

A Reconfigurable Camera Add-On for High Dynamic Range, Multispectral, Polarization, and Light-Field Imaging

Alkhazur Manakov^{1*} John F. Restrepo^{1*} Oliver Klehm² Ramon Hegedüs²
 Elmar Eisemann³ Hans-Peter Seidel² Ivo Ihrke^{1,2,4}

¹Saarland University ²MPI Informatik ³Delft University of Technology ⁴INRIA

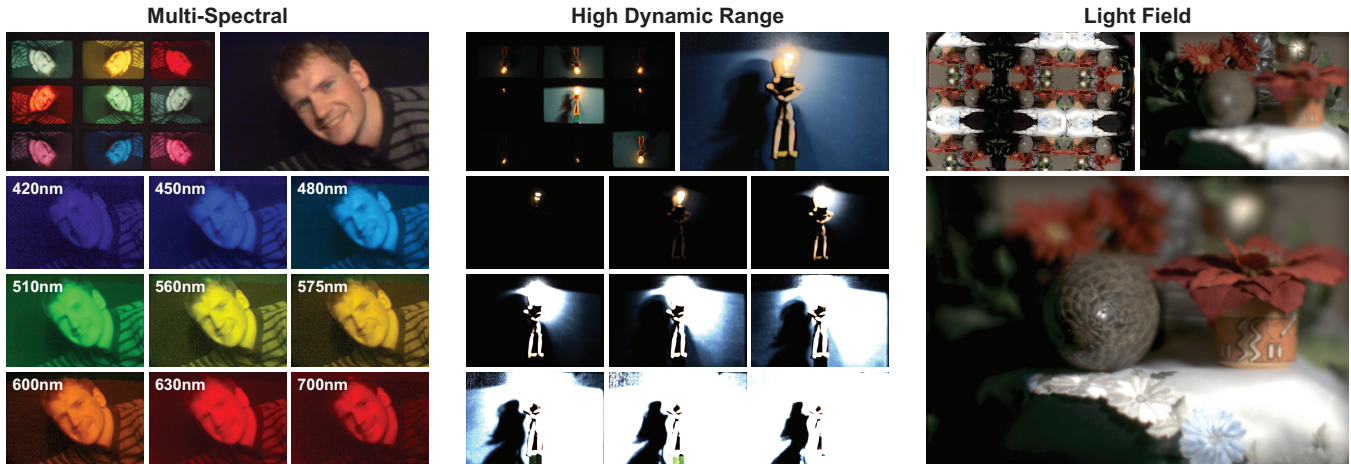


Figure 1: Exemplary applications of our camera add-on: Left: Multispectral imaging: unprocessed output (top left), a spectral stack of images after processing (bottom), a neutral image relit with a flat spectrum (top right). Middle: High Dynamic Range Imaging: unprocessed output (top left), a simulated exposure sweep of contrast ratio 100 : 1 (bottom), a tone-mapped version of the HDR image (top right). Right: Light-field imaging: unprocessed output (top left), virtual refocussing on foreground (bottom) and background (top right).

Abstract

We propose a non-permanent add-on that enables plenoptic imaging with standard cameras. Our design is based on a physical copying mechanism that multiplies a sensor image into a number of identical copies that still carry the plenoptic information of interest. Via different optical filters, we can then recover the desired information. A minor modification of the design also allows for aperture subsampling and, hence, light-field imaging. As the filters in our design are exchangeable, a reconfiguration for different imaging purposes is possible. We show in a prototype setup that high dynamic range, multispectral, polarization, and light-field imaging can be achieved with our design.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture

Keywords: computational photography, computational optics

Links: [DL](#) [PDF](#) [WEB](#) [VIDEO](#) [DATA](#)

*equal contribution

ACM Reference Format

Manakov, A., Restrepo, J., Klehm, O., Hegedüs, R., Eisemann, E., Seidel, H., Ihrke, I. 2013. A Reconfigurable Camera Add-On for High Dynamic Range, Multispectral, Polarization, and Light-Field Imaging. ACM Trans. Graph. 32, 4, Article 47 (July 2013), 13 pages. DOI = 10.1145/2461912.2461937 <http://doi.acm.org/10.1145/2461912.2461937>

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
 Copyright © ACM 0730-0301/13/07-ART47 \$15.00.
 DOI: <http://doi.acm.org/10.1145/2461912.2461937>

1 Introduction

Imaging dimensions of the plenoptic function [Adelson and Bergen 1991] has been a long-standing goal of the imaging community even before its formal definition by Adelson and Bergen. Access to the full properties of the incident light on a sensor has a large number of applications in scientific imaging, industrial quality control, remote sensing, computer vision, and computer graphics.

Numerous specialized devices, ranging from space-borne imagers to microscope cameras, exist for classic multi-spectral and polarization imaging. More recently, dynamic range restrictions of sensors (high dynamic range imaging) and directional variation of light (light-field capture) have become a major focus in computer graphics. In order to gain access to these physical dimensions of an image, the light integration has to be adapted. The three major approaches are *i*) temporal multiplexing where an image stack is recorded and different filters are placed in the light path of different exposures. This approach can only be applied to static or quasi-static scenes. The latter requires a registration of the individual images which is a difficult problem in itself. The second approach is *ii*) hardware parallel acquisition, where the optical image is multiplied by means of a beam-splitter arrangement and projected onto different sensor units that are spatially de-localized. Different optical pre-filters can be inserted into the different optical light paths. This arrangement allows for dynamic scenes to be imaged. It comes, however, at the price of large, expensive, and bulky setups that have to be custom built. Further, synchronization and radiometric calibration of the different sensors with respect to each other is another problematic aspect. The third approach is *iii*) spatial multiplexing. Here, a single sensor unit is being employed where every pixel is associated with a different optical pre-filter. This design avoids syn-

chronization and calibration issues and allows for single-exposure (snapshot) retrieval. Its most familiar application is color imaging via a color filter array. The concept has, however, been extended, to cover all plenoptic dimensions and even devices that do not exploit physical per-pixel filters such as lenslet and mask-based light-field cameras can be interpreted this way [Ihrke et al. 2010]. The major drawback of this solution are custom sensor designs or the permanent modification of standard sensors.

We aim at lifting these constraints by enabling snapshot imaging of the plenoptic dimensions of an image as in *ii*) and *iii*), however, avoiding the necessity of custom camera designs, but aiming at utilizing a single sensor as in *iii*). Moreover, we wish to design an optical element that can reversibly be attached to a standard camera, but can easily be removed for standard operation.

Our device physically multiplies an optical image by means of repeated mirror operations, which makes it suitable for time-crucial shots. The resulting array of *identical* images maintains the plenoptic properties of the physical image. These copies are projected into a plane where they are optically pre-filtered and then imaged by a 1:1 imaging unit onto the sensor of a standard camera. Our optical element can easily (and reversibly) be inserted between the lens and body of a standard camera, converting it into a high dynamic range, multi-spectral, polarization, or light-field imaging unit, depending on the configuration of our proposed optical element. We demonstrate this design in an optical table setup, but discuss how the construction could be miniaturized.

Our design features several advantages. In particular, our contributions are an optical design for a removable camera add-on for snapshot plenoptic imaging, that:

- features *identical* sub-images without parallax for filter-based imaging by employing a diffuser architecture, enabling scene-independent, pre-calibrated acquisition,
- allows for an increased photographic control of plenoptic imagery as compared to existing approaches, by avoiding the use of sub-apertures for HDR, multi-spectral, and polarization imaging,
- enables a light-field camera design which allows for the control of depth-of-field vs. light throughput, and a spatial/angular resolution trade-off that facilitates high spatial resolution (full HD) recording and refocusing.

2 Related Work

Our device allows for snapshot plenoptic imaging and we shortly discuss the most related literature. General Plenoptic image acquisition is described in two excellent articles; Wetzstein et al. [2011] discuss related work by plenoptic dimension whereas Zhou et al. [2011] concentrate on a classification by coding approaches.

Snapshot imaging captures differently-filtered images in a single exposure. A simple hardware-parallel setup is to place several cameras close to each other (e.g. [Wilburn et al. 2005]), each with a different pre-filter or with modified recording parameters. Nonetheless, if light-field [Levoy and Hanrahan 1996; Gortler et al. 1996] imaging is not intended (as is the case for all plenoptic dimensions other than direction), the different viewpoints of the cameras need to be compensated for. Optical flow [Horn and Schunck 1981] can be applied, but in practice, it is of limited utility since it is based on the brightness constancy assumption which is necessarily violated when different optical pre-filters are being used. Similar considerations hold for light-field based plenoptic imaging solutions [Horstmeyer et al. 2009; Georgiev et al. 2011].

For a *monocular* snapshot, the scene is to be observed through a common main lens to avoid parallax between the differently-filtered images. A hardware-parallel solution is to use optical splitting trees [McGuire et al. 2007]. The light path is separated into different branches by a system of beam splitters. At the end of each light path, a standard sensor is installed and a separate optical filter can be introduced in each one of them. This mode of imaging is employed in a number of practical systems for HDR imaging [Tocci et al. 2011], polarization imaging [Pezzaniti et al. 2008], and multispectral imaging [Spiering 1999; McGuire et al. 2007].

Single-sensor setups often make use of multiplexing, which comes in two flavors, *direct multiplexing*, where the optically pre-filtered images are re-distributed on the sensor and directly imaged, and *computational multiplexing*, where super-positioned plenoptic quantities are observed. In the latter case, a computational inversion has to recover the plenoptic quantities of interest. Often, this inversion step is ill-posed and has to be regularized using prior information about the signal. The different approaches are most easily exemplified in specific application areas, presented hereafter.

Multispectral imaging Capturing different color primaries has manifold applications in computer graphics and vision (e.g., re-lighting beyond white point adjustment, or improved tracking and segmentation accuracy [Park et al. 2007; Cao et al. 2011]) and in general scientific imaging (e.g., satellite-based remote sensing, fluorescence microscopy, optical quality control).

For color photography, the Bayer filter [1976] is often used, which is placed directly in front of the sensors. In principle, various primaries or other plenoptic dimensions [Narasimhan and Nayar 2005] can be captured in this way. Other direct multiplexing approaches use a dispersing element that is inserted into the optical path, while light is simultaneously re-directed or blocked from reaching nearby sensor pixels. Examples include the use of mirror arrays, so-called image splicers [Bonnet et al. 2004], or blocking masks [Du et al. 2009; Cao et al. 2011]. The resulting output is typically of low spatial resolution. A recent trend is to capture a high-resolution RGB image in conjunction with the low-resolution multispectral image and to transfer the spectral information to the high-resolution RGB image [Rump and Klein 2010; Cao et al. 2011].

Alternatively, computational multiplexing approaches have been investigated. Computed Tomography Imaging Spectrometry [Okamoto and Yamaguchi 1991; Descour and Dereniak 1995], inserts a diffraction grating into the optical path, effectively generating multiple copies of the scene that are smeared according to wavelength. Via a tomographic reconstruction the spectral scene content can be computed. Unfortunately, due to the weak transmission of higher order diffraction signals, the method requires HDR imaging to be successfully employed [Habel et al. 2012]. Coded Aperture Snapshot Spectral Imaging [Gehm et al. 2007] employs an optical design that allows for a random modulation of the spectral image content. The computational scheme then employs sparse recovery methods. Other solutions combine Fourier transform imaging and aperture splitting [Kudenov and Dereniak 2011]. All these approaches yield low-resolution images dominated by regularizing information.

High Dynamic Range HDR imaging [Mann and Picard 1995; Debevec and Malik 1997; Reinhard et al. 2010] circumvents the dynamic-range restriction of standard imaging sensors. It can be considered a variant of multispectral imaging since it employs spectral filters with a flat response and the previously-discussed approaches could, theoretically, be applied. In practice, exposure sequences [Mann and Picard 1995; Debevec and Malik 1997] are predominantly used. Snapshots are possible with varying pixel

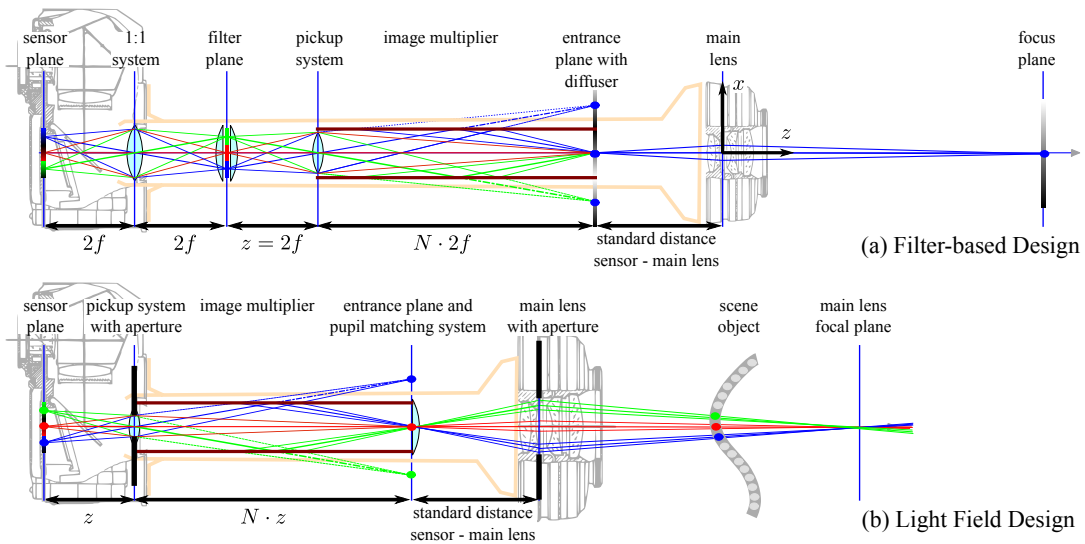


Figure 2: Our camera add-on for filter-based (a) and light-field imaging (b). The illustrations are described in detail in the text. In case (a), differently-colored light paths image differently-filtered copies of the same scene point. In case (b), differently-colored light paths indicate different viewing directions in the entrance plane of the system, which corresponds to the micro-lens plane in lenslet-based light-field cameras.

exposures [Nayar and Mitsunaga 2000], but this approach shares the previously-mentioned drawbacks. Recently, a beam splitter tree [McGuire et al. 2007] has been used to construct an HDR video camera [Tocci et al. 2011].

Polarization imaging The polarization of the field reflected off a scene object is affected by material type, roughness, surface orientation, the polarization state of the illumination, as well as scattering properties of the surrounding medium. Polarization, hence, carries important information about material and scene properties, enabling, e.g., dehazing [Schechner et al. 2001], or the removal and editing of specularities [Neumann et al. 2008].

While for multispectral imaging relatively simple physical wavelength splitting mechanisms via dispersion or diffraction are possible, polarization does not share such properties (birefringence is only a comparatively small effect). Hence, fast temporal multiplexing by filter wheels or beam splitting [Pezzaniti et al. 2008] is applied. Alternatively, the generalized mosaic of Schechner and Nayar [Schechner and Nayar 2005] can be used, which has also been applied to other plenoptic dimensions.

Light-field imaging Light-fields represent the angular variation of incident light at every sensor element. Even though the principles have been invented a century ago [Ives 1903; Lippmann 1908], only the advent of digital camera technology made them practical. Virtual refocussing and view point change [Isaksen et al. 2000; Ng 2005] within the limits of the main camera lens, but also depth estimation [Wanner and Goldluecke 2012a; Wanner and Goldluecke 2012b] are possible applications.

The primary means to record monocular, i.e. in-camera, light-fields is to employ lenslet arrays [Adelson and Wang 1992; Ng et al. 2005], but also attenuation masks placed closely to the sensor can be used [Veeraraghavan et al. 2007; Lanman et al. 2008].

A camera aperture exceeding light-field was achieved by Levoy et al. [2004] who use a mirror array for synthetic aperture imaging and projection. Most related to our approach, Han et al. [2003] obtain hemispherical imaging of flat objects via a kaleidoscopic mirror. This arrangement to measure BTFs was extended to three-

dimensional objects by Reshetouski et al. [2011] and subsequently used for reflectance scanning [Ihrke et al. 2012].

We present a novel optical design using a direct spatial multiplexing scheme for monocular snapshot plenoptic imaging. It also relies on inter-reflections within mirror systems, but can be added non-permanently to any standard camera and be reconfigured to image different plenoptic dimensions in a unified design. This aspect is one of the key contributions of this article. We demonstrate the feasibility of our setup with a large-scale prototype that is used in a number of application areas, ranging from HDR, multispectral, and polarization imaging to light-field acquisition. In contrast to other approaches, our light-field design is most suitable for imaging at a low directional, but high spatial resolution. We demonstrate that high-quality refocussing and view-point changes are possible and that unprecedented large virtual apertures at a high spatial resolution can be achieved.

3 Overview

Our goal is the reversible modification of standard camera hardware to flexibly acquire several plenoptic dimensions, such as wavelength, polarization and direction, via a simple optical design placed between a camera body and a standard objective lens.

We propose two variants of our design, both built on a physical copy mechanism that exploits the inter-reflections in a kaleidoscopic system of mirrors. The following discussion refers to Fig. 2 (a) and describes the first variant, which allows for capturing improved dynamic range, multi-spectral information, and polarization without parallax between the different copies of a scene.

The *original image* of the standard camera lens arrives on a *diffuser* screen that is placed in the location that would usually be occupied by the camera's sensor. This diffuser is observed through a mirror arrangement that we refer to as *image multiplier*, which produces a number of $N \times N$ *identical* copies of the original image that still carry the physical information of the plenoptic function. Our design then employs a *pickup imaging system* that images the information exiting the mirror system to a *filter plane*. The image at the filter plane has the dimensions of the original sensor, but contains




Figure 3: The maximum observation angle α at the entrance plane is obtained when observing a point at the image boundary of one of the outermost mirrored copies. The maximum angle is formed by the ray exiting the pickup-system aperture on the opposite side.

$N \times N$ spatially separated identical copies of the original image. These copies can be individually modulated by optical filters placed in the filter plane, hereby, enabling, among others, snapshot high dynamic range, multi-spectral, and polarization imaging (Sec. 6.1). One could also have placed a custom sensor in this plane with the respective filters attached to its surface, however, since we aim at a reversible add-on, we choose to image the filtered results onto the original camera sensor by employing a *1:1* imaging system.

The second variant of our setup, Fig. 2 (b), allows for aperture sub-sampling of the main camera lens and enables light-field imaging, Sec. 4.2. As we will show, this can be achieved by replacing the diffuser by a correcting lens. The *1:1* imaging system is not needed in this design. The details of both designs are explained in the following section.

4 Optical Design

4.1 Filter-Based Design

Our basic design for filter-based imaging is shown in Fig. 2 (a). We describe the individual components, their function and design parameters by moving from the scene towards the sensor, i.e. from the right towards the left.

The diffuser The main lens is imaging the scene onto a plane that would typically contain the camera sensor. In our design, a diffuser is placed at this location. Its size matches what the main optics are optimized for, as important imaging characteristics like the field-of-view directly depend on it. The diffuser acts as a rear-projection screen, i.e., observing it from the left shows the image that would be observed by a sensor at this location. The key role of the diffuser is to remove the directional dependency in this image and thus the generation of an identical appearance when viewed from different directions. Otherwise all physical attributes of the plenoptic function are maintained. The operation of our copying mechanism can be regarded as a multiplexing scheme that transfers identical image content into the directional component of the plenoptic function. We use a polarization-preserving diffuser, for details see Sec. 6.1.

It is important that the diffuser lobe is wide enough to accommodate the different viewing directions that create the image copies, see Fig. 2 (a), otherwise vignetting occurs. However, if the lobe is too wide, stray light is spread into the system. The diffuser scattering profile should therefore be adapted to the maximum observation angle α , see Fig. 3, for best performance and light efficiency of the system. In addition, a pupil matching (not shown) lens should be used to adapt the entrance pupil of the image multiplier to the exit pupil of the main lens. We will discuss this component in the

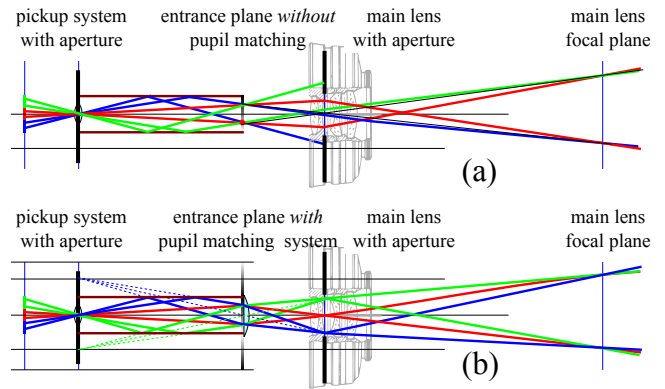


Figure 4: For light-field imaging, the exit pupil of the main lens and the entrance pupil of the image multiplier have to be matched. a) without pupil matching b) with pupil matching.

light-field imaging case, Sec. 4.2, where it is essential for proper operation. For the current discussion, this lens homogenizes the illumination picked up from the entrance plane in the case of a weak diffuser, for which strong and directionally varying illumination may otherwise manifest itself in a non-uniform transmission of the system.

The image multiplier Once the image is visible on the diffuser screen, we aim at copying it by means of mirror reflections via the image multiplier. A suitable choice is a kaleidoscope with parallel walls, resulting in a virtual plane of image copies.

The width and the height of the image multiplier are defined by the sensor size. Therefore, the only variable is its length along the optical axis. This length is determined by the $1:N$ minification that the pickup imaging system is designed to produce, and by its focal length f_{ps} . The effect of the pickup imaging system is that $N \times N$ views of the diffuser are compressed to the size of a standard sensor image that are made accessible as a real image in the filter plane. Following geometrical optics, the relation between image multiplier length $N \cdot z$, number of image copies N , and focal length of the pickup system f_{ps} is approximately given by the thin lens equation

$$\frac{1}{z} + \frac{1}{N \cdot z} = \frac{1}{f_{ps}}; \quad \frac{N}{N+1}z = f_{ps}. \quad (1)$$

In practice, this means that a short focal length f_{ps} and a low image multiplication factor N lead to short lengths of the image multiplier. Another aspect of the design is the aperture of the pickup lens. In conjunction with the diffuser lobe, it determines the light efficiency of the system. Hence, it should be chosen as large as possible. Note that a large aperture for the pickup system does not involve a loss of image quality, since we are imaging a planar object (i.e. the entrance plane). In general, it is however difficult to obtain large aperture for short focal length lenses, as they become bulky and have a strong curvature, leading to significant deviations from the geometric optics model. Due to the competing demands for the aperture and the focal length of the pickup system, a compromise between them, and, therefore, the image multiplier length, has to be found. An additional effect of the length parameter is the observation angle under which the different copies of the entrance plane are seen. A larger length leads to smaller observation angles and therefore to weaker diffusion requirements. More specifically, the maximum observation angle α is given by

$$\alpha = \tan^{-1}\left(\frac{(N/2 \cdot l_f + a_{ps}/2)}{l_{im}}\right),$$




Figure 5: Effect of the light-field pupil-matching system (simulation): without pupil-matching lens (left), with pupil-matching lens (right). Without adapting the pupils of the main lens and image multiplier, large parts of the dragons are vignetted out.

where l_f is the original sensor (and therefore the diffuser) size, N is the targeted number of image copies, a_{ps} the aperture of the pickup lens, and l_{im} the length of the image multiplier (Fig. 3). It can be reduced by a longer image multiplier, a low number of image copies, a smaller sensor size, and to a minor effect by reducing the aperture of the pickup system.

For the best optical quality and geometric accuracy, the multiplier can be made from glass, utilizing the effect of total internal reflection to create the mirror images. In this case, its length is approximately multiplied by the refractive index of the glass, which can be derived by considering two planar air/glass – glass/air interfaces. The condition on the maximum observation angle does not change; since the diffusion lobe refracts into the image multiplier, it narrows by the same amount as the maximum observation angle.

The filter plane The system discussed so far generates a real image of $N \times N$ copies of the physical image that a standard camera would have captured. Further, it makes these copies accessible in the filter plane where an array of optical filters allows us to gain access to the different plenoptic dimensions.

A slight complication occurs due to the fact that the image in the filter plane is diverging in the direction of the sensor. Severe vignetting occurs without proper compensation. To avoid this problem, the exit pupil of the image multiplier system has to be adapted to the entrance pupil of the 1:1 imaging system. For this purpose, we insert a pair of plano-convex lenses at the filter plane that together form an additional 1:1 imaging system between the aperture plane of the pickup system and that of the 1:1 imaging system.

The 1:1 imaging system The task of the 1:1 imaging system is to project the $N \times N$ optically pre-filtered copies of the diffuser-plane image onto the sensor that integrates the incoming photons. Since 1:1 imaging occurs at two focal lengths, we choose to dimension our system with respect to the focal length f of the 1:1 imaging lens. The choice of placing the pickup system at a distance of $2f$ from the filter plane is determined by keeping all imaging planes of the system equal in size to the original sensor dimensions. The overall length of our system is therefore $(6 + 2N) \cdot f$ and the individual lens components have focal lengths of $2f$ for the pair of the plano-convex lenses and $2N/(N + 1) \cdot f$ for the pickup lens.

4.2 Light-Field Design

With minor modifications, Fig. 2 (b), our system can be used for light-field imaging. For this case, we can omit the filter plane and the 1:1 imaging system of the filter-based design and image the output of the pickup system directly onto the sensor (although one could also combine both systems to obtain optically-filtered light-field images as in [Horstmeyer et al. 2009]). The light-field imaging



Figure 6: The aperture of the pickup system determines the depth of field of the light-field views (shown is the center view, but other views behave similarly). With open pickup aperture, depth of field effects can be clearly observed (left). With a closed pickup aperture, the views are sharply imaged over the full depth range (right).

effect is based on a subsampling of the main lens aperture [Ng et al. 2005]. In the following, we explain our setup and discuss the resulting imaging geometry.

Aperture Sub-Sampling By removing the diffuser component, we preserve the directionality of the plenoptic function on the entrance plane and can sample it in the sub-images. A difficulty is the divergent nature of the image cast by the main lens onto the entrance plane, see Fig. 4 (a). For clarity, a pinhole-like aperture of the pickup system is being shown. We show two scene points that are imaged from three directions each. One sees that severe vignetting occurs; the blue view is completely blocked for the upper scene point whereas the green view is missing for the lower scene point. The problem is caused by a mismatch between the exit pupil of the main lens and the entrance pupil of the pickup system. The pickup system images regions of the main lens aperture far outside its physical limits (Fig. 5, top).

We circumvent the problem by introducing a *pupil-matching lens* that images the aperture plane of the main lens onto the aperture plane of the pickup system. The mirror operation introduced by the image multiplier generates virtual viewpoints through the mirrored pickup apertures. These, in turn, are imaged onto stable regions of the main lens aperture. As shown in Fig. 4 (b), this matches the exit pupil of the main lens to the entrance pupil of the pickup system. Now the condition for non-vignetted imaging only depends on the choice of the maximum observation angle α . It should be designed to match the maximal angle of the exit pupil of the main lens. A visual example is shown in Fig. 5 (bottom).

Please notice that this pupil-matching lens was already mentioned for the filter-based design since it allows us to select the optimal diffuser scattering profile (Sec. 4.1).

Pickup-System Aperture An additional modification is to equip the pickup system with an aperture. This way, the global *depth-of-field* of the light-field views can be regulated at the expense of light efficiency. This feature allows for an increased photographic control of light-field imagery that was previously unavailable. As an example, in a lenslet-based light-field camera [Ng et al. 2005] the depth-of-field of the sub-views is fixed since the micro-lenses do not feature controllable apertures. Fig. 6 (bottom) shows the improvements in depth-of-field that can be obtained by stopping down the pickup aperture.

Imaging Geometry The resulting imaging geometry is similar to the focused plenoptic camera of Georgiev and Lumsdaine [2011]. We obtain a set of virtual cameras in the aperture plane of the main lens. These virtual cameras are arranged in a symmetric configu-




Figure 7: Photograph of our prototype setup. The inset shows a closeup on the filter plane assembly. Note that Fig. 2 shows a schematic design highlighting functional properties whereas our practical implementation uses lens assemblies to reduce the aberrations of the complete system.

ration as prescribed by the image-multiplier geometry which is the generator of the set of view points. The virtual cameras have view points that converge in the focal plane of the main lens. In this plane, there is zero parallax between the views. Out-of-plane objects are imaged with positive or negative parallax depending on whether they are located in the front or in the back of the focal plane, see also Fig. 2 (b).

5 Prototype & Preprocessing

In this section, we discuss the construction of our practical prototype, Fig. 7, design decisions, and alternatives to the setup proposed in Secs. 4.1 and 4.2.

5.1 Construction details

We built an optical table setup of the proposed design in order to evaluate its effectiveness. The camera in our system is a Canon EOS 5D mark II with 22 Mpx resolution, featuring a full-frame sensor. We designed our system to generate 3×3 copies of the scene for both our filter-based design, and for the light-field imaging case. This choice leads to 9 different exposure settings for HDR imaging, Sec. 6.1, 27 different spectral measurements for the multispectral imaging case, Sec. 6.1, since the 9 filters are imaged by three color channels each, 9 channels for polarization measurements, Sec. 6.1, and 3×3 different views for the light-field imaging case, Sec. 6.2.

Due to experimental constraints and the limited availability of suitable stock parts, our practical implementation differs slightly from the ideal case discussed in Sec. 4.

The diffuser Our diffuser has a thickness of $\approx 1 \text{ mm}$ and polarization preserving properties since it was designed for polarization-based 3D rear-projection screens (ScreenTech GmbH, material type “ST-Professional-DCF”). The diffuser scattering profile falls to 50% transmittance at about 20° off-axis, which is well above the maximum observation angle of our system (12.95° for 15 mm pickup lens aperture). The requirements for the elimination of the directional light variation are thus met. For a more detailed discussion, please refer to the supplemental material.

The image multiplier To create the copies of the imaged scene, we employed a rectangular kaleidoscope, $36 \text{ mm} \times 24 \text{ mm} \times 300 \text{ mm}$ in size. It was made from optical front-surface mirrors and constructed by an artisan kaleidoscope maker (Karleidoskope GmbH). The quality of the mirror system that can be achieved by manual construction is not sufficient for the imaging purposes proposed in this article. Since an individual pixel covers about $\approx 18 \mu\text{m}$ of diffuser surface, a perfectly parallel arrangement of the mirrors is necessary. Due to inadequate alignment, our kalei-

doscope suffers from some imaging imperfections that most prominently show in the corner views of the kaleidoscope. In a practical implementation, a rectangular prism utilizing total internal reflection can be used as an accurate image multiplier.

Pupil matching in the filter plane While the ideal design, Sec. 4 features two plano-convex lenses with the filter array being placed in the aperture of the resulting effective bi-convex lens, in practice this arrangement is more easily implemented by a single bi-convex lens at a small distance to the filter array. Shifting the filter array out of the aperture has the additional benefit of masking imperfections in the optical filters themselves: manual construction of the filter array results in unavoidable scratches and other minor deviations from a perfectly planar optical filter of zero width. If the filter array was placed directly into the aperture of the filter-plane pupil matching arrangement, these imperfections would readily become apparent in the recorded images, while they are now blurred and less noticeable.

The 1:1 imaging system We implement the 1:1 imaging by using a Canon 100 mm, $f/2.8$ macro lens. This results in a distance of about 300 mm between the lens and the filter plane. In addition, the distance between the pickup system and the filter plane has to be adjusted to this length to ensure 1:1 imaging, preserving the overall width and height of the imaging system to match that of a full-frame sensor, i.e. $36 \text{ mm} \times 24 \text{ mm}$. Overall, this leads to an overall system length of about 1000 mm including the camera and the main lens for our prototype system.

Light Field Design For practical light-field imaging, we reuse the filter-based optical system, i.e., we do not remove the filter-plane and 1:1 imaging optical systems, we only remove the diffuser and the filters in the filter plane. The pupil matching lens was chosen as a 2" diameter, $f = 70 \text{ mm}$ plano-convex lens.

5.2 Pre-Processing

The pre-processing procedure consists in registering the 3×3 sub-images that are recorded by the sensor with one another. Since the images we are interested in are located in the entrance plane and are coincident, a single geometric calibration procedure suffices for all applications presented below. Moreover, this calibration is *scene-independent* since only two-dimensional optical distortions in the entrance plane need to be compensated for. The sub-images suffer from misregistration on the sensor primarily for two reasons: an imperfect arrangement of the mirror planes due to manual construction and geometric/chromatic aberrations induced by our optical setup.

We address these imperfections in two steps. While keeping the diffuser in place and removing the main lens, we use a transparency slide with a checkerboard pattern that we place in close distance to the diffuser. We then illuminate the slide with a far-away point light source, hereby projecting the pattern onto the diffuser. Taking an image of this pattern indicates the perspective distortions introduced by misalignments of the mirrors. We note that the corner images of our 3×3 matrix of views encounter two levels of reflection. These images show a noticeable disagreement along their diagonals. Thus, we compensate each half of these images separately. This first compensation is performed by estimating homographies between the outer and the central views and aligning all images to the central view.

Residual registration imperfections are caused by geometrical and chromatic aberrations. We measure these using a wavelet noise pattern [Cook and DeRose 2005] as a target that provides dense features. The distortions are then estimated via optical flow [Horn



Figure 8: Tone-mapped results of three example scenes captured with our high dynamic range setup. The maximum dynamic range is about 18 f-stops.

and Schunck 1981] with a strong smoothness prior. To register the outer views to the central image, we use the composite of all displacements (homography-based warping and residual optical flow compensation). In the filter-based system, this process matches the images. In the light-field imaging case, the views are brought into agreement, as expected by the geometrical construction; horizontally-neighboring views show only horizontal and vertically-neighboring views only vertical parallax, diagonally-displaced views show combinations thereof.

Our prototypical system also shows radiometric distortions, i.e. vignetting is observable throughout the geometrically-registered images. To measure the effect, we re-introduce the main lens into the system and add a strong diffuser, which is illuminated from a far-away diffuse light source. The resulting image is used to divide out vignetting effects caused by the optical system.

As a result of these pre-processing steps, we end up with a stack of images I_i that are optically pre-filtered as if taken in a time-sequential manner:

$$s^j = \int_{\omega, \lambda, t} l_\lambda^j(x, y, \omega, \lambda, t) f_i(\omega, \lambda) d\omega d\lambda dt, \quad j = 0 \dots 3,$$

$$I_i = \left\{ [1 \ 0 \ 0 \ 0] \cdot \mathbf{M}_i \cdot [s^0 \ s^1 \ s^2 \ s^3]^T \right\}_0^1. \quad (2)$$

Since we are dealing with polarized light in Sec. 6.1, the formulation includes polarization parameters. We formulate the plenoptic function l_λ as consisting of four parts; the four Stokes parameters s^j with the following definitions: $l_\lambda^1 = E_X E_X^* + E_Y E_Y^*$, $l_\lambda^2 = E_X E_X^* - E_Y E_Y^*$, $l_\lambda^3 = E_X E_Y^* + E_Y E_X^*$, $l_\lambda^4 = i(E_X E_Y^* - E_Y E_X^*)$, where E_X and E_Y are the two orthogonal plane wave components of the electric field E and * signifies complex conjugation. The optical filters are denoted by sets $\{\mathbf{M}_i, f_i\}$, consisting of a standard optical filter f_i and a Mueller matrix \mathbf{M}_i . For the plenoptic dimensions, we denote wavelength as λ , directions as ω , and time as t . Multiplication by the $[1 \ 0 \ 0 \ 0]$ vector extracts the irradiance measurement that is being registered by the sensor. The clamping operation $\{\cdot\}_0^1$ models the saturation limit imposed by a real sensor. We note that not all of the filter dimensions (wavelength, polarization, and direction) are used simultaneously in the following. Rather, each of the application areas that we describe uses one dimension at a time.

6 Applications

6.1 Filter-based Applications

High Dynamic Range Imaging

For HDR imaging, we arrange our filter array to consist of 3×3 neutral density filters and the optical filters in Eq. 2 become $\{\mathbf{1}, c_i\}$, $i = 1 \dots 9$ with a unit Mueller matrix and constant spectral filters $f_i(\lambda) = c_i$. We choose a set with transmittance values of $\{1.0, 0.5, 0.25, 0.126, 0.063, 0.032, 0.016, 0.008, 0.004\}$,




Figure 9: Ground truth evaluation of our multispectral imaging pipeline. The RMS error is 1%, the maximum error is about 20% (in the far blue regions that have a poor spectral sensitivity). The other patches show similar performance.

yielding a dynamic range improvement of about 8 f-stops over that of the sensor dynamic range. These images have a verified linear response and can be merged by standard mechanisms [Debevec and Malik 1997]. For video operation, the camera applies an adaptive response curve out of our control. We therefore estimate the radiometric response by a variant of Mitsunaga and Nayar's [1999] polynomial technique that estimates the response from a series of photographs of a MacBeth color checker while enforcing curvature constraints on the final curve [Ihrke 2012]. Examples captured with the resulting HDR imaging system are shown in Fig. 8.

Multispectral Imaging

For multispectral imaging, we equip our system with 3×3 broadband spectral filters as manufactured by Rosco Labs (Roscolux swatchbook). The filters in Eq. 2 become $\{\mathbf{1}, f_i(\lambda)\}$, $i = 1 \dots 9$. We chose to use broadband spectral filters because our imaging system features a color filter array. Imaging 9 spectral filters through three different Bayer filters each results in an overall measurement of 27 broadband filtered images. Using narrow band filters would only yield 9 different measurements since the Bayer filters are largely orthogonal. The filters in our array are {Cyan #4360, Yellow #4590, Red #26, Orange #23, Green #89, Blue-Green #93, Lavender #4960, Blue #80, Magenta #4760}. We measured their spectral response using a spectrometer (Thorlabs CCS 200).

For spectral calibration of the Bayer filters, we illuminated a scene containing a MacBeth color checker with a high pressure mercury vapor lamp with a previously acquired spectrum s_{mv} . In the multispectral imaging case, Eq. 2 can be simplified to

$$I_i^{r|g|b}(x, y) = \int_{\lambda} l_\lambda(x, y, \lambda) f_i^{r|g|b}(\lambda) f_i(\lambda) s(\lambda) d\lambda, \quad (3)$$

where $f_i^{r|g|b}(\lambda)$ denotes the spectral sensitivity of the camera for the R, G, and B channels, $f_i(\lambda)$ are the known spectra of the Roscolux filters, and s is the spectrum of the light source. In this case, the plenoptic function $l_\lambda(x, y, \lambda)$ only depends on the spectral scene reflectance whose spectrum $l_\lambda(x, y, \lambda)$ is known through collected measurements obtained from http://www.babelcolor.com/main_level/ColorChecker.htm. The spectrum of the light source is s_{mv} . Therefore, all components of the integral in Eq. 3 except for the Bayer filter responses $f_i^{r|g|b}(\lambda)$ are known and can be estimated by an expansion into basis functions similar to Toyooka and Hayasaka [1997]. We chose a set of 50 overlapping Gaussians distributed in the range between 400 and 700 nm as our basis. The optimization problem uses images through all 116 Roscolux

filters and enforces a non-negativity constraint via quadratic programming.




Figure 10: Multispectral-imaging applications. Spectral relighting with ideal Planckian black body light sources of different temperature, and relighting with natural day light and a high-pressure sodium street light (top row). The street light leads to the familiar green loss. The scene is metameric for low color temperatures, but under daylight illumination the plastic and the two real flowers are clearly distinguishable. Simulation of color-deficient vision for a color checker and a natural scene (bottom rows). Protanomalous vision shifts the response of the red sensitive cones towards green by $\approx 10\text{nm}$, whereas deutanomalous vision shifts the green cone response by $\approx 10\text{nm}$ towards the red.

Once the spectral response of the sensor is known, arbitrary scenes can be imaged. If the spectrum of the light source $s(\lambda)$ is known, a neutral reflectance spectrum can be recovered, otherwise, only the product $l_\lambda(x, y, \lambda)s(\lambda)$ is accessible. The scene spectra are recovered similar to spectral calibration of the sensor, except that now the spectral sensitivities $f^{r|g|b}(\lambda)$ of the sensor are known whereas the scene spectrum $l_\lambda(x, y, \lambda)$ or its product with the illumination spectrum $l_\lambda(x, y, \lambda)s(\lambda)$ are estimated. In this case, spectral white balancing, similar to RGB white balancing can be performed by dividing all spectra by the spectrum of a known white scene patch. In contrast to the spectral calibration step, image spectra have to be estimated for every pixel and quadratic programming becomes too costly. Instead, we drop the non-negativity constraint, solve a least squares problem per-pixel and clamp negative values to zero. For improved regularization, we use a PCA basis as in [Toyooka and Hayasaka 1997]

We verified the performance of our multispectral imaging pipeline by imaging a Gretag Macbeth color checker under known illumination. The reconstructed spectral reflectance agrees well with collected data (babelcolor), see Fig. 9. An application of our system configured for multispectral imaging is shown in Fig. 10.

Polarization Imaging

We can also equip our filter-based system with polarization filters. Hereby, the camera is made sensitive to the polarization state of

light and acts as a pixel-by-pixel polarization state analyzer. To this end at least three independent measurements have to be carried out and four if the full polarization state that also includes the circular polarization component is to be retrieved [Goldstein 2003]. In the case of polarization imaging, it is essential to use a polarization-preserving diffuser. This type of diffuser is commonly used in screens for polarization-based stereo projection. Regular diffusers, on the other hand, act as depolarizers and are unsuitable.

In this article, we restrict our scope to linear polarization imaging, since, apart from some special cases of, e.g., circular dichroism and circularly-polarized luminescence, significant levels of circular polarization are rarely encountered in nature [Hegedus et al. 2006]. For this purpose, we placed five linear sheet polarizers with differently orientated transmission axes $\approx \{0^\circ, 36^\circ, 72^\circ, 108^\circ, 144^\circ\}$ into the filter array of our system. In terms of Eq. 2, our filters become $\{\mathbf{M}_i, 1\}$, $i = 0 \dots 4$, where 1 denotes an optical filter $f_i = 1$. The four corners of the array were left empty and the corresponding sub-images were ignored. The setup still provides more measurements per pixel than needed and we avoid those images produced by second-order reflections, which are more prone to optical aberrations and complex polarization modulation.

When only linear polarization is measured, we can retrieve the first three Stokes components $s^j, j = 0 \dots 2$, and the fourth circular component s^3 , if any, is considered as part of the unpolarized component s^0 of the light. Correspondingly, we employ 3×3 Mueller matrices, which is a common procedure in linear polarimetry [Neumann et al. 2008]. To determine the Stokes vectors, we first construct the 3×5 matrix \mathbf{W} whose consecutive rows are identical to the upper row of the respective Mueller matrices $\mathbf{M}_i, i = 0 \dots 4$. For each pixel, we store the measured intensities through the five polarization filters in a vector \mathbf{p} , the Stokes vector $\mathbf{s} = (s^0, s^1, s^2)$ is obtained by least-squares regression

$$\mathbf{s} = (\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T \mathbf{p}. \quad (4)$$

Some additional care is needed because the filter array is placed inside our optical system, whose reflections and scattering affect the polarization state of light. The total influence of the system including that of the polarization filters can be characterized by an effective Mueller matrix $\mathbf{M}_{sys}(x, y)$, which is spatially dependent. The most prominent effect is caused by the mirrors of the image multiplier. We determine this pixel-wise Mueller matrix by a calibration procedure that uses a ground truth polarimeter to obtain the Stokes vectors of 6 scenes with homogenous (i.e. spatially non-dependent) polarization states and relating these values to the ones observed by our system. The linear relation $\mathbf{s}_{sys}^{(i)}(x, y) = \mathbf{M}_{sys}(x, y)\mathbf{s}_{gt}^{(i)}$, $i = 1..6$ is then solved in a least-squares sense for \mathbf{M}_{sys} . Here, $\mathbf{s}_{sys}^{(i)}$ are the Stokes parameters measured by our system, whereas $\mathbf{s}_{gt}^{(i)}$ are the Stokes parameters measured by the ground truth polarimeter. In practice, we use 30 different polarizer/analyzer pair images to perform the polarization calibration. For more details, please refer to the supplementary document. Applications of our snapshot polarimeter configuration are shown in Fig. 11.

6.2 Light-Field Imaging

The low angular resolution of our 3×3 light-fields necessitates an angular up-sampling scheme in order to perform convincing refocussing and view-point changes at a reasonable distance outside the focal plane of our system. In practice, the observed parallax in our system can exceed 100 pixels. The spatial resolution of the images returned by our system is, however, large with a resolution of approx. 1800×1200 pixels for each sub-view. We address the angular interpolation problem by first performing a depth estimate, and a parallax-based morphing operation. This morphing

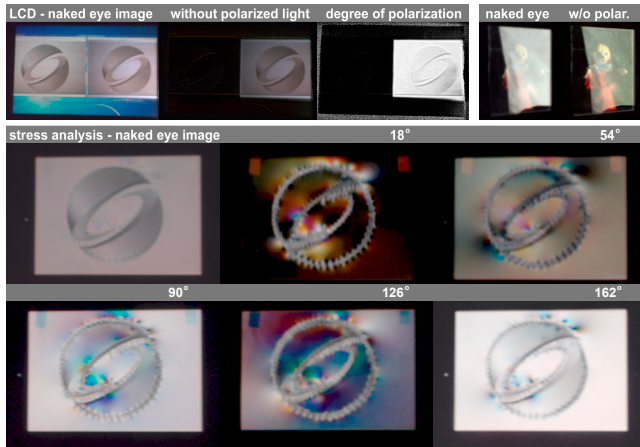


Figure 11: Polarization imaging (top row): the logo on the left is shown on the screen, whereas the logo on the right is printed on paper. The paper acts as a diffuser and destroys the LCD’s light polarization. We show the image perceived by the eye (left), a linear-polarizer simulation that removes the LCD’s light (middle) and the degree of polarization (right, white indicates an unpolarized state). The doll example illustrates reflection removal by a virtual polarization filter. Polarization stress analysis (bottom rows): a transparent acrylic glass window is attached to an LCD screen acting as a source for polarized light. The window contains drilled holes in form of the logo’s outline that induce stress in the material. The images show different simulations from our system output, making stress patterns appear in the form of color fringes. The upper left image shows the logo displayed on the LCD screen as it would be perceived by an unaided human eye.

also makes view extrapolation possible, which enables an aperture synthesis beyond the limits of the main lens.

Depth Estimation

The recent literature on depth estimation from light-fields [Wanner and Goldluecke 2012a; Wanner and Goldluecke 2012b] is not applicable to our data. In [Wanner and Goldluecke 2012a], the authors estimate depth based on epipolar plane images. An angular resolution of 3×3 is too low to use this approach. A different version, based on depth maps [Wanner and Goldluecke 2012b] requires at least 5×5 views because derivatives of the directional light-field components have to be estimated. Instead, we adapt an optical-flow technique for our purpose [Horn and Schunck 1981].

Our modifications consist in introducing a coupling between the flow variables of different views. It is well known that optical flow suffers from the so-called aperture problem, i.e. two variables are sought at every image location, but only a single constraint is available. In the case of light-field imaging, the flow is known to be constrained to the directions of the epipolar lines between views. Moreover, the structure of these epipolar lines is very regular due to the fixed spacing between the virtual views. The optical-flow vectors are therefore replaced by depth estimates $d(x, y)$ that couple the flow estimates in all surrounding light-field views via the depth-induced parallax $d(x, y) \cdot [u_i, v_i]$, where the vector $[u_i, v_i]$ is a constant for every view I_i and describes the slope of the epipolar lines. Due to the constraints of our setup, we can safely assume the epipolar lines to be parallel in every sub-view.

This depth estimation is based on the optical flow brightness constancy assumption and, therefore, does not estimate the real scene

depth. However, it computes an estimate of *apparent depth*. Since we are not interested in depth per se, but in its view interpolating properties, this approach is reasonable for angular light-field up-sampling.

View Interpolation and Extrapolation

We estimate a depth map for each of the sub-views, which allows us to generate a new view by morphing the sub-views I_i according to the parallax displacement $d \cdot [u_i, v_i]^T$. The main challenges for a high-quality interpolation are a proper handling of the occlusion boundaries, the handling of multiple pixels of an input view mapping to the same destination pixel, and the avoidance of interpolation holes by forward warping. Our morphing uses forward and backward warping steps followed by a blending procedure.

Each of the nine sub-views may contain exclusive information not available in any other sub-view but valuable for the interpolated view. However, warping all views can lead to blur because the depth estimation is only approximate. Using the four neighboring views of the *interpolated* position on the main lens’ aperture is a good tradeoff. A similar scheme can be used for *extrapolation*; using the two (for extrapolation in u or v) or the one closest view (for extrapolation in u and v).

Applications of our design configured for light-field imaging are shown in Fig. 12. We illustrate virtual refocussing, stereo-pair reconstruction, and depth-based stylization. The figure shows that our extrapolation solution allow us to virtually extend the aperture of the main lens to generate increased parallax and extremely shallow depth-of-field effects.

7 Discussion

7.1 Potential for Miniaturization

As discussed in Secs. 4.1 and 4.2, the size of the overall system is determined by the distance between the sensor and the first imaging element of our optical design. In SLR type cameras, this distance is bounded from below by the moving mirror of these cameras and can be assumed to be around 50 mm for a full-frame sensor. In the filter-based design, this requirement determines the focal length f of the 1:1 imaging system and with it the overall length of the optical system as $(6 + 2N) \cdot f$ if $N \times N$ copies are to be imaged. The focal length f is therefore fixed to $\approx 25 \text{ mm}$. With $N = 3$ (9 sub-images), a length of $\approx 300 \text{ mm}$ is needed. The diameters of the 1:1 imaging lens and the pickup lens determine the maximum pickup aperture and are therefore uncritical.

For the light-field design, we can remove the 1:1 imaging and the filter plane optics. The minimum distance of our system is now determined by the closest position that the pickup lens can assume. Given these considerations, z in Eq. 1 equals 50 mm and the overall system length is $4 \cdot z = 200 \text{ mm}$ for 3×3 copies. A light-field unit could, therefore, be produced at a shorter length.

Overall, the system is suitable for imaging a low number of copies with its size increasing linearly for a larger number. The system size also scales linearly with the sensor size of the camera being employed. Smaller units could thus be designed for smaller sensors. Future cameras may also remove the mirror in SLR cameras since an optical view finder is not strictly necessary for computational cameras. It might therefore be possible to miniaturize the design even further.



Figure 12: Examples of light-field imaging applications. Topmost row: strong refocussing with a virtual aperture matching the real Canon EF 50 mm f/1.4 USM lens, which was used to record the scene. The all-in-focus view is an example of an interpolated light-field view. Second/Third row: Refocussing in a challenging setting (specular, mirroring and refracting objects). The second row shows refocussing within the physical limits of the real aperture, the third uses extrapolation to virtually extend the aperture by one full size to each side, leading to a virtual f/0.7. This mode enables macro-photography effects for standard lenses and wide fields of view. Please, compare the quality of the in-focus regions; no significant loss in sharpness is observed for the extended aperture case, validating the suitability of our depth-estimation and view-morphing procedures. Fourth row: Virtual stereo images (in form of anaglyphs) with an adaptable base line of three light-field scenes. Fifth row: Depth-based stylization examples: from left to right, painterly/real style, bloom+contrast+coloring/real, and sharpening/blur.

7.2 Filter-Based Design

Our implementation of the filter-based design suffers from a number of problems whose combination leads to a loss in optical resolution. First, the manual construction of the table setup and the use of stock parts limit the quality of the optical system, resulting in an inherent optical point spread function (PSF) of the system. In addition, the diffuser is quite thick and shows a substantial amount of multiple scattering, resulting in glare around bright features. The spatial PSF of the diffuser convolves with the optical PSF of the remainder of the system. The resulting PSF varies significantly over the image plane, resulting in a non-uniform distortion of the individual sub-images, which creates a slight offset between the registered copies. Since all applications depend on a computational recombination of the registered images, this creates an additional *numerical* PSF that again convolves with the one of the combined diffuser/optical system. This three-fold combination results in a considerable loss of resolution. However, an improvement in any of the components will improve the complete system and we believe that considerable headroom is available.

The nominal output resolution of our system is $\approx 1800 \times 1200$ pixels, however, due to the mentioned difficulties this number reduces to about 600×400 pixels at an acceptable quality. Even though this number appears low, in comparison to other snapshot multispectral [Habel et al. 2012; Du et al. 2009] (typical resolution below 200×200 pixels) or polarization imaging techniques [Pezzaniti et al. 2008] (resolution 782×582), these results are competitive with a large margin for improvements.

Compared to aperture-splitting approaches [Aggarwal and Ahuja 2001; Bando et al. 2008; Horstmeyer et al. 2009; Georgiev et al. 2011], our design offers two main advantages: first, the filtered sub-images have an identical view point, avoiding the necessity of scene-dependent registration. Second, the aperture of the main lens maintains its full functionality whereas in aperture-splitting approaches the sub-views become vignetted of the main lens aperture is stopped down. This way, the choice between the light efficiency and the depth-of-field of the system remains under control of the photographer without compromising the acquired information.

7.3 Light-Field Design

Our light-field prototype has a sufficient quality to render refocused images at the full resolution of 1800×1200 pixels. Since the depth map allows for an extrapolation beyond the physical limits of the main lens' aperture, strong out-of-focus and shallow depth-of-field effects can be simulated convincingly. Further, other applications like distance-based stylization become possible by using the derived depth (Fig. 12).

Comparing to the established lenslet-based design [Ng et al. 2005], our proposed solution explores an orthogonal end of the design space. While lenslet-based light-field cameras offer a high angular resolution, their spatial resolution is limited and has to be interpolated to enable an acceptable image size. Contrary, our design is most suitable for imaging at a high spatial resolution with a low number of angular views. In our case, the angular resolution has to be interpolated to enable convincing refocussing effects. These trade-offs are an inherent feature of the two designs. *Decreasing* the angular resolution to gain spatial resolution is difficult for the lenslet-based design since the size of the microlenses represents a limiting factor. In addition, color imaging with color-filter arrays would become unfeasible when approaching low angular resolutions because parallax effects would become noticeable. In contrast, *increasing* the number of angular views in our design makes the device size less feasible. A unique advantage of our design is

that the depth of field of the light-field views can be regulated by the aperture of the pickup lens. Light-field aliasing and light efficiency of the system can therefore be controlled by the user unlike in existing light-field camera designs. The maximum achievable parallax is in both cases limited by the aperture size of the main lens and this performance parameter is the same for both designs.

8 Conclusions and Future Work

We presented an optical design suitable for implementation as an in-between optical element that enables plenoptic image acquisition with standard cameras. We have implemented the design in a prototypical system and illustrated high dynamic range, multispectral, polarization, and light-field imaging applications. Our light-field design offers the unique advantages of high spatial resolution imaging with control over the depth of field of the light-field sub-views. We achieve an unprecedented spatial resolution in refocussing and virtual view point adjustment applications. The design potentially lends itself to miniaturization to practical scales.

In the future, we will concentrate on designing a small-scale implementation. The numerical part of our system PSF is determined by the algorithmic components of the computational-optical device. Improvements in depth-estimation and view-interpolation schemes will reduce its size. In particular, we plan to investigate adaptations of more modern optical-flow variants to light-field depth estimation. For the filter-based design, we plan a principled investigation of different diffuser types to bring the quality of the HDR, multispectral, and polarization images closer to that of the light-field case.

Acknowledgements

This work was supported by the German Research Foundation (DFG) through an Emmy-Noether fellowship IH 114/1-1 and by the Intel Visual Computing Institute at Saarland University.

References

- ADELSON, E. H., AND BERGEN, J. R. 1991. The Plenoptic Function and the Elements of Early Vision. In *Computational Models of Visual Processing*, MIT Press, 3–20.
- ADELSON, E., AND WANG, J. 1992. Single Lens Stereo with a Plenoptic Camera. *IEEE Trans. PAMI* 14, 2, 99–106.
- AGGARWAL, M., AND AHUJA, N. 2001. Split aperture imaging for high dynamic range. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, vol. 2, 10–17 vol.2.
- BANDO, Y., CHEN, B.-Y., AND NISHITA, T. 2008. Extracting Depth and Matte using a Color-Filtered Aperture. *ACM TOG* 27, 5, 134.
- BAYER, B. E., 1976. Color Imaging Array. US Patent 3,971,065.
- BONNET, H., ABUTER, R., BAKER, A., ET AL. 2004. First Light of SINFONI at the VLT. *ESO Messenger* 117, 17, 17–24.
- CAO, X., TONG, X., DAI, Q., AND LIN, S. 2011. High-Resolution Multi-Spectral Video Capture with a Hybrid Camera System. In *Proc. CVPR*, 297 – 304.
- COOK, R. L., AND DEROSE, T. 2005. Wavelet Noise. *ACM TOG* 24, 3, 735–744.

- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. SIGGRAPH*, 369–378.
- DESCOUR, M., AND DERENIAK, E. 1995. Computed-tomography Imaging Spectrometer: Experimental Calibration and Reconstruction Results. *Appl. Optics* 34, 22, 4817–4826.
- DU, H., TONG, X., CAO, X., AND LIN, S. 2009. A Prism-Based System for Multispectral Video Acquisition. In *Proc. ICCV*, 175–182.
- GEHM, M. E., JOHN, R., BRADY, D. J., WILLET, R. M., AND SCHULZ, T. J. 2007. Single-Shot Compressive Spectral Imaging with a Dual-Disperser Architecture. *Optics Exp.* 15, 21, 14013–14027.
- GEORGIEV, T., LUMSDAINE, A., AND CHUNEV, G. 2011. Using Focused Plenoptic Cameras for Rich Image Capture. *IEEE CG&A* 31, 1, 62–73.
- GOLDSTEIN, D. H. 2003. *Polarized Light, 2nd edition*. CRC Press, New York, US.
- GORTLER, S., GRZESZCZUK, R., SZELINSKI, R., AND COHEN, M. 1996. The Lumigraph. In *Proc. SIGGRAPH*, 43–54.
- HABEL, R., KUDENOV, M., AND WIMMER, M. 2012. Practical Spectral Photography. *CGF* 31, 2 (May), 449–458.
- HAN, J. Y., AND PERLIN, K. 2003. Measuring Bidirectional Texture Reflectance with a Kaleidoscope. In *Proc. SIGGRAPH*, 741–748.
- HEGEDUS, R., SZEL, G., AND HORVATH, G. 2006. Imaging polarimetry of the circularly polarizing cuticle of scarab beetles (Coleoptera: Rutelidae, Cetoniidae). *Vision Research* 46, 2786–2797.
- HORN, B. K. P., AND SCHUNCK, B. G. 1981. Determining Optical Flow. *Artif. Intell.* 17, (1-3), 185–203.
- HORSTMEYER, R., EULISS, G., ATHALE, R., AND LEVOY, M. 2009. Flexible Multimodal Camera Using a Light Field Architecture. In *Proc. ICCP*, 1–8.
- IHRKE, I., WETZSTEIN, G., AND HEIDRICH, W. 2010. A Theory of Plenoptic Multiplexing. In *Proc. CVPR*, 1–8.
- IHRKE, I., RESHETOUSKI, I., MANAKOV, A., TEVS, A., WAND, M., AND SEIDEL, H.-P. 2012. A Kaleidoscopic Approach to Surround Geometry and Reflectance Acquisition. In *Proceedings of IEEE International Workshop on Computational Cameras and Displays*, 1–8.
- IHRKE, I., 2012. Color Calibration Toolbox for MATLAB v2.0. <http://giana.mmci.uni-saarland.de/software.html>.
- ISAKSEN, A., MCMILLAN, L., AND GORTLER, S. J. 2000. Dynamically Reparameterized Light Fields. In *Proc. SIGGRAPH*, 297–306.
- IVES, H., 1903. Parallax Stereogram and Process of Making Same. US patent 725,567.
- KUDENOV, M., AND DERENIAK, E. 2011. Compact Snapshot Real-Time Imaging Spectrometer. In *SPIE Conf. on Elec.-Opt. Rem. Sens., Phot. Tech., and Appl. V*, 81860W–1–81860W–12.
- LANMAN, D., RASKAR, R., AGRAWAL, A., AND TAUBIN, G. 2008. Shield Fields: Modeling and Capturing 3D Occluders. *ACM TOG* 27, 5, 131.
- LEVOY, M., AND HANRAHAN, P. 1996. Light Field Rendering. In *Proc. SIGGRAPH*, 31–42.
- LEVOY, M., CHEN, B., VAISH, V., HOROWITZ, M., MCDOWALL, I., AND BOLAS, M. 2004. Synthetic Aperture Confocal Imaging. *ACM TOG* 23, 3, 825–834.
- LIPPMANN, G. 1908. La Photographie Intégrale. *Academie des Sciences* 146, 446–451.
- MANN, S., AND PICARD, R. W. 1995. Being 'Undigital' with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures. In *Proc. IS&T*, 442–448.
- MCGUIRE, M., MATUSIK, W., PFISTER, H., CHEN, B., HUGHES, J. F., AND NAYAR, S. K. 2007. Optical Splitting Trees for High-Precision Monocular Imaging. *IEEE CG&A* 27, 2, 32–42.
- MITSUNAGA, T., AND NAYAR, S. K. 1999. Radiometric Self Calibration. In *Proc. CVPR*, 374–380.
- NARASIMHAN, S., AND NAYAR, S. 2005. Enhancing Resolution along Multiple Imaging Dimensions using Assorted Pixels. *IEEE Trans. PAMI* 27, 4, 518–530.
- NAYAR, S., AND MITSUNAGA, T. 2000. High Dynamic Range Imaging: Spatially Varying Pixel Exposures. In *Proc. CVPR*, vol. 1, 472–479.
- NEUMANN, L., HEGEDUS, R., HORVATH, G., AND GARCIA, R. 2008. Applications of High Precision Imaging Polarimetry. In *Proc. Computational Aesthetics in Graphics, Visualization and Imaging*, 89–97.
- NG, R., LEVOY, M., BRÉDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light Field Photography with a Hand-Held Plenoptic Camera. Tech. Rep. Computer Science CSTR 2005-02, Stanford University.
- NG, R. 2005. Fourier Slice Photography. *ACM TOG* 24, 3, 735–744.
- OKAMOTO, T., AND YAMAGUCHI, I. 1991. Simultaneous Acquisition of Spectral Image Information. *Optics Lett.* 16, 16, 1277–1279.
- PARK, J.-I., LEE, M.-H., GROSSBERG, M. D., AND NAYAR, S. K. 2007. Multispectral Imaging Using Multiplexed Illumination. In *Proc. ICCV*, 1–8.
- PEZZANITI, J. L., CHENAULT, D., ROCHE, M., REINHARDT, J., PEZZANITI, J. P., AND SCHULTZ, H. 2008. Four Camera Complete Stokes Imaging Polarimeter. In *Proc. SPIE 6972, Polarization: Measurement, Analysis, and Remote Sensing VIII*, 69720J–1–69720J–12.
- REINHARD, E., WARD, G., DEBEVEC, P., PATTANAIK, S., HEIDRICH, W., AND MYSZKOWSKI, K. 2010. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann Publishers.
- RESHETOUSKI, I., MANAKOV, A., SEIDEL, H.-P., AND IHRKE, I. 2011. Three-Dimensional Kaleidoscopic Imaging. In *Proc. CVPR*, 353–360.
- RUMP, M., AND KLEIN, R. 2010. Spectralization: Reconstructing spectra from sparse data. In *Proc. EGSR*, 1347–1354.
- SCHECHNER, Y., AND NAYAR, S. 2005. Generalized Mosaicing: Polarization Panorama. *IEEE Trans. PAMI* 27, 4, 631–636.

- SCHECHNER, Y., NARASIMHAN, S. G., AND NAYAR, S. K. 2001. Instant Dehazing of Images using Polarization. In *Proc. CVPR*, 325–332.
- SPIERING, B. A., 1999. Multispectral Imaging System. US Patent 5,900,942.
- TOCCI, M. D., KISER, C., TOCCI, N., AND SEN, P. 2011. A Versatile HDR Video Production System. *ACM TOG* 30, 4.
- TOYOOKA, S., AND HAYASAKA, N. 1997. Two-Dimensional Spectral Analysis using Broad-Band Filters. *Optical Communications* 137 (Apr), 22–26.
- VEERARAGHAVAN, A., RASKAR, R., AGRAWAL, A., MOHAN, A., AND TUMBLIN, J. 2007. Dappled Photography: Mask Enhanced Cameras For Heterodyned Light Fields and Coded Aperture Refocussing. *ACM TOG* 26, 3, 69.
- WANNER, S., AND GOLDLUECKE, B. 2012. Globally Consistent Depth Labeling of 4D Lightfields. In *Proc. CVPR*, 41 – 48.
- WANNER, S., AND GOLDLUECKE, B. 2012. Spatial and Angular Variational Super-Resolution of 4D Light Fields. In *Proc. ECCV*, 608 – 621.
- WETZSTEIN, G., IHRKE, I., LANMAN, D., AND HEIDRICH, W. 2011. Computational Plenoptic Imaging. *CGF* 30, 8, 2397–2426.
- WILBURN, B., JOSHI, N., VAISH, V., ET AL. 2005. High Performance Imaging using Large Camera Arrays. *ACM TOG* 24, 3, 765–776.
- ZHOU, C., AND NAYAR, S. 2011. Computational Cameras: Convergence of Optics and Processing. *IEEE Trans. IP* 20, 12 (Dec), 3322–3340.

