4], known as Plenoptic 1.0 in photography. However, the ideas and methods presented here could potentially be extended to other alternative optical designs, such as Plenoptic 2.0 [6] and Fourier LFM [7]. The main novelty of our work includes the following aspects:

- A generalized framework based on shift-invariant subspace modelling is proposed for system discretization. Unlike the discretization approach used in conventional

reconstruction methods, our framework allows diverse sampling densities and template functions. The canonical discretization approach is a particular case of our framework.
- The LFM system is modelled using filter banks and the corresponding forward model is simplified using singular value decomposition (SVD). Our approach allows faster computation compared to conventional modelling.
- An ADMM-based optimization algorithm is developed for artifact-free volume reconstruction. The proposed approach is more effective than conventional Richardson-Lucy-based methods and faster due to the efficient computation of the forward model.

The rest of the paper is organized as follows: in Section II we summarize previous work in the area. In Section III, we present a formal description of the system based on the optic model proposed by Broxton et al [8]. In Section IV, we perform a novel analysis of the forward model that comprises its discretization and its simplification. In Section V, we propose volume reconstruction techniques based on our analysis of the forward model. Finally, in Section VI, we show reconstruction results on synthetic and real data and present our conclusions in Section VII.

II. Previous Work

The light field is a function that maps each ray in the space into a non-negative value known as the radiance (the amount of light of the ray). As mentioned in the previous section, the concept of light field was first proposed for photography by Levoy and Hanrahan [1], and by Gortler et al. [9] in 1996. However, the idea of plenoptic cameras comes from even earlier work; the pioneer was Lippman in 1908 with his contribution of the idea of "Integral Photography", as summarized in [10]. Ren Ng et al. reported the first successful implementation of a modern light field camera using microlens arrays in [10]. Now, this camera design is known as Plenoptic 1.0. The concept of light field became attractive because capturing the light field function allows the perfect recovery of the original 3D scene. A review of the main principles of light field photography can be found in [11].

The same optical configurations designed for photography can be adapted for microscopy. In [4], Levoy et al. first proposed this idea by implementing the Plenoptic 1.0 configuration in a standard microscope: A microlens array is placed at the native imaging plane of a 4F system with a camera placed at the back focal plane of the microlens array. In general, LFM is a technique that aims to recover a 3D image from a single light field image using computational reconstruction strategies.

In its early years, two methods were proposed for reconstruction in LFM. In their pioneering work, Levoy et al. introduced an algorithm that follows two simple steps. It first refocuses the light field to obtain a raw volume or a focal stack, and then it sharpens the stack by using 3D deconvolution. Even though this method is fast, the reconstruction has limited quality, and the size of the microlens pitch determines the lateral sampling interval of the reconstructed volume. Specifically, the lateral sampling interval is always the same as the sampling interval of the light field function, which in turn is restricted by the existing trade-off between spatial and angular sampling of the light field [4]. In LFM, a single camera sensor and a microlens array are used to capture both the position and direction of the light rays or, equivalently, both spatial and angular information of the light field. This optical arrangement imposes physical limitations on the sampling rates of the light field function: the spatial sampling interval is proportional to the microlens pitch, whereas the angular sampling interval is inversely proportional to the number of pixels under each microlens [4]. Thus, if the light field function is sampled with a high spatial sampling density, it is necessarily sampled with a low angular sampling density.

Later, Broxton et al. [8] proposed computing the measurement matrix of the system by using a wave-optics model. This model is used to state a linear inverse problem solved using the Richardson-Lucy (RL) algorithm. Furthermore, this method achieves better performance than the refocusing approach in terms of quality, and it does not limit the sampling interval of the reconstructed volume to be the same as the microlens pitch. However, this method has high computational complexity and suffers from square-like artifacts at specific depths.

Since the introduction of LFM, multiple alternative optical designs have been proposed to make the impulse response of the system less redundant. For instance, among others, Cohen et al. [12] proposed to place phase masks in the optical path; Guo et al. [7], proposed the Fourier Light Field Microscopy, He et al. [13] proposed the Snapshot Multifocal Light Field Microscopy, and recently Zhang et al. [14] proposed the Confocal Light Field Microscopy.

Compared with the multiple improvements proposed for the optical system, the computational strategy relies mainly on a RL strategy. In [15], Prevedel et al. proposed a RL-like algorithm called Iterative Image Space Reconstruction Algorithm (ISRA) [16]. When applied to time series extraction from neuronal cells expressing fluorescent activity reporters, Nöbauer et al. [5] included a reconstruction step using total variation and sparse priors using a modified version of ISRA. More recently, Lu et al. [17] proposed a Phase-Space deconvolution method that avoids reconstruction artifacts by first upscaling the light field image and then using one ISRA for each light field view. Stefanoiu et al. [18] proposed an approach that includes a filtering step to remove artifacts after each RL iteration. There also exist solutions that avoid the high computational complexity by reconstructing only point-like sources, as proposed in [19]. Finally, a few studies that exploit pure deep learning have recently emerged, e.g. [20], [21], [22]. However, deep learning approaches suffer from a lack of ground-truth data and are not always effective in this context [23]. In this work, we focus on the theoretical analysis of the system and iterative reconstructions methods, which we show to be useful in practice.

III. System Description

A light field microscope can be described as an operator that transforms a 3D scalar field into a 2D discrete image called a "light field", as in [8]. For simplicity, we ignore one lateral

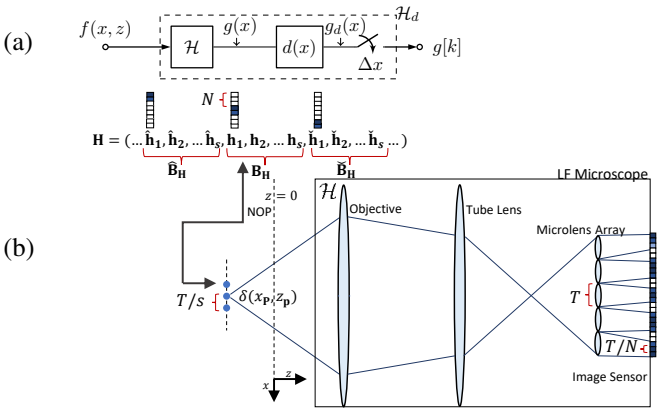

Fig. 1. Description of a light field system. (a) Basic block diagram of a light field microscope $\mathcal{H}_d$. The system is composed of a continuous ideal microscope $\mathcal{H}$ followed by a discretization step represented by a convolution with kernel $d(x)$ followed by sampling. (b) Conventional discretization approach and optical diagram. The conventional approach computes the impulse response for different input locations, which leads to a block-circulant-like matrix for each z-location. The upsampling factor s and the number of pixels under each microlens N define the block sizes of $\mathbf{H}$.

dimension and assume unit magnification since extending the analysis to the original problem is straightforward. Under these assumptions, a monochromatic light field microscope can be modelled as an operator $\mathcal{H}_d$ that transforms an intensity distribution $f(x, z)$ defined for any lateral and axial coordinates x, z into a pixel value $g[k]$ for any index k as follows:

$$g[k] = \mathcal{H}_d\{f(x, z)\}. \tag{1}$$

The functions $f(\cdot, \cdot) \in L^2 : \mathbb{R}^2 \mapsto \mathbb{R}^+$ and $g[\cdot] \in \ell^2 : \mathbb{Z} \mapsto \mathbb{R}^+$ represent the volumetric input and the light field, respectively. Moreover, Table I summarizes the main symbols used in this paper. To characterize the light field microscope $\mathcal{H}_d$, we divide it into a cascade of two blocks, as shown in Figure 1(a). The first block describes a continuous ideal microscope $\mathcal{H}$ that is usually described using wave optics. The second block represents a discretization process of the light field. This process is performed intrinsically by the pixels in the image sensor, and it can be modelled by a convolution followed by sampling.

For fluorescent microscopy imaging, it is realistic to assume that light emitted from two different sources is incoherent and that effects of occlusion or non-constant refractive indexes can be safely omitted. Based on these assumptions, the operator $\mathcal{H}$ that maps the volumetric input to the output light intensity (magnitude square of the complex field) is linear. Hence, for a volumetric input $f(\mathbf{p})$, the intensity $g(x)$ observed before sampling at the image sensor can be described with a superposition integral as follows:

$$g(x) = \int h(x, \mathbf{p}) f(\mathbf{p}) d\mathbf{p}, \tag{2}$$

where the function $h(x, \mathbf{p})$ is the impulse response of the system $\mathcal{H}$, or, equivalently, the light intensity at the output when the input is a point source located at $\mathbf{p} = (x_\mathbf{p}, z_\mathbf{p}) \in \mathbb{R}^2$.

Our work assumes that $h(x, \mathbf{p})$ is computed using the optical analysis first proposed by Broxton et al. [8].

TABLE I
TABLE OF SYMBOLS

Symbol	Description
$\mathcal{H}_d$	LFM operator
$\mathcal{H}$	LFM operator with continuous output
$h(x, \mathbf{p})$	Impulse response of $\mathcal{H}$
$\mathbf{H}_\delta, \mathbf{H}$	Conventional LFM measurement matrix
$\mathbf{H}_i$	LFM measurement matrix for the i-th depth
$\mathbf{B_H}$	Elementary matrix that circularly repeated forms $\mathbf{H}_i$
$\mathbf{S}_i$	Operator that extracts the depth i
$\varphi(x, z)$	Synthesis filter
$\tilde{\varphi}(x, z)$	Dual basis function of $\varphi(x, z)$
$d(x)$	Output sampling filter
$\mathbf{S}_\varphi$	Synthesis operator with $\varphi[\cdot]$ as template
$\mathbf{A}_\varphi$	Analysis operator with $\varphi[\cdot]$ as template
$\mathbf{D}_m$	Downsampling matrix of factor m
$\mathbf{U}_m$	Upsampling matrix of factor m
$f(x, z)$	Continuous volumetric input signal
$g(x)$	Continuous LF output signal
$f[k, l], f_{k,l}$	Discretized volumetric input signal
$g[k]$	Discrete LF output signal
$\mathbf{f}$	Vectorized input volume
$\mathbf{g}$	Vectorized light field
Δx_1	Lateral Sampling interval for $f(x, z)$
Δz	Axial sampling interval for $f(x, z)$
Δx_2	Sampling interval for $g(x)$
T	Microlens Pitch
N	Number of pixels under each microlens
s	Upsampling factor

Thus, the impulse response $h(x, \mathbf{p})$ is described by independently studying three parts of the system: (a) propagation from the source to the microlens array, (b) modulation due to the microlenses, and (c) propagation from the microlens array to the image sensor plane. However, as mentioned previously, an additional process must model the light field discretization performed by the image sensor. This process is simply a convolution with a kernel $d(x)$ followed by sampling. The final output of the system is a discrete image $g[k]$:

$$g[k] = g_d(k\Delta x), \tag{3}$$

where $g_d(x) = g(x) * d(x)$, and Δx is the sampling interval. In many practical cases, it holds that $\Delta x = T/N$, where T is the microlens pitch and N is the number of pixels under each microlens. See Figure 1 (b).

As explained in [8], the system $\mathcal{H}$ is not shift-invariant. However, it satisfies an important property: periodic shift-invariance. Namely, ignoring the finite extent of the device, for any $x, z \in \mathbb{R}$, if the input is shifted by multiples of T along the lateral axis, then the output is also shifted by multiples of T. Hence, the following relationship holds:

$$g(x - nT) = \mathcal{H}\{f(x - nT, z)\}, \forall n. \tag{4}$$

This means that for any depth z, this equation describes a periodically shift-invariant behaviour. An equivalent form to characterize this property is to state that the impulse response $h(x, x_\mathbf{p}, z_\mathbf{p})$ of the system $\mathcal{H}$ is periodic:

$$h(x, x_\mathbf{p}, z_\mathbf{p}) = h(x - T, x_\mathbf{p} - T, z_\mathbf{p}). \tag{5}$$

Finally, the entire system is discretized such that the inversion can be computed numerically. In [8], a standard discretization approach was proposed: first, a Dirac delta is

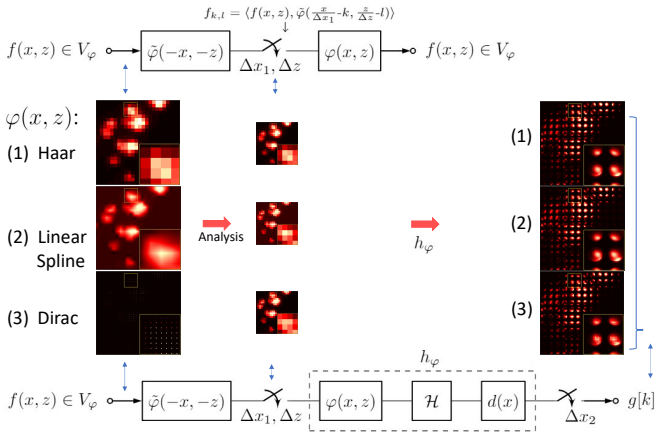

Fig. 2. Discretization of the light field system using shift-invariant subspaces. We show the discretization of a volume f. If f belongs to the shift-invariant subspace defined by template function $\varphi(x,z)$, the discretization can be inverted after filtering with $\varphi(x,z)$, as shown in the block diagram on the top of the figure. This allows performing a non-standard discretization that generalizes the conventional discretization method in LFM, as indicated on the bottom of the figure. Notice that the input discretization depends on the sampling density and template function $\varphi(x,z)$, while the output is discretized intrinsically by the device depending on the physics of the pixel sensors. In our example we show three possible choices for the input filters (Haar, Linear Spline and Dirac delta).

shifted laterally and axially; then, each shifted Dirac is used individually as an input to compute the corresponding impulse response, which is then stored in a matrix. The shifts of the Dirac delta are defined by an axial sampling interval Δz and a lateral sampling interval $\frac{T}{s}$, where the integer s is the upsampling factor, as shown in Figure 1(b). Since the system is not shift-invariant, multiple impulse responses for different positions must be stored. The final discrete system can be described by:

$$\mathbf{g} = \mathbf{H}\mathbf{f}, \tag{6}$$

where matrix $\mathbf{H} \in \mathbb{R}^{m \times n}$ maps a vectorized volume $\mathbf{f} \in \mathbb{R}^n$ into a light field $\mathbf{g} \in \mathbb{R}^m$. The number n of voxels of the volume is usually much larger than the number m of pixels of the light field.

In [8], the inverse problem derived from Equation (6) is solved using RL. Since property (4) leads to matrix $\mathbf{H}$ that is block-circulant for each z position, it is feasible to solve the problem by using RL [8]. However, this can be highly computationally demanding. Furthermore, the recovered volume suffers from artifacts near the native object plane (NOP): the plane where a point source is in focus at the microlens array, as described in [8] (conventionally, the plane $z = 0$).

IV. Forward Model Analysis

In this section, we study the description of the light field system. First, we analyze a framework that generalizes the conventional discretization approach by using shift-invariant subspaces for sampling. Then, we describe the system using filter banks, which in turn suggests a model simplification based on SVD that can be used to accelerate the computation of the forward model.

A. A general discretization framework

As mentioned previously, the standard approach presented in [8] discretizes the system by computing the impulse response for different input positions. However, a more general strategy is to discretize both the input and output using the assumption that the volume $f(x,z)$ belongs to a shift-invariant subspace (SIS) defined by a template function $\varphi(x,z)$, which was proposed in [24] for convolution operators in digital holography. If a function $f(x,z)$ belongs to a SIS V_φ generated by $\varphi(x,z)$, it can be written as follows:

$$f(x,z) = \sum_k \sum_l f_{k,l} \varphi\left(\frac{x}{\Delta x_1} - k, \frac{z}{\Delta z} - l\right), \tag{7}$$

where $f_{k,l} = \langle f(x,z), \tilde{\varphi}(\frac{x}{\Delta x_1} - k, \frac{z}{\Delta z} - l)\rangle$ and $\tilde{\varphi}(x,z)$ is the dual basis of $\varphi(x,z)$. As shown in Figure 2 (top), the relationship shown in Equation (7) can be represented with a simple pipeline: convolution, sampling, convolution. Note that any input $f(x,z)$ belonging to the subspace V_φ can be perfectly reconstructed from its samples $f_{k,l}$. Furthermore, the samples $f_{k,l}$ can be computed by filtering $f(x,z)$ with $\tilde{\varphi}(x,z)$ and then sampling with sampling interval Δx_1, Δz. This idea is exemplified in Figure 2 (top) by using three types of filters: a Haar, a linear spline and a Dirac delta.

If $f(x,z)$ is the input to the light field microscope, the model in Equation (7) suggests an alternative discretization process illustrated in Figure 2 (bottom). The input is sampled using the template function $\varphi(x,z)$ with sampling intervals Δx_1, Δz. The output is sampled using the template function $d(x)$ with a sampling interval Δx_2. Notice that the sampling of the output is performed intrinsically by the light field microscope, whereas the sampling of the input depends on how $f(x,z)$ is modelled. For example, if we assume $f(x,z)$ is a bandlimited signal, $\varphi(x,z)$ is the sinc function. If we assume $f(x,z)$ is a uniform spline, then $\varphi(x,z)$ is a spline. In contrast, $d(x)$ is usually a fixed box function (Haar filter) that models the effect of each pixel in the camera sensor [16]. Since the filter $d(x)$ should have support equal to the pixel size to simulate the effect of the pixel, a box function is usually enough to model the system properly. See Figure 2.

The discretization process leads to a new system defined by the impulse response $h_\varphi(x, x_\mathbf{p}, z_\mathbf{p})$. This system is equivalent to a cascade of three systems, (a) a 2D convolution with kernel $\varphi(x,z)$, (b) the system $\mathcal{H}$, and (c) a 1D convolution with kernel $d(x)$. See Figure 2 (b) for clarification. Therefore, by construction, periodicity still holds for the impulse response $h_\varphi(x, x_\mathbf{p}, z_\mathbf{p})$:

$$h_\varphi(x, x_\mathbf{p}, z_\mathbf{p}) = h_\varphi(x - T, x_\mathbf{p} - T, z_\mathbf{p}). \tag{8}$$

The corresponding discretized impulse response can be found from the impulse response $h_\varphi(x, x_\mathbf{p}, z_\mathbf{p})$ by sampling it along each dimension. Therefore, the discrete system is defined by:

$$h_\varphi[k, k_\mathbf{p}, l_\mathbf{p}] = h_\varphi(k\Delta x_2, k_\mathbf{p}\Delta x_1, l_\mathbf{p}\Delta z). \tag{9}$$

Furthermore, using Equation (8) and Equation (9) it is possible to show that a periodicity property also holds for the discrete impulse response $h_\varphi[\cdot]$:

$$h_\varphi[k, k_\mathbf{p}, l_\mathbf{p}] = h_\varphi[k - rq, k_\mathbf{p} - ts, l_\mathbf{p}], \tag{10}$$

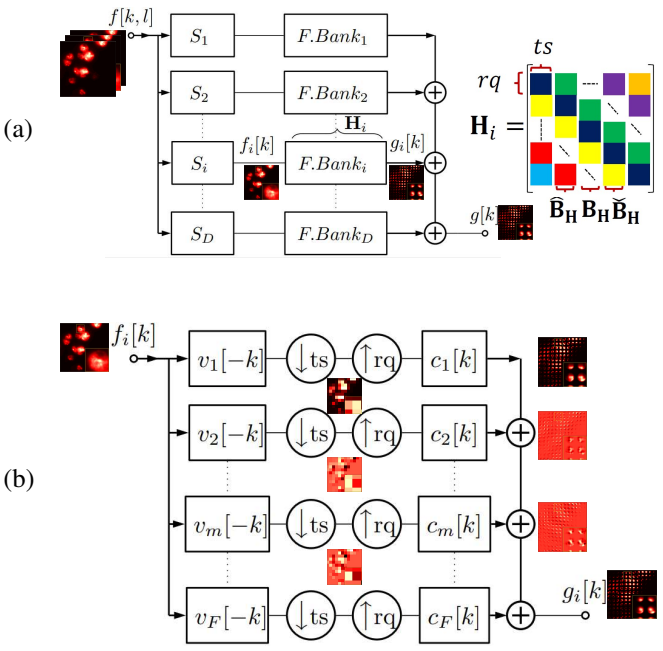

Fig. 3. Filter Bank Representation. In (a), a block diagram of light field microscope is depicted. For each depth i, the linear system S_i performs a slicing operation which chooses the respective i-th slice of the volume $f_i[k] = f[k,i]$. Then, each slice is passed through a filter bank to obtain a light field $g_i[k]$ per depth i. Finally, the light field $g[k]$ is the summation of all $g_i[k]$. (b) For each depth, the filter bank can be approximated by reducing the number of branches of the filter bank from ts to F. The sampling filter $v_m[k]$ defines the circulant matrix $\mathbf{V_m}$ and the filter $c_m[k] = \sigma_m u_m[k]$ defines the circulant matrix $\mathbf{C_m}$. Filters $v_m[k]$, $u_m[k]$ are the left and right singular vectors and coefficients σ_m are the singular values coming from a singular value decomposition. Notice that when $F = ts$, the approximation becomes equality.

where we assumed that integers q, r, s and t exist such that $\frac{\Delta x_1}{T} = \frac{q}{s}$ and $\frac{\Delta x_2}{T} = \frac{t}{r}$ and where both fractions are irreducible.

Finally, we remark that the conventional discretization procedure can be understood as a particular case of this more general framework. Note that if we allow $\varphi(x,z)$ to be a Dirac delta, $q=1$ and $t=1$, our framework reduces to the standard discretization approach. In this case, s is the upsampling factor, r is the number of pixels under each microlens N, and $h_\varphi[k, k_\mathbf{p}, l_\mathbf{p}]$ is the standard discrete impulse response, as described in [8].

B. Model Simplification

Equation (10) indicates that periodically-shift invariance holds for each depth of the discrete system. Therefore, the measurement matrix $\mathbf{H}$ derived from the discretized impulse response $h_\varphi[k, k_\mathbf{p}, l_\mathbf{p}]$ can be written as

$$\mathbf{H} = \sum_{i=1}^{D} \mathbf{H}_i \mathbf{S}_i, \tag{11}$$

where $\mathbf{S}_i$ is an auxiliary matrix that selects the depth $l_\mathbf{p} = i$ from the discrete volumetric input, D is the number of depths, and each matrix $\mathbf{H}_i$ follows a block-circulant structure due to

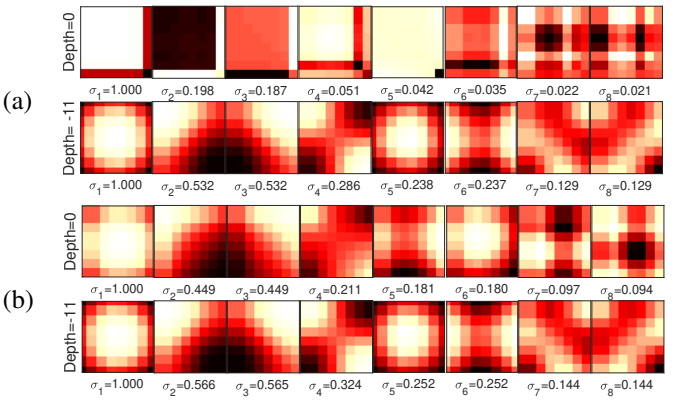

Fig. 4. Optimal analysis filters. This figure shows $F = 8$ optimal filters found for two different depths ($z = -11\ \mu m$, $z = 0\ \mu m$) from a maximum of $ts = 19 \times 19$ filters. For each depth, these filters are the 2-D version of the 1-D filters named $\{v_m\}_{m=1}^{F}$. In (a) we use a dirac Delta as the template function φ and in (b) we use a spline. The shown singular values σ_m are normalized to be between 0 and 1.

Equation (10). Furthermore, as depicted in Figure 3, the periodic shift-invariant property allows each $\mathbf{H}_i$ to be represented as a filter bank that performs a set of convolutions with the input. This accelerates the computation of the forward model. Assuming that each $\mathbf{H}_i$ is full rank, $2 \times ts \times D$ convolutions are needed to describe the whole $\mathbf{H}$. This is because the number of branches needed to represent each $\mathbf{H}_i$ equals ts, and 2 convolutions are required for each branch, as shown in Figure 3 (b). However, these computational requirements may be excessive in practice since the number ts of branches can be very high.

To mitigate this issue, we propose a simplification of the forward model. The computation can be accelerated by reducing the number of branches of each filter bank to F by optimally choosing the corresponding filters. First, we analyze the case $F = 1$, and then we generalize for arbitrary F. Thus, to approximate each $\mathbf{H}_i$ using a filter bank with a single branch, we state the following optimization problem:

$$\min_{\mathbf{c},\mathbf{v}} \quad \|\mathbf{H}_i - \mathbf{C}\mathbf{U}_{rq}\mathbf{D}_{ts}\mathbf{V}^\mathsf{T}\|_2^2, \tag{12}$$

where the matrix $\mathbf{D}_{ts}$ is a downsampling matrix of factor ts and the matrix $\mathbf{U}_{rq}$ models upsampling by rq. The matrix $\mathbf{V}$ is circulant, and it is completely defined by one column $\mathbf{v}$, which corresponds to the input filter of the filter bank. Similarly, the matrix $\mathbf{C}$ is circulant, and one column $\mathbf{c}$ corresponds to the output filter of the filter bank. The proper selection of $\mathbf{c}$ and $\mathbf{v}$ will allow the best approximation of $\mathbf{H}_i$ in the least square sense. The norm used here is the Frobenius norm.

To solve this problem, we limit the analysis to $\mathbf{v}$ with compact support ts. Since both $\mathbf{H}_i$ and $\mathbf{C}\mathbf{U}_{rq}\mathbf{D}_{ts}\mathbf{V}^\mathsf{T}$ have a similar structure, they can be expressed as follows:

$$\mathbf{H}_i = (...\hat{\mathbf{B}}_\mathrm{H}, \mathbf{B}_\mathrm{H}, \check{\mathbf{B}}_\mathrm{H}...), \tag{13}$$

and

$$\mathbf{C}\mathbf{U}_{rq}\mathbf{D}_{ts}\mathbf{V}^\mathsf{T} = (...\hat{\mathbf{B}}_\mathrm{CV}, \mathbf{B}_\mathrm{CV}, \check{\mathbf{B}}_\mathrm{CV}...). \tag{14}$$

The notation $\hat{\mathbf{B}}_\mathrm{H}$ means that each column of $\mathbf{B}_\mathrm{H}$ is shifted upward by rq elements. Similarly, $\check{\mathbf{B}}_\mathrm{H}$ indicates a downward shift by rq elements. Similar notations apply to $\mathbf{B}_\mathrm{CV}$. See

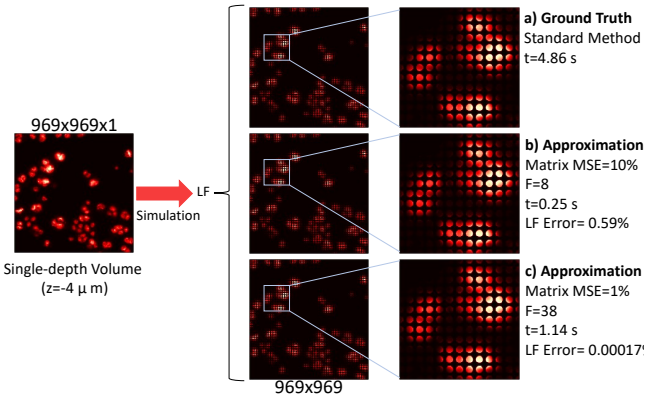

Fig. 5. Model simplification example. A synthetic light field was computed from a single depth volume. In (a) we show the simulated light field image using the standard model without any approximation. Furthermore, we show results using our approximation method by setting $F = 8$ and $F = 38$, and this leads to a matrix error (mean square error) of 10% in (b) and 1% in (c). The rightmost column shows the average computational time and the error between the approximated light field image and the ground truth. The single-depth volume is a sample of Drosophila melanogaster Kc167 cells taken from a publicly available library [25]. All the methods were tested in a CPU (Intel Core i7-6700, 16 GB RAM) using MATLAB R2018b.

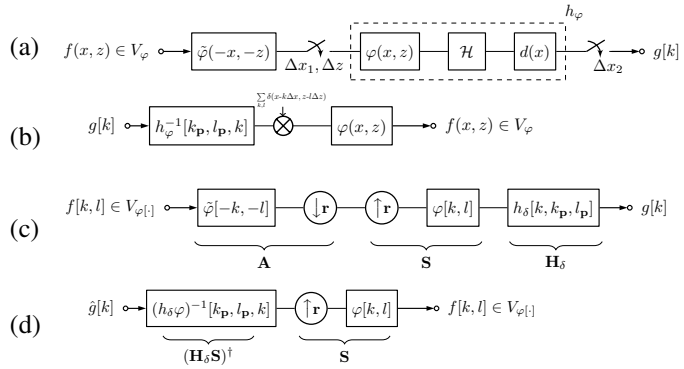

Fig. 6. Shift-invariant subspace assumption. In (a), we show a method to discretize the system using a shift-invariant subspace (SIS) V_φ. An arbitrary volumetric signal $f(x, z) \notin V_\varphi$ is discretized with a template function $\varphi(x, z)$ at an arbitrary sampling density. In contrast, the output is discretized intrinsically by the microscope with the kernel $d(x)$ at a sampling density defined by the pixel sensor to obtain a discrete output $g[k]$. In (b) we show how to perform perfect reconstruction from $g[k]$ under the assumption that the volume $f(x, z)$ belongs to the assumed SIS V_φ and the discrete system $h_\varphi[k, k_\mathbf{p}, l_\mathbf{p}]$ is invertible. In (c), we show the discrete version of the shift-invariant subspace discretization shown in (a). In (d), we show how to perform perfect reconstruction in the discrete space, under the assumption that $\mathbf{H}_\delta\mathbf{S}$ is full rank and the discrete volume $f[x, z]$ belongs to the assumed discrete SIS $V_{\varphi[\cdot]}$.

Figure 1 (b) and Figure 3 (a) for clarification. Hence, using Equation (13) and (14), expression (12) reduces to the minimization of $\|\mathbf{B_H} - \mathbf{B_{CV}}\|_2^2$. Based on the compact support assumption for $\mathbf{v}$, it holds that $\mathbf{B}_{\mathrm{CV}} = \mathbf{c}\mathbf{v}^\mathsf{T}$. Therefore, $\mathbf{c}$ and $\mathbf{v}$ should be chosen such that:

$$\min_{\mathbf{c},\mathbf{v}} \quad \|\mathbf{B_H} - \mathbf{c}\mathbf{v}^\mathsf{T}\|_2^2. \tag{15}$$

This minimization is a rank-1 approximation easily solved by using the SVD of $\mathbf{B_H}$. More importantly, computing the approximated forward model $\mathbf{C}\mathbf{U}_{rq}\mathbf{D}_{ts}\mathbf{V}^\mathsf{T}$ involves only two convolutions, implying much less computational complexity than computing the whole model.

As a generalization to the above rank-1 (very coarse) approximation of $\mathbf{H}_i$ by solving problem (12), the following rank-F approximation that keeps the largest F singular values of $\mathbf{B_H}$ is proposed:

$$\mathbf{H}_i \approx \sum_{m=1}^{F} \mathbf{C}_m\mathbf{U}_{rq}\mathbf{D}_{ts}\mathbf{V}_m^\mathsf{T}. \tag{16}$$

For each matrix $\mathbf{V}_m$, the respective $\mathbf{v}_m$ is a right singular vector of $\mathbf{B_H}$ and the $\mathbf{c}_m$ of $\mathbf{C}_m$ is a left singular vector $\mathbf{u}_m$ multiplied by the respective singular value σ_m: $\mathbf{c}_m = \sigma_m\mathbf{u}_m$. Notice that, since this approximation comes from an SVD, the summation progressively improves the approximation of the matrix $\mathbf{H}_i$. In particular, when F equals ts, the approximation turns into an equality. We also note that Equation (16) can be interpreted as a filter bank with F branches. Thus, for each depth, the approximation of $\mathbf{H}_i$ with F terms reduces the number of branches of the filter bank from a maximum of ts to F (see Figure 3).

Finally, given a volumetric input with an arbitrary number of depths D, the rank-F approximation strategy can be applied to approximate each matrix $\mathbf{H}_i$ for each depth i independently.

Approximation examples. To demonstrate the proposed model simplification strategy, Figure 4 shows $F = 8$ optimal filters computed using our SVD approximation for two depths. For these experiments, the input sampling intervals are set as $\Delta x_1 = T/N_1$ for both lateral axes and $\Delta z = T$, while the output sampling intervals are set to $\Delta x_2 = T/N_2$ for both lateral axes, where $N_1 = 8$, $N_2 = 19$, $T = 125\ \mu m$, and the microscope magnifies the volume by a factor of $M = 25$. In part (a), we set $\varphi(x, z)$ to be a Dirac delta and a linear spline in part (b). The shown filters are the equivalent 2D versions of the filters $\{v_m\}_{m=1}^F$ of the filter bank shown in Figure 3. One can observe that the shape of each filter changes with the depth. In part (a), for $z = 0\ \mu m$, the filter for singular value σ_1 is close to a square, which means that the system averages a region of the size of the microlens area, and projects this region into the light field space. Furthermore, the singular value σ_2 is almost 0.09 times the singular value σ_1, which shows the rank deficiency of the matrix $\mathbf{B_H}$ for this depth. However, changing the template function to a linear spline modifies the measurement matrix of the system, as shown in Equation (9). We observe in part (b) that the optimal filter corresponding to σ_1 for $z = 0\ \mu m$ is not a square anymore, which shows how the rank deficiency of the matrix $\mathbf{B_H}$ was reduced due to the integration along the axial dimension that occurs when a spline is selected as a template function.

As explained previously, truncating the number of branches of each filter bank of the model accelerates its computation. To show the utility of this approximation, we take a single-depth volume ($z = -4\ \mu m$) and simulate the forward model using both the standard method and our approximation. We use the same downsampling and upsampling factor $ts = rq = 19$ for both lateral axes. Figure 5 shows visual results for $F = 8$ and $F = 38$, which correspond to a matrix approximation error of approximately 10% and 1%, respectively. Notice that the number of filters F is much lower than a maximum of $ts = 19 \times 19$. Furthermore, for both cases, the approximated

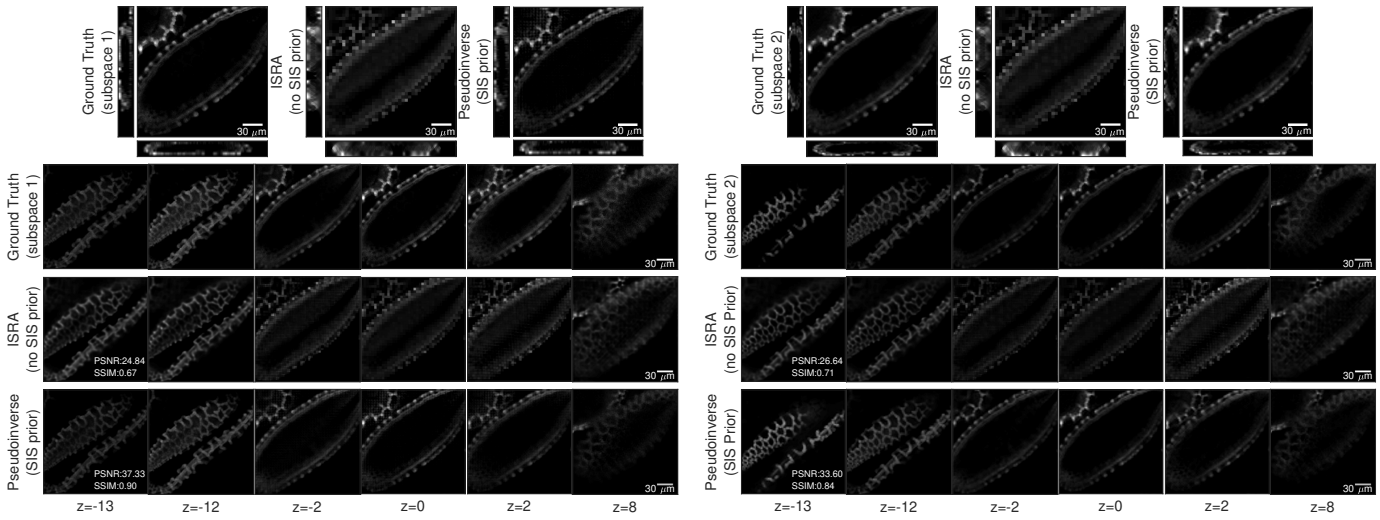

Fig. 7. Ideal reconstruction using the pseudoinverse on synthetic data. In this experiment, the ground truth volume satisfies the SIS assumption. We show two volumes lying into 2 different subspaces as ground truth, subspace 1 (a) and subspace 2 (b). Then, we show the respective reconstruction without any prior assumption (using ISRA) and the reconstruction using the pseudoinverse with the prior assumption. Top, we show the in-focus plane, one xz, and one yz plane. Below, we show additional slices for different depths. All the distances are measured in μm. The volumetric data was taken from Lilium Longiflorum Pollen [25], while the light field images were simulated from the 3D pollen volumes.

light field is visually indistinguishable from the ground truth, the error measured from the light field is less than 6% for $F = 8$, and it is almost zero when $F = 38$. The latter is because the system matrix is inherently low rank. Furthermore, in this experiment, our simplified model is almost 19 times faster to compute when $F = 8$ and four times faster when $F = 38$ than the standard approach.

V. Volume Reconstruction

This section analyzes the problem of reconstructing a 3D volume from a single light field image. In particular, we study model-based reconstructions to perform this task. First, we analyze an ideal case where the volumetric signal lies in a SIS; then, we study the more general scenario.

A. Ideal reconstruction under shift-invariant-subspace assumption

In Section IV-A, we discussed an alternative discretization process based on the assumption that the volumetric signal $f(x, z)$ lies in a shift-invariant subspace generated by $\varphi(x, z)$. This is again depicted in Figure 6 (a). Under this assumption, the image formation process is described by the discrete filter $h_\varphi[k, k_\mathbf{p}, l_\mathbf{p}]$ and, if this filter is invertible, reconstruction is achieved by first filtering the discrete light field image $g[k]$ with $h_\varphi^{-1}[k, k_\mathbf{p}, l_\mathbf{p}]$ and then by resynthesizing the original volume using a continuous filter $\varphi(x, z)$, see Figure 6 (b). In other words, perfect reconstruction (aliasing-free) of the continuous signal $f(x, z)$ is always achievable if two conditions hold: (a) $h_\varphi[\cdot]$ is invertible and (b) the volume being imaged genuinely belongs to the SIS generated by $\varphi(x, z)$ e.g. Linear splines, box functions, sinc functions (band-limited functions). As an analogy, in the problem of reconstructing a signal from its discrete samples, the named conditions are analog to the requirement for perfect reconstruction of the input. In Nyquist theorem, when the input signal genuinely belongs to the space of band-limited functions when the cutoff frequency equal to half the sampling frequency, the signal can be perfectly reconstructed from its samples. This observation suggests a way to find volumes that can be reconstructed from the light field image.

In practice, it is difficult to find a close form for h_φ; we therefore first discretize both $h(x, x_\mathbf{p}, z_\mathbf{p})$ and $f(x, x_\mathbf{p}, z)$ as in [8], but at a much finer resolution than the native lateral resolution of the LFM and then apply the shift-invariant model to the discretized f, which we denote as $f[k, l]$. This leads to the alternative model shown in Figure 6 (c), which can be described in matrix vector form as follows:

$$\mathbf{g} = \mathbf{H}_\delta \mathbf{S}_\varphi \mathbf{A}_\varphi \mathbf{f}. \tag{17}$$

Here, vector $\mathbf{f} \in \mathbb{R}^n$ is the vectorized volumetric input, vector $\mathbf{g} \in \mathbb{R}^m$ is the light field image, matrix $\mathbf{A}_\varphi \in \mathbb{R}^{\frac{n}{\mathrm{r}_1 \mathrm{r}_2} \times n}$ represents the convolution with the filter $\tilde{\varphi}(x, z)$ followed by downsampling, matrix $\mathbf{S}_\varphi \in \mathbb{R}^{n \times \frac{n}{\mathrm{r}_1 \mathrm{r}_2}}$ is the upsampling followed by convolution with the filter $\varphi(x, z)$, the downsampling and upsampling are defined by the factor $\mathbf{r} = (\mathrm{r}_1, \mathrm{r}_2)$, and $\mathbf{H}_\delta \in \mathbb{R}^{m \times n}$ is the forward model discretized using [8] (see Figure 6 (c)).

To recover the original volume $\mathbf{f}$, one can choose the downsampling factor $\mathbf{r}$ and filter $\varphi[\cdot]$ in Figure 6 (c) in a way that $\mathbf{H}_\delta \mathbf{S}_\varphi$ is full column rank and then the pseudoinverse can be used as follows:

$$\mathbf{f} = \mathbf{S}_\varphi (\mathbf{H}_\delta \mathbf{S}_\varphi)^\dagger \mathbf{g}. \tag{18}$$

Notice that this approach allows perfect reconstruction if matrix $\mathbf{H}_\delta \mathbf{S}_\varphi$ is full rank, and the input belongs to the assumed SIS. This process is depicted in Figure 6 (d).

Algorithm 1: Volume Reconstruction

Input : Light field image $\mathbf{g}$, matrix block $\mathbf{B_H}$ for every depth, learning rate α, importance of derivatives $\{\rho_x, \rho_y, \rho_z\}$, non-negativity importance ρ_+, soft-threshold parameter λ, number of iterations K.

Output: Reconstructed Volume $\mathbf{f}^K$

$k = 0$

Initialize $\mathbf{f}^k$ and auxiliar variables $\{\mathbf{u}_i^k\}_{i=1}^4$, $\{\mathbf{z}_i^k\}_{i=1}^4$ to zero.

for $k < K$ **do**

$\mathbf{z}_1^{k+1} = prox_+(\mathbf{f}^k + \mathbf{u}_1^k)$

$\mathbf{z}_3^{k+1} = prox_{|\cdot|}(\mathbf{D}_y^n\mathbf{f}^k + \mathbf{u}_3^k; \lambda)$

$\mathbf{z}_2^{k+1} = prox_{|\cdot|}(\mathbf{D}_x^n\mathbf{f}^k + \mathbf{u}_2^k; \lambda)$

$\mathbf{z}_4^{k+1} = prox_{|\cdot|}(\mathbf{D}_z^n\mathbf{f}^k + \mathbf{u}_4^k; \lambda)$

$\mathbf{f}^{k+1} \leftarrow$ solve Equation (19) ($\mathbf{H}\mathbf{f}^{k+1}$ is computed from $\mathbf{B_H}$, possibly using our fast approximation)

$$\begin{aligned}(\mathbf{I}+\rho_+\mathbf{I}+\rho_x\mathbf{D}_x^{n\mathsf{T}}\mathbf{D}_x^n+\rho_y\mathbf{D}_y^{n\mathsf{T}}\mathbf{D}_y^n+\rho_z\mathbf{D}_z^{n\mathsf{T}}\mathbf{D}_z^n)\mathbf{f}^{k+1}\\ = \mathbf{f}^k - \alpha\mathbf{H}^\mathsf{T}(\mathbf{H}\mathbf{f}^k - \mathbf{g}) + \rho_+(\mathbf{z}_1^k - \mathbf{u}_1^k)\\ + \rho_x\mathbf{D}_x^{n\mathsf{T}}(\mathbf{z}_2^k - \mathbf{u}_2^k) + \rho_y\mathbf{D}_y^{n\mathsf{T}}(\mathbf{z}_3^k - \mathbf{u}_3^k)\\ + \rho_z\mathbf{D}_z^{n\mathsf{T}}(\mathbf{z}_4^k - \mathbf{u}_4^k) \quad (19)\end{aligned}$$

$\mathbf{u}_1^{k+1} = \mathbf{u}_1^k + \mathbf{f}^{k+1} - \mathbf{z}_1^{k+1}$

$\mathbf{u}_2^{k+1} = \mathbf{u}_2^k + \mathbf{D}_x^n\mathbf{f}^{k+1} - \mathbf{z}_2^{k+1}$

$\mathbf{u}_3^{k+1} = \mathbf{u}_3^k + \mathbf{D}_y^n\mathbf{f}^{k+1} - \mathbf{z}_3^{k+1}$

$\mathbf{u}_4^{k+1} = \mathbf{u}_4^k + \mathbf{D}_z^n\mathbf{f}^{k+1} - \mathbf{z}_4^{k+1}$

$k = k + 1$

end

However, we highlight that there are limitations that make $\mathbf{H}_\delta$ rank deficient that cannot be overcome by the selection of a proper SIS. These limitations arise from the inherent lack of diversity in the impulse response for different input positions, e.g., the impulse response for in-focus points. Thus, the rank of $\mathbf{H}_\delta$ imposes a limit on the maximum number of samples that can be recovered. In particular, if the lateral sampling interval Δx_1 is decreased, the axial sampling interval Δz needs to be increased. This implies that it is not feasible to achieve high resolution simultaneously in both axes. Finally, one needs to choose the template function $\varphi[\cdot]$ and downsampling factor $\mathbf{r}$ such that the inversion of $\mathbf{H}_\delta\mathbf{S}_\varphi$ is as stable as possible, which gives a degree of freedom to investigate in practice.

Our experiments demonstrate this framework by modelling the SIS using separable 3D filters formed by one-dimensional linear splines. For the subspace used in the first scenario (scenario 1), we set a downsampling factor $\mathbf{r} = (4, 4, 16)$ which means a factor of 4 for both lateral dimensions and 16 for the axial dimension. In the second case, scenario 2, we use the same filter type but a downsampling factor $\mathbf{r} = (8, 8, 8)$. For this experiment, we use publicly available 3D data: the Lilium Longiflorum Pollen [25]. This volume is projected into the named subspaces to ensure the assumption holds. Figure 7 (a) shows that the ground truth volume for scenario 1 has a visually slightly higher lateral resolution than for scenario 2 due to the different downsampling factors $\mathbf{r}$ used for each subspace.

To compute $\mathbf{H}_\delta$ we use lateral and axial sampling intervals, Δx_1 and Δz, equal to $T/16$, and the output sampling interval Δx_2 defined by the pixel size of the sensor is equal to $T/19$, where T is the microlens pitch. Notice that our framework allows diverse sampling densities for both input and output, as opposed to [8]. Then, we compute the light field image using matrix $\mathbf{H}_\delta$ and attempt to recover the original volume using Equation (17). We clarify that the pseudoinverse was approximated iteratively since the closed-form expression of the pseudoinverse is not accessible.

Figure 7 shows the comparison between the proposed approach and the conventional reconstruction using ISRA. We show both lateral and axial slices of the volume. The reconstructed volume using the pseudoinverse and the SIS assumption matches the ground-truth volume very accurately, which is verified by the high Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). Moreover, the reconstruction does not suffer from heavy square-like artifacts near $z = 0$ as in the standard reconstruction approach (ISRA). Note that the volume was reconstructed at the maximum sampling density $16/T$ for both scenarios. However, for scenario 1, the volumetric signal lies in a SIS with a shift equal to 0.25 times the microlens pitch, and for scenario 2, the volume lies in a SIS with a shift equal to 0.5 times the microlens pitch. This experiment shows that, for the proposed ideal scenarios, exploiting the underlying low-dimensionality of the volume allows artifact-free reconstruction with much higher accuracy than ISRA in terms of two metrics, SSIM and PSNR.

B. General Scenario and Additional priors

As mentioned previously, when no prior assumption is made about the input, the reconstruction suffers from square-like artifacts near the plane $z = 0$. This is explained in work [8] and also verified in our experiments in the previous section. The reason is that a solution found using the pseudoinverse of $\mathbf{H}_\delta$ lies in the row space of the matrix $\mathbf{H}_\delta$, which only contains volumes with artifacts, due to the rank deficiency and the particular matrix structure for this depth.

Furthermore, during the volumetric reconstruction, commonly used algorithms such as ISRA and RL only impose mild priors (e.g., non-negativity and noise distribution), which are insufficient to regularize the solution. Therefore, we suggest incorporating more advanced priors as additional regularization to enforce the solution to live in a richer space far away from the row space of $\mathbf{H}_\delta$.

In particular, we propose two techniques for volume reconstruction. First, we propose to solve the following optimization problem:

$$\begin{aligned}\min_{\mathbf{f}} \quad & \|\mathbf{H}_\delta\mathbf{f} - \mathbf{g}\|_2^2 + \|\mathbf{D}_x^n\mathbf{f}\|_1 + \|\mathbf{D}_y^n\mathbf{f}\|_1 + \|\mathbf{D}_z^n\mathbf{f}\|_1 \\ \text{s.t.} \quad & \mathbf{f} \geq \mathbf{0},\end{aligned} \quad (20)$$

where $\mathbf{f}$ is the desired volume, $\mathbf{g}$ is the light field image, and $\mathbf{D}_x^n$, $\mathbf{D}_y^n$ and $\mathbf{D}_z^n$ are the n-th order derivatives along each axis. The parameter n can be adjusted experimentally since the derivative order does not imply too much additional

computation. Notice that if $n = 1$, the regularizer is an anisotropic total variation. However, this whole optimization is computationally demanding due to the high dimension of matrix $\mathbf{H}_\delta$. To alleviate this issue, we propose to use our approximation method based on SVD to compute the forward model.

As shown in the previous section, a SIS spanned by linear splines is able to reproduce typical structures that appear in the LFM image of a pollen volume. Considering this, we state the following optimization problem that exploits the SIS assumption to impose an additional constraint on the reconstructed volume:

$$\begin{aligned} \min_{\mathbf{f}} \quad & \|\mathbf{H}_\delta \mathbf{S}_\varphi \mathbf{A}_\varphi \mathbf{f} - \mathbf{g}\|_2^2 + \|\mathbf{D}_x^n \mathbf{f}\|_1 + \|\mathbf{D}_y^n \mathbf{f}\|_1 + \|\mathbf{D}_z^n \mathbf{f}\|_1 \\ \text{s.t.} \quad & \mathbf{f} \geq \mathbf{0}, \end{aligned} \tag{21}$$

where matrix $\mathbf{A}_\varphi$ and matrix $\mathbf{S}_\varphi$ have been defined before. Note that if the projection $\mathbf{S}_\varphi \mathbf{A}_\varphi$ is orthogonal, the transpose of the matrix $\mathbf{H}_\delta \mathbf{S}_\varphi \mathbf{A}_\varphi$ satisfies:

$$(\mathbf{H}_\delta \mathbf{S}_\varphi \mathbf{A}_\varphi)^\mathsf{T} = (\mathbf{S}_\varphi \mathbf{A}_\varphi)^\mathsf{T} \mathbf{H}_\delta{}^\mathsf{T} = \mathbf{S}_\varphi \mathbf{A}_\varphi \mathbf{H}_\delta{}^\mathsf{T}. \tag{22}$$

Therefore, the rows of the new measurement matrix span volumes that always lie in the assumed subspace. This characteristic, together with the sparsity imposed into the n-th order derivatives, helps to reduce typical reconstructions artifacts. Furthermore, the dimension of the matrix $\mathbf{H}_\delta \mathbf{S}_\varphi$ is smaller than that of $\mathbf{H}_\delta$. This allows solving the problem by exploiting the underlying low-dimensional 3D space where $\mathbf{f}$ lives, thereby contributing to efficient computation.

Finally, to solve optimization problems (20) and (21), we require an algorithm that avoids time-consuming inner loops and does not introduce strong artifacts from the computation of the transpose. To this end, we propose to adapt the inner-loop-free ADMM proposed in [26] for cryogenic electron microscopy to the light field problem, as shown in Algorithm 1. This algorithm can be implemented efficiently since it only needs the computation of the forward model at each iteration, the adjoint (transpose), convolutions, and point-wise non-linear operators (soft-thresholding and rectifier linear units).

VI. Experiments and Results

In this section, we show experimental results using synthetic and real data. For synthetic data, we measure reconstruction performance using the peak signal to noise ratio (PSNR) and the structural similarity index measure (SSIM) since the ground truth is available in this case. For real data, we only show reconstructed results for visual comparison. If not stated otherwise, the settings of the light field microscope used for both the simulated and the real data are set as: numerical aperture = 1, refractive index = 1.33, wavelength = 490 nm, magnification (M) = 25, microlens pitch (T) = 125 μm, microlens focal length = 1250 μm, tube lens focal length = 0.18 m, pixels under each microlens =19. Furthermore, to compute $\mathbf{H}_\delta$ we set the lateral and axial sampling intervals, Δx_1 and Δz to $T/16$, and set the output sampling interval Δx_2 defined by the pixel size of the sensor to $T/19$, where T is the microlens pitch.

TABLE II
Performance of model-based methods on synthetic data.

	PSNR	SSIM	Time/Iteration (s)
	Pollen		
ISRA [15]	26.05	0.54	470
ISRA (Total Variation) [5]	26.18	0.53	520
ISRA (Artifact-free) [18]	25.10	0.54	485
ADMM	27.22	0.65	79
ADMM (SIS)	**27.70**	**0.66**	**32**
	Neurons		
ISRA [15]	30.81	0.94	471
ISRA (Total Variation) [5]	28.38	0.91	530
ISRA (Artifact-free) [18]	32.20	0.92	481
ADMM	33.83	0.96	68
ADMM (SIS)	**39.67**	**0.98**	**26**

A. Synthetic Data

To validate our analysis, we perform experiments using two types of volumes. First, from the pollen structure [25], we simulate a light field image using matrix $\mathbf{H}_\delta$. The size of the synthetic light field is 779×779 pixels, while the original and reconstructed volumes are of size $656 \times 656 \times 96$ voxels (notice that $659 = 779 \times 16/19$). Then, from the light field image, we try to reconstruct the volume using five methods: (a) conventional ISRA [5], (b) ISRA with total variation, (c) artifact-free ISRA [18], (d) ADMM with sparsity regularization on the derivatives, as shown in Equation (20), and (e) ADMM with SIS assumption as well as sparsity regularization on the derivatives, as shown in Equation (21). For the latter case, we choose splines as template functions. Specifically, we choose the settings named 'scenario 1' in the previous section.The order of the derivative n in Equation (20) and (21) is set to 3. For all the three methods, the number of iterations is set to 24. In the supplemental material we discuss the use of a different stopping criterion.

We clarify that we use the same algorithm to solve both optimization problems (20) and (21). However, for the first case, our approximation method (16) is used to make the optimization computationally tractable, while for the second case, the approximation is not needed since $\mathbf{H}_\delta \mathbf{S}_\varphi$ has a smaller size than $\mathbf{H}_\delta$. In these experiments, we choose F in Equation (16) to achieve a matrix approximation error of 10%. Furthermore, note that in Algorithm 1, the additional computation performed includes convolutions and the adjoint operator of the system. The convolutions can be computed efficiently, and the adjoint operator can be found easily from the original measurement matrix using reflections and permutations. Therefore, Algorithm 1 can be implemented efficiently. Finally, the parameters α, ρ_x, ρ_y,ρ_z, ρ_+ and λ of Algorithm 1 can be found manually. In this section, we show that different choices of values have a limited effect on the performance of the method.

Figure 8 (a) shows the 3D structures recovered from a simulated light field image using the ISRA approaches and our method. Both the standard ISRA and the ISRA with total variation [5] introduce artifacts at planes near $z = 0$. In contrast, our ADMM methods achieve artifact-free reconstruction. Even though the ISRA version proposed in [18] also achieves artifact-free reconstruction, the solution is over-smoothed near

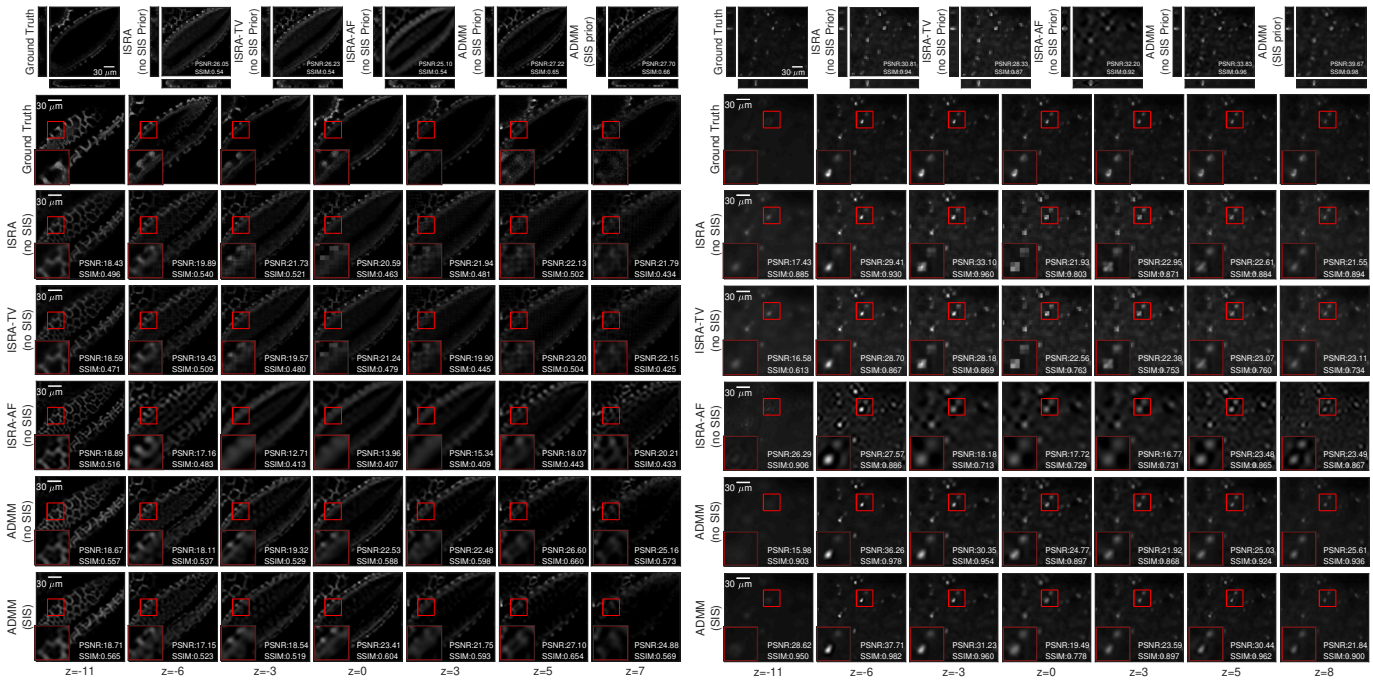

Fig. 8. Reconstruction using synthetic light field data. Top, we show the in-focus plane, one xz, and one yz slice of the ground truth volume, the reconstruction using ISRA, ISRA with total variation [5], artifact-free (AF) ISRA [18], ADMM without SIS assumption, and ADMM with SIS assumption. The shown PSNR and SSIM correspond to the whole volume. Furthermore, below, we show additional slices for different depths. The shown PSNR and SSIM correspond to each slice. In this case, the ground truth volume does not satisfy the SIS assumption. We use the Lilium Longiflorum Pollen [25] dataset in (a), and a brain slice taken from an EGFP tagged triple transgenic mouse line in (b). The light field images were created synthetically from wide field volumes. All distances are measured in μm.

TABLE III
PERFORMANCE OF ITERATIVE METHODS WITH POISSON NOISE.

α_0	48	40	32	24	16	8	48	40	32	24	16	8
	Pollen											
	PSNR						SSIM					
ISRA [15]	27.17	26.89	26.35	27.16	27.07	26.73	0.62	0.63	0.54	0.57	0.56	0.55
ISRA (Total Variation) [5]	26.96	26.19	26.48	26.28	26.41	25.84	0.56	0.54	0.54	0.54	0.54	0.53
ISRA (Artifact free) [18]	26.77	25.16	25.92	26.03	25.50	25.71	0.55	0.51	0.52	0.53	0.51	0.52
ADMM	26.83	26.65	**27.11**	27.04	27.21	26.79	0.63	0.63	0.64	0.64	0.64	0.64
ADMM (SIS)	**27.82**	**26.97**	26.61	**27.27**	**27.80**	**27.67**	**0.64**	**0.64**	**0.66**	**0.65**	**0.66**	**0.66**
	Neurons											
	PSNR						SSIM					
ISRA [15]	**34.20**	27.93	32.51	32.61	29.31	33.79	0.88	0.80	0.88	0.90	0.88	0.94
ISRA (Total Variation) [5]	29.50	25.92	28.13	29.77	27.71	28.60	0.84	0.77	0.83	0.86	0.84	0.86
ISRA (Artifact-free) [18]	30.05	30.82	31.34	32.15	31.58	31.88	0.86	0.87	0.88	0.90	0.90	0.91
ADMM	33.02	**31.50**	**34.25**	**34.69**	**32.70**	34.36	**0.94**	**0.93**	**0.96**	**0.96**	**0.95**	**0.96**
ADMM (SIS)	31.96	29.68	32.88	34.60	32.12	**34.37**	0.90	0.88	0.92	0.95	0.94	**0.96**

$z = 0$. Furthermore, our ADMM approaches give higher PSNR and SSIM than all the ISRA approaches. The ADMM with the SIS prior gives slightly higher PSNR and SSIM than the ADMM with no SIS assumption, as shown in Table II.

For the second experiment, we imaged a 50-um-thick brain slice from a mouse line Ai90 [27] expressing an Enhanced Green Fluorescent Protein (EGFP) with a wide-field microscope (Figure 8 (b)). Then, we simulated light field images using the captured volume as ground truth data. For this type of data, we obtain similar results as in the previous case. Both ADMM methods achieve artifact-free reconstruction, while the standard ISRA and the ISRA with total variation [5] reconstructs squares near the plane $z = 0$, as shown in Figures 8 (b). The artifact-free ISRA [18] removes squares, but in this case the solution is distorted. Quantitatively, as shown in Table II, the resulting PSNR and SSIM using the SIS assumption are slightly higher than those without the SIS assumption. Both ADMM methods outperform ISRA approaches in terms of PSNR and SSIM.

Furthermore, for the given settings, the ADMM methods are faster than the conventional methods. Each iteration of all the ISRA methods takes more than 470 seconds on average, whereas the ADMM without SIS prior takes 74 seconds, and ADMM with the SIS prior takes 29 seconds on average, as shown in Table II. All the methods were tested in a CPU (Intel Core i7-6700, 16 GB RAM) using MATLAB

TABLE IV
HYPER-PARAMETER STUDY(PSNR).

		α	ρ_+	ρ_x, ρ_y	ρ_z	λ
		$[0.3, 3] \times 10^{-6}$	$[1.2, 12]$	$[250, 2500]$	$[0, 100]$	$[0.15, 15] \times 10^{-4}$
ADMM	Min.	24.91	24.32	27.32	25.84	27.31
	Max.	27.48	27.47	27.47	27.47	27.48
ADMM(SIS)	Min.	25.67	25.98	27.47	27.50	27.62
	Max.	27.82	27.86	27.71	27.71	27.71

R2018b. Further results with synthetic light field created from mouse blood vessels [21] and mouse neurons co-expressing the functional calcium indicator GCaMP8f, and the structural marker tdTomato [28] are shown in the supplement.

Performance under noise. To further analyze the performance of our method we included Poisson noise in the system as follows:

$$\mathbf{y} \sim \alpha_0 \text{Pois}(\frac{\mathbf{H}_\delta \mathbf{x}}{\alpha_0}), \tag{23}$$

which means that $\mathbf{y}$ is drawn from a Poisson distribution with mean $\frac{\mathbf{H}_\delta \mathbf{x}}{\alpha_0}$. The scalar α_0 controls the amount of noise (larger α_0 means more noise).

We use the same two datasets used in the previous section for this experiment. We synthetically created LF images with different noise levels by varying the scalar α_0 from 8 to 48, as shown in Table III. We found that both ADMM methods can manage adequately low to moderate noise conditions. Furthermore, the ADMM methods achieve the highest PSNR and SSIM for most noise conditions. Specifically, for the pollen data, the SIS prior consistently gives the best performance in most of the cases, as shown in Table III.

Similarly, for the neuron data, the ADMM without the SIS assumption achieves the best performance in most cases. Moreover, for this dataset, the SIS assumption only provides the best PSNR for minimum noise levels ($\alpha_0 = 8$). Visual results related to these simulations are shown in the supplemental material. Furthermore, the following section shows that our methods also exhibit satisfactory results for real noisy acquisition conditions.

Ablation Study. To evaluate the robustness of our algorithm, we evaluated the PSNR achieved with our ADMM methods when varying each hyper-parameter in Algorithm 1 independently, while the rest of the parameters were kept fixed using the values shown in Table V. As mentioned previously, these hyper-parameters were found manually. We use the Pollen dataset in these experiments. Note that the last column in Table V shows the PSNR achieved by the corresponding configuration.

TABLE V
ADMM HYPER-PARAMETERS.

	α	ρ_+	ρ_x, ρ_y	ρ_z	λ	PSNR
(No SIS)	2×10^{-6}	1.21	1000	100	5×10^{-4}	27.22
(SIS)	2.5×10^{-6}	1.21	2000	0	5×10^{-4}	27.70

We tested the 6 parameters of our algorithm: learning rate α, importance of derivatives $\{\rho_x, \rho_y, \rho_z\}$, non-negativity importance ρ_+ and the threshold λ for soft-thresholding. Note that we fixed ρ_x and ρ_y to be the same.

In our experiments, we found that the maximum PSNR achieved by the ADMM with the SIS prior is always larger than the maximum PSNR achieved by the ADMM without the SIS prior, as shown in Table IV. These results coincide with the results in Table II. Furthermore, the PSNR shows no significant variations when the hyperparameters are modified within the ranges in Table IV. We also note that, the non-negativity importance ρ_+ and the learning rate α have a greater impact on the reconstruction.

B. Real Data

In this section, we evaluate the proposed approaches using real light field data. The experimental settings we use for the ADMM algorithms are the same as those in the previous section. The size of the real input light field image is 2033×2033 pixels. From this single image, we reconstruct a volume of size $1712 \times 1712 \times 96$ voxels covering a depth range $-15\mu m < z < 15\mu m$.

The light field image is captured from the same 50-μm-thick mouse brain slice of the previous section. We show the light field image in Figure 9 (a). We also capture the corresponding volume with the microscope in wide-field modality by scanning it along the z-axis without the microlens array. This wide-field image stack provides an approximation of the target volume we aim to reconstruct. The in-focus plane, together with two axial slices, are shown in Figure 9 (a). Furthermore, different lateral slices corresponding to seven depths are shown in Figure 9 (b). Notice that the non-ideal optical system and the specimen type cause increased noise in the images. To avoid noise amplification, the number of iterations used for reconstruction must be chosen properly. As proposed in previous works [5], [18], [15], a typical empirical number of iterations used for ISRA is between 8 and 10. We fixed this value to 10 for ISRA, and keep 24 iterations for our ADMM methods.

We compare the conventional method ISRA with the ADMM methods with and without SIS prior. In this case, we only present visual results since there is no ground truth available. As shown in Figure 9, both ADMM methods give a better reconstruction performance than ISRA in terms of artifacts reduction. In particular, these artifacts are strongly present at $z = 0$ μm, but also at other slices close to the in-focus plane, such as $z = 3$ μm and $z = -3$ μm. In general, comparable performance is achieved for depths far from $z = 0$ μm, but artifacts near the native object plane are clearly removed when using the proposed ADMM methods.

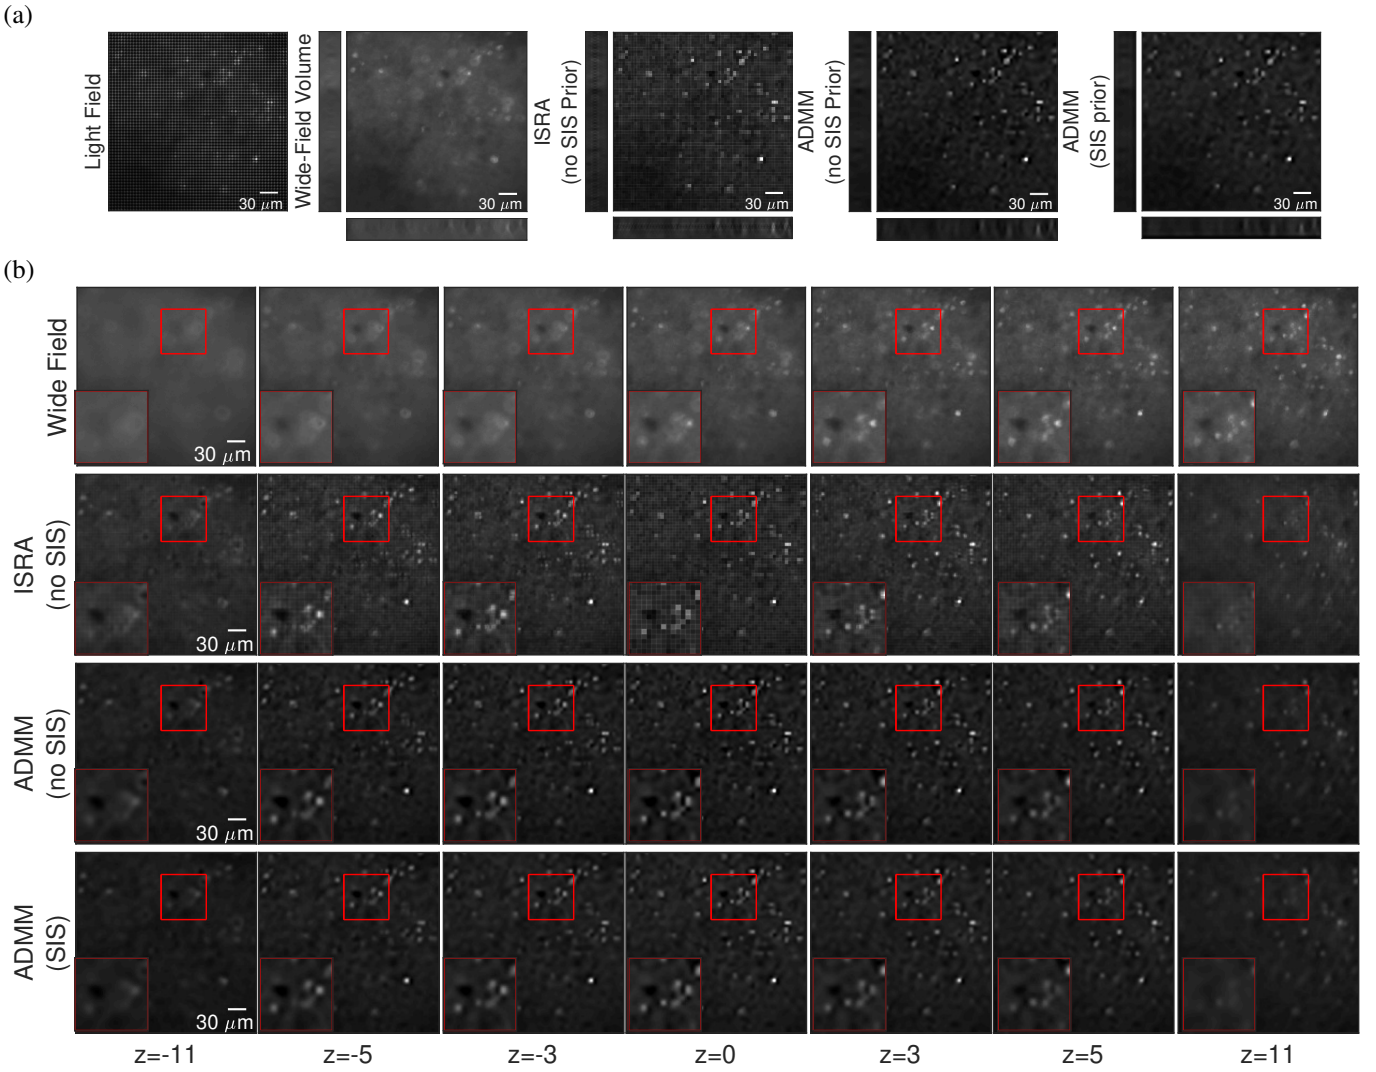

Fig. 9. Reconstruction using real light field data from a brain slice (50 μm thick) imaged from an EGFP tagged triple transgenic mouse line. In part (a), from left to right, we show the original light field image; then, for comparison, we show a wide-field image taken without the microlens array; finally, the reconstruction using ISRA, our ADMM without and with the SIS assumption. We show the in-focus plane and two xz and yz slices below and to the left of each lateral slice, respectively. In part (b), we show additional slices for different depths. All the distances are measured in μm. The settings used to capture both the light field image, and the wide-field image are specified in Section VI.

For the second experiment, we use a light field image captured from acute mouse brain slices expressing the calcium indicator NIR-GECO2G [29], [30]. Furthermore, the size of the input light field is 1881×1881 pixels. From this single image, we reconstruct a volume of size $1584 \times 1584 \times 96$ voxels covering a depth range $-30\mu m < z < 30\mu m$.

Figure 10 (a) shows the light field image and the corresponding wide-field image stack (in-focus plane with two axial slices) captured from the same volume. Furthermore, different lateral slices corresponding to seven depths of the captured volume are shown in Figure 10 (b). In this case, we find similar results as in the previous example. We observe that both ADMM methods behave similarly but remarkably remove artifacts near the native object plane, specifically at depths $z = -2\ \mu m$, $z = 0\ \mu m$ and $z = 2\ \mu m$. ISRA instead produces strong square-like artifacts in these planes.

Finally, we show how our method adapts to different types of data by performing reconstruction using the light field data provided in [22]. For this experiment, part of the previous microscope settings is modified. According to [22], the wavelength is set to 488 nm, the numerical aperture is 0.5, the magnification is 22.5, the microlens focal length is 3125 μm, and the tube lens focal length is 0.3 m. From a single light field of size 931×931 pixels, we reconstruct a volume of size $752 \times 752 \times 96$ in the range $-50\mu m < z < 50\mu m$. The image corresponds to a group of beads, as shown in Figure 11 (a). Furthermore, in this case, a high-quality image stack of the volume obtained using selective plane illumination microscopy (SPIM) is available. Additionally, the reconstruction obtained using the deep network hylfm-net from [22] is also available.

The results obtained using light field data from [22] are consistent with our previous conclusions. Figure 11 (a) and 11 (b), show that ISRA produces square-like artifacts near the native object plane while both ADMM methods do not introduce noticeable artifacts. In addition, this figure show that the

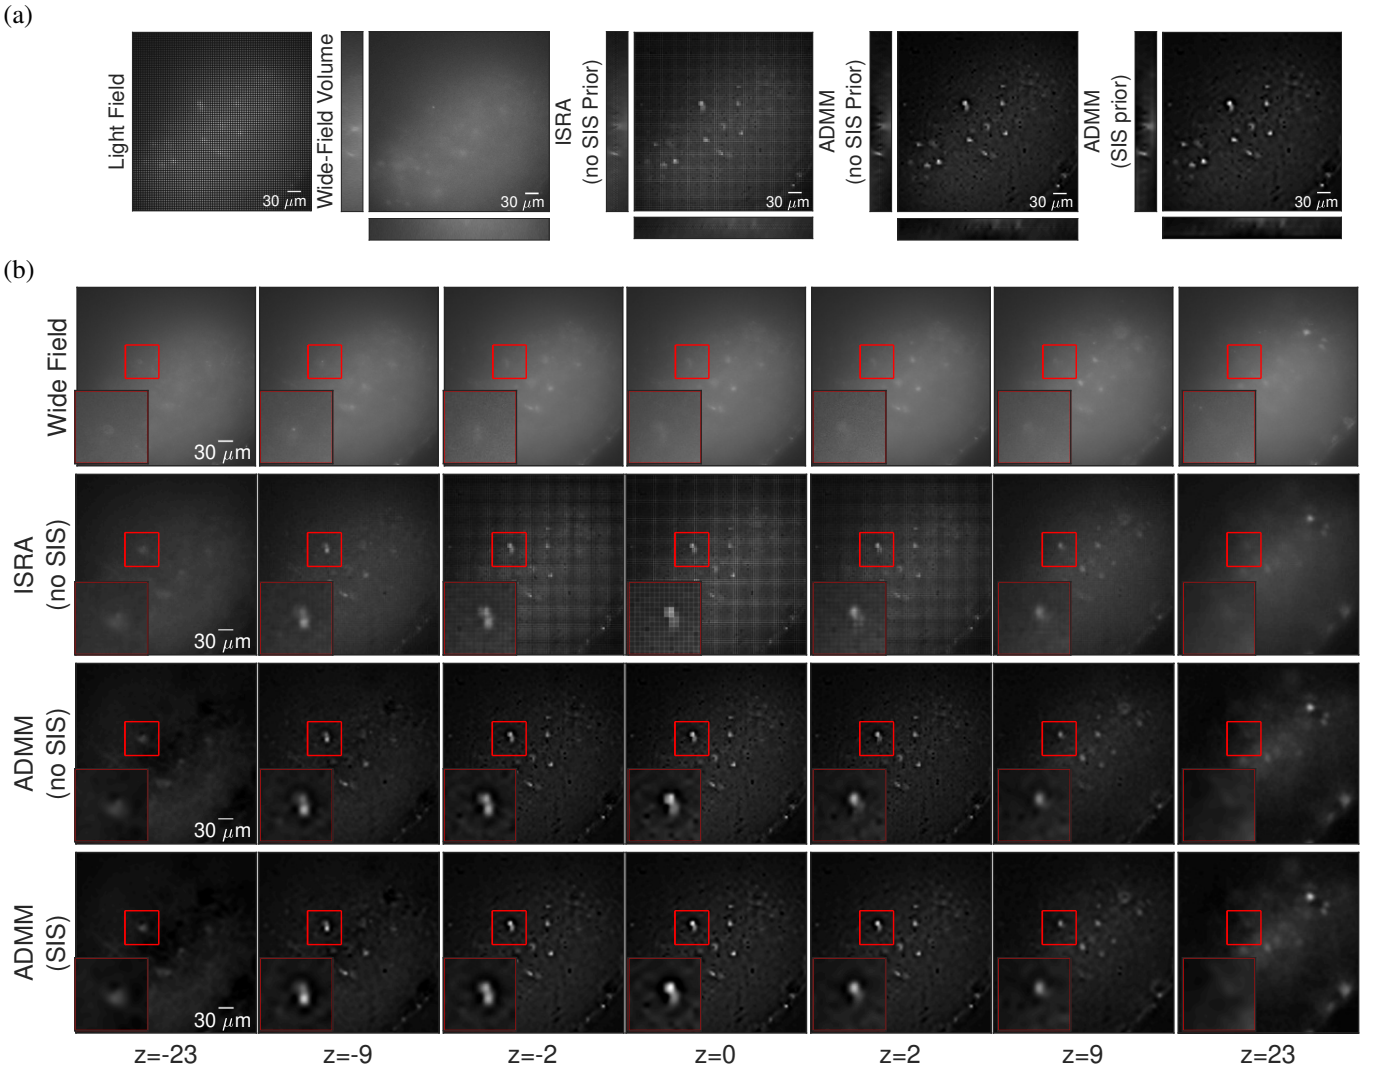

Fig. 10. Reconstruction using real light field data from acute mouse brain slices expressing the calcium indicator NIR-GECO2G [29], [30]. In part (a) from left to right, we show the original light field image, a wide-field image taken with the microscope, the reconstruction using ISRA, and our ADMM without and with the SIS assumption. We show the in-focus plane and two xz and yz slices below and to the left of each lateral slice, respectively. In part (b), we show additional slices for different depths. All the distances are measured in μm. The settings used to capture both the light field image, and the wide-field image are specified in Section VI

TABLE VI
PERFORMANCE OF LEARNING METHODS ON SYNTHETIC DATA.

	PSNR	SSIM
	Pollen	
VCDNet	22.76	0.49
hyLFMNet	22.84	0.57
	Neurons	
VCDNet	23.03	0.23
hyLFMNet	22.54	0.10

hyLFM-net can reconstruct sharp beads that are visually similar to the SPIM images because SPIM volumes were used as labels in the training stage. However, when the trained hyLFM-net [22] and VCDNet [20] were tested on synthetic light field images of pollens and neurons (shown in Figure 8), the performance of both networks degraded dramatically, as shown by the poor PSNR and SSIM scores shown in Table VI. Even though both learning approaches can perform well if tested with similar data as used in the training stage [22], [20], producing thousands of labelled data pairs for each type of sample is challenging and even unfeasible under many realistic conditions. Furthermore, in our scenario, a network trained on synthetic data will not perform well on real data without proper modelling of the background noise and light scattering, as opposed to [20]. This highlights the relevance of model-based methods in scenarios like those discussed in this paper.

Deep-Learning prior. In this section, we perform additional experiments using a self-supervised learning technique that aims to solve problems related to the lack of data for training.

Specifically, we use a deep-learning prior for reconstruction [31]. In Section V, we proposed to use a customized ADMM algorithm to perform volume reconstruction. However, ADMM is a general tool that can also be used to impose a deep prior. Therefore, another approach to solve the LFM problem is to replace the proposed SIS prior with a

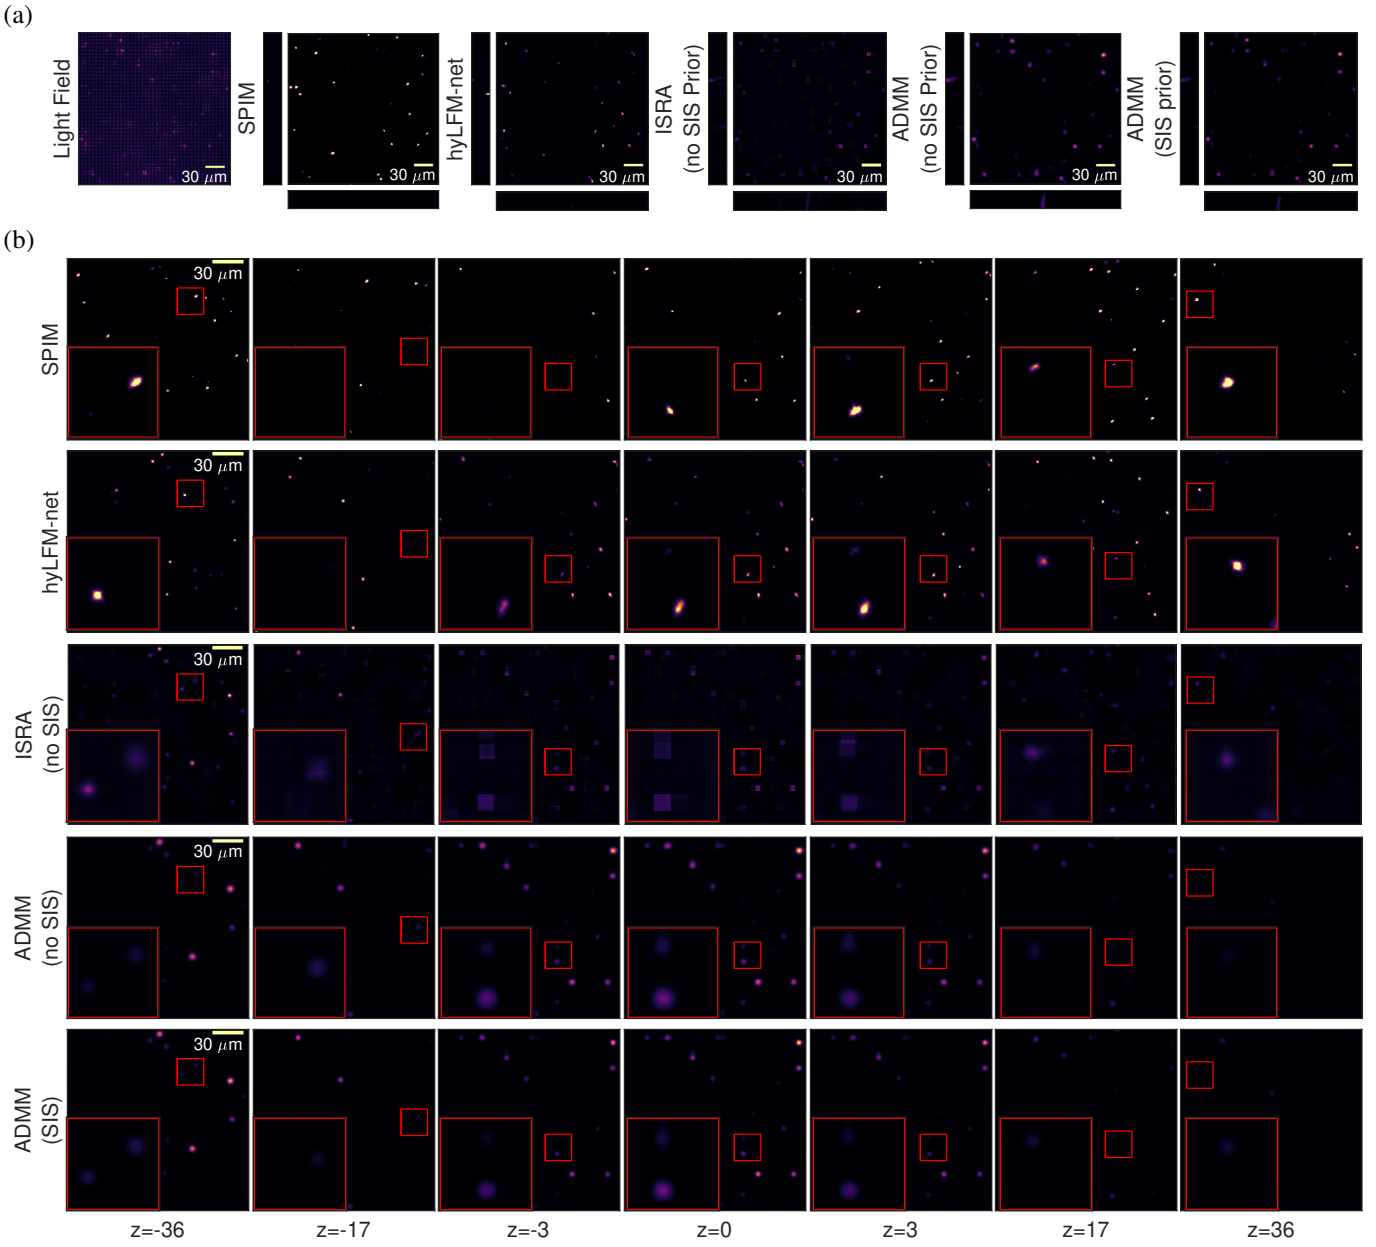

Fig. 11. Reconstruction using real light field data from beads taken from [22]. In part (a), from left to right, we show the original light field image, a SPIM image, the reconstruction using the HyLFM-Net[22], the reconstruction using ISRA, and finally, our ADMM without and with the SIS assumption. We show the in-focus plane and two xz and yz slices below and to the left of each lateral slice, respectively. In part (b), we show additional lateral slices for different depths. All the distances are measured in μm. The settings used to capture the light field image are specified in Section VI according to [22].

deep-learning prior and use the ADMM algorithm for reconstruction, as proposed in [32] for medical image restoration. Thus, in this method, we aim to solve the inverse problem by assuming that the reconstructed volume is the output of a neural network with input $\mathbf{z}$. Furthermore, an additional isotropic total variation prior is imposed as follows:

$$\begin{aligned} \min_{\theta} \quad & \|\mathbf{H}_\delta f_\theta(\mathbf{z}) - \mathbf{g}\|_2^2 + \|\mathbf{t}\|_2, \\ \text{s.t.} \quad & \mathbf{D} f_\theta(\mathbf{z}) = \mathbf{t}, \end{aligned} \tag{24}$$

where $f_\theta(\mathbf{z})$ is a neural network with learnable parameters θ, $\mathbf{z}$ is a fixed random vector and $\mathbf{D}$ is the discrete gradient operator. To solve this problem we use ADMM, as proposed in [32]. The following sub-problems must be solved for each iteration k:

$$\begin{aligned} \theta^{k+1} \in \arg\min_{\theta} \; & \frac{1}{2}\|\mathbf{H}_\delta f_\theta(\mathbf{z}) - \mathbf{g}\|_2^2 + \\ & + \frac{\beta}{2}\|\mathbf{D} f_\theta(\mathbf{z}) - \mathbf{t}^k + \frac{\lambda^k}{\beta}\|_2^2, \end{aligned} \tag{25}$$

$$\begin{aligned} \mathbf{t}^{k+1} = \arg\min_{\mathbf{t}} \; & \|\mathbf{t}\|_2^2 + \\ & + \frac{\beta}{2}\|\mathbf{t} - (\mathbf{D} f_{\theta^{k+1}}(\mathbf{z}) + \frac{\lambda^k}{\beta})\|_2^2, \end{aligned} \tag{26}$$

$$\lambda^{k+1} = \lambda^k + \beta(\mathbf{D} f_{\theta^{k+1}}(\mathbf{z}) - \mathbf{t}^k), \tag{27}$$

where Equation (25) is solved inexactly by applying a prefixed number of iterations of a gradient-based method using Adam optimizer. Equation (26) is solved directly since it is

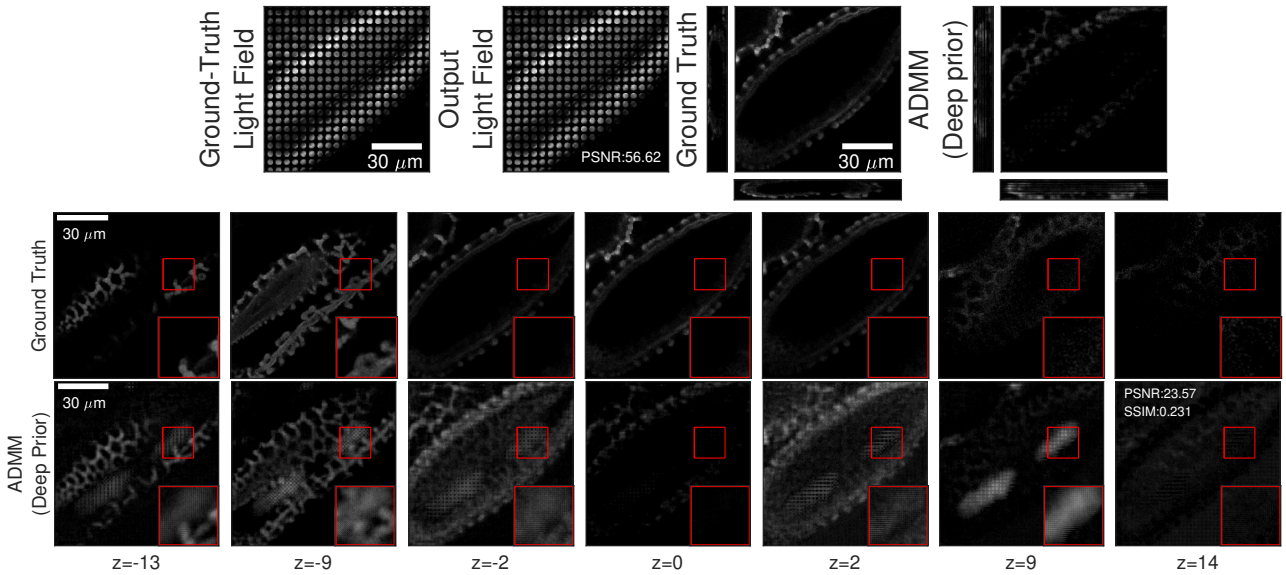

Fig. 12. Reconstruction using ADMM with deep prior. Top, we show the in-focus plane, one xz, and one yz slice of the ground truth volume, the reconstruction using ADMM with deep prior. All distances are measured in μm. Furthermore, below, we show additional slices for different depths. The shown PSNR and SSIM correspond to the whole volume. We use the Lilium Longiflorum Pollen [25] dataset. The light field images were created synthetically from the 3D volumes.

a least-square minimization, and the last equation (Eq. (27)) simply updates the dual variable λ.

The deep-learning-prior approach is evaluated using the pollen dataset described in the previous section. For this experiment, $\mathbf{z}$ is a random vector sampled from a uniform distribution. The structure of the network $f_\theta(\cdot)$ is a 2D-UNet, as proposed in previous works [21], [20] for supervised learning for LFM.

In Figure 12, we show the results after solving Equation (24) iteratively. The recovered volume has limited quality (PSNR=23.47) and suffers from reconstruction artifacts at different depths, for instance, at $z = -2$ and $z = 9$. These artifacts are not the classic square-like artifacts close to the native object plane; they appear to be due to the structure of the network used. However, part of the volume structure is recovered at specific depths, for instance, at $z = -9$ and $z = -13$. As shown in the supplemental material, in the same conditions our model-based ADMM achieves PSNR 31.32. Furthermore, note in Figure 12 that the light field image re-synthesised from the reconstruction is almost identical to the input light field image (PSNR=56.62) since the problem is highly ill-posed. Hence, the proposed priors and settings are not enough to avoid obtaining undesired solutions for this particular experiment.

In addition, we would like to emphasize that this approach is highly demanding in terms of computation since one artificial neural network must be "trained" for each reconstruction. Specifically, to obtain the result shown in Figure 12, we fixed the number of inner iterations to 100, and we ran 1000 iterations for the outer loop [32]. As mentioned in the introduction, LFM usually seeks to reconstruct volume time series. Therefore, reconstruction approaches for LFM should take computational complexity into account. In our approach, we successfully achieved a balance between complexity and reconstruction performance.

VII. Conclusion

We have proposed a novel discretization method that uses more diverse sampling densities and shift-invariant subspaces to sample the impulse response of the light field system. In an ideal scenario, perfect reconstruction can be achieved by using the pseudoinverse; however, this depends on the proper selection of the template functions and shifts defining the shift-invariant subspace. Moreover, we propose to model the system using filter banks. This model leads to a method that accelerates the computation of the forward model by using SVD.

Finally, we propose a new 3D reconstruction algorithm for light field microscopy with improved PSNR and SSIM, fewer artifacts, and faster speed (lower computational complexity) than conventional methods. This is achieved by incorporating additional priors and using a specific algorithm based on ADMM. The improvement in computational time is due to the exploitation of the underlying low-rank property of the measurement matrix by either using a specific SVD or assuming that the reconstructed volume lies in a SIS. Numerical results on both synthetic and real data also demonstrate that our approach consistently outperforms conventional volume reconstruction approaches.

Acknowledgment

The authors would like to thank Gerald Moore for preparing the EGFP-tagged brain sample, as well as Robert Campbell for the calcium indicator NIR-GECO2G.

References

[1] M. Levoy and P. Hanrahan, "Light field rendering," in *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '96. New York, NY, USA: ACM, 1996, pp. 31–42. [Online]. Available: http://doi.acm.org/10.1145/237170.237199

[2] D. G. Dansereau, "Plenoptic signal processing for robust vision in field robotics," Ph.D. dissertation, 2014.

[3] R. Farrugia and C. Guillemot, "Light field super-resolution using a low-rank prior and deep convolutional neural networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, p. 1, Jan 18, 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8620368

[4] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz, "Light field microscopy," *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 924–34, 2006. [Online]. Available: http://dx.doi.org/10.1145/1141911.1141976

[5] T. Nöbauer, O. Skocek, A. P.-A. J., L. Weilguny, F. M. Traub, M. I. Molodtsov, and A. Vaziri, "Video rate volumetric ca2+ imaging across cortex using seeded iterative demixing (sid) microscopy," *Nature Methods*, vol. 14, p. 811, 2017. [Online]. Available: https://doi.org/10.1038/nmeth.4341

[6] H. Li, C. Guo, and S. Jia, "High-resolution light-field microscopy," in *Frontiers in Optics 2017*. Optical Society of America, 2017, p. FW6D.3. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=FiO-2017-FW6D.3

[7] C. Guo, W. Liu, X. Hua, H. Li, and S. Jia, "Fourier light-field microscopy," *Opt. Express*, vol. 27, no. 18, pp. 25 573–25 594, Sep 2019. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-27-18-25573

[8] M. Broxton, L. Grosenick, S. Yang, A. A. N. Cohen and, K. Deisseroth, and M. Levoy, "Wave optics theory and 3-d deconvolution for the light field microscope," *Optics express*, vol. 21, no. 21, pp. 25 418–25 439, 2013.

[9] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, "The lumigraph," in *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '96. New York, NY, USA: ACM, 1996, pp. 43–54. [Online]. Available: http://doi.acm.org/10.1145/237170.237200

[10] R. Ng, M. Levoy, M. B. 'edif, G. Duval, M. Horowitz, and P. Hanrahan, "Light field photography with a hand-held plenoptic camera," Tech. Rep., apr 2005. [Online]. Available: http://graphics.stanford.edu/papers/lfcamera/

[11] I. Ihrke, J. Restrepo, and L. Mignard-Debise, "Principles of light field imaging: Briefly revisiting 25 years of research," *IEEE Signal Processing Magazine*, vol. 33, no. 5, pp. 59–69, 2016.

[12] N. Cohen, S. Yang, A. Andalman, M. Broxton, L. Grosenick, K. Deisseroth, M. Horowitz, and M. Levoy, "Enhancing the performance of the light field microscope using wavefront coding," *Optics Express*, vol. 22, no. 20, pp. 24 817–24 839, 2014, j2: Opt. Express. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-22-20-24817

[13] K. He, X. Wang, Z. W. Wang, H. Yi, N. F. Scherer, A. K. Katsaggelos, and O. Cossairt, "Snapshot multifocal light field microscopy," *Opt. Express*, vol. 28, no. 8, pp. 12 108–12 120, Apr 2020. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-28-8-12108

[14] Z. Zhang, L. Bai, L. Cong, P. Yu, T. Zhang, W. Shi, F. Li, J. Du, and K. Wang, "Imaging volumetric dynamics at high speed in mouse and zebrafish brain with confocal light field microscopy," *Nature biotechnology*, vol. 39, no. 1, pp. 74–83, 2021, iD: Zhang2021. [Online]. Available: https://doi.org/10.1038/s41587-020-0628-7

[15] R. Prevedel, Y. Yoon, M. Hoffmann, N. Pak, G. Wetzstein, S. Kato, T. Schrödel, R. Raskar, M. Zimmer, E. S. Boyden, and A. Vaziri, "Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy," *Nature methods*, vol. 11, no. 7, p. 727, 2014.

[16] M. E. Daube-Witherspoon and G. Muehllehner, "An iterative image space reconstruction algorthm suitable for volume ect," *IEEE transactions on medical imaging*, vol. 5, no. 2, pp. 61–66, 1986.

[17] Z. Lu, J. Wu, H. Qiao, Y. Zhou, T. Yan, Z. Zhou, X. Zhang, J. Fan, and Q. Dai, "Phase-space deconvolution for light field microscopy," *Opt.Express*, vol. 27, no. 13, pp. 18 131–18 145, Jun 2019. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-27-13-18131

[18] A. Stefanoiu, J. Page, P. Symvoulidis, G. G. Westmeyer, and T. Lasser, "Artifact-free deconvolution in light field microscopy," *Opt. Express*, vol. 27, no. 22, pp. 31 644–31 666, Oct 2019. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-27-22-31644

[19] P. Song, H. Verinaz-Jadan, C. L. Howe, P. Quicke, A. J. Foust, and P. L. Dragotti, "3D localization for light-field microscopy via convolutional sparse coding on epipolar images," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1017–1032, 2020.

[20] Z. Wang, L. Zhu, H. Zhang, G. Li, C. Yi, Y. Li, Y. Yang, Y. Ding, M. Zhen, S. Gao *et al.*, "Real-time volumetric reconstruction of biological dynamics with light-field microscopy and deep learning," *Nature Methods*, vol. 18, no. 5, pp. 551–556, 2021.

[21] J. Page, F. Saltarin, Y. Belyaev, R. Lyck, and P. Favaro, "Learning to reconstruct confocal microscope stacks from single light field images," 2020.

[22] N. Wagner, F. Beuttenmueller, N. Norlin, J. Gierten, J. C. Boffi, J. Wittbrodt, M. Weigert, L. Hufnagel, R. Prevedel, and A. Kreshuk, "Deep learning-enhanced light-field imaging with continuous validation," *Nature Methods*, vol. 18, no. 5, pp. 557–563, 2021.

[23] C. Belthangady and L. A. Royer, "Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction," *Nature methods*, vol. 16, no. 12, pp. 1215–1225, 2019.

[24] N. Chacko, M. Liebling, and T. Blu, "Discretization of continuous convolution operators for accurate modeling of wave propagation in digital holography," *JOSA A*, vol. 30, no. 10, pp. 2012–2020, 2013.

[25] A. Carpenter, "Drosophila melanogaster kc167 cell." [Online]. Available: http://cellimagelibrary.org/images/21766

[26] L. Donati, E. Soubies, and M. Unser, "Inner-loop-free ADMM for cryo-em," in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, 2019, pp. 307–311.

[27] L. Madisen, A. R. Garner, D. Shimaoka, A. S. Chuong, N. C. Klapoetke, L. Li, A. Van Der Bourg, Y. Niino, L. Egolf, C. Monetti *et al.*, "Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance," *Neuron*, vol. 85, no. 5, pp. 942–958, 2015.

[28] Y. Zhang, M. Rózsa, D. Bushey, J. Zheng, D. Reep, G. Broussard, A. Tsang, G. Tsegaye, R. Patel, S. Narayan *et al.*, "jgcamp8 fast genetically encoded calcium indicators," *Janelia Research Campus*, vol. 10, p. 13148243, 2020.

[29] Y. Qian, D. M. O. Cosio, K. D. Piatkevich, S. Aufmkolk, W.-C. Su, O. T. Celiker, A. Schohl, M. H. Murdock, A. Aggarwal, Y.-F. Chang *et al.*, "Improved genetically encoded near-infrared fluorescent calcium ion indicators for in vivo imaging," *PLoS Biology*, vol. 18, no. 11, p. e3000965, 2020.

[30] Y. Qian, K. D. Piatkevich, B. Mc Larney, A. S. Abdelfattah, S. Mehta, M. H. Murdock, S. Gottschalk, R. S. Molina, W. Zhang, Y. Chen *et al.*, "A genetically encoded near-infrared fluorescent calcium ion indicator," *Nature methods*, vol. 16, no. 2, pp. 171–174, 2019.

[31] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9446–9454.

[32] P. Cascarano, A. Sebastiani, M. C. Comes, G. Franchini, and F. Porta, "Combining weighted total variation and deep image prior for natural and medical image restoration via admm," *arXiv preprint arXiv:2009.11380*, 2020.