



# 3-D visualization of transparent fluid flows from snapshot light field data

Martin Eberhart<sup>1</sup> · Stefan Loehle<sup>1</sup> · Philipp Offenhäuser<sup>2</sup>

Received: 28 January 2021 / Revised: 23 June 2021 / Accepted: 27 June 2021 / Published online: 15 July 2021  
© The Author(s) 2021

## Abstract

This paper presents the use of light field data, recorded in a snapshot from a single plenoptic camera, for 3-D visualization of transparent fluid flows. We demonstrate the transfer of light field deconvolution, a method so far used only in microscopy, to macroscopic scales with a photographic setup. This technique is suitable for optically thin media without any additional particles or tracers and allows volumetric investigation of non-stationary flows with a simple single camera setup. An experimental technique for the determination of the shift-variant point spread functions is presented, which is a key for applications using a photographic optical system. The paper shows results from different test cases with increasing complexity. Reconstruction of the 3-D positions of randomly distributed light points demonstrates the achievable high accuracy of the technique. Gas flames and droplets of a fluorescent liquid show the feasibility of the proposed method for the visualization of transparent, luminous flows. The visualizations exhibit high quality and resolution in low-contrast flows, where standard plenoptic software based on computer vision fails. Axial resolution depends on the data and is about an order of magnitude lower than the lateral resolution for simple point objects. The technique also allows the time-resolved analysis of flow structures and the generation of 3D3C-velocity fields from a sequence of exposures.

## 1 Introduction

Fluid flows involve complex three-dimensional structures and interactions on different spatial scales, and their volumetric investigation and visualization are of utmost interest for the understanding of flow physics. Modern techniques are often based on tomography, where the volume is reconstructed from multiple projections under different viewing angles. Exploiting flow parameters such as refractive index, luminosity or the positions of added tracers allows to derive volumetric results from a set of 2D measurements (Cai et al. 2013). Elsinga et al. (2006) brought particle imaging velocimetry (PIV) to the third dimension by photographing particles within the flow from a number of viewpoints. The 3-D particle positions were determined using an algebraic tomography algorithm, and the shifts between subsequent

frames of an image sequence led to a three-dimensional, three-component velocity field. Anikin et al. (2010) also followed a tomographic approach to investigate the chemiluminescence emission within flames and to analyze the distribution of excited  $\text{OH}^*$  molecules. In our group, Hermann et al. (2016) applied tomography to the measured light intensity of a plasma flow, which revealed a pronounced asymmetry in the volumetric radiation field at short time scales. Halls et al. (2016) and Upton et al. (2011) investigated 3-D turbulent structures in complex flows by an arrangement of high-resolution cameras in combination with a tomographic reconstruction method. Atcheson et al. (2008) exploited the relation between density and refractive index in the hot plume of a flame: They used a background oriented schlieren technique (BOS) to monitor the apparent 2-D shifts of a background pattern, viewed through the flow by a set of cameras, and determined the time-resolved 3-D density distribution with the aid of a tomographic reconstruction. Tracers are used in the technique described by Lynch and Thurow (2012), where small oil droplets are introduced into a turbulent jet. The authors do not use tomography, but instead sweep a laser sheet rapidly through the volume, and record the illuminated droplets with a high-speed camera to yield an almost instantaneous visualization of the flow field.

✉ Martin Eberhart  
eberhart@irs.uni-stuttgart.de

<sup>1</sup> High Enthalpy Flow Diagnostics Group (HEFDiG), Institute of Space Systems, University of Stuttgart, Pfaffenwaldring 29, 70569 Stuttgart, Germany

<sup>2</sup> Hewlett PackardEnterprise (HPE), Herrenberger Straße 140, 71034 Böblingen, Germany

All these techniques require a complex measurement setup including several cameras, an arrangement of mirrors or an elaborated lighting system. This is critical if optical access to the flow is limited. In such cases, a technique is favorable which operates on data recorded from a single viewpoint, and non-stationary flows additionally demand a snapshot measurement in a single exposure. An instrument for the instantaneous acquisition of 3-D data that has raised considerable interest in the field of fluid flow investigations is the plenoptic camera (Adelson and Wang 1992; Ng 2006). Compared to a conventional photographic camera, it features an additional array of microlenses close to the image sensor. Incoming light rays are sorted onto different pixels, depending on their direction, and this additional directional information is contained in the raw image. Appropriate decoding of the spatio-angular data allows to extract information on the depth coordinates within the scene, which permits to derive, e.g., 3-D positions of particles within a flow.

In recent years, an increasing number of articles have been published that combine three-dimensional plenoptic imaging with PIV or particle tracking velocimetry (PTV) to yield 3-D, three-component (3D3C) velocity fields of particle-loaded flows from single camera data. Lillo et al. (2014) used a commercial Raytrix camera and the associated software RxLive to analyze the evolution of the 3-D structure of fuel sprays from an automobile injector, however, without deriving any velocity data. This was shown, e.g., by Chen and Sick (2017), who used the commercial Raytrix software RxFlow to perform a PTV analysis of the flow field within a combustion engine and computed 3D3C values from a plenoptic image sequence. Tan et al. (2020) developed a modular plenoptic imaging system that can be mounted on various cameras and used it for high-speed recording of the flow induced by the propelling motion of marine jellyfish in a water tank. Tracer particles allowed to derive the three-dimensional flow field. Fahringer and Thurow (2018) considerably increased the depth resolution of PIV measurements in a particle-seeded flow by using a set of multiple plenoptic cameras.

All these publications deal with flows carrying tracer particles, which present regions of high local contrast in the image. In such cases, algorithms from the field of computer vision are very effective in computing depth coordinates, and the commercial software RxLive allows to generate depth maps of a scene in almost real time. Things are different without any particles or other sources of contrast like edges, surface textures or patterns. This includes fluid flows where no tracers have been added, luminous plasmas or flames (without soot particles) in combustion research or mixing of fluorescent liquids. Standard correlation-based algorithms perform very poorly (Greene and Sick 2013) for such low-contrast problems, and alternative approaches are required to reconstruct the volume of the flow. In this paper, we put

forward light field deconvolution for the volumetric visualization of particle-free, transparent, and luminous flows. This method has been developed for microscopy applications on very small spatial scales, where the special optical properties of a microscope prove to be beneficial. We transfer light field deconvolution to macroscopic scales using a photographic setup (Eberhart and Loehle 2021) and present an experimental calibration approach, which is required to characterize the complex optical system. As a scanless technique operating on a single snapshot recording, light field deconvolution allows the volumetric investigation of non-stationary flows under conditions of limited optical access. Measurements are performed with a single camera, which eliminates the need for complex alignment, image registration and precise triggering required for setups with multiple instruments.

The article is organized as follows: After a brief introduction to plenoptic cameras and light field deconvolution in Sect. 2, we present our approach for the experimental determination of the point spread function (PSF) in Sect. 3, which is a key for the calibration of the optical system. Using the acquired PSF, we perform light field deconvolution on test cases recorded by a plenoptic camera: A first test on a simple planar stripe pattern shows the robustness of the method throughout the field of view. Reconstruction of the 3-D positions of a point cloud builds intuition about the achievable depth resolution, before testing the method on real transparent, luminous flows, an ensemble of gas flames and a fluorescent liquid in a water tank. Results are presented in Sect. 4, before closing the article by considering computational aspects in Sect. 1.

## 2 Background

### 2.1 Plenoptic imaging

A digital sensor at the image plane of a conventional photographic camera measures the intensity distribution of light transferred by the main lens, and each pixel integrates the incident radiation over a certain solid angle. Directional information on the light rays is lost as a result of this integration. A plenoptic camera features an additional array of microlenses (MLA), placed at some distance in front of the sensor, which allows to additionally capture this lost information (Lippmann 1908): Depending on its direction, the microlenses distribute incoming light to different locations on the image sensor beneath. Even though the image is still recorded by a flat, two-dimensional sensor, it now holds spatio-angular data, the so-called light field, which defines the transport of light energy along rays in space (Ng 2006). Information about the third coordinate of the scene is coded into the camera's pixel values. Appropriate decoding of the recorded data allows to extract the light field of the object

space and to reconstruct its volume. The MLA can be understood as a multiplexer (Wetzstein et al. 2013), distributing the depth content of the scene within the camera image, or, alternatively, as a transformer (Wender et al. 2015) that structures the light field before its acquisition by the sensor. An example of a raw image of such a system is given in Fig. 1 on the left. Each lenslet creates a circular microimage, which is clearly visible in the zoomed-in inset. The overall image is structured by an overlay of two coordinate systems, an  $(s, t)$ -system defining the position of the microimages, and a  $(u, v)$ -system of the pixel locations within them.

In the present work, we used a commercial camera Raytrix R29 (monochrome, type hr29050MFLCPC), a camera of the so-called focused type (Lumbsdaine and Georgiev 2009), often termed plenoptic 2.0, with a schematic sketched on the right of Fig. 1. The main lens forms a miniaturized, three-dimensional image of the object world within the camera, which is picked up by the microlenses. They act like a large number of miniaturized cameras with varying viewing angles and form sharp sub-images of the object on the sensor plane. Lenslet focal length  $f$  and stand-off  $b$  match the distance  $a$  according to the thin-lens equation. The parameter  $a$  is a function of the object distance, and so the camera's total depth of field can be increased by using an MLA composed of lenslets with different focal lengths (Georgiev and Lumbsdaine 2012). A Raytrix R29 is such a multifocus plenoptic camera with three lenslet types arranged in a hexagonal grid (Perwaß and Wietzke 2012). Details of the optical setup used in this work are summarized in Table 1.

The plenoptic camera Raytrix R29 was equipped with a Nikkor 200 mm f/4 Micro main lens at a working distance of 250 mm, which yields a 1:1 magnification. The f-number of the microlenses has to match the image-sided f-number of the main lens, defined as the distance between its rear principal plane and the camera's imaging plane (Ng 2006; Perwaß and Wietzke 2012). This distance changes with focus settings, and the aperture of the main lens has to be adjusted manually such that the single microimages touch without overlapping.

From a photographer's point of view, the extra angular information recorded by a plenoptic camera opens the intriguing possibility to render 2-D images with different

**Table 1** Details of used optical setup

Camera	Raytrix R29 (hr29050MFLCPC)
Interface	Camera Link
Sensor type	CCD, monochrome
Sensor size	full frame 35mm
Pixel resolution	6576 x 4384 (29MP)
Pixel pitch	5.5 $\mu\text{m}$
Bit depth	8 bit
Shutter	electronic, global
MLA type	hexagonal, multi-focus (3 lens types)
MLA f-number	7
MLA pitch	31×31 pixels
Framerate	max. 6 fps
Main lens	AF Micro-Nikkor 200 mm 1:4D IF-ED
Working distance	250 mm
Magnification	1:1
f-number	7-8, matching MLA and working distance
Volume sizes	max. 36×24×20 mm (x,y,z)


perspectives (Levoy and Hanrahan 1996) or different focal planes (Isaksen et al. 2000) from a single snapshot. However, there is an obvious caveat of this technique, as the total information that can be captured by an imaging sensor is limited: Irrespective of the type of plenoptic camera, the directional content has to be paid for by trading in effective lateral resolution of the sensor for additional depth information.

## 2.2 Volume reconstruction by light field deconvolution

Light field deconvolution has been first introduced by Broxton et al. (2013) for volumetric investigations under a microscope. For an in-depth reading, we refer to the original publication and give a brief overview of the technique in the following.

In general, deconvolution methods are based on a linear image formation model according to

**Fig. 1** *Left:* Exemplary raw image of a plenoptic camera, here a Lytro Illum with color sensor. *Right:* Schematic sketch of a plenoptic camera of the focused type, used in Raytrix models



$$f(x_i, y_i) = g(x, y, z) * h + n \quad (1)$$

Here, the recorded image  $f$ , with coordinates  $x_i$  and  $y_i$ , is a convolution of the volume  $g$  with the spatial impulse response  $h$  under consideration of an additional noise term  $n$ . Lateral object space dimensions are denoted  $x$  and  $y$ , and  $z$  is the depth coordinate along the optical axis. If we discretize the problem to be solved on a computer, the image  $f$  is divided into  $N_p = n_{x_i} \times n_{y_i}$  pixels and the volume  $g$  is made up of  $N_v = n_x \times n_y \times n_z$  voxels.

The spatial impulse response is commonly termed Point Spread Function (PSF) and defines the spreading of light, emitted by a point source, through the optical system onto the image plane. It models the blur observed in out-of-focus regions, as well as the imaging quality within the focal plane. For diffraction-limited systems like ideal microscopes, the PSF takes the form of the Airy-pattern, in photographic systems, it is furthermore affected by various sources of optical aberrations (Gross et al. 2005). The MLA within a plenoptic camera picks up the PSF of the main lens and generates a distinct spot pattern on the sensor beneath, which typically spans several microimages. This pattern represents the light field PSF and carries information on the 3-D position of an isotropic point emitter within the volume (Broxton et al. 2013).

The PSF of a standard microscope without an MLA is shift-invariant, i.e., it does not depend on the position of the point source. This is due to the system's telecentric design, but this property is lost when an MLA is inserted into the optical path: The pattern of the light field PSF is a function of the 3-D position of the point source, so that an emitter at  $(x, y, z)$  within the volume produces an explicit intensity distribution at the  $(x_i, y_i)$  image plane. As a result, the single PSF  $h$  has to be replaced by a PSF matrix  $H$ . For every object space position, it holds the generated  $(x_i, y_i)$ -intensity distribution:

$$H = H(x_i, y_i, x, y, z) \quad (2)$$

The image formation of Eq. 1 may be written in a discretized form, with the convolution operation expressed via a matrix-vector multiplication:

$$\mathbf{f} = H\mathbf{g} \quad (3)$$

Here, the recorded noisy image has been rearranged into a column vector  $\mathbf{f}$  with  $N_p$  pixels, the volume is contained in a column vector  $\mathbf{g}$  having  $N_v$  voxels and the PSF matrix  $H$  is rearranged to have dimensions  $N_p \times N_v$ .

Deconvolution seeks to revert the image formation process given in Eq. 3, effectively trying to use a known PSF as a tool to recover the volumetric intensity distribution  $\mathbf{g}$  from a measured image  $\mathbf{f}$ . However, because the images are noisy, the problem is ill-conditioned and a simple inversion of Eq. 1 fails,

as it would result in a severe noise amplification. A solution for the volume  $\mathbf{g}$  is therefore computed by means of iterative deconvolution methods, and an overview of common algorithms can be found, e.g., in (Sage et al. 2017). Light field deconvolution has so far relied mostly on the classical Richardson–Lucy scheme, which assumes Poisson-distributed noise within the measured image. Its iterative update scheme in matrix-vector notation, with  $k$  being the iteration index, reads

$$\mathbf{g}^{k+1} = \frac{\mathbf{g}^k}{H^T \mathbf{1}} \left( H^T \frac{\mathbf{f}}{H \mathbf{g}^k} \right) \quad (4)$$

The matrix  $H^T$  has the interesting property of defining a back-projection of single image pixels into the volume (Broxton et al. 2013), while a forward projection of the object space onto the image plane is defined by  $H^T$ .

In summary, the algorithm computes an error quotient by comparing the measured image  $\mathbf{f}$  to the forward projection of the current volume estimate  $H\mathbf{g}^k$  and then backprojects this error by means of  $H^T$  to update the volumetric intensity distribution. Richardson–Lucy is a maximum-likelihood based method and as such tends to over-fit noise at the point of convergence (Sage et al. 2017). This is counterbalanced by stopping the algorithm early, which acts as a pseudo-regularization. In our experiments, we found eight iterations to produce a good balance between deconvolution, noise amplification, and computing time.

Due to the regular arrangement of the lenslets in the MLA, the pattern of the light field PSF is periodically repeating upon lateral shifts of the point source. This means that only a representative sub-matrix of  $H$  has to be determined and stored in memory, which drastically reduces both the computational burden of the reconstruction and the effort required for the identification of  $H$ . The observed periodicity is, however, only given in combination with a shift-invariant PSF of the main lens (or the objective in a microscope). This does not hold for typical photographic systems, including plenoptic cameras, and we will deal with the effect of this imperfection as part of the discussed test cases in Sect. 4.1. From the preceding paragraph, it is obvious that a PSF matrix  $H$ , which accurately represents the optical system, is key for successful light field deconvolution. In the following, we turn our attention to the acquisition of such a matrix for the case of a photographic camera, which is optically less well defined than a microscope, and where the optical path of the experimental setup may contain additional elements.

### 3 Experimental PSF acquisition

Light field deconvolution in the context of microscopy has so far relied on simulated PSFs, computed on the basis of wave optical theories (Broxton et al. 2013; Prevedel



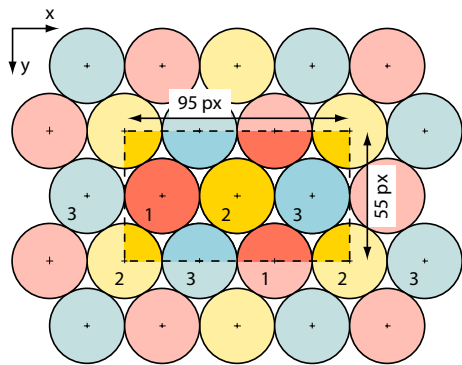
et al. 2014; Lu et al. 2019; Stefanoiu et al. 2019). This is a meaningful approach for two reasons: First, the optical system of a microscope is precisely defined and can be modeled accurately on a computer. And second, an experimental calibration of the matrix  $H$  would require to position sub-resolution-sized light sources with nanometer accuracy within the measurement volume. The situation is reversed for a photographic camera: As we do not work with a high magnification main lens, the spatial accuracy required for positioning the point source is drastically reduced. Photographic lenses, on the other hand, are a complex ensemble of multiple optical elements, some of them shifting relative to the others to allow focusing, including non-circular apertures. A photographic setup often involves additional filters or windows, and it is very difficult to represent all such influences in a simulation. For a plenoptic camera, we therefore propose an experimental approach for the determination of the matrix  $H$ , which is detailed in the following.

### 3.1 Definition of an elementary cell

Experimental PSF calibration of a plenoptic camera means imaging a point source at a high number of voxel positions. As outlined in the preceding section, only a representative part of the volume needs to be calibrated due to the periodicity of the light field patterns. We call this part an elementary cell, and its size and shape are determined by the arrangement of the microlenses in the MLA. A sketch of the lenslet array used in a Raytrix R29 is given in Fig. 2. It is composed of lenses with three different focal lengths, indicated by color, which are arranged in a hexagonal grid. The entire sensor area can be tiled by repeating copies of the elementary cell, indicated by a dashed rectangle. With the given pixel pitch of the camera, the cell has a size of 95×55 pixels in the  $x$ - and  $y$ -direction, respectively. This means that for each  $z$ -plane, the camera’s response to point sources at the corresponding 5225 voxel positions has to be determined.


The geometrical boundaries of the elementary cell have to be found experimentally in a pre-calibration step.

Here, the point source is shifted in sub-voxel steps, and we compare the respective patterns with the one recorded at the origin via a cross-correlation. The left of Fig. 3 gives the result  $r_{0a}^x$  for the horizontal  $x$ -direction, where maxima at around  $x = 0.5$  mm represent borders of the elementary cell, where patterns begin to repeat. The camera’s main lens is not telecentric, and so the field of view is extending cone-like into object space. The graph includes measurements for a number of different  $z$ -values, and cell borders obviously get wider with increasing distance from the camera. Analysis yields the lateral  $x$ -width of the elementary cell  $\Delta x$  in object space as a function of the depth coordinate  $z$ , which is plotted on the right of the figure and shows a linear dependence. The lateral voxel size is obtained by dividing the cell width by the number of positions as sketched in Fig. 2, which also defines the single shifts of the point source during calibration.



**Fig. 2** Hexagonal layout of the microlens array in an R29 camera. Colors indicate three lenslet types with different focal length. The rectangle defines the elementary cell used for PSF calibration, with a total of 95×55 positions

**Fig. 3** *Left:* Cross-correlation (normalized) of recorded point source PSF patterns along the horizontal  $x$ -axis and the pattern recorded at the origin. Colors indicate  $z$ -depth. *Right:*  $x$ -position of peak cross-correlation as a function of  $z$ -depth




### 3.2 Calibration stage

Discretization of the volume and the sensor's pixel pitch is tied together by the magnification of the main lens, which means that the voxel size is a function of the chosen focal length and the working distance. In the used setup with the parameters of Table 1, the magnification is 1, and so the pixel pitch translates to a lateral voxel size of  $5.5 \mu\text{m}$ . This does not directly correspond to the lateral resolution, which is in the order of  $100\text{--}200 \mu\text{m}$  and is explained in Sect. 4.2. Voxel sizing in the z-direction, i.e., the depth discretization, can be chosen freely and determines the number of individual z-planes that must be considered in the calibration runs.

Positioning has to be very precise in all three dimensions and is accomplished in this study by a custom-made translation stage shown in Fig. 4. It is designed from three individual linear micrometer stages, each driven by computer-controlled stepping motors which a positioning accuracy in the order of  $1\text{--}2 \mu\text{m}$ , well below the voxel size. The orientation of the three object space axes is sketched additionally, with the origin at the intersection of the optical axis and the native object plane (NOP), the focal plane of the main lens.

To represent a spatial impulse, the size of the point source has to be below the pixel resolution of the camera. It is realized by a small circular aperture with a diameter of  $5 \mu\text{m}$ , illuminated from the back by a diffuse LED.

Despite working with a monochrome camera, the wavelength of light must not be neglected: Due to chromatic aberrations at the main lens and especially at the uncorrected MLA, different wavelengths generate slightly different



**Fig. 4** Computerized three-axes translation stage used for PSF calibration and definition of the coordinate system

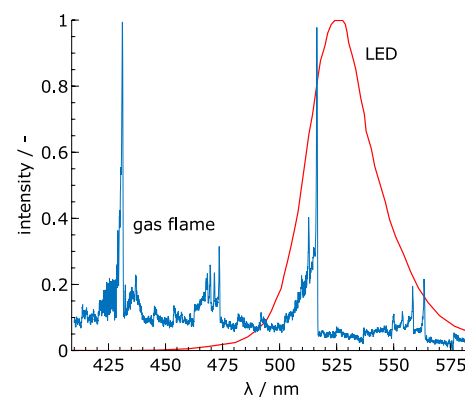
patterns at the image sensor. This means that the emission spectrum of the light source used during calibration should ideally match the spectrum of the measurement. The present study used gas flames and fluorescent droplets as test objects, and analysis of the flame spectrum using an Echelle spectrometer (LTB Aryelle 150) revealed a prominent emission peak at around  $516 \text{ nm}$ , with a second blueish peak at  $430 \text{ nm}$ . Figure 5 shows the measured spectrum together with the emission profile of the chosen green LED (Nichia NSPG300D), which was used to illuminate the point source aperture. The fluorescent dye (fluorescein sodium) for the droplet test case has an emission maximum at around  $515 \text{ nm}$  and matches well the LED.

### 3.3 Building the PSF $H$

While modern digital camera sensors have several millions of pixels, the image of a point source is almost entirely black, except for the nonzero pixel values of the light field PSF patterns. This is to our benefit, as we do not have to store the entire images in a huge matrix  $H$ , but instead take rectangular cutouts around the patterns. This is illustrated in Fig. 6 on the left, showing a sample raw image of a point source centered on the optical axis. The size of the cutout region has to be chosen such that no information within the pattern is cropped. From step to step, both the point source and the cutout have to be shifted likewise by a pixel- and a voxel-width, respectively.

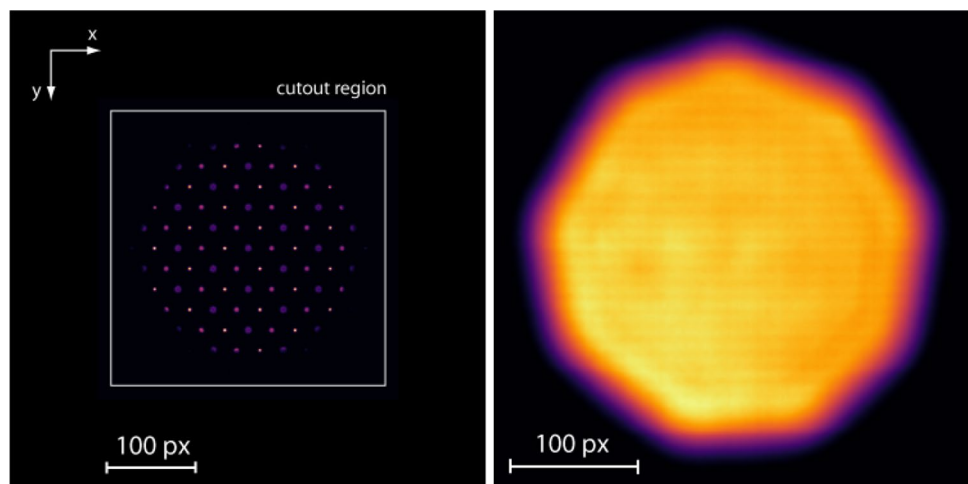
The matrix  $H$  is assembled from all images recorded by the camera, with the point source located at the voxel positions relating to the elementary cell. It is convenient to replace the  $N_p \times N_v$ -matrix used in Eq. 3 by the five-dimensional equivalent given in Eq. 2: Every  $(x_i, y_i)$ -slice contains the cutout patterns measured at source location  $(x, y, z)$ .

It is interesting to examine the image that forms when we sum up all the  $(x_i, y_i)$ -slices of  $H$  that belong to a single z-plane. This is shown on the right of Fig. 6, and we notice



**Fig. 5** Emission spectra of gas flame and of the LED used for calibration

**Fig. 6** *Left:* Spot pattern generated by a point source, defining a single PSF within  $H$ . Raw image is monochrome with applied color map. *Right:* Summation of all single PSFs within one  $z$ -slice of the matrix  $H$ , revealing the nonagon shape of the main lens aperture



that the result is not circular: It is in fact a nonagonal structure and corresponds to the aperture of the main lens, which is designed with nine movable blades. This showcases the importance of an experimental calibration in the case of a photographic system, where the arrangement of the optical elements is not precisely known.

With the PSF in its memory-saving, five-dimensional form, constructed from the nonzero cut out patterns, the backprojection array  $H^T$  is no longer a simple transpose of  $H$ , but requires a more complex procedure. A high pixel count of the image sensor and a hexagonal MLA arrangement lead to a high computational effort with the codes published in (Prevedel et al. 2014; Lu et al. 2019; Stefanou et al. 2019). The computing time could be cut drastically by a new algorithm based on a pure re-arrangement of the array elements (Eberhart 2020). Figure 7 shows two resulting slices of  $H^T$ . It is the intensity distribution in object space, generated by a single pixel being backprojected through the optical system, on the left side for the center pixel, on the right side for an off-axis pixel. Notice the fringe pattern due to diffraction by one of the oblique


aperture blades, which highlights the importance of an experimental calibration.

Realizing high axial resolutions is one of the main challenges in 3-D imaging. Deciding on the number of  $z$ -planes for the calibration is a trade-off between the achievable resolution gain on the one hand, which is subject to the used optical system, and the experimental effort and memory requirements on the other hand. For computations presented in this paper, the object space was discretized with 20 depth layers and a spacing of 1 mm. With the size of the elementary cell sketched in Fig. 2, this required to record the point source patterns at a total of  $95 \times 55 \times 20$  positions.

### 3.4 Relation to plenoptic tomography

The background section on plenoptic imaging mentioned the possibility to render multiple 2D images with varying view-points from single-exposure plenoptic data. This establishes a close link to tomography, where such series of projections under different angles are used as input for 3-D reconstruction. Both classes of tomographic algorithms, analytical and

**Fig. 7** Two elements of the matrix  $H^T$ , defining a back-projection of a pixel value into object space. The right image reveals diffraction patterns at the main lens aperture



iterative, are based on one or several backprojections of the recorded images into the volume, and it is obvious that data from a plenoptic camera lend itself to tomographic methods. In fact, a number of publications, notably those by Fahringer and Thurow (2012) in the framework of PIV analysis of particle loaded flows, use tomography to determine three-dimensional particle positions (Fahringer et al. 2015). Here, the authors favor an iterative MART algorithm, which at the first glance seems to be much more straightforward than a deconvolution approach, as no complex PSF matrix has to be determined. However, as demonstrated, e.g., by Levoy et al. (2006), deconvolution and tomographic reconstruction with a limited number of view angles are formally equivalent, and so these techniques also share similar requirements. In tomography, a matrix of weighting coefficients takes the place of the PSF and expresses the influence of object space voxels on the different image pixels. In a simple form, they are found by tracing a ray of light emerging from a single pixel throughout the optical system and determining its overlap with the volume's voxels. With a photographic setup, recorded images are no simple parallel projections, and the weighting matrix has to be found under consideration of the optical system. This can be done using simplifying assumptions, e.g., replacement of the main lens by a single thin lens (Fahringer et al. 2015). This neglects aberrations present in the real setup, which can be improved by performing an additional calibration step using known targets (Hall et al. 2018; Fahringer and Thurow 2018). In any case, the weighting matrix is huge, as it links all voxels to all pixels and does not exploit the periodicity found in the light field PSFs.

Given the formal similarity of limited angle tomography and deconvolution, it could be expected that a tomographic approach on plenoptic data may reach a reconstruction quality comparable to the method presented in this study, with similar efforts for calibration and computation. We favor light field deconvolution in combination with an experimental PSF due to its precise representation of all actually used optical elements and potential aberrations within the measurement setup.

## 4 Results

We used a number of test cases with increasing complexity to show light field deconvolution for volumetric visualization on macroscopic scales. The first is a cloud of discrete luminous points in space, intended to build intuition about the achievable spatial resolution of the reconstruction. The second case is an arrangement of multiple stationary, transparent gas flames, followed by a fluorescent droplet falling into a water pool, which created a time-series of light field data. All volume reconstructions were carried out using the MATLAB code published by Prevedel et al. (2014), which

is an implementation of the Richardson–Lucy-based deconvolution method proposed by Broxton et al. (2013).


It has been pointed out by several authors that sampling of the light field by a plenoptic camera is not uniform throughout the volume, but changes significantly with the depth coordinate (Bishop and Favaro 2009; Broxton et al. 2013; Stefanoiu et al. 2019). Sampling is especially low in close vicinity of the main lens focal plane, which leads to artifacts in the reconstruction. We adopted a method and the associated MATLAB code published by Stefanoiu et al. (2019), where an additional smoothing step within the deconvolution algorithm effectively reduces such artifacts. The calculation of the employed filter kernels requires knowledge of optical parameters of the main lens, such as the position of the principal planes and the effective focal length, which were determined by means of a raytracing simulation using Zemax OpticStudio. Before going into details of the respective results, we need to address an important limiting aspect of the used photographic setup, the shift-variance of the point spread function, which is discussed in the following.

### 4.1 Shift-variant main lens PSF


As outlined in the preceding paragraphs, light field deconvolution was introduced for microscopic imaging, where it takes advantage of the special properties of microscope optics. Due to their telecentric design, all ray bundles that are captured by the lens have chief rays parallel to the optical axis. In consequence, a microscope produces strictly orthographic views, and a lateral shifting of the object yields no parallax (Levoy and Hanrahan 1996): It is not possible to examine the object's side faces by sliding it toward the outer limits of the field of view. This has the important implication that the PSF of the optical system, without any additional microlenses, is shift-invariant and does not depend on the lateral  $x/y$ -position in object space. This is why the patterns of the light field PSF show the discussed periodicity. Even though photographic cameras can also be equipped with telecentric main lenses, this is reserved to special applications, as such lenses are extremely bulky and expensive. Using a standard photographic lens, it is therefore important to assess the effect of the shift-variance of the PSF on volume reconstructions by means of light field deconvolution.

Figure 3 demonstrates that the recorded patterns were self-repeating after shifting the point source over a certain distance, defining the object space boundaries of the elementary cell. An extension of such measurements toward larger shifts along the horizontal  $x$ -axis is given in Fig. 8, again plotting the cross-correlation with the pattern at the central position. The dashed vertical lines mark multiples of the width of the elementary cell, defined by the position of the first maximum. We observe two things: The overall agreement between the patterns slightly decreases with growing





**Fig. 8** Normalized cross-correlation of point source patterns along the lateral  $x$ -axis and pattern recorded at the origin. With a wider  $x$ -range than in Fig. 3, it shows a slight decrease in the correlation toward the rims due to a shift-variant main lens PSF. Vertical lines indicate multiples of the determined width of the elementary cell



**Fig. 9** Reconstruction of a planar, regular stripe pattern, spanning the entire field of view, to probe the influence of a shift-variant main lens PSF. *Top*: Computed result. *Bottom*: Intensity profile along the horizontal  $x$ -axis (red line), showing no significant falloff toward the rims

distance from the center, indicated by the red line, and the location of the maxima shifts relative to the dashed lines. Both effects are due to the variation in the main lens PSF with respect to the lateral position.


How do these variations affect the reconstruction quality? This was addressed by imaging a planar test pattern consisting of regular, oblique black and white stripes, which spanned the complete field of view of the camera ( $36 \times 24$  mm in object space). The result of the reconstruction, using the experimentally acquired PSF matrix  $H$  and eight iterations of light field deconvolution, is given on the top

of Fig. 9. An intensity profile was taken along the red line, which is plotted on the lower part of the figure and ranges over the entire horizontal width of the data. Obviously, the intensity does not drop toward the outer rims, and the peaks are regularly spaced throughout the plot. This means that the reconstruction method is robust concerning the shift-variant nature of the main lens PSF, which also demonstrates the feasibility of light field deconvolution with our photographic setup. This is clearly not a general statement, but is valid for the combination of camera, main lens and working distance that has been used in the present work. Other optical systems may well suffer from a more pronounced effect, which should be investigated in dedicated examinations.

### 4.2 Point cloud

A cloud of luminous points in space was created by using the motorized stage to position the  $5 \mu\text{m}$  point source, which had already been used during calibration, at 200 random coordinates within the depth range of the PSF matrix, and subsequent linear superpositioning of the recorded raw images. Such a cloud has the benefit of accurately known point locations. The volume was again reconstructed by eight iterations of light field deconvolution and is illustrated in Fig. 10. On the left, the top row shows  $x/y$  slices at different  $z$  depths, the bottom shows a  $z/y$  section of the same volume, and the right presents a 3-D rendering of the reconstructed light points. The different intensities of neighboring points in the lower left image are due to their different  $x$ -position, so that some have not been sliced through their center. The origin of the object space coordinate system is again located in the focal plane of the main lens, the native object plane (NOP), with the  $z$ -axis pointing along the viewing direction of the camera. The right of the figure shows a perspective 3-D visualization of the reconstructed volume. Here, little red balls mark the actual position of the light points in space, which are matched well in all three coordinates.

**Fig. 10** Reconstructed volume of a point cloud, generated by superposition of single raw recordings of a light point at 200 coordinates in 3-D space. *Left, top*: Lateral  $xy$ -slices at different depth positions. *Left, bottom*: 2-D slice along the depth coordinate  $z$ , showing a rough estimate of the depth resolution. *Right*: 3-D visualization of the point cloud. Little red balls mark the precisely know, actual positions of the points




It is apparent that both lateral and depth resolution decrease with increasing distance from the NOP. A closer examination is given in Fig. 11: The left shows lateral intensity profiles of reconstructed points at three different z-depths, with a shape close to Gaussian. Parts of the deconvolution algorithm are carried out in the Fourier domain, and results are prone to ringing in regions with high local contrast, which can be recognized at the wings of the blue and green profiles. The non-negativity constraint within the Richardson–Lucy scheme reflects negative values onto the positive side.

The right of the figure plots these profiles along the depth coordinate. Spacing of the data points is here determined by the PSF calibration, where a discretization of 1 mm was used in z-direction. The dashed lines represent Gaussian profiles that have been fitted to the data in order to derive FWHM values. Figure 12 shows lateral and axial resolution as a function of the depth coordinate, where the FWHM values of all points within the respective z-planes were averaged. The lateral FWHM, on the left of the figure, is below 100 μm in the rear part of the volume and then rises to about 200 μm toward the front boundary. The FWHM along the optical axis, plotted on the right, is 1 mm close to the NOP and increases to about 3.5 mm in the front part.

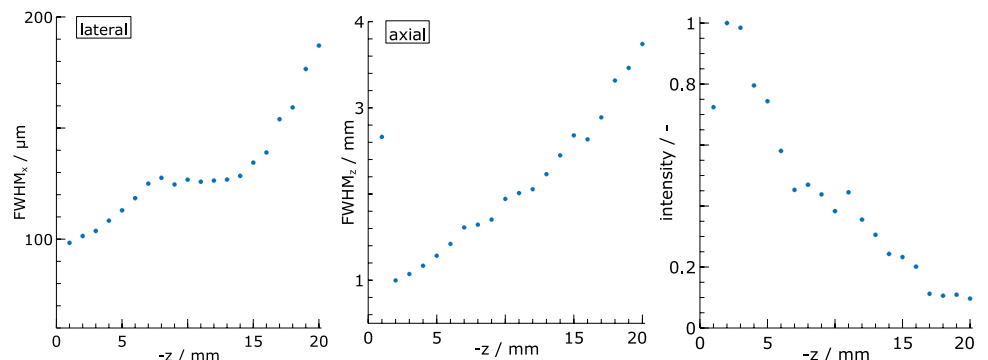
A note on the lateral resolution that can be expected within a reconstructed volume: With a 1:1 magnification of the optical setup, the lateral voxel size is equivalent to the pixel pitch, which is 5.5 μm for the used camera. The sensor of a plenoptic camera, however, does not only

sample spatial data, but also angular information, which drastically reduces the lateral resolution. In focused plenoptic cameras, like the used Raytrix model, spatial and directional sampling is intertwined and changes with scene depth. In the worst case, the lenslet pitch of the MLA dictates the maximum achievable lateral resolution, which in our case translates to about 170 μm. Light field deconvolution is closely related to superresolution approaches (Bishop and Favaro 2012), but we cannot expect to recover details on the scale of the sensor’s pixel resolution. Compared to the lateral resolution, the measured axial resolution is lower by a factor of 10–20. This results in a visible elongation of the light points in Fig. 10 along the z-axis. A poorer depth resolution by comparison with the lateral direction is, e.g., also observed in standard microscopes, due to the three-dimensional structure of the diffraction patterns of the optical systems. Light field deconvolution of microscope data achieved lateral/axial resolution ratios of about 2–6 (Cohen et al. 2014; Prevedel et al. 2014). The lower axial resolution in the present experiment is believed to be by reason of the complex and less precise photographic setup, which prevents from using theoretical PSFs, but requires an experimental acquisition, which adds additional uncertainties and noise. In contrast to pure optical measurements, computational imaging and reconstruction also depend on the recorded data itself. Very shallow gradients and noisy measurements can downgrade the achievable resolution. In this respect, the present test case benefits from the clear structure of the

**Fig. 11** *Left:* Intensity profiles of reconstructed points along their x-axis at different z-positions. Profiles closely follow Gaussian shapes, and resolution is defined as their full width at half maximum (FWHM). *Right:* Intensity profiles of the points along the depth coordinate z, additionally showing fitted Gaussian curves



**Fig. 12** Lateral (*left*) and depth (*middle*) resolution of reconstructed points along the z-axis. Note the different orders of magnitude. *Right:* Normalized intensities of reconstructed points as a function of depth



singular light points, and the axial resolution is superior to the results from the more complex cases, which are discussed in the next sections.

Even though the brightness of the light points during the measurement was constant throughout the volume, the intensity of their reconstruction is also subject to the axial position, which is evident from the right of Fig. 10. A closer examination is given in the rightmost plot of Fig. 12. The intensity drops with increasing distance from the NOP by a factor of about ten. This issue needs to be addressed in upcoming investigations. We speculate that this is due to a reduced fidelity of the raw images toward higher z-coordinates: A larger distance from the NOP means that the light from the point source is spread over a larger area and a higher number of microimages, with reduced pixel intensities. This reduces the signal-to-noise ratio, which also applies to the recording of the PSF during calibration. A remedy could be an increased bit depth within the raw images, which requires a re-programming of the camera.

The reconstruction of a volume filled with discrete points or particles is not the prime application of a deconvolution approach. Here, algorithms based on stereoscopic principles, like the commercial light field software RxLive or RxFlow, provide much faster and presumably more accurate results. The intention of this test is to demonstrate the achievable lateral- and depth resolution under ideal conditions at a precisely defined object. It shows the general feasibility of the approach and a resolution sufficient for the analysis of small-scaled flow structures, before moving on to more realistic cases: transparent, luminous flows that do not carry any particles or other tracers.

### 4.3 Gas flames


A propane/butane gas burner, shown on the left of Fig. 13, was chosen as a test object. The head of the burner has a diameter of 20 mm. It is covered by a slightly convex,

perforated plate, where the single holes have diameters of 1 mm are spaced by 2 mm in a hexagonal arrangement and generate multiple small flamelets. A part of the burner head was covered, as sketched in the figure, so that the flames extended over a depth range of about 15 mm. The burner was positioned within the z-coordinates of the calibration, and the flames were recorded in a single exposure of the plenoptic camera (shutter speed 50 ms, gain 10) from a side view. The right of Fig. 13 shows the raw image, and the zoomed-in region reveals the microimages formed by the MLA. This luminous, transparent flow shows only little local contrast, so that standard computer vision algorithms largely fail (Greene and Sick 2013).

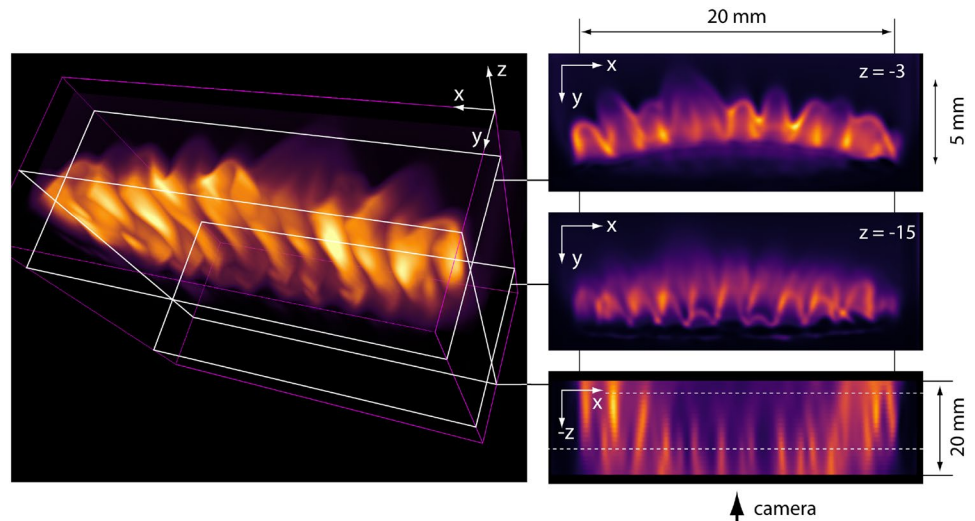
Light field deconvolution (eight iterations) was carried out on the same raw image and resulted in a computed volumetric emission which is visualized in Fig. 14. The right hand side shows, from top to bottom, two lateral xy-slices through the volume at different depths and an xz-slice. The axes have been scaled for better representation, and values at intermediate, non-integer z-positions are interpolated. The left of the figure presents a perspective 3-D visualization in the form of a maximum intensity projection, additionally showing position and orientation of the respective slices.

The two xy-slices confirm that overlapping flamelets at different z-depths clearly can be discriminated, which indicates a true three-dimensional reconstruction. The dark, curved region, also visible in the xz-slice, is due to the convex grid of the burner head. The single flamelets have a shape resembling a hollow cone, so that a planar section should appear circular or elliptical. The xz-slice of Fig. 14 shows curved side faces of the little flames which, however, do not form a closed oval. We speculate that this is owed to the fact that the intensity distribution within the front and back faces are very homogeneous with extremely shallow gradients, which is challenging to reconstruct. A certain amount of intensity variation, e.g., due to turbulence, is therefore beneficial for volumetric visualization by light field

**Fig. 13** *Left:* Transparent flames of a camping stove, used as a test case. As sketched, part of the burner head was covered during the experiments. *Right:* Raw image recorded by the plenoptic camera. The zoomed-in inset shows the circular microimages in a hexagonal arrangement



**Fig. 14** *Left:* 3-D visualization (maximum intensity projection) of transparent gas flames above a camping stove. *Right:* Two-dimensional slices through the volume



deconvolution. The axial resolution of the visualization is reduced compared to the point cloud test case, and the flamelets appear smeared over several millimeters in the depth direction. This lowers the possible optical sectioning of the volume. Light contributions from out-of-focus regions are to some extent still present within selected depth planes, which can be regarded as artifacts of the reconstruction. Using appropriate spectral band-pass filters during measurement and calibration could lead to improved 3-D results due to an exactly matching PSF. This also opens the possibility to selectively investigate the volumetric distribution of single radiating species within a flame.

It should be noted that the flames, though assumed to be optically thin, may still affect the path of light through their volume, caused by refraction due to gradients of the refractive index of the hot air. This also alters the PSF of the system which can lead to a lower reconstruction quality. We expect to mitigate this effect by incorporating an appropriate physical model into the deconvolution process, which is part of ongoing research.

#### 4.4 Fluorescent droplet

A fluorescent liquid was prepared by mixing a fluorescent dye (fluorescein-sodium) with water and dripping the solution into a small water pool. This pool was made from blackened aluminum with a transparent acrylic front pane and a depth of 2 cm along the optical axis of the camera and was positioned within the depth coordinates of the PSF calibration. Illumination with a blueish UV-LED caused fluorescence with an emission peak at around 515 nm. Figure 15 shows a photograph of the resulting luminous turbulent flow, taken by a standard color DSLR camera.

A time resolved light field sequence of the flow was recorded by the plenoptic camera at a frame rate of 2.5 fps.

Shutter speed and gain were set to 10 ms and 10, respectively. Similar to gas flames, the fluorescent flow is rather foggy and does not provide enough local contrast for standard reconstruction algorithms. Each raw image was used for a volumetric reconstruction by eight iterations of light field deconvolution, and results are given in Fig. 16. Here, the left column shows a 3-D visualization (maximum intensity projection) of the temporal evolution of the flow, for five selected frames at the given points in time. The little gray dots in the back are not artifacts but show gas bubbles at the rear wall. The straight features that form in the rear after 10.8 s show fluorescent dye accumulating at the physical edges of the water pool. Despite the low contrast,



**Fig. 15** Exemplary photograph (standard color DSLR camera) of a flow induced by dropping fluorescent liquid into water




deconvolution resolves the volumetric nature of the turbulent structures with high resolution.

The second column of Fig. 16 shows  $xy$ -slices through the computed volume at a fixed  $z$ -depth for respective time points.

Time resolved three-dimensional results in principle also allow to derive 3D3C-velocity data of the flow. Particle tracking velocimetry (PTV) or particle image velocimetry (PIV) techniques determine displacements of tracer particles within the flow on a frame-by-frame basis to derive velocity vectors. The commercial software RxFlow, e.g., performs a PTV analysis of light field data recorded by Raytrix cameras. The underlying feature matching algorithms rely on particle flows and require a suitably high contrast within the single frames.

To explore the potential of a 3D3C flow field analysis in the given low-contrast case of the fluorescent droplet, we performed an optical flow computation on two consecutive frames of the reconstructed volume. In general, the optical flow can be derived for 3-D data, which yields a three-dimensional field of displacement vectors, which can be interpreted as velocities (Mustafa 2016). In the present case, the axial and lateral resolution within the volume is very different, and we chose to compute 2D optical flows, but at various axial depths. The used method, implemented in MATLAB, is based on a modified Horn–Schunck algorithm (Sun et al. 2010). A pre-processing step was used




**Fig. 16** Time resolved sequence of a fluorescent droplet in water. *Left:* 3-D visualization (maximum intensity projection). *Middle/Right:* 2-D slices at two different fixed  $z$ -depths

which served to alleviate large intensity differences within the images by applying a  $\log$ -transform on the pixel values (Zhuk et al. 2017). Results of the optical flow computation are shown in Fig. 17, where displacement magnitude is coded in both color and length of the arrows, plotted over monochrome 2D intensity maps of the volume. Velocities were calculated based on the pixel displacements, converted to object space dimensions, which were divided by the time interval between the two frames.

The figure presents two  $xy$ -slices, one close to the front boundary (top), and the other in the middle of the depth range (bottom). They show clear differences in the flow fields in both magnitude and orientation. Regions with a source/sink-like structure, which would be unphysical in a planar flow, are possible in the volumetric case due to mass flow in the axial direction.

The setup of the droplet flow is a challenging test case for the robustness of the deconvolution approach: Compared to the camera calibration, it contains two additional media with different refractive indices and their interfaces. Light is emitted in water and propagates through acrylic and air toward the main lens, which alters the PSF of the optical system. For best results, calibration would have to be done under the same conditions, with the point source positioned within



**Fig. 17** 2D velocity fields derived from the optical flow between two consecutive frames of the light field sequence at depth  $z = -10$  (top) and  $-19$  (bottom)

the water pool. But even with a basic calibration, light field deconvolution generates volumetric visualizations of such complex transparent fluid flows which are not possible with standard algorithms based on computer vision.

## 5 Conclusion

In this paper, we have demonstrated the application of light field deconvolution for the 3-D visualization of transparent fluid flows on macroscopic scales. A key contribution is the development of an experimental calibration technique, which allows to acquire the system's PSF matrix in complex optical setups, involving photographic main lenses and other possible elements. As the time-consuming, computationally expensive volume reconstruction is done as a post-processing step, the temporal resolution of the method is limited alone by the frame rate of the recording system, and three-dimensional intensity distributions can be determined from snapshot data of a single plenoptic camera. We have presented results from different test cases, luminous flows produced by transparent gas flames and fluorescent droplets, which show the capability of the proposed method. The achieved lateral object space resolution was in the order of 100–200  $\mu\text{m}$  in the best case, whereas the axial resolution was lower by a factor of 10–20. This creates a visible elongation along the depth direction, which has to be taken into account in the interpretation of the results. Without additional tracers or other sources of high local contrast, reconstruction of light field sequences and analysis of the optical flow between consecutive frames even allows to derive information on the velocity field within the volume, even though no particles have been added as in common PIV/PTV techniques.

A strength of plenoptic light field deconvolution is its simple setup requiring only a single snapshot exposure from a single camera, which is highly beneficial for measurements with limited optical access like combustion chambers or arc jet facilities. This eliminates the need of precise triggering of multiple instruments, and the technique is applicable to the investigation of fast three-dimensional processes like, e.g., turbulent fluid structures.

## Appendix: Computational aspects

Light field deconvolution is a computationally intensive technique, in terms of both memory requirement and computation time. The PSF array  $H$  used in the present study holds  $95 \times 55 \times 20$  single PSFs with about  $150 \times 150$  pixels in single precision, which requires approximately 9.5 GB. The same amount of memory is occupied by the array  $H^T$ . This high number of single PSFs is dictated by the MLA design

of the R29 camera, with three different lenslet types in a hexagonal arrangement. Usage of custom made cameras with a dedicated layout of the MLA may alleviate this requirement.

The code published by Prevedel et al. (2014), written in MATLAB, takes advantage from parallel computation on GPUs and requires a CUDA compatible graphics card. The volume of the gas flames example contains  $1705 \times 4275 \times 20$  voxels and was computed on a desktop machine with an Intel i7-6700K CPU, 48 GB memory and a Nvidia GeForce 980 Ti GPU having 6 GB of memory. One iteration of light field deconvolution required 2700 s.

A single frame of the fluorescent drop case is made up of  $4235 \times 3039 \times 20$  voxels, and volume reconstruction was done on the Vulcan cluster at the High-Performance Computing Center Stuttgart (HLRS) of the University of Stuttgart. We used a dual socket node with 256 GB of main memory and a Nvidia GPU accelerator. Each socket is equipped with an 8-core Intel Xeon E5-2667v4 processor with a base frequency of 3.2 GHz. The accelerator is a Nvidia Tesla P100 with 12 GB of memory. Here, the computation took 2750 s per iteration per frame. Computational efforts could also be reduced by adjusting the discretization of the volume to the spatial resolution than can be achieved in practice. As already pointed out by Broxton et al. and others, sampling of the light field is not uniform along the optical axis, which could be taken into account by sophisticated reconstruction and calibration procedures.

Visualization of 3-D intensities was done using 3D Slicer, a software developed for medical image processing, and ParaView was used for the optical flow vectors.

**Acknowledgements** This work was funded by the German Research Foundation under grant LO 1772/4-1. We appreciate continuous discussions and ideas within the High Enthalpy Flow Diagnostics Group. This work was partially executed within the frame of the German SiVeGCS project. The authors gratefully acknowledge access to the high-performance computing facility "Vulcan" at HLRS and would like to thank the team of HLRS for their kind support.

**Author Contributions** ME and SL conceptualized and designed the research presented in this paper. ME developed hardware components and conducted experiments. Implementation of computer code for data analysis and modification of existing algorithms was done by ME and PO. Visualizations were created by ME. The initial manuscript was written by ME and SL. All authors read and approved the final manuscript. SL acquired funding required for this project.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated

otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Adelson E, Wang J (1992) Single lens stereo with a plenoptic camera. *IEEE Trans Pattern Anal Mach Intell* 14(2):99–106. <https://doi.org/10.1109/34.121783>
- Anikin NB, Suntz R, Bockhorn H (2010) Tomographic reconstruction of the OH\*-chemiluminescence distribution in premixed and diffusion flames. *Appl Phys B* 100(3):675–694. <https://doi.org/10.1007/s00340-010-4051-5>
- Atcheson B, Ihrke I, Heidrich W, Tevs A, Bradley D, Magnor M, Seidel HP (2008) Time-resolved 3D capture of non-stationary gas flows. *ACM Trans Graphics* 27(5):1–9. <https://doi.org/10.1145/1409060.1409085>
- Bishop T, Favaro P (2009) Plenoptic depth estimation from multiple aliased views. In: 12th international conference on computer vision, pp 1622–1629. <https://doi.org/10.1109/ICCVW.2009.5457420>
- Bishop T, Favaro P (2012) The light field camera: extended depth of field, aliasing, and superresolution. *IEEE Trans Pattern Anal Mach Intell* 34(5):972–986. <https://doi.org/10.1109/TPAMI.2011.168>
- Broxton M, Grosenick L, Yang S, Cohen N, Andalman A, Deisseroth K, Levoy M (2013) Wave optics theory and 3-D deconvolution for the light field microscope. *Opt Express* 21(21):25418–25439. <https://doi.org/10.1364/OE.21.025418>
- Cai W, Li X, Ma L (2013) Practical aspects of implementing three-dimensional tomography inversion for volumetric flame imaging. *Appl Opt* 52(33):8106–8116. <https://doi.org/10.1364/ao.52.008106>
- Chen H, Sick V (2017) Three-dimensional three-component air flow visualization in a steady-state engine flow bench using a plenoptic camera. *SAE Int J Engines* 10(2):625–635. <https://doi.org/10.4271/2017-01-0614>
- Cohen N, Yang S, Andalman A, Broxton M, Grosenick L, Deisseroth K, Horowitz M, Levoy M (2014) Enhancing the performance of the light field microscope using wavefront coding. *Opt Exp* 22(20):24817–24839. <https://doi.org/10.1364/OE.22.024817>
- Eberhart M (2021) Efficient computation of backprojection arrays for 3D light field deconvolution. *Opt Exp*. Accepted for publication
- Eberhart M, Loehle S (2021) Three-dimensional analysis of transparent flames by light field deconvolution. *J Thermophys Heat Transfer* 35(1):200–205. <https://doi.org/10.2514/1.T6046>
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006) Tomographic particle image velocimetry. *Exp Fluids* 41:933–947. <https://doi.org/10.1007/s00348-006-0212-z>
- Fahringer T, Thurow BS (2012) Tomographic reconstruction of a 3-D flow field using a plenoptic camera. In: 42nd AIAA fluid dynamics conference, AIAA, pp 2012–2826. <https://doi.org/10.2514/6.2012-2826>
- Fahringer TW, Thurow BS (2018) Plenoptic particle image velocimetry with multiple plenoptic cameras. *Meas Sci Technol* 29(7). <https://doi.org/10.1088/1361-6501/aabe1d>
- Fahringer TW, Lynch KP, Thurow BS (2015) Volumetric particle image velocimetry with a single plenoptic camera. *Meas Sci Technol* 26(11). <https://doi.org/10.1088/0957-0233/26/11/115201>
- Georgiev T, Lumsdaine A (2012) The multifocus plenoptic camera. *Proc SPIE* 8299:829908–829908–11. <https://doi.org/10.1117/12.908667>
- Greene ML, Sick V (2013) Volume-resolved flame chemiluminescence and laser-induced fluorescence imaging. *Appl Phys B* 113:87–92. <https://doi.org/10.1007/s00340-013-5664-2>
- Gross H, Singer W, Totzeck M (2005) *Handbook of Optical Systems: Physical Image Formation*, vol 2. Wiley-VCH. <https://doi.org/10.1002/3527606688>
- Hall EM, Fahringer TW, Guildenbecher DR, Thurow BS (2018) Volumetric calibration of a plenoptic camera. *Appl Opt* 57(4):914–923. <https://doi.org/10.1364/AO.57.000914>
- Halls BR, Gord JR, Jiang N, Splichenko M, Roy S, Meyer TR (2016) High-speed three-dimensional tomographic measurements for combustion systems. In: 32nd AIAA aerodynamic measurement technology and ground testing conference, AIAA. <https://doi.org/10.2514/6.2016-4027>
- Hermann T, Loehle S, Fasoulas S, Andrianatos A (2016) Tomographic optical emission spectroscopy for plasma wind tunnel testing. *Appl Opt* 55(36):10290–10298. <https://doi.org/10.2514/6.2016-3203>
- Isaksen A, McMillan L, Gortler SJ (2000) Dynamically reparameterized light fields. In: Proceedings of the 27th annual conference on computer graphics and interactive techniques, ACM Press/Addison-Wesley Publishing Co., pp 297–306. <https://doi.org/10.1145/344779.344929>
- Levoy M, Hanrahan P (1996) Light field rendering. In: Proceedings of the 23rd annual conference on computer graphics and interactive techniques, association for computing machinery, SIGGRAPH '96, pp 31–42. <https://doi.org/10.1145/237170.237199>
- Levoy M, Ng R, Adams A, Footer M, Horowitz M (2006) Light field microscopy. *ACM Trans Graph* 25(3):924–934. <https://doi.org/10.1145/1141911.1141976>
- Lillo PM, Greene ML, Sick V (2014) Plenoptic single-shot 3D imaging of in-cylinder fuel spray geometry. *Z Phys Chem* 229(4). <https://doi.org/10.1515/zpch-2014-0601>
- Lippmann G (1908) Épreuves réversibles donnant la sensation du relief. *Journal de Physique Théorique et Appliquée* 7. <https://doi.org/10.1051/jphysap:019080070082100>
- Lu Z, Wu J, Qiao H, Zhou Y, Yan T, Zhou Z, Zhang X, Fan J, Dai Q (2019) Phase-space deconvolution for light field microscopy. *Opt Exp* 27(13):18131–8145. <https://doi.org/10.1364/OE.27.018131>
- Lumsdaine A, Georgiev T (2009) The focused plenoptic camera. In: 2009 IEEE international conference on computational photography (ICCP), pp 1–8. <https://doi.org/10.1109/ICCPHOT.2009.5559008>
- Lynch KP, Thurow BS (2012) 3-D flow visualization of axisymmetric jets at reynolds number 6,700 and 10,200. *J Visualization* 15:309–319. <https://doi.org/10.1007/s12650-012-0141-2>
- Mustafa M (2016) A data-driven learning approach to image registration. PhD thesis, University of Nottingham
- Ng R (2006) Digital light field photography. PhD thesis, Stanford University
- Perwaß C, Wietzke L (2012) Single lens 3D-camera with extended depth-of-field. *Human Vision and Electronic Imaging XVII, SPIE Proceedings* 8291:45–59. <https://doi.org/10.1117/12.909882>
- Prevedel R, Yoon YG, Hoffmann M, Pak N, Wetzstein G, Kato S, Schrödel T, Raskar R, Zimmer M, Boyden ES, Vaziri A (2014) Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy. *Nat Methods* 11:727–730. <https://doi.org/10.1038/nmeth.2964>
- Sage D, Donati L, Soulez F, Fortun D, Schmit G, Seitz A, Guiet R, Vonesch C, Unser M (2017) DeconvolutionLab2: an open-source software for deconvolution microscopy. *Methods* 115:28–41. <https://doi.org/10.1016/j.jymeth.2016.12.015>

- Stefanoiu A, Page J, Symvoulidis P, Westmeyer G, Lasser T (2019) Artifact-free deconvolution in light field microscopy. *Opt Exp* 27(22):31644–31666. <https://doi.org/10.1364/OE.27.031644>
- Sun D, Roth S, Black MJ (2010) Secrets of optical flow estimation and their principles. In: *IEEE conference on computer vision and pattern recognition*, IEEE. <https://doi.org/10.1109/CVPR.2010.5539939>. [http://cs.brown.edu/~dqsun/code/cvpr10\\_flow\\_code.zip](http://cs.brown.edu/~dqsun/code/cvpr10_flow_code.zip)
- Tan ZP, Alarcon R, Allen J, Thurow BS, Moss A (2020) Development of a high-speed plenoptic imaging system and its application to marine biology PIV. *Meas Sci Technol* 31(5). <https://doi.org/10.1088/1361-6501/ab553c>
- Upton TD, Verhoeven DD, Hudgins DE (2011) High-resolution computed tomography of a turbulent reacting flow. *Exp Fluids* 50:125–134. <https://doi.org/10.1007/s00348-010-0900-6>
- Wender A, Iseringhausen J, Goldluecke B, Fuchs M, Hullin MB (2015) Light field imaging through household optics. In: *Vision, modeling & visualization*, The Eurographics Association. <https://doi.org/10.2312/vmv.20151271>
- Wetzstein G, Ihrke I, Heidrich W (2013) On plenoptic multiplexing and reconstruction. *Int J Comput Vis* 101(2):384–400. <https://doi.org/10.1007/s11263-012-0585-9>
- Zhuk S, Tchrakian T, Akhriev A, Lu S, Hamann H (2017) Dynamic cloud motion forecasting from satellite images. In: *IEEE 56th annual conference on decision and control (CDC)*. <https://doi.org/10.1109/CDC.2017.8264113>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.