

MicroLens array grid estimation, light field decoding, and calibration

Maximilian Schambach and Fernando Puente León, *Senior Member, IEEE*
 Karlsruhe Institute of Technology, Institute of Industrial Information Technology
 Hertzstr. 16, 76187 Karlsruhe, Germany
 {schambach, puente}@kit.edu

Abstract—We quantitatively investigate multiple algorithms for microlens array grid estimation for microlens array-based light field cameras. Explicitly taking into account natural and mechanical vignetting effects, we propose a new method for microlens array grid estimation that outperforms the ones previously discussed in the literature. To quantify the performance of the algorithms, we propose an evaluation pipeline utilizing application-specific ray-traced white images with known microlens positions. Using a large dataset of synthesized white images, we thoroughly compare the performance of the different estimation algorithms. As an example, we apply our results to the decoding and calibration of light fields taken with a Lytro Illum camera. We observe that decoding as well as calibration benefit from a more accurate, vignetting-aware grid estimation, especially in peripheral subapertures of the light field.

I. INTRODUCTION

Computational cameras, that is, cameras utilizing combined optical and digital image processing techniques, have been gaining attention both in consumer applications, such as multi-lens camera systems in mobile devices, as well as scientific and industrial applications, such as light field cameras [1], [16] or snapshot hyperspectral cameras [2]. Computational imaging systems can usually well be described using the so-called *4D light field* $L_{\lambda,t}(u, v, a, b)$, where λ describes a wavelength, t the time, and the coordinates (u, v, a, b) correspond to a certain parametrization of the spatio-angular dependency of the light field of which there are numerous. For computational cameras, one usually uses the plane-plane parametrization: a light ray inside a camera is uniquely described by the intersection points $\mathbf{u} = (u, v)$ and $\mathbf{a} = (a, b)$ of two parallel planes, e.g. the main lens plane and the sensor plane.

In particular, microlens arrays (MLAs) are used in computational imaging applications allowing for a complex coding (or multiplexing) scheme of the light field onto an imaging sensor. Most prominently, MLAs are used in compact MLA-based light field cameras [16], [13], but also in other applications such as multi- or hyperspectral imaging [19], [20]. As usual, a certain camera model is then used to calibrate the intrinsic (and extrinsic) parameters of the camera to relate the image-side to the object-side light field. The model is then evaluated, e.g. by using the ray re-projection error. In most light field camera models [7], [6], [22], [23], the calibration is performed using image-side light fields that have been decoded from the sensor lenslet images, while others perform calibration utilizing the raw images directly [4]. In either case, this includes multiple

non-trivial pre-processing steps, such as the detection of the projected microlens (ML) centers and estimation of a regular grid approximating the centers, alignment of the lenslet image with the sensor, slicing the image into a light field and, in the case of hexagonal MLAs, resampling the light field onto a rectangular grid. These steps have a non-negligible impact on the quality of the decoded light field and camera calibration. Hence, a quantitative evaluation is necessary where possible. Here, we will focus on the estimation of the MLA grid parameters (to which we refer to as *pre-calibration*), which is the basis for all decoding and calibration schemes found in the literature. In spite of the importance of the pre-calibration pipeline, the literature focuses mostly on the camera models and decoding but pays little or no attention to the necessary details emerging in the pre-calibration, most importantly non trivial effects such as mechanical and natural vignetting. While for a correct pre-calibration, the detection of the perspective projected ML centers is necessary, all methods proposed in the literature rely on estimating the center of each ML image brightness distribution, approximating the orthogonally projected centers. Due to natural and mechanical vignetting, this results in severe deviations from the true projected centers, in particular in off-center MLs. This is the main scope of this article. In particular, our contributions are as follows:

- We propose a camera model and ray tracer implementation to synthesize application-specific white images with known ML centers as reference data.
- We propose a new pre-calibration algorithm motivated by our physical camera model, taking into account natural and mechanical vignetting effects.
- We present detailed accuracy requirements that the pre-calibration pipeline has to fulfill and show that the proposed algorithm, in case the of a Lytro Illum camera, fulfills these requirements. We compare our algorithm to different schemes proposed in the literature (which we show fail the accuracy requirements).
- We evaluate the full light field decoding pipeline for the different pre-calibration algorithms for simulated as well as real light field data.
- We investigate the influence of the quality of the pre-calibration on the full calibration of a Lytro Illum light field camera for different calibration methods.

Since the ML grid parameters are application-specific, we make the source code for a full evaluation pipeline (image

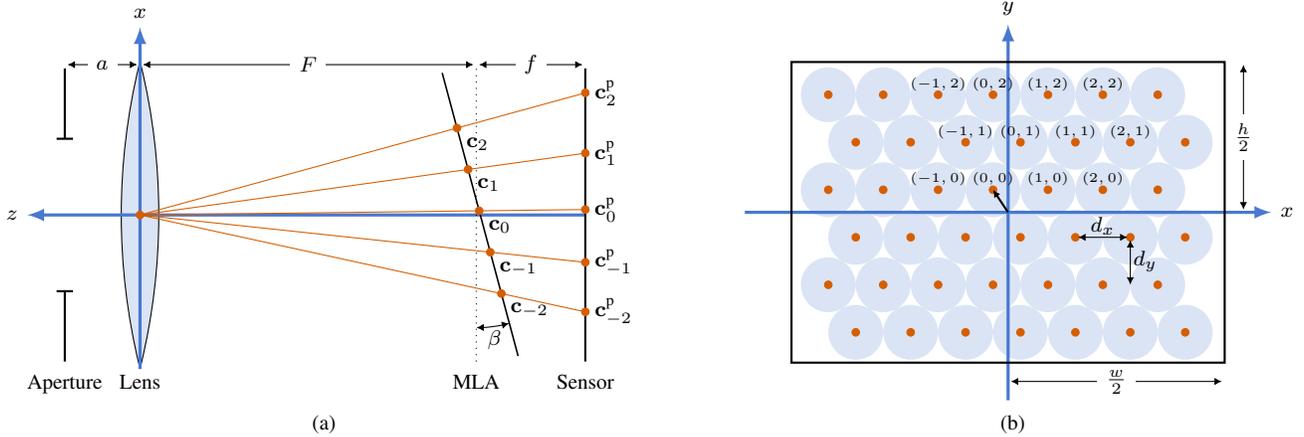


Figure 1. Schematic drawings of the used camera model of the unfocused plenoptic camera with exaggerated ML size. (a) A 2D section ($y = 0$) of the used camera model. The coordinates $\mathbf{c}_{\pm i}$ denote the centers of the ML (which are not explicitly depicted) and their perspective projections $\mathbf{c}_{\pm i}^p$. (b) The used MLA model before rotation and tilt. The ML centers \mathbf{c}_{ij} are depicted in orange with corresponding index labels (i, j) .

synthesis, ML grid estimation, and decoding) freely available [10] to be used in scientific research for any kind of MLA-based application. This includes the release of an open source Python framework for light field decoding and analysis as well as for applications in hyperspectral imaging. Even though the following presentation relies on a light field camera in the unfocused design [16], the proposed pre-calibration method is equally applicable to cameras in the focused design [13].

The remainder of this paper is organized as follows. We introduce the used camera model and chosen camera parameters in Section II. In Section III, we review different methods for the estimation of ML grids and formulate precise accuracy requirements which the algorithms ought to fulfill. Furthermore, we introduce a new estimation method which we thoroughly motivate using the physical camera model. All estimation methods are then quantitatively evaluated. In the remaining Sections IV and V, we investigate the influence of the ML grid estimation accuracy on light field decoding and calibration, respectively.

II. CAMERA MODEL AND REFERENCE DATA

The pre-calibration of MLA-based cameras is usually performed using so-called white images (WIs)—images of a white scene, for example taken using an optical diffuser. In order to quantitatively evaluate the performance of the ML grid estimation algorithms, appropriate reference data is needed. Of course, real WIs, as for example provided by the Lytro cameras, are unsuited since the actual ML centers are unknown. Therefore, reference data has to be synthesized. Previously, Hog et al. [9] used a simple addition of three 2D cosine waves to synthesize a WI with known parameters, but the results are too crude for a precise evaluation of the estimation algorithms. In particular, they neither account for natural, nor mechanical vignetting of the main lens and the MLs. Here, we use a self-developed ray tracer [17], [10], which we have extended by the following camera model to render a multitude of white images with precisely known ML centers.

A. Camera model

The camera model used is depicted in Fig. 1. In our model, the camera consists of a main lens and a collection of MLs, arranged in a hexagonal grid, which may be rotated (not depicted in the figure) and tilted. All lenses are modeled as thin lenses. As is usual in the focused design, f-number matching of main lens and MLs is assumed. To simulate irregularities of the grid, we add independent uncorrelated Gaussian noise ϵ to the ideal grid point's x - and y -coordinates. Natural vignetting is implemented by using the $\cos^4 \Phi$ law and the ray's incident angle Φ . Finally, an object-side aperture with variable entrance pupil is placed at distance a to the main lens to account for mechanical vignetting effects. Note that we do not model systematic, non-rigid deformations of the MLA as considered in [18]. We argue that these irregularities should be eradicated in the manufacturing process of high-quality MLAs, as they introduce irreducible blur in the light field (on which we will elaborate in Section II-B).

The ideal, unrotated, untilted and unshifted ML center coordinates are given by

$$\mathbf{c}_{\pm i \pm j}^{\text{id}} = \mathbf{o}_g + \begin{pmatrix} (\pm i \pm \frac{1}{2}j \bmod 2) d_x \\ \pm j d_y \\ 0 \end{pmatrix} + \epsilon_{\pm i \pm j}, \quad (1)$$

for $(i, j) \in \mathbb{N}^2$. Here, d_x, d_y denote the ideal grid spacing, $\mathbf{o}_g = (o_{g,x}, o_{g,y}, 0)^T$ the grid offset, and $\epsilon = (\epsilon, \epsilon, 0)^T$ the grid noise with variance σ_g^2 . The ideal hexagonal grid is determined by a single grid spacing d via $d_x = d, d_y = \sqrt{3}d/2$, where the ML radius is given by $r = d/2$. The ideal grid points are then rotated in the xy -plane by α , rotated around the y -axis by β , rotated around the x -axis by γ , and shifted to $z = -F$, where F denotes the main lens focal length (or image distance). Hence, we obtain the final grid point coordinates

$$\mathbf{c}_{\pm i \pm j} = \mathbf{R}_{x,\gamma} \mathbf{R}_{y,\beta} \mathbf{R}_{z,\alpha} \mathbf{c}_{\pm i \pm j}^{\text{id}} + (0, 0, -F)^T. \quad (2)$$

We will at times refer to them simply by \mathbf{c}_k for $k \in \mathbb{Z}$, where we do not need to specify the re-indexing $(\pm i, \pm j) \mapsto k$.

The size (w, h) of the MLA is chosen such that the projection of the grid, after rotation α and tilt (β, γ) , covers the full sensor of size (s_x, s_y) , i.e. it can be calculated via

$$\mathbf{R}_{x,\gamma} \mathbf{R}_{y,\beta} \mathbf{R}_{z,\alpha} \begin{pmatrix} w \\ h \\ 0 \end{pmatrix} = \begin{pmatrix} s_x \\ s_y \\ z \end{pmatrix}, \quad (3)$$

where z is arbitrary.

The perspective projection of the ML centers (2) from the center $(0, 0, 0)^T$ of the exit pupil onto the sensor is given by

$$\mathbf{c}_{\pm i \pm j}^p = \lambda_{\pm i \pm j} \mathbf{c}_{\pm i \pm j} \quad (4)$$

with scaling factor $\lambda_{\pm i \pm j}$ such that

$$(\mathbf{c}_{\pm i \pm j}^p)_z = (\lambda_{\pm i \pm j} \mathbf{c}_{\pm i \pm j})_z = -F - f, \quad (5)$$

where f denotes the ideal ML focal length. Therefore, using (2), we find

$$\lambda_{\pm i \pm j} = \frac{-F - f}{(\mathbf{R}_{x,\gamma} \mathbf{R}_{y,\beta} \mathbf{R}_{z,\alpha} \mathbf{c}_{\pm i \pm j}^{\text{id}})_z - F}. \quad (6)$$

The orthogonally projected centers $\mathbf{c}_{\pm i}^o$ are simply obtained from the $\mathbf{c}_{\pm i \pm j}$ by setting their z -value to $(-F - f)$.

B. MLA accuracy estimates

In order to simplify some of the parameters, consider the following estimates. Assuming an ideal grid, $0 = \alpha = \beta = \gamma$, $\mathbf{0} = \epsilon_k$, the focal length f_k of a ML has to be accurate within

$$\Delta_f < pf/d \quad (7)$$

in order for the disk of confusion to lie within a pixel [16] with pixel pitch p . Deviations from this constraint will lead to blur in the decoded image which cannot be compensated. Following the same argument, the rotation α and tilt (β, γ) have to be constrained such that the maximum change in distance Δ_z to the sensor fulfills the same restriction. To estimate this, we use the point $(w/2, h/2, 0)^T$ and rotate and tilt it using $\mathbf{R}_{x,\gamma} \mathbf{R}_{y,\beta} \mathbf{R}_{z,\alpha}$. The resulting z -component then yields the maximum change of distance of the MLA to the sensor. Using (3), we find

$$\Delta_z = s_x (\tan \beta / \cos \gamma) + s_y \tan \gamma < pf/d. \quad (8)$$

Note that the result does not depend on the rotation α . To obtain a common upper bound Δ_δ for the accuracies of the tilt angles, we perform a Taylor series expansion in $\beta, \gamma = 0$,

$$\Delta_z \approx s_x (\beta + \beta^3/3 + \beta\gamma^2/2) + s_y (\gamma + \gamma^3/3), \quad (9)$$

set $\beta = \gamma \equiv \Delta_\delta$ and solve

$$\Delta_z \approx \Delta_\delta (s_x + s_y) + \Delta_\delta^3 (5s_x/6 + s_y/3) < pf/d. \quad (10)$$

Additionally, the tilt introduces geometric distortion in the perspectively projected grid. That is, the ideally regular grid $\{\mathbf{c}_k\}_k$ with constant grid spacing d will be projected onto an irregular grid $\{\mathbf{c}_k^p\}_k$ with a local grid spacing

$$\begin{aligned} d_{x,\pm i \pm j} &= \|\mathbf{c}_{\pm i \pm j}^p - \mathbf{c}_{\pm(i-1)\pm j}^p\|, \\ d_{y,\pm i \pm j} &= \sqrt{3} \|\mathbf{c}_{\pm i \pm j}^p - \mathbf{c}_{\pm i \pm (j-1)}^p\|/2. \end{aligned} \quad (11)$$

The maximum difference in local grid spacing is given using the largest ML indices $i_{\max} = \lceil s_x/2d_x \rceil, j_{\max} = \lceil s_y/2d_y \rceil$ by

$$\begin{aligned} \Delta_{d,x,\max} &= |d_{x,i_{\max}j_{\max}} - d_{x,-i_{\max}-j_{\max}}|, \\ \Delta_{d,y,\max} &= |d_{y,i_{\max}j_{\max}} - d_{y,-i_{\max}-j_{\max}}|. \end{aligned} \quad (12)$$

This formula can be used to estimate whether the tilt (if within the constraint (10)) is detectable in the ML image, as we exemplarily examine in the next section in the case of the Lytro Illum camera. For all mathematical and numerical estimates made above, we provide a simple Python script which helps checking the estimates in a given application.

C. Reference data parameter choice

In the remainder, all parameters are chosen according to a Lytro Illum light field camera. That is, we use a sensor of size 7728×5368 px with a pixel pitch of $p = 1.4 \mu\text{m}$, a resolution of 10 bit and a gamma factor of 0.4. The MLs have an approximate diameter of $d = 20 \mu\text{m}$ and a fixed f-number of $f/2$, hence an ideal focal length of $f = 40 \mu\text{m}$. The MLs are arranged in a hexagonal grid with an estimated grid noise standard deviation of 0.1% of the ML diameter, i.e. $\sigma_g = 0.0143$ px. Unfortunately, it was not possible to find manufacturer specifications on the grid spacing accuracies so they had to be estimated. Of course, they can be adapted in the simulation when known in a specific application. The main lens of the Lytro Illum camera is a zoom lens with a focal length of 30 mm to 250 mm. The Lytro Illum camera provides a set of 33 different white images, taken at 10 different zoom settings and different focus settings. In order to be able to compare the synthetic results to actual white images, we choose four main focal lengths for which a corresponding white image is provided by the camera. In particular, we choose focal lengths F of 30 mm, 47 mm, 117 mm, and 249 mm. For every white image, we simulate three different aperture settings, ranging from no mechanical vignetting to strong vignetting, where the object-side aperture is chosen such that the resulting vignetting effect is visually comparable with the Lytro white image of the corresponding zoom setting and a focus setting showing the strongest vignetting.

Following (7), the ML focal lengths have to be accurate within $\Delta_f < 2.8 \mu\text{m}$. Furthermore, according to (10), assuming ideal ML focal lengths, the tilt has to be accurate within $\Delta_\delta < 0.0088^\circ$. The accuracy will have to be even higher in order for the combined $\Delta_f + \Delta_z$ to fulfill the constraint. For the maximum geometrical distortion within these constraints, following (12), we find $\Delta_{d,x,\max} \approx \Delta_{d,y,\max} = 0.0022$ px assuming a 30 mm main lens for which the scaling factor λ and distortion effects are the largest. The geometric distortion hence is negligibly small. Therefore, all images are synthesized with zero tilt, $\beta = 0 = \gamma$. When, in a given application, the geometric distortion is non-negligible, it can be included in the simulation and has to be estimated and corrected before estimating a regular ML grid from the WI. The remaining parameters, such as the grid rotation α and offset \mathbf{o}_g , are varied to obtain a collection of different WIs in order to increase the statistical significance of the evaluation.

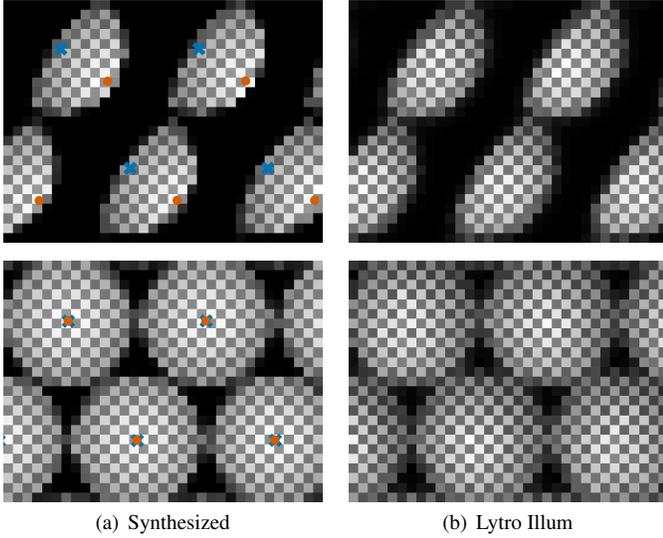


Figure 2. Comparison of a raw synthetic and Lytro Illum WI in the case of a 30 mm main lens and strong mechanical vignetting. Crops are from the top left corner (top row) and center of the image (bottom row). The perspective-projected reference centers \mathbf{c}_k^p are depicted as blue crosses, the orthogonally projected \mathbf{c}_k^o as red dots.

We ray-trace a total of 240 white images. The synthetic WIs are then mosaiced using a Bayer pattern with color response according to a Lytro Illum camera. Furthermore, we add Gaussian image noise with standard deviation σ_n of four different levels to the WIs to investigate the robustness of the grid estimation algorithms with respect to image noise. Hence, we evaluate a total of 960 different WIs. A comparison of a synthesized and a Lytro Illum WI is shown in Fig. 2. The synthesized image incorporates all characteristics of the real one, in particular natural vignetting, which causes off-center brightest pixels, and mechanical vignetting resulting in the characteristic cat eye shape of the projected microlens images close to the sensor edges.

III. MICROLENS GRID ESTIMATION

The main purpose of pre-calibration is to estimate a regular grid approximating the perspective-projected ML centers \mathbf{c}_k^p , which correspond to the coordinates of the central rays of the target light field. In the further decoding pipeline, the estimated grid is used to align the lenslet image with the sensor and slice it to a 4D light field (see Section IV).

Usually, the ML grid is estimated by detecting the ML centers from the corresponding WI [4] and building a regular grid best approximating the detected centers [7] or by directly estimating regular grid from the WI [6]. Challenges in the detection are versatile: on the one hand, the sheer amount of MLs in MLAs used in practice limits the algorithm’s complexity. On the other hand, the geometry of the MLA is not trivial and usually slightly irregular. Furthermore, main lens and ML vignetting influences the (local) shape and brightness of the ML images, particularly of those that are close to the sensor edge: while ML images near the sensor center are circular and brightest in the center, ML images closer to the sensor edge are cat-eye-shaped and show an off-center brightest pixel

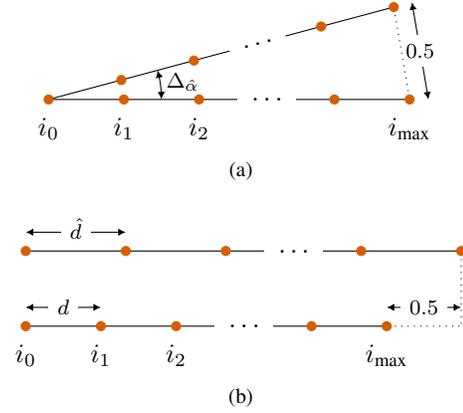


Figure 3. Exaggerated sketch for the estimated spacing and rotation accuracy requirements. (a) Grid rotation accuracy. (b) Grid spacing accuracy.

(see Fig. 2). There are mainly two methods proposed in the literature: Cho et al. [6] first compensate the rotation using an estimate obtained from the Fourier transform of the WI. In the spatial domain, they perform a grayscale erosion and clustering of the demosaiced WI. To estimate the ML centers, they use a parabolic least squares (LS) regression of the clustered MLs. In the decoding pipeline by Dansereau et al. [7], implemented as the de-facto light field decoding standard in the *MATLAB Light Field Toolbox*, the raw WI is convolved with a disk kernel. The ML centers are then estimated by finding the local maxima in the filtered image. This does not result in subpixel precision. However, in the succeeding pre-calibration, the grid parameters are estimated with subpixel precision. Furthermore, Hog et al. [9] present a Fourier-based estimation algorithm. But since it does not yield results that are significantly different from those of Dansereau et al. [7], we do not re-evaluate it here. For completeness, we will also evaluate the ML center detection used by Bok et al. [4], which is not presented in their paper but implemented in the corresponding *MATLAB* reference implementation.

None of the mentioned algorithms consider vignetting effects. Taking into account the natural and mechanical vignetting by estimating the ML grid parameters and coordinates in the spatial domain of the WI is extremely challenging. Local circle search algorithms have been proposed [15] which show good performance in the image’s central regions and in cases of strong vignetting close to the sensor edge, but mediocre performance in cases of only slight mechanical vignetting. Lytro supposedly uses a similar local arc fitting to account for mechanical vignetting [11], but, since their software is closed-source, this is speculative.

A. Grid estimation accuracy requirements

The grid parameters (grid spacing, rotation, and offset) have to be estimated with very high accuracy. Assuming that the maximum deviation from a grid point of the estimated regular grid to a true grid point may not exceed 0.5 px, we can estimate upper bounds of the individual accuracy requirements (compare Fig. 3). Assuming that the grid offset is estimated perfectly and the grid matches the real grid at the sensor center, in order for the grid points furthest from the center to be

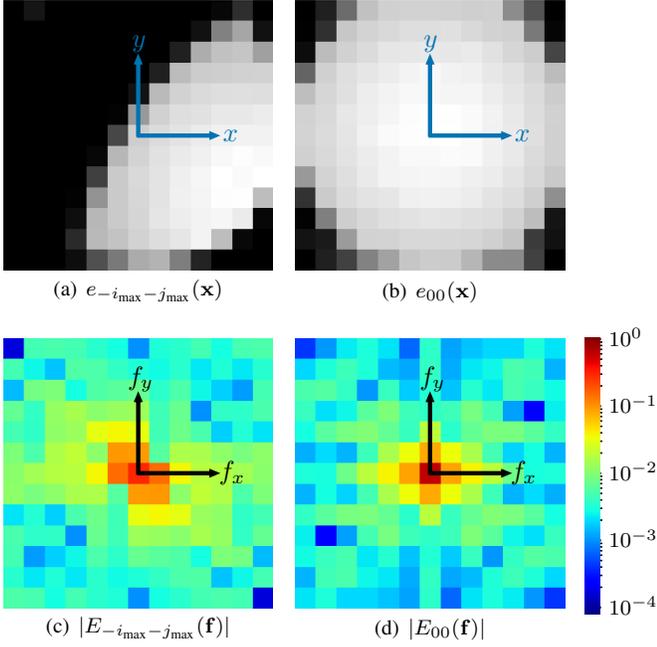


Figure 4. Comparison of central and upper left texel cropped from Fig. 2 in the spatial domain (demaosaiced, upper row) and in the Fourier domain (bottom row, normalized).

within 0.5 px of the true grid centers, the accuracy $\Delta_{\hat{d}}$ of the estimated grid spacing \hat{d} has to be within the upper bound estimate

$$|\Delta_{\hat{d}}| < \frac{0.5}{l_{\max}} = 0.0018 \text{ px} \quad (13)$$

in the case of the Lytro Illum camera. Here,

$$l_{\max} = \max\{i_{\max}, j_{\max}\} = \max\{\lceil 2s_x/d_x \rceil, \lceil 2s_y/d_y \rceil\} \quad (14)$$

is determined by the longer side of the sensor. Following a similar argument, the accuracy $\Delta_{\hat{\alpha}}$ of the estimated grid rotation $\hat{\alpha}$ has to satisfy

$$\sin|\Delta_{\hat{\alpha}}| < \frac{0.5 \cdot p}{i_{\max} \cdot d} \implies |\Delta_{\hat{\alpha}}| < \arcsin \frac{0.5 \cdot p}{i_{\max} \cdot d} = 0.0074^\circ. \quad (15)$$

Finally, the grid offset leads to a global shift of the estimated grid, and should hence at least be accurate within

$$|\Delta_{\hat{\delta}}| < 0.5 \text{ px}. \quad (16)$$

B. Proposed pre-calibration

The accuracy estimates made above pose challenging requirements on the estimation algorithms, in particular on the estimation of the grid spacing. We propose a novel algorithm which operates in the Fourier domain to estimate the grid spacing and rotation, and in the spatial domain to estimate the grid offset. The estimation takes into account the natural and mechanical vignetting present in the white images.

1) *Grid rotation and spacing estimation:* A white image can be interpreted as an (approximately) regular structure: ignoring natural and mechanical vignetting, a white image is made up of a regular texel $e(\mathbf{x})$ which is arranged in a grid spanned by the vectors \mathbf{b}_1 and \mathbf{b}_2 . For example,

$\mathbf{b}_1 = (0, 2r)$, $\mathbf{b}_2 = (\sqrt{3} \cdot r, r)$ for a perfect hexagonal grid with a ML radius of $r = d/2$. Using 2D convolution (denoted by $**$), this results in the (continuous) white image

$$g_{\text{ideal}}(\mathbf{x}) = e(\mathbf{x}) ** \sum_{i,j \in \mathbb{Z}} \delta(\mathbf{x} - i\mathbf{b}_1 - j\mathbf{b}_2) \quad (17)$$

$$\Downarrow$$

$$G_{\text{ideal}}(\mathbf{f}) \propto E(\mathbf{f}) \cdot \sum_{i,j \in \mathbb{Z}} \delta(\mathbf{f} - i\mathbf{f}_1 - j\mathbf{f}_2), \quad (18)$$

where $E(\mathbf{f})$ denotes the Fourier transform of $e(\mathbf{x})$. For the frequency basis vectors $\mathbf{f}_k = (f_{k,x}, f_{k,y})$, it holds [3]

$$\begin{pmatrix} b_{1,x} & b_{2,x} \\ b_{1,y} & b_{2,y} \end{pmatrix} = \begin{pmatrix} f_{1,x} & f_{1,y} \\ f_{2,x} & f_{2,y} \end{pmatrix}^{-1}, \quad (19)$$

where $b_{k,x}, b_{k,y}$ are the components of the vectors \mathbf{b}_k . Ideally, we can estimate the grid spacing and rotation by detecting the peaks corresponding to $\mathbf{f}_1, \mathbf{f}_2$ (and their multiples) in the absolute value of the Fourier transform $E(\mathbf{f})$.

Now, introducing natural and mechanical vignetting will not change the grid vectors but instead modulate the (now local) texel: in the image center, the texel will not be altered, but deviating from the center, the brightness distribution of the texel will change due to natural vignetting. Furthermore, due to mechanical vignetting, some pixels of the texel will be blocked. Therefore, the texel $e_{ij}(\mathbf{x})$ is different at every grid position (i, j) and we write

$$g(\mathbf{x}) = \sum_{i,j \in \mathbb{Z}} e_{ij}(\mathbf{x}) ** \delta(\mathbf{x} - i\mathbf{b}_1 - j\mathbf{b}_2). \quad (20)$$

As an example, two local (discrete) texels are shown in Fig. 4. We calculate the Fourier transform of (20) to

$$G(\mathbf{f}) = \sum_{i,j \in \mathbb{Z}} E_{ij}(\mathbf{f}) \cdot e^{-2\pi i (i\mathbf{b}_1 + j\mathbf{b}_2) \cdot \mathbf{f}}, \quad (21)$$

where $E_{ij}(\mathbf{f})$ denotes the Fourier transform of $e_{ij}(\mathbf{x})$. Since every texel is different, (21) cannot directly be written as a Dirac comb like in the ideal case. We now model the local texels as

$$e_{ij}(\mathbf{x}) = e(\mathbf{x}) \cdot m_{ij}^{\text{nv}}(\mathbf{x}) \cdot m_{ij}^{\text{mv}}(\mathbf{x}) \quad (22)$$

where $e(\mathbf{x})$ is a binary circular mask with ML radius r , $m_{ij}^{\text{nv}}(\mathbf{x})$ is the modulation due to natural vignetting, whose shape we do not have to specify more explicitly but could for example model as a wide Gaussian bell, and $m_{ij}^{\text{mv}}(\mathbf{x})$ describes the modulation due to mechanical vignetting, which can be modeled again as a binary circular mask of large radius with non-zero center. In more detail, the natural vignetting can be written as

$$m_{ij}^{\text{nv}}(\mathbf{x}) = m_{00}^{\text{nv}}(\mathbf{x} - \mathbf{o}_{ij}) \quad (23)$$

where \mathbf{o}_{ij} is the distance of the perspective projected to the orthogonally projected ML center on the sensor plane, since natural vignetting causes the brightest pixel to be at the orthogonally projected center, but the modulation shape does not change otherwise. A schematic drawing of the local texel

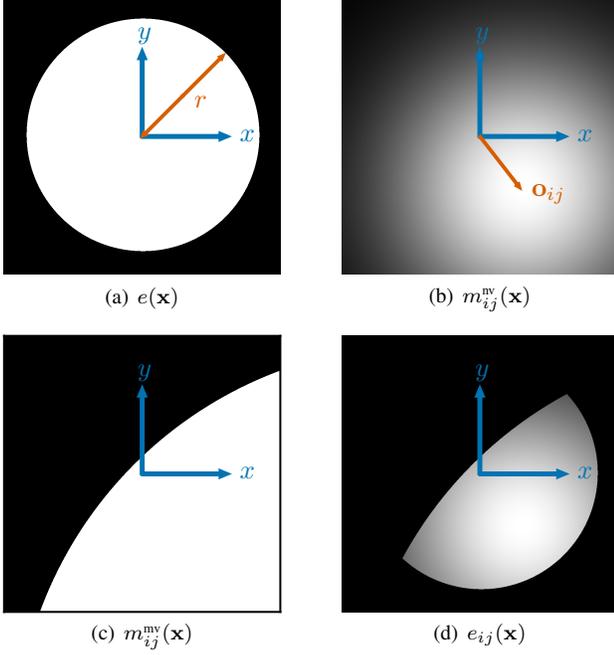


Figure 5. Schematic drawing of the continuous local texel model according to (22) in the case of strong mechanical vignetting

model is shown in Fig. 5. We calculate the Fourier transform of $e_{ij}(\mathbf{x})$ using (22) and (23):

$$\begin{aligned} E_{ij}(\mathbf{f}) &= E(\mathbf{f}) ** M_{ij}^{mv}(\mathbf{f}) ** M_{ij}^{mv}(\mathbf{f}) \\ &= E(\mathbf{f}) ** M_{00}^{mv}(\mathbf{f}) \cdot e^{-2\pi i \mathbf{o}_{ij} \cdot \mathbf{f}} ** M_{ij}^{mv}(\mathbf{f}), \end{aligned} \quad (24)$$

where both the Fourier transform $E(\mathbf{f})$ and $M_{ij}^{mv}(\mathbf{f})$ are Airy discs of different widths (and phase).

We observed that natural vignetting will cause the periodic peaks in the Fourier transform of the WI to shift. This is likely due to the underlying periodic structure of the modulation itself which manifests itself in the phase factor: the natural vignetting of the overall WI can be seen as a regular texture which is arranged in the hexagonal grid of the orthogonally projected ML centers (instead of the perspectively projected ones). To eliminate this effect, we perform a strong gamma compression to effectively eliminate the modulation $m_{00}^{mv}(\mathbf{x})$. That is, we calculate

$$\begin{aligned} \tilde{g}_\gamma(\mathbf{x}) &= g^\gamma(\mathbf{x}) \\ &\approx \sum_{i,j \in \mathbb{Z}} e(\mathbf{x}) \cdot m_{ij}^{mv}(\mathbf{x}) ** \delta(\mathbf{x} - i \mathbf{b}_1 - j \mathbf{b}_2) \end{aligned} \quad (25)$$

$$\tilde{G}_\gamma(\mathbf{f}) \approx \sum_{i,j \in \mathbb{Z}} E(\mathbf{f}) ** M_{ij}^{mv}(\mathbf{f}) \cdot e^{-2\pi i (i \mathbf{b}_1 + j \mathbf{b}_2) \cdot \mathbf{f}}, \quad (26)$$

for $\gamma \ll 1$ such that

$$(m_{00}^{mv}(\mathbf{x}))^\gamma \approx 1 \iff \delta(\mathbf{f}). \quad (27)$$

Note that since $e(\mathbf{x})$ and $m_{ij}^{mv}(\mathbf{x})$ are binary, they are unchanged from the gamma compression. This gamma compression has the additional advantage that the algorithm can equivalently operate on the raw, mosaiced WI, since the

compression will push all gray values to one, mitigating the possible effects introduced by demosaicing algorithms, which are particularly severe in the case of microlens images [8]. For the evaluation, we will hence apply our method to the raw WIs.

In practice, the texels will deviate from our texel model (22), in particular, the $e(\mathbf{x})$ and $m_{ij}^{mv}(\mathbf{x})$ will not be exactly binary. The gamma compression then would have no effect as all non-zero pixels would be mapped to one. To make sure that the texel model (22) is acceptable, we perform contrast stretching prior to gamma compression. That is, for some $q \in (0, 1)$, we linearly map the value range $[q, 0.99]$ to $[0, 1]$ and clip all values below q to zero and above 0.99 to 1. The value of q depends on the white image and the camera parameters. The resulting white image is denoted by $\tilde{g}_{\gamma,q}$. Furthermore, since the mechanical vignetting only has an effect on off-center texels $(i, j) > (I, J)$ for some $I, J \in \mathbb{Z}$, we perform windowing using a rotationally symmetric Gaussian window $w_\sigma(\mathbf{x})$ prior to calculating the Fourier transform. This suppresses off-center texels which are distorting the ideal spectrum due to mechanical vignetting. The standard deviation σ of the Gaussian window is chosen such that

$$\begin{aligned} \tilde{g}_{\sigma,\gamma,q}(\mathbf{x}) &= w_\sigma(\mathbf{x}) \cdot \tilde{g}_{\gamma,q}(\mathbf{x}) \\ &\approx w_\sigma(\mathbf{x}) \cdot \left(e(\mathbf{x}) ** \sum_{i,j \in \mathbb{Z}} \delta(\mathbf{x} - i \mathbf{b}_1 - j \mathbf{b}_2) \right). \end{aligned} \quad (28)$$

Hence, the Fourier transform of can be approximated as

$$\begin{aligned} \tilde{G}_{\sigma,\gamma,q}(\mathbf{f}) &\approx W_\rho(\mathbf{f}) ** \left(E(\mathbf{f}) \cdot \sum_{i,j \in \mathbb{Z}} \delta(\mathbf{f} - i \mathbf{f}_1 - j \mathbf{f}_2) \right) \\ &= \sum_{i,j \in \mathbb{Z}} W_\rho(\mathbf{f} - i \mathbf{f}_1 - j \mathbf{f}_2) \cdot E(i \mathbf{f}_1 + j \mathbf{f}_2), \end{aligned} \quad (29)$$

where the \mathbf{f}_i are given by (19) and the Fourier transform $W_\rho(\mathbf{f})$ of the Gaussian window is again a Gaussian window with standard deviation $\rho = 1/(2\pi\sigma)$. Therefore, the Fourier transform of $\tilde{G}_{\sigma,\gamma,q}(\mathbf{f})$ is given by a sum of shifted Gaussians centered at linear combinations of the grid frequency vectors \mathbf{f}_i . Since the standard deviation σ of the window in the spatial domain is much larger than the grid spacing, the standard deviation ρ of the window in the Fourier domain is much smaller than the grid frequency spacing. More specifically, in the case of the Lytro Illum camera, using a standard deviation of $\sigma = 100$ px and a hexagonal grid spacing of 15 px, we find a standard deviation of $\rho = 0.0016$ px⁻¹ and a smallest distance of frequency basis vectors of $d_f = 0.0770$ px⁻¹. Therefore, the center of each Gaussian in (29) lies outside the 51σ neighborhood of the closest neighboring Gaussian. The local maximum of each Gaussian is accordingly virtually undisturbed by neighboring ones. Hence, the peaks in $\tilde{G}_{\sigma,\gamma,q}(\mathbf{f})$ approximate well the linear combinations of the grid frequency spacing vectors \mathbf{f}_i .

In the discrete case we apply a rotationally symmetric Hann window whose width is determined by the length of the smaller dimension of the white image to reduce spectral leakage.

By estimating the peaks in the spectrum $\hat{G}_{\sigma,\gamma,q}(\mathbf{f})$ of the contrast-stretched, gamma-compressed and windowed white

image, we can estimate the grid spacing of the underlying perspective project microlens centers via (19). We estimate these basis vectors \mathbf{f}_i in the Fourier domain by finding the local maxima in the magnitude of the Fourier-transformed WI that correspond to the first n multiples of the frequency basis vectors $n\mathbf{f}_i$. The number n of detected maxima depends on the application. In the case of the Lytro Illum camera, we are able to find $n = 5$ values per frequency basis vector. We use zero padding and a centroid calculation to estimate the (sub-pixel) coordinate of those frequency vectors, i.e. for every cluster $\mathcal{C}_{i,n}$ around a peak corresponding to $n\mathbf{f}_i$ we calculate

$$\hat{\mathbf{f}}_{i,n} = \frac{\sum_{j \in \mathcal{C}_{i,n}} \mathbf{f}_j \tilde{G}_{\sigma,\gamma,q}(\mathbf{f}_j)}{\sum_{k \in \mathcal{C}_{i,n}} \tilde{G}_{\sigma,\gamma,q}(\mathbf{f}_k)}. \quad (30)$$

In total, we have introduced three hyperparameters (σ, γ and q) which our proposed method depends on. In order to determine these parameters appropriately (i.e., such that the made approximations hold), we need a certain measure which does not depend on any prior knowledge (e.g. the ML centers or the underlying grid spacing). To this end, we propose the following: from the estimated frequencies $\hat{\mathbf{f}}_{i,n}$, we calculate their distances

$$\hat{d}_{i,n} = \|\hat{\mathbf{f}}_{i,n+1} - \hat{\mathbf{f}}_{i,n}\| \quad (31)$$

which, ideally, should be constant in the case of a regular grid for $i = 1, 2$, respectively. Hence, the estimated standard deviation \hat{s}_i of the samples $\hat{d}_{i,n}$ should ideally be zero. We can therefore use this estimated variance in the Fourier grid spacing to optimize the hyperparameters γ, q , and σ : the optimal parameters should show a minimal variance. For a specific task, e.g. given a light field camera with a fixed prime lens, this minimization could be done manually. But for the Lytro camera, the white images corresponding to the different focal lengths show different features in terms of vignetting, grid spacing, and brightness distribution. Therefore, and to obtain an automated calibration process, we use differential evolution on a predefined search space of the hyperparameters to minimize the estimated standard deviation \hat{s}_i of the Fourier grid spacing automatically.

Using the final estimated grid frequency vectors $\hat{\mathbf{f}}_k$ with the corresponding estimated grid spacing vectors $\hat{\mathbf{b}}_k$, we straightforwardly obtain estimates of the grid spacing and rotation of the perspective projected ML grid. An overview of the proposed grid estimation is shown in Fig. 6.

2) *Grid offset estimation*: Having estimated the grid rotation and spacing, we are left to estimate the overall grid offset, which is done in the spatial domain. Dansereau et al. [7] estimate the offset by building an initial regular grid, using the previously estimated grid spacing and rotation, and measuring the median distance of the regular grid points to the previously detected ML centers.

Here, we propose a refinement of this method. First, we estimate the ML centers only in the central region of the image where the expected difference of perspective to orthogonally projected centers is less than 0.5 px. This increases the accuracy of the detection since the orthogonally projected centers

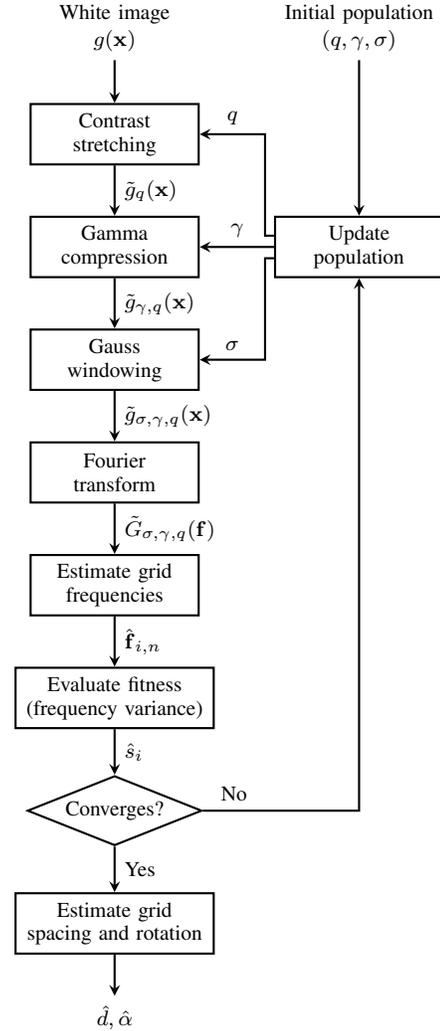


Figure 6. Flowchart of the proposed grid spacing and rotation estimation algorithm. The detailed steps for the minimization of \hat{s}_i via differential evolution have been omitted for clarity.

are more easy to detect due to natural vignetting. That is, we are restricting the detection region, denoted by Λ , such that

$$\|\mathbf{c}_{\pm i_{\max} 0}^p - \mathbf{c}_{\pm i_{\max} 0}^o\| = d \cdot i_{\max} (\lambda_{\pm i_{\max} 0} - 1) < 0.5. \quad (32)$$

Since $d_x > d_y$, this will suffice to fulfill the analogous constraint in the y -direction as well. We will use a rough estimate of the factor $\lambda \approx (F + f)/F$ obtained from (6), given a rough estimate of the image distance F and the ML focal length f , and the previously estimated MLA spacing \hat{d} . The restricted region Λ will be comparatively small, depending on the factor λ and hence on the main lens focal length, consisting of as few as 50×50 MLs in the case of a 30mm main lens. While this small region is not suited to estimate the grid spacing with high accuracy, estimating the overall offset can be done with much fewer measurements, i.e. fewer available ML centers. Having restricted the detection region, we convolve the image with a Gaussian kernel to reduce image noise. In a second step, the image is clustered using local thresholding (by local Gaussian weighted mean with a block size of 17 px), to find areas around local peaks, and a standard cluster labeling

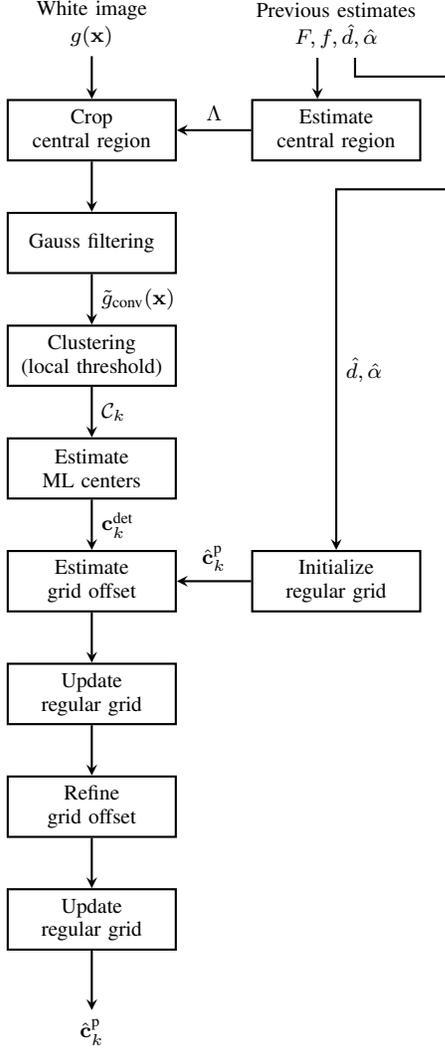


Figure 7. Flowchart of the proposed full grid estimation algorithm with offset refinement.

algorithm. Each cluster represents exactly one ML. Finally, we estimate the ML centers from the detected clusters. That is, for each detected cluster C_k , we calculate the center of mass, analogously to the calculation (30) in the Fourier domain:

$$\mathbf{c}_k^{\text{det}} = \frac{\sum_{m \in C_k} \mathbf{x}_m g_{\text{conv}}(\mathbf{x}_m)}{\sum_{n \in C_k} g_{\text{conv}}(\mathbf{x}_n)}. \quad (33)$$

To estimate the grid offset, analogously to Dansereau et al. [7], we first calculate the median distance of the initialized regular grid points $\hat{\mathbf{c}}_k^{\text{p}}$ to the detected ML centers $\mathbf{c}_k^{\text{det}}$. Additionally, in a refinement step, we will then calculate a *weighted* median distance of the updated regular grid point to detected ML centers, assigning a higher weight to those ML centers that are more central (using a symmetric Gaussian window). Since the detection inaccuracies due to natural vignetting will be smaller in the image center, this should yield a more reliable final result of the estimated grid offset. An overview of the proposed offset estimation method is depicted in Fig. 7.

Using the estimated grid spacing, rotation, and offset, we calculate a final, estimated, regular hexagonal grid $\{\hat{\mathbf{c}}_k^{\text{p}}\}_k$ approximating the perspective projected ML centers \mathbf{c}_k^{p} .

C. Quality measures

To quantitatively measure the performance of the grid estimation algorithms, we use the following quality measures. We measure the overall grid estimation accuracy Q_g by calculating the root mean square distance of estimated to true grid points:

$$Q_g = \sqrt{\frac{1}{M} \sum_{k=1}^M \|\hat{\mathbf{c}}_k^{\text{p}} - \mathbf{c}_k^{\text{p}}\|^2}. \quad (34)$$

Here, M denotes the number of grid points in the estimated grid. When grid noise has been added to the ideal grid points, we will measure higher values of Q_g . Ideally, we would estimate the perfect regular grid that the grid points are derived from, i.e. we obtain

$$Q_{g, \text{ideal}} = \sqrt{\frac{1}{M} \sum_{k=1}^M \epsilon_k^2} = \frac{\sigma_g}{\sqrt{M}} \sqrt{\sum_{k=1}^M \left(\frac{\epsilon_k}{\sigma_g}\right)^2} =: \frac{\sigma_g}{\sqrt{M}} X \quad (35)$$

where $\epsilon_k \sim \mathcal{N}(0, \sigma_g)$ and hence X is distributed according to the chi distribution with M degrees of freedom. We find the expected value

$$\mathbf{E}[Q_{g, \text{ideal}}] = \bar{Q}_{g, \text{ideal}} = \sigma_g \cdot \sqrt{\frac{2}{M}} \cdot \frac{\Gamma((M+1)/2)}{\Gamma(M/2)}, \quad (36)$$

where Γ denotes the gamma function. Using the identity

$$\lim_{n \rightarrow \infty} \frac{\Gamma(n+\gamma)}{\Gamma(n)n^\gamma} = 1, \quad \text{for all } \gamma \in \mathbb{C}, \quad (37)$$

with $n = M/2$ and $\gamma = 1/2$, we find the approximation

$$\bar{Q}_{g, \text{ideal}} \approx \sigma_g. \quad (38)$$

We will view this as the ideal mean grid estimation accuracy. To gain further insight into the grid estimation performance, we calculate the mean absolute difference of estimated to true grid rotation α as well as the mean absolute difference of estimated to true grid spacing d ,

$$Q_s = |\hat{d} - d|, \quad Q_r = |\hat{\alpha} - \alpha|. \quad (39)$$

Furthermore, we measure the runtime of each grid estimation algorithm. The evaluation was carried out on 8 cores of an AMD EPYC 7351 CPU@2.40 GHz utilizing multithreading where possible.

D. Results

For all 960 WIs, a regular grid is estimated with the different estimation algorithms and the overall grid accuracy Q_g as well as the spacing and rotation accuracies Q_s, Q_r are calculated using the ground truth ML centers and grid parameters. A detailed comparison of the overall grid estimation performances is shown in Fig. 8. Note that, since the calibration by Bok et al. [4] does not utilize a regular grid but the individually detected centers, the grid spacing and rotation accuracies Q_s, Q_r cannot be specified in this case. We observe that only the proposed algorithm satisfies the accuracy requirements as stated in Section III-A. While the algorithm by Dansereau et al. performs very well in estimating the grid rotation, its spacing estimation is robust but limited in

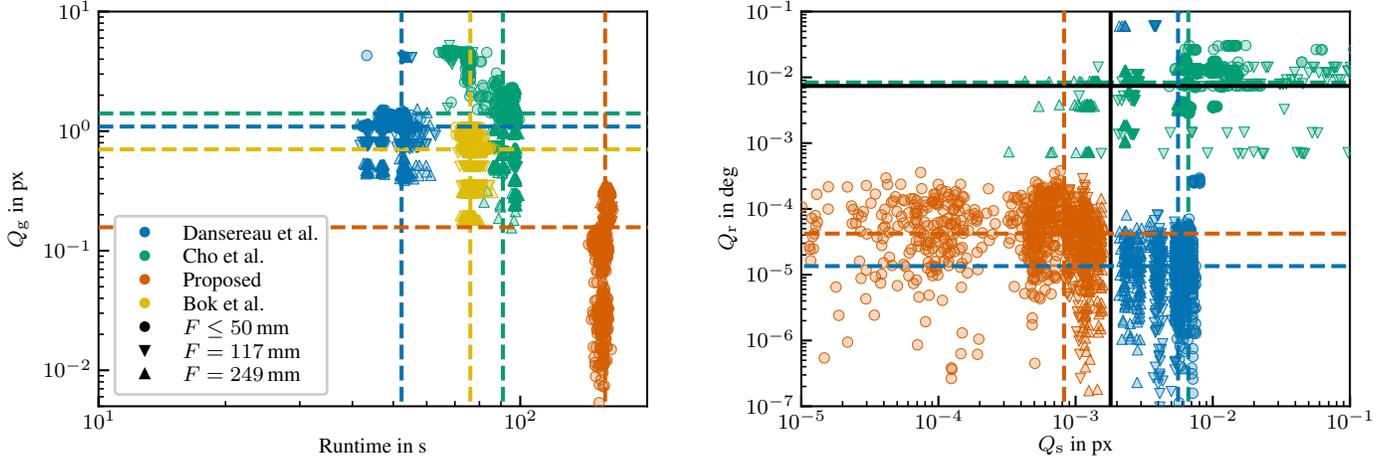


Figure 8. Performance comparison of the different ML grid estimation algorithms. The drawn colored lines show the median values of the corresponding datasets. The solid black lines denote the accuracy requirements as stated in Section III-A.

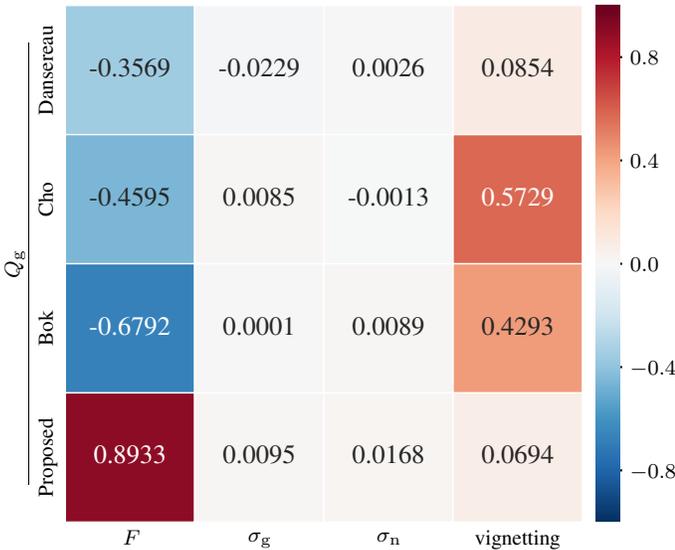


Figure 9. Pearson correlation of the grid estimation accuracy Q_g with the image distance F , grid noise σ_g , image noise σ_n and (mechanical) vignetting for the different estimation algorithms.

accuracy, in particular for short main lens focal lengths. The proposed method outperforms the others in the overall grid estimation accuracy. The improved accuracy stems from the more accurate grid spacing estimation while the performance regarding the rotation accuracy is slightly worse than the algorithm by Dansereau et al. [7]. Still, the rotation estimation is performed with very high accuracy of about 10^{-4} deg. Conversely, the method proposed by Cho et al. does not yield a robust estimation of the grid spacing and rotation. In terms of the grid accuracy Q_g , we find that the proposed method yields results of about one order of magnitude better than the other methods while occasionally showing extreme results performing even two orders of magnitude better. We observe that the grid estimation proposed by Bok et al. [4] slightly outperforms the one by Dansereau et al. [7].

Even though the proposed method has a longer runtime, with an average of about 160 s per WI compared to an average

Table I
MEAN GRID ESTIMATION ACCURACY FOR THE DIFFERENT ALGORITHMS FOR DIFFERENT IMAGE DISTANCES F (IN MM). ALL OTHER QUANTITIES IN PIXEL.

F	σ_g	$\bar{Q}_{g, \text{ideal}}$	Q_g			
			Dansereau	Cho	Bok	Proposed
30	0	0	1.2850	2.4631	0.9724	0.0865
	0.0143	0.0143	1.2855	2.6117	0.9723	0.0881
47	0	0	1.1323	1.8162	0.7124	0.0498
	0.0143	0.0143	1.1075	1.6608	0.7126	0.0561
117	0	0	0.9418	2.8216	0.5056	0.1973
	0.0143	0.0143	0.9420	2.8486	0.5057	0.1990
249	0	0	0.8238	0.6398	0.4339	0.2949
	0.0143	0.0143	0.7613	0.6369	0.4340	0.2913

of about 75 s in the case of the method by Bok et al. and about 50 s in the case of the method by Dansereau et al., we argue that this is still feasible, as the calibration usually only has to be executed once per camera. Furthermore, the runtime will be shorter in practice using a desktop PC with a clock speed higher than the used 2.4 GHz.

Investigating the results in more detail, as shown in Fig. 9, we find that all methods are unsusceptible to image and grid noise. While the method proposed by Cho et al. and the method by Bok et al. show a strong dependency on the mechanical vignetting present in the white image, the proposed method and the method by Dansereau et al. do not show such correlation. On the other hand, there seems to be a strong dependency on the image distance for all pre-calibration methods. While the methods by Bok et al., Cho et al., and Dansereau et al. perform increasingly better with larger image distances, the accuracy of the proposed method decreases. This is further analyzed in Tab. I. Since the scaling factor λ converges to 1 when the image distance increases, the influence of natural vignetting in the white image decreases. That is, with larger image distances, the orthogonally projected centers and the perspectively projected centers coincide. Hence, the methods relying on the local brightness distribution of every ML, such as the method by

Bok et al. or Dansereau et al., show an increase in accuracy. On the other hand, the proposed method shows extremely accurate estimates, close to the expected ideal mean accuracy $\bar{Q}_{g, \text{ideal}}$, in the case of a 30 mm and 47 mm main lens but a decreasing performance in the case of the 117 mm and the 249 mm lens. This is likely due to the characteristics of the mechanical vignetting. While for shorter main lens focal lengths, the mechanical vignetting only influences MLs very close to the sensor edge but with a sharp cutoff, the vignetting is more spread out in the case of a longer focal length. Therefore, the proposed algorithm is likely to use a smaller window size which decreases the estimation accuracy, since the effective resolution of the Fourier-transformed image is decreased. Still, the performance in those cases is better than the estimation accuracy reached by Bok et al., Dansereau et al. or Cho et al. Also, using a high-quality main lens for very long focal lengths should mitigate the effects of the mechanical vignetting and lead to higher estimation accuracies.

Overall, the proposed grid estimation algorithm outperforms the ones by Dansereau et al. [7], Cho et al. [6], and Bok et al. [4]. As the method by Cho et al. could not provide reliable results, we exclude it from the remaining evaluation. Next, we will evaluate how the increased grid estimation accuracy influences the light field decoding and calibration quality.

IV. LIGHT FIELD DECODING

As an example MLA-based application, we decode the raw sensor images of a Lytro Illum light field camera (and its ray tracer implementation) and investigate the decoding quality for the different ML centers and grid estimation algorithms. For the camera pre-calibration, an ideal grid is estimated as described in Section III for each of the 34 WIs that are provided by the camera. Each white image is taken at a different main lens focal length and focus distance. The used decoding pipeline mostly follows Dansereau et al. [7] and consists of the following: black and white level correction as specified in the raw image’s meta data file, devignetting and color correction, using the white image corresponding to the zoom and focus settings of the raw image, and demosaicing, here using the method by Malvar et al. [14]. Using the estimated grid parameters, the sensor image is aligned with the grid by rotating, scaling, and translating the image, such that the ideal grid point coordinates fall on pixel centers and rotation is compensated. Finally, the sensor images can be sliced to a light field and resampled to a rectangular grid. For the last step, we use a gradient-guided interpolation technique, deviating from the 1D- or Delaunay-based interpolations that are used in the *MATLAB Light Field Toolbox*: depending on the gradients of the image in x -, y - or xy -direction, we either perform 1D vertical, horizontal, or 2D bilinear interpolation to resample the light field.

The proposed decoding pipeline is included in a new Python framework that we release under an open-source license for the scientific community [10]. All shown results, if not explicitly stated otherwise, are obtained using the proposed framework for comparability. All presented images were taken with a Lytro Illum camera, or its ray tracer implementation, set to a 30 mm main lens focused at infinity.

A. Results

Two peripheral subapertures of two decoded example light fields, one synthesized using the proposed camera model and ray tracer implementation, and one taken with a Lytro Illum camera, for different MLA grid estimation algorithms together with their respective estimated ML centers and ideal grid point centers, are shown in Fig. 10.

Similarly to the results of the previous section, we observe that the proposed ML grid estimation performs best: at the image center as well as the sensor edges, the proposed method approximates the ground truth ML centers with high accuracy while the method by Dansereau et al. slightly underestimates the grid spacing and hence the ML centers at the sensor edge are not correctly approximated by the regular grid (see top row of Fig. 10). A similar observation can be made for the real WI of a Lytro Illum camera. Even though the true ML centers are unknown, the method by Dansereau et al. again seems to slightly underestimate the grid spacing compared to the proposed method (see third row of Fig. 10). The results of our framework, using the method by Dansereau et al., agrees with the results obtained with the *MATLAB Light Field Toolbox*, validating our re-implementation of Dansereau’s methods.

Regarding the decoding quality, we observe a slight improvement by using the proposed MLA estimation: in peripheral subapertures, ghosting artifacts appear less dominant (see last row of Fig. 10). The effect can be observed for both the synthetic as well as the real lenslet images. The effect in the synthesized images seems to be stronger, probably due to a stronger simulated mechanical vignetting.

Overall, we can conclude that the MLA grid estimation accuracy influences the light field decoding quality, even though the influence is subtle and mostly significant for peripheral subapertures. Nevertheless, even a slight improvement could make more subapertures available for further light field analysis, such as depth estimation and refocusing. In principle, the decoding quality could be further improved, e.g. by using an optimized demosaicing scheme [8] or a different interpolation method for optimal hexagonal to rectangular resampling [6].

V. LIGHT FIELD CALIBRATION

To further quantitatively evaluate the influence of the ML grid estimation accuracy, we perform a calibration of a Lytro Illum camera, again set to a focal length of 30 mm and focused at infinity. For the calibration, we use two different methods that have been proposed in the literature and for which the source code is publicly available. Namely the calibration by Dansereau et al. [7], using corner features in the decoded light fields, and the calibration by Bok et al. [4] which directly utilizes the raw lenslet images, using line features. Since the different methods rely on different features during the calibration, we use a different, optimized dataset for each. While the calibration by Dansereau et al. profits from many corners being present in the calibration images (and hence from smaller grid sizes), the method by Bok et al. performs better with larger grid sizes that show more line features in the raw lenslet images.

For the calibration using line features, we created a dataset containing 10 images of a checker grid with a baseline of

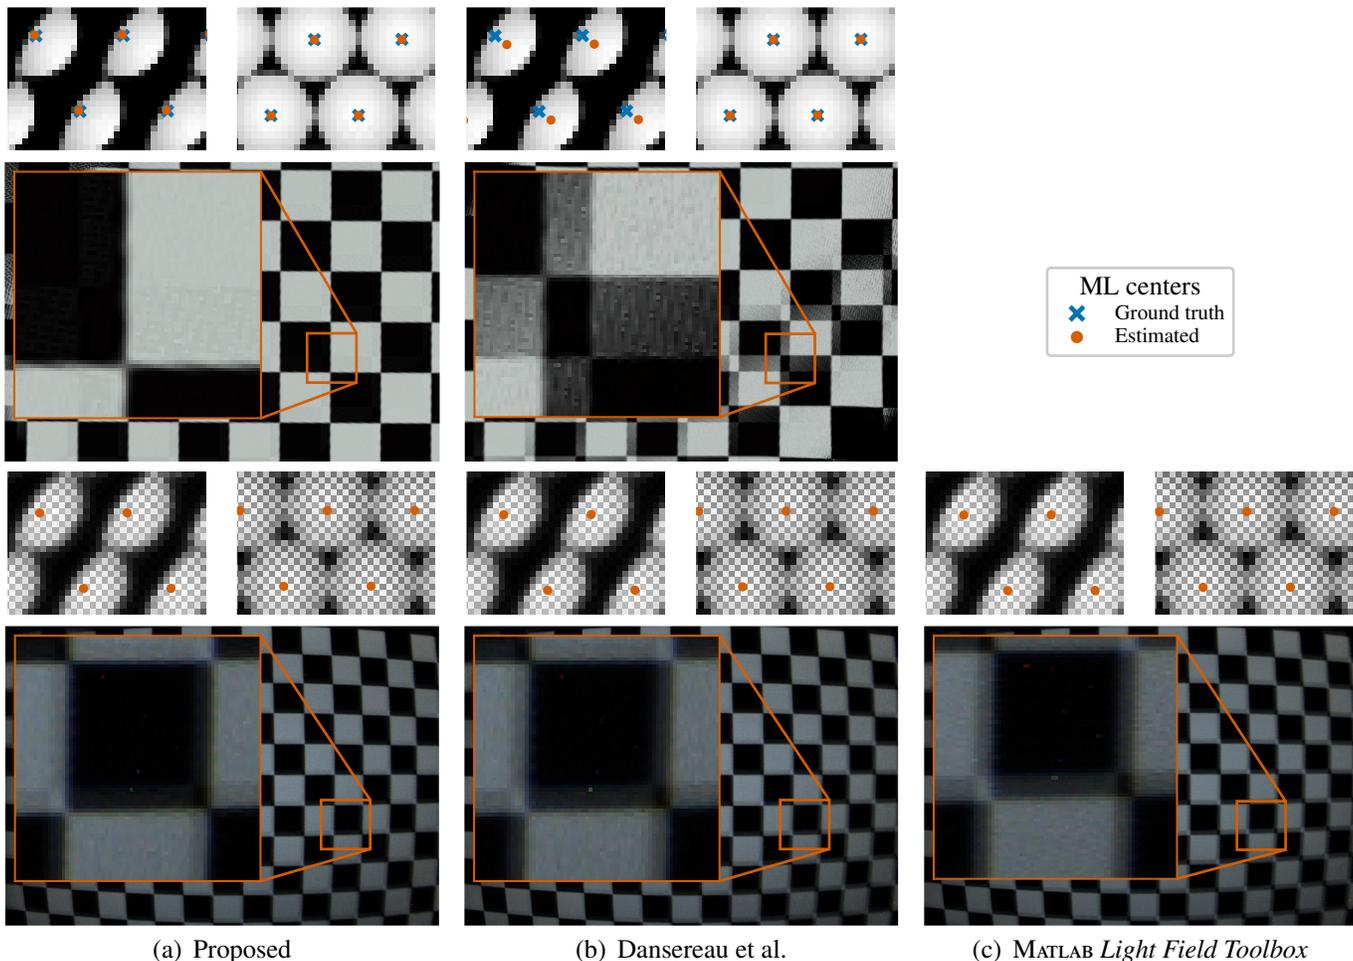


Figure 10. Comparison of peripheral subapertures decoded from the raw lenslet images and corresponding white image sections using different MLA grid estimation methods. Upper two rows: from synthetic raw images with known ground truth ML centers. Lower two rows: from images taken with a Lytro Illum camera.

15.57 mm. The calibration was performed using the MATLAB code provided by Bok et al. [4] which we have modified to be able to run the calibration with our previously estimated ML centers. We have performed the calibration using the unmodified code by Bok et al. as well as with the estimated ML grid points using the method by Dansereau et al. and our proposed method. Across all images and subapertures, the ray re-projection root mean square error (RMSE) as well as the projection RMSE is calculated.

For the calibration using corner features, we again created a dataset containing 10 images of a checker grid, now with a baseline of 6.23 mm. The calibration was performed using the MATLAB *Light Field Toolbox* by Dansereau. We have slightly modified the code such that we could include the different MLA grid estimation results. Again, we calculate the ray re-projection RMSE across all images and subapertures.

A. Results

The calibration results are shown in Tab. II and, in the case of the calibration using line features, in more detail in Fig. 11. We observe that both calibration methods, the calibration using line features as well as the calibration using corner features,

profit from the improved accuracy of the proposed MLA grid estimation algorithm. Depending on the used calibration, we are able to improve the overall ray re-projection RMSE by about 0.02 mm to 0.04 mm which is an improvement by about 15% to 20% compared to previous methods. Similarly to the results from Section IV, the gain in accuracy is larger in peripheral subapertures (as shown in Fig. 11). Again, this reflects the more accurate grid spacing estimation, which we observed in Section III, and can make further light field analysis more robust and accurate.

Additionally, in the case of the calibration using line features, we observe that the methods utilizing a regular grid to estimate the ML centers (namely the method by Dansereau et al. as well as the proposed method) result in a higher calibration accuracy. To some extent, this is surprising, as the results in Section III-D showed a slightly better ML center estimation performance in the case of the algorithm used by Bok et al. [4]. However, note that in this instance, no regular grid approximating the centers in estimated. This suggests two conclusions: first, the ML centers are more robustly estimated when approximated by a regular grid. When estimating a regular grid, a multitude of measurements are fused to obtain

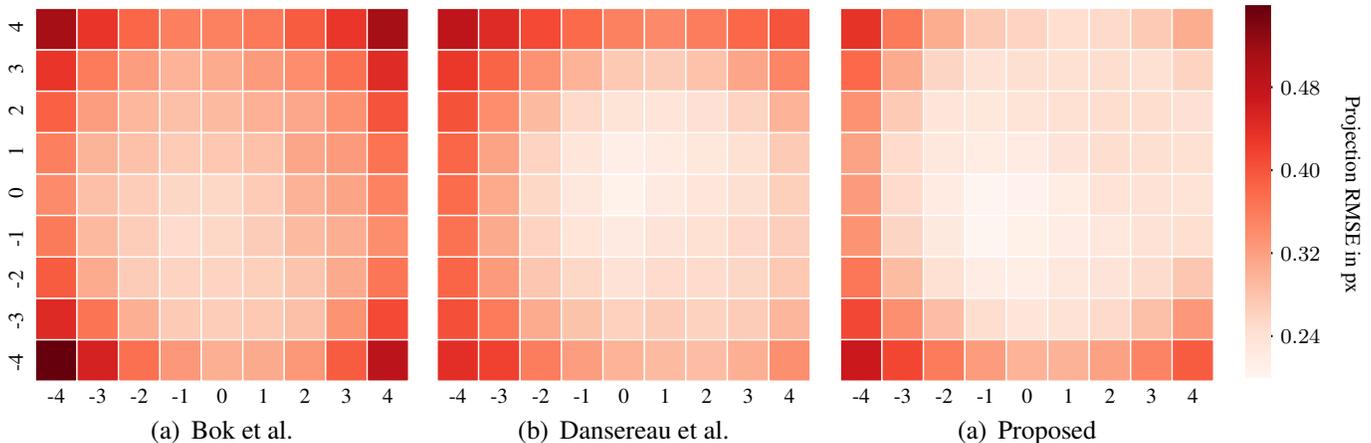


Figure 11. Projection RMSE across all calibration images for the different subaperture indices (u, v) for different ML estimation algorithms in the case of calibration using line features for a Lytro Illum camera.

Table II
ROOT MEAN SQUARE ERRORS ACROSS ALL SUBAPERTURES OF ALL CALIBRATION IMAGES FOR THE DIFFERENT CALIBRATION AND GRID ESTIMATION METHODS FOR A LYTRO ILLUM CAMERA.

Calibration method	ML estimation method	Ray reprojection RMSE in mm	Projection RMSE in px
Line features	Bok et al.	0.1409	0.4747
	Dansereau et al.	0.1263	0.4208
	Proposed	0.1118	0.3719
Corner features	Dansereau et al.	0.2583	-
	Proposed	0.2196	-

an accurate and robust result. Secondly, systematic, non-rigid deformations of the MLA are likely negligible (in the case of the used Lytro Illum camera). These irregularities should well be detected in the algorithm by Bok et al. which however shows a worse calibration result. This reinforces our decision to not include these deformations in our camera model.

VI. CONCLUSIONS

We have thoroughly investigated multiple algorithms for MLA grid estimation and the influence of the estimation accuracy on light field decoding and calibration of MLA-based light field cameras using the Lytro Illum camera as an example. We proposed a physical camera model and ray tracer implementation to synthesize application-specific white images with known ML centers that we used to quantify the performances of the different algorithms. Based on this camera model, we proposed a new MLA estimation algorithm which explicitly takes into account the natural and mechanical vignetting present in the WI and show that the proposed method outperforms the ones previously discussed in the literature. By releasing the MLA grid estimation pipeline, from WI synthesis to the evaluation of the measured accuracies and making it publicly available, we provide an easy way for the research community to find the optimal estimation parameters for their specific MLA-based application.

We have shown the importance of accurate MLA grid estimation by decoding the raw sensor images of a Lytro Illum

camera as well as its corresponding ray tracer implementation for the different estimation algorithms. With more accurate grid estimates, we find a higher quality in the decoded light fields, mostly in peripheral subaperture images. Additionally, the more accurate grid estimates contribute positively to the camera calibration, lowering the ray reprojection error by up to 20 %, which we have shown for a calibration using line features [4] as well as a calibration using corner features [7]. Investigations by Bok et al. [4] suggest that increased calibration performance positively effects disparity estimation as well as 3D reconstruction of the calibrated light fields. Furthermore, we assume that measurement application with high accuracy demands, such as light field deflectometry [21] or structured light-based methods [5], [12], can potentially profit from this improved calibration performance. To investigate the impact of calibration performance on a specific measurement task, additional thorough evaluation is necessary.

For the pre-calibration and decoding of light fields, we introduced a new Python package that we make publicly available [10]. The framework offers an easy-to-use user interface as well as a modern modular and object-oriented approach to light field decoding and calibration which is easy to extend for new applications (i.e. new computational cameras). Furthermore, the framework provides modules for general light field and hyperspectral image analysis, such as depth estimation, refocusing, color conversion, and more.

VII. ACKNOWLEDGMENTS

The authors acknowledge support by the state of Baden-Württemberg through bwHPC, a massively parallel computer. The authors would like to thank David Uhlig for many helpful discussions on the topic.

REFERENCES

- [1] E. H. Adelson and J. Y. Wang. Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):99–106, 1992. 1

- [2] G. R. Arce, D. J. Brady, L. Carin, H. Arguello, and D. S. Kittle. Compressive coded aperture spectral imaging: An introduction. *IEEE Signal Processing Magazine*, 31(1):105–115, 2014. 1
- [3] J. Beyeler, F. Puente León, and C. Frese. *Machine Vision – Automated Visual Inspection: Theory, Practice and Applications*. Springer, Berlin Heidelberg, 2016. 5
- [4] Y. Bok, H.-G. Jeon, and I. S. Kweon. Geometric calibration of microlens-based light field cameras using line features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(2):287–300, 2017. 1, 4, 8, 9, 10, 11, 12
- [5] Z. Cai, X. Liu, X. Peng, Y. Yin, A. Li, J. Wu, and B. Z. Gao. Structured light field 3D imaging. *Optics express*, 24(18):20324–20334, 2016. 12
- [6] D. Cho, M. Lee, S. Kim, and Y.-W. Tai. Modeling the calibration pipeline of the Lytro camera for high quality light-field image reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3280–3287, 2013. 1, 4, 10
- [7] D. G. Dansereau, O. Pizarro, and S. B. Williams. Decoding, calibration and rectification for lenselet-based plenoptic cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1027–1034, 2013. 1, 4, 7, 8, 9, 10, 12
- [8] P. David, M. Le Pendu, and C. Guillemot. White lenslet image guided demosaicing for plenoptic cameras. In *IEEE 19th International Workshop on Multimedia Signal Processing*, pages 1–6, 2017. 6, 10
- [9] M. Hog, N. Sabater, B. Vandame, and V. Drazic. An image rendering pipeline for focused plenoptic cameras. *IEEE Transactions on Computational Imaging*, 3(4):811–821, 2017. 2, 4
- [10] Institute of Industrial Information Technology, Karlsruhe Institute of Technology. Public GitLab repositories. <https://gitlab.com/iit-public>, GNU GPLv3 License, 2019. 2, 10, 12
- [11] C.-K. Liang and Z. Wang. Calibration of light-field camera geometry via robust fitting, Aug. 16 2016. US Patent 9,420,276. 4
- [12] X. Liu, Z. Cai, X. Peng, and B. Z. Gao. Fringe projection profilometry performed with a light field camera. In *Digital Holography and Three-Dimensional Imaging*, pages DTh4B–5. Optical Society of America, 2018. 12
- [13] A. Lumsdaine and T. Georgiev. The focused plenoptic camera. In *IEEE International Conference on Computational Photography*, pages 1–8. IEEE, 2009. 1, 2
- [14] H. S. Malvar, L.-W. He, and R. Cutler. High-quality linear interpolation for demosaicing of Bayer-patterned color images. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 3, pages iii–485. IEEE, 2004. 10
- [15] L. Mignard-Debise. *Tools for the paraxial optical design of light field imaging systems*. PhD thesis, Université de Bordeaux, 2018. 4
- [16] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a hand-held plenoptic camera. *Computer Science Technical Report CSTR*, 2(11):1–11, 2005. 1, 2, 3
- [17] T. Nürnberg, M. Schambach, D. Uhlig, M. Heizmann, and F. Puente León. A simulation framework for the design and evaluation of computational cameras. In *Automated Visual Inspection and Machine Vision III*, volume 11061 of *Proceedings of SPIE*, Bellingham, WA, 2019. 2
- [18] C. Pitts, T. J. Knight, C.-K. Liang, and Y.-R. Ng. Compensating for variation in microlens position during light-field image processing, Sept. 9 2014. US Patent 8,831,377. 2
- [19] R. Shogenji, Y. Kitamura, K. Yamada, S. Miyatake, and J. Tanida. Multispectral imaging using compact compound optics. *Optics Express*, 12(8):1643–1655, 2004. 1
- [20] R. P. Stanley, A. Chebira, A. Ghasemi, and A. L. Dunbar. Hyperspectral imaging using a commercial light-field camera (conference presentation). *Proc. SPIE*, 2017. 1
- [21] D. Uhlig and M. Heizmann. Multi-stereo deflectometry with a light-field camera. *tm – Technisches Messen*, 85(S1):59–65, 2018. 12
- [22] C. Zhang, Z. Ji, and Q. Wang. Decoding and calibration method on focused plenoptic camera. *Computational Visual Media*, 2(1):57–69, 2016. 1
- [23] Q. Zhang, C. Zhang, J. Ling, Q. Wang, and J. Yu. A generic multi-projection-center model and calibration method for light field cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1