DOI: 10.1111/j.1467-8659.2011.02073.x

COMPUTER GRAPHICS *forum*
Volume 30 (2011), number 8 pp. 2397–2426

Computational Plenoptic Imaging

Gordon Wetzstein[1], Ivo Ihrke[2], Douglas Lanman[3] and Wolfgang Heidrich[1]

[1]University of British Columbia, Canada
gordonw@media.mit.edu and heidrich@cs.ubc.ca
[2]Universität des Saarlandes / MPI Informatik, Germany
ihrke@mpi-inf.mpg.de
[3]MIT Media Lab, Cambridge, MA, USA
dlanman@media.mit.edu

Abstract

The plenoptic function is a ray-based model for light that includes the colour spectrum as well as spatial, temporal and directional variation. Although digital light sensors have greatly evolved in the last years, one fundamental limitation remains: all standard CCD and CMOS sensors integrate over the dimensions of the plenoptic function as they convert photons into electrons; in the process, all visual information is irreversibly lost, except for a two-dimensional, spatially varying subset—the common photograph. In this state-of-the-art report, we review approaches that optically encode the dimensions of the plenoptic function transcending those captured by traditional photography and reconstruct the recorded information computationally.

Keywords: Computational photography, the plenoptic function, image acquisition

ACM CCS: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture— I.4.5 [Image Processing and Computer Vision]: Reconstruction

1. Introduction

Evolution has resulted in the natural development of a variety of highly specialized visual systems among animals. The mantis shrimp retina, for instance, contains 16 different types of photoreceptors [MO99]. The extraordinary anatomy of their eyes not only allows the mantis shrimp to see 12 different colour channels, ranging from ultraviolet to infrared, and distinguish between shades of linear and circular polarization, but it also allows the shrimp to perceive depth using trinocular vision with each eye. Other creatures of the sea, such as cephalopods [MSH09], are also known to use their ability to perceive polarization for communication and unveiling transparency of their prey. Although the compound eyes found in flying insects have a lower spatial resolution compared to mammalian single-lens eyes, their temporal resolving power is far superior to the human visual system. Flys, for instance, have a flicker fusion rate of more than 200 Hz [Ruc61], which is an order of magnitude higher than that of the human visual system.

Traditionally, cameras have been designed to capture what a single human eye can perceive: a two-dimensional trichromatic image. Inspired by the natural diversity of perceptual systems and fuelled by advances of digital camera technology, computational processing and optical fabrication, image processing has begun to transcend limitations of film-based analogue photography. Applications for the computerized acquisition of images with a high spatial, temporal, spectral and directional resolution are manifold; medical imaging, remote sensing, shape reconstruction, surveillance and automated fabrication are only a few examples. In particular, the computer graphics and vision communities benefit from computational plenoptic imaging. Not only do the techniques discussed in this survey allow highly detailed visual information to be captured, which is essential for the acquisition of geometry, scene reflectance, materials and refractive index properties, but they can also be directly used for image-based rendering and lighting, increasing the realism of synthetically generated content.

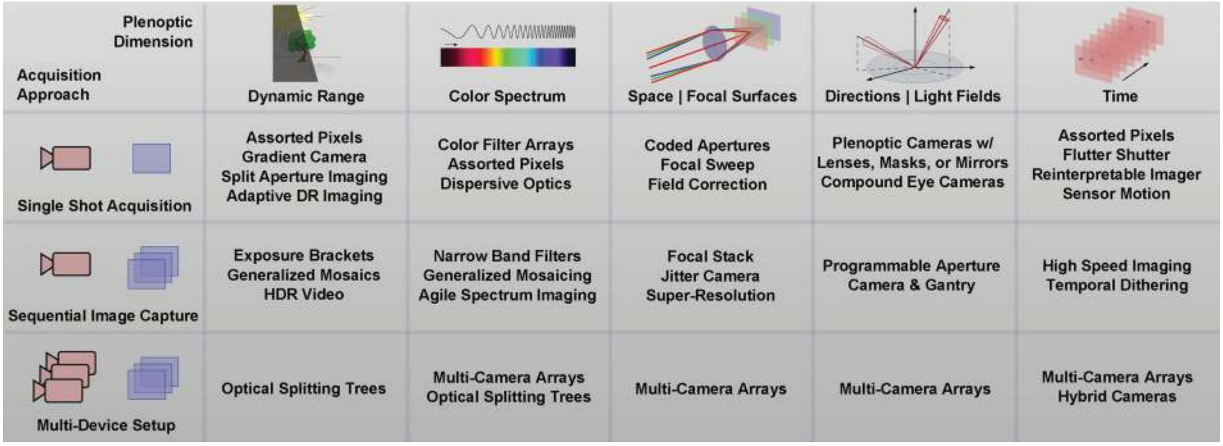

Figure 1: *Taxonomy and overview of plenoptic image acquisition approaches.*

The plenoptic function [AB91] provides a ray-based model of light encompassing properties that are of interest for image acquisition. Most of these properties, however, are irreversibly lost by standard sensors integrating over the plenoptic dimensions during image acquisition. As illustrated in Figure 1, the plenoptic dimensions include the colour spectrum as well as spatial, temporal and directional light variation. We also consider dynamic range a desirable property, as common sensors have a limited dynamic range.

In addition to the plenoptic dimensions, we further categorize plenoptic image acquisition techniques according to their hardware configuration (see Figure 1). While single-device, multi-shot approaches are usually the preferred method for capturing plenoptic properties of static scenes, either multiple devices or single-image multiplexing are required to record dynamic phenomena. Using multiple devices is often the most expensive solution, whereas multiplexing commonly reduces the spatial resolution of captured content in favour of increased plenoptic resolution. The optimal acquisition approach for a given problem is, therefore, dependent on the properties of the photographed scene and the available hardware.

1.1. Computational photography and plenoptic imaging

What makes plenoptic imaging different than general computational photography? Plenoptic imaging considers a subset of computational photography approaches; specifically, those that aim at acquiring the dimensions of the plenoptic function with combined optical light modulation and computational reconstruction. Computational photography has grown tremendously in the last years with dozens of published papers per year in a variety of graphics, vision and optics venues. The dramatic rise in publications in this interdisciplinary field, spanning optics, sensor technology, image processing and illumination, has made it difficult to encompass all research in a single survey.

We provide a structured review of the subset of research that has recently been shown to be closely related in terms of optical encoding and especially in terms of reconstruction algorithms [IWH10]. Additionally, our report serves as a resource for interested parties by providing a categorization of recent research and is intended to aid in the identification of unexplored areas in the field.

1.2. Overview and definition of scope

In this report, we review the state of the art in joint optical light modulation and computational reconstruction approaches for acquiring the dimensions of the plenoptic function. Specifically, we discuss the acquisition of high dynamic range (HDR) imagery (Section 2), the colour spectrum (Section 3), light fields and directional variation (Section 4), spatial super-resolution and focal surfaces (Section 5), as well as high-speed events (Section 6). We also outline the acquisition of light properties that are not directly included in the plenoptic function, but related, such as polarization, phase imaging and time-of-flight (Section 7) and point the reader to more comprehensive literature on these topics. Conclusions and possible future avenues of research are discussed in Section 8.

Due to the fact that modern, digital acquisition approaches are often closely related to their analogue predecessors, we outline these whenever applicable. For each of the plenoptic dimensions we also discuss practical applications of the acquired data. As there is an abundance of work in this field, we focus on imaging techniques that are designed for standard planar 2D sensors. We will only highlight examples of modified sensor hardware for direct capture of plenoptic image information. We do not cover pure image processing

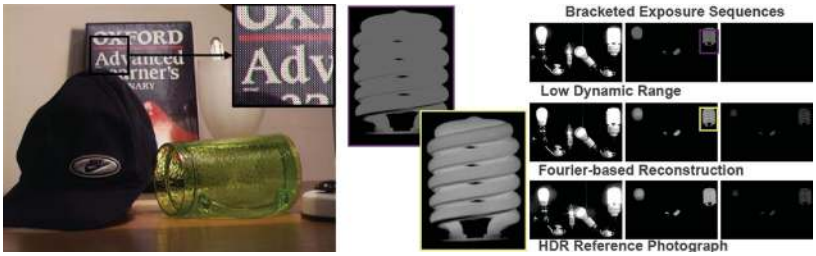

Figure 2: *Sensor image captured with an array of ND filters [NM00] (left). Exposure brackets and magnifications for Fourier-based HDR reconstruction from multiplexed sensor images [WIH10] (right).*

techniques, such as tone-reproduction, dynamic range compression and tone-mapping [RWD*10], or the reconstruction of geometry [IKL*10], BSDFs and reflectance fields.

2. High Dynamic Range Imaging

HDR image acquisition has been a very active area of research for more than a decade. The dynamic range of an imaging system is commonly defined as the ratio of largest and smallest possible value in the range, as opposed to the domain, of a recorded signal. Unfortunately, standard sensors have a limited dynamic range, which often results in clipping of bright and dark parts of a photographed scene. HDR imaging is important for many computer vision applications, including image-based geometry, material and lighting reconstruction. Applications for HDR imagery in computer graphics include physically based rendering and lighting [Deb02], image editing, digital photography and cinema, perceptual difference metrics based on absolute luminance [MDMS05, MKRH11], virtual reality and computer games. With the introduction of the HDR display prototype [SHS*04] and its successor models becoming consumer products today, high-contrast photographic material is required for a mass market. For a comprehensive overview of HDR imaging, including applications, radiometry, perception, data formats, tone reproduction and display, the reader is referred to the textbook by Reinhard *et al*. [RWD*10]. In this section, we provide a detailed and up-to-date list of approaches for the acquisition of HDR imagery.

2.1. Single-shot acquisition

According to DxOMark (www.dxomark.com), the latest high-end digital SLR cameras are equipped with CMOS sensors that have a measured dynamic range of up to 13.5 f-stops, which translates to a contrast of 11 000:1. This is comparable to that of colour negative films [RWD*10]. In the future, we can expect digital sensors to perform equally well as negative film in terms of dynamic range, but this is not the case for most sensors today.

Specialized sensors, which allow HDR content to be captured, have been commercially available for a few years. These include professional movie cameras, such as Grass Valley's Viper [Val10] or Panavision's Genesis [Pan10]. The SpheroCam HDR [Sph10] is able to capture full spherical 360-degree images with 26 f-stops and 50 megapixels in a single scan. A technology that allows per-pixel exposure control on the sensor, thereby enabling adaptive HDR capture, was introduced by Pixim [Pix10]. This level of control is achieved by including an analogue-to-digital converter for each pixel on the sensor.

Capturing image gradients rather than actual pixel intensities was shown to increase the dynamic range of recorded content [TAR05]. In order to reconstruct intensity values, a computationally expensive Poisson solver needs to be applied to the measured data. While a *Gradient Camera* is an interesting theoretical concept, to the knowledge of the authors this camera has never actually been built.

The maximum intensity that can be resolved with standard ND filter arrays is limited by the lowest transmission of the employed ND filters. Large, completely saturated regions in the sensor image are usually filled with data interpolated from neighbouring unsaturated regions [NM00]. An analysis of sensor saturation in multiplexed imaging along with a Fourier-based reconstruction technique that boosts the dynamic range of captured images beyond the previous limits was recently proposed [WIH10]. Figure 2 shows an example image that is captured with an ND filter array on the left and a Fourier-based reconstruction of multiplexed data on the right.

An alternative to mounting a fixed set of ND filters in front of the sensor is an aligned spatial light modulator, such as a digital micromirror device (DMD). This concept was explored as *Programmable Imaging* [NBB04, NBB06] and allows for adaptive control over the exposure of each pixel. Unfortunately, it is rather difficult to align a DMD with a sensor on a pixel-precise basis, partly due to the required additional relay optics; for procedures to precisely calibrate such a system please consult [RFMM06]. Although a transmissive spatial light modulator can, alternatively, be mounted near the aperture plane of the camera, as proposed by Nayar and Branzoi [NB03], this *Adaptive Dynamic Range Imaging* approach only allows lower spatial frequencies in the image to be modulated. The most practical approach to adaptive exposures is a per-pixel control of the readout in software, as implemented by the Pixim camera [Pix10]. This has also been simulated for the specific case of CMOS sensors with rolling shutters [GHMN10], but only on a per-scanline basis. The next version of the *Frankencamera* [ATP*10] is planned to provide non-destructive sensor readout for small image regions of interest [Lev10], which would be close to the desired per-pixel exposure control.

Rouf *et al.* [RMH*11] propose to encode both saturated highlights and low-dynamic range (LDR) content in a single sensor image using cross-screen filters. Computerized tomographic reconstruction techniques are employed to estimate the saturated regions from glare created by the optical filters.

2.2. Multi-sensor and multi-exposure techniques

The most straightforward way of acquiring HDR images is to sequentially capture multiple photographs with different exposure times and merge them into a single, high-contrast image [MP95, DM97, MN99, RBS99]. Some of these approaches simultaneously compute the non-linear camera response function from the image sequence [DM97, MN99, RBS99]. Extensions to these techniques also allow HDR video [KUWS03]. Here, successive frames in the video are captured with varying exposure times and aligned using optical flow algorithms. Today, all of these methods are well established and discussed in the textbook by Reinhard *et al.* [RWD*10].

In addition to capturing multiple exposures, a static filter with varying transmissivity, termed *Generalized Mosaicing* [SN03a], can be mounted in front of the camera but also requires multiple photographs to be captured. Alternatively, the optical path of an imaging device can be divided using prisms [AA04] (*Split Aperture Imaging*) or beam-splitters [TKTS11, MMP*07], so that multiple sensors capture the same scene with different exposure times. While these approaches allow dynamic content to be recorded, the additional optical elements and sensor hardware make them more expensive and increase the form factor of the device.

2.3. Analysis and tradeoffs

Given a camera with known response function and dynamic range, Grossberg and Nayar [GN03] analyse the best possible set of actual exposure values for a LDR image sequence used to compute an HDR photograph. By also considering variable ISO settings, Hasinoff *et al.* [HDF10] provide the optimal choice of parameters for HDR acquisition with minimal noise. Granados *et al.* [GAW*10] analyse how to optimally combine a set of different exposures.

3. Spectral Imaging

Imaging of the electromagnetic spectrum comes in a number of flavours. For photographs or movies, the goal is typically to capture the colours perceived by the human visual system. Since the human visual system is based on three types of colour sensing cells (the cones), three colour bands are sufficient to form a natural colour impression. This discovery is usually credited to Maxwell [Max60].

In this report we are mainly concerned with methods for capturing the physical properties of light in contrast to their perceptual counterparts that are dealt with in the areas of Applied Perception and Colour Sciences. For readers interested in issues of colour perception, we refer to standard literature: Wyszeski and Stiles [WS82] provide raw data for many perceptual experiments. Fairchild's book [Fai05] is a higher-level treatise focusing on models for perceptual effects as, for instance, adaptation issues. Hunt's books [Hun91, Hun04] deal with measurement and reproduction of colour for human observers (e.g. in digital imaging, film, print, and television). Reinhard *et al.* [RKAJ08] discuss colour imaging from a computer graphics perspective.

In this section we discuss spectral imaging from a radiometric, i.e. physical, perspective. To simplify the discussion we first introduce some terminology as used in this subfield of plenoptic imaging.

3.1. Glossary of terms

Spectral radiance is the physical quantity emitted by light sources or reflected by objects. The symbol is L_λ and its unit is $[W/m^2 \cdot sr \cdot nm]$. Spectral radiance is constant along a ray. It is the quantity returned by the plenoptic function.

Spectral filters selectively attenuate parts of the electromagnetic spectrum. There are two principles of operation, *absorptive* spectral filters remove parts of the spectrum by converting photons into kinetic energy of the atoms constituting the material. Interference-based filters, also referred to as *dichroic* filters, consist of a transparent substrate which is coated with thin layers that selectively reflect light, reinforcing and attenuating different wavelengths in different ways. The number and thicknesses of the layers determine the spectral reflection profile. Absorptive filters have a better angular constancy, but heating may be an issue for narrow-band filters. Interference-based filters have the advantage that the spectral filter curve can be designed within certain limits by choosing the parameters of the coatings. However, the angular variation of these filters is significant. In general, filters are available both for transmission and reflection modes of operation.

Narrow-band filters have a small support in the wavelength domain.

Broad-band filters have a large support in the wavelength domain. They are also known as *panchromatic* filters.

The **spectral response curve** of a sensor is a function that describes its quantum efficiency with respect to photons of different wavelengths. A higher value means a better response of the sensor to photons of a particular wavelength, i.e. more electrons are freed due to the photo-electric effect.

Colour is the perceptual interpretation of a given electromagnetic spectrum.

The **gamut** of an imaging or colour reproduction system is the range of correctly reproducible colours.

Multi-spectral images typically consist of a low number of spectral bands. They often include a near infrared (NIR) band. The bands typically do not form a full spectrum, there can be missing regions [Vag07].

Hyper-spectral images contain 30 to several hundred spectral bands which are approximations to the full spectrum [Vag07]. The different spectral bands do not necessarily have the same spatial resolution. In this report, we will use the term multi-spectral to refer both to multi-spectral and hyper-spectral image acquisition methods.

A **multi-spectral data cube** is a stack of images taken at different wavelength bands.

3.2. Colour imaging

In a limited sense, the most common application of multi-spectral imaging is the acquisition of colour images for human observers. In principle, three spectral bands mimicking the human tri-stimulus system are sufficient to capture colour images. This principle was first demonstrated by Maxwell performing colour photography by time-sequential acquisition of three images using different band-pass filters (see Figure 3) . Display was achieved by super-imposed spectrally filtered black-and-white projections using the same filters as used for capture. This acquisition principle was in use for quite some time until practical film-based colour photography was invented. One of the earliest collections of colour photographs was assembled by the Russian photographer Sergej Mikhailovich Prokudin-Gorskij [PG12]. Time-sequential imaging through different filters is still one of the main modes of capturing multi-spectral images (see Section 3.3).

In the digital age, colour films have been replaced by electronic CMOS or CCD sensors. The two technologies to capture an instantaneous colour image are optical splitting trees employing dichroic beam-splitter prisms [Opt11], as used in three-CCD cameras, and spatial multiplexing [NN05, IWH10], trading spatial resolution for colour information. The spatially varying spectral filters in multiplexing applications are also known as colour filter arrays (CFAs). A different principle, based on volumetric or layered measurements, is employed by the Foveon sensor [Fov10], which captures tri-chromatic images at full spatial resolution.

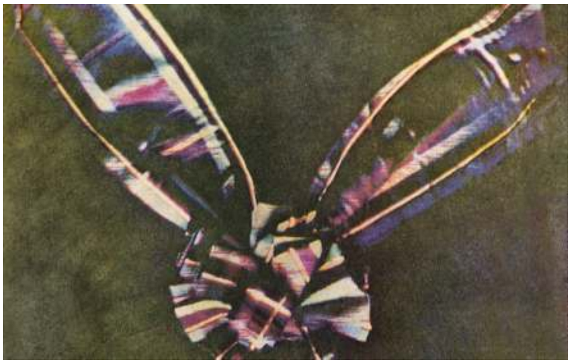

Figure 3: *Tartan Ribbon, considered to be the world's first colour photograph, taken by Thomas Sutton for James Clerk Maxwell in 1861 by successively placing three colour filters in front of the camera's main lens and taking three monochromatic photographs (Wikimedia Commons).*

The most popular spatial multiplexing pattern is the well-known Bayer pattern [Bay76]. It is used in most single-sensor digital colour cameras. The associated problem of reconstructing a full-resolution colour image is generally referred to as demosaicing. An overview of demosaicing techniques is given in [RSBS02, GGA*05, LGZ08]. Li *et al.* [LGZ08] present a classification scheme of demosaicing techniques depending on the prior model being used (explicitly or implicitly) and an evaluation of different classes of algorithms. An interesting result is that the common constant-hue assumption seems to be less valid for modern imagery with a wider gamut than the classical Kodak photo CD test set [Eas], which was scanned from film and has predominantly been used for evaluating demosaicing algorithms. Mostly, demosaicing is evaluated through simulation. However, in a realistic setting, including camera noise, Hirakawa and Parks [HP06] have shown that demosaicing on noisy images performs poorly and that subsequent denoising is affected by demosaicing artefacts. They propose a joint denoising and demosaicing framework that can be used with different demosaicing and denoising algorithms.

In recent years, a large number of alternative CFAs have been explored by camera manufacturers, some of which are already being used in consumer products. Examples and many references to alternative CFA designs can be found in Hirakawa and Wolfe's work [HW08]. Traditionally, imaging through CFAs and reconstruction of the signal have been seen as sub-sampling and up-sampling operations, respectively. Recent research in the analysis of these multiplexing patterns has, however, produced a new view of multiplexing as a projection onto a linear subspace of basis functions (the spectral responses of the filters in this case), i.e. of multiplexing as a coding operation [LV09, IWH10]. Correspondingly, in this view, the reconstruction process is seen as a recovery of the subspace, or a decoding of the signal. This view originated in Fourier analysis of CFAs [ASH05], stimulated by the desire to apply digital signal processing methodology to the colour multiplexing problem. Being a linear framework, it allows for the optimization of the subspace onto which colour is projected [HW07, HW08, LV09]. Practical realizations are alternative CFA designs that suffer from less aliasing than their *ad hoc*, heuristically designed counterparts. While in [HW07, HW08, LV09] a fixed number of primary colour response functions are assumed which can be linearly mixed to optimize the CFA [PR06, PR10] and optimize the spectral response functions themselves in order to improve CFA design. More recently, [SMH*11] proposed to use shiftable layers of CFAs; this design allows the colour primaries to be switched dynamically and provides an optimal SNR in different lighting conditions. Wetzstein *et al*. [WIH10] explore CFA designs that, in conjunction with a demosaicing-like optimization algorithm, allow trichromatic HDR images to be captured.

Generalizing CFAs, Narasimhan and Nayar [NN05] proposed the *Assorted Pixels* framework, where individual pixels can be modulated by arbitrary plenoptic filters, yielding an image mosaic that has to be interpolated to obtain the full-resolution multi-channel image. In subsequent work [YMIN10], aliasing within the *Assorted Pixels* framework was minimized. Ihrke *et al*. [IWH10] have shown how this (and other) approaches that are tiling the image in a superpixel fashion can be interpreted as belonging to one group of imaging systems that share common analysis and reconstruction approaches.

3.3. Multi-spectral imaging

As for the other plenoptic dimensions, the three basic approaches of Figure 1, single-shot capture, sequential image acquisition and multiple device set-ups are valid alternatives for multi-spectral imaging and have been investigated intensely.

3.3.1. *Spectrometers*

Traditionally, spectroscopy has been carried out for single rays entering an instrument referred to as spectrometer. It was invented by Joseph von Fraunhofer in 1814 and used to discover the missing lines in the solar spectrum bearing his name. Typically the ray is split into its constituent wavelengths which are displaced spatially. This is achieved by placing either dispersive or diffractive elements into the light path, where the latter come both in transmissive and reflective variants. If dispersion is used to split the ray, typically a prism is employed. The separation of wavelengths is caused by the wavelength-dependent refractive index of the prism material. The function mapping wavelength to refractive index is typically decreasing with increasing wavelength, but usually in a non-linear fashion. Under certain conditions, it can even have an inverted slope (anomalous dispersion) [Hec02]. Diffractive elements are usually gratings, where maxima of the diffraction pattern are spatially shifted according to the grating equation [Hec02]. After the light path is split by some means, the light is brought onto a photo-detector which can, for instance, be a CCD. Here, relative radiance of the individual wavelengths is measured. Spectrometers have to be calibrated in two ways: first, the mapping of wavelengths to pixels has to be determined. This is usually done using light sources with few and very narrow emission lines of known wavelength, the pixel positions of other wavelengths are then interpolated [GWHD09]. The second step establishes the relative irradiance measured for every wavelength. This is done by measuring a surface of known flat reflectance, for example Spectralon, which is illuminated with a known broad-band spectrum. The relative inhomogeneities imposed by the device are then divided out [GWHD09]. Spectrometers that are designed to image more than one ray are referred to as *imaging spectrometers*.

3.3.2. *Scanning imaging spectrometers*

Traditional devices are usually based on some form of scanning. Either a full two-dimensional image is acquired with changed band-pass filters, effectively performing a spectral scan, or a pushbroom scan is performed where the two-dimensional CCD images a spatial dimension on one axis of the image and the spectral dimension on the other. The full multi-spectral data cube is then obtained by scanning the remaining spatial dimension.

Spectral scanning can be performed in a variety of ways. Most of them involve either a filter wheel (e.g. [WH04]) or electronically tunable filters. The former method usually employs narrow-band filters such that spectral bands are imaged directly. The disadvantage is a low light throughput. Toyooka and Hayasaka [TH97] present a system based on broad-band filters with computational inversion. Whether or not this is advantageous depends on the camera noise [IWH10]. Electronically tunable filters are programmable devices that can exhibit varying filter curves depending on control voltages applied to the device. Several incarnations exist; the most well known include Liquid Crystal Tunable Filters (LCTFs) [ci09], which are based on a cascade of Lyot-filter stages [Lyo44], acousto-optical tunable filters, where an acoustically excited crystal serves as a variable diffraction grating, and interferometer-based systems, where the spectrum is projected into the Fourier basis. In the latter, the spectral scan is performed in a multiplexing manner: by varying the position of the mirror in one arm of an interferometer, for instance a Michelson-type device, different phase shifts are induced for every wavelength. The resulting spectral modulation is in the form of a sinusoid. Thus, effectively, the measurements are performed in the Fourier basis, similar to the Dappled Photography technique [VRA*07] for light fields (see Section 4). The spectrogram is obtained by taking an inverse Fourier transform. A good overview of these technologies is given in [Gat00].

A more flexible way of programmable wavelength modulation is presented by Mohan *et al*. [MRT08]. They modulate the spectrum of a entire image by first diffracting the light and placing an attenuating mask into the ‘rainbow plane’ of the imaging system. However, the authors do not recover multiplexed spectra but only demonstrate modulated spectrum imaging. Usually, the scanning and the imaging process have to be synchronized, i.e. the camera should only take an image when the filter in front of the camera is set to a known value. Schechner and Nayar [SN04] introduce a technique to computationally synchronize video streams taken with a periodically moving spectral filter.

All previously discussed techniques attempt to recover scene spectra passively. An alternative technique using active spectral illumination in a time-sequential manner is presented by Park *et al*. [PLGN07]. The scene, which can include ambient lighting, is imaged under different additional spectral lighting. The acquired images allow for reasonably accurate per-pixel spectra to be recovered.

Spatial scanning has been widely employed in satellite-based remote sensing. Two technologies are commonly used: pushbroom and whiskbroom scanning. Whereas pushbroom scanning uses a two-dimensional sensor and can thus recover one spectral and one spatial dimension per position of the satellite, whiskbroom systems employ a one-dimensional sensor, imaging the spectrum of a single point which is then scanned to obtain a full scan-line with a rotating mirror. The main idea is that a static scene can be imaged multiple times using different spectral bands and thus a full multi-spectral data cube can be assembled. A good overview of space-borne remote sensing, and more generally, multi-spectral imaging techniques is given in [Vag07]

In computer vision, a similar concept, called *Generalized Mosaicing*, has been introduced by Schechner and Nayar [SN01]. Here, a spatially varying filter is mounted in front of the main lens, filtering each column of the acquired image differently. By moving the camera and registering the images, a full multi-spectral data cube can be recovered [SN02].

3.3.3. *Single-shot imaging spectrometers*

To enable the spectral acquisition of fast-moving objects, it is necessary to have single-shot methods available. Indeed, this appears to be the focus of research in recent years. We can differentiate between three major modes of operation. The first is a trade of spatial for spectral resolution. Optical devices are implemented that provide empty space on the sensor which can, with a subsequent dispersion step through which a scene ray is split into its wavelength constituents, be filled with spectral information. The second option is multi-device set-ups which operate mostly like their spectral scanning counterparts, replacing sequential imaging by additional hardware. The third class of devices employs computational imaging, i.e. computational inversion of an image formation process where spectral information is recorded in a super-imposed manner.

Spatial multiplexing of the spectrum, in general, uses a dispersive or diffractive element in conjunction with some optics redirecting rays from the scene onto parts of the sensor surrounded by void regions. The void regions are then filled with spectral information. All these techniques take advantage of the high resolution of current digital cameras. Examples using custom manufactured redirecting mirrors include [HFHG05, GKT09, GFHH10]. These techniques achieve imaging of up to 25 spectral bands in real-time and keep the optical axis of the different slices of the multi-spectral data cube constant. Bodkin *et al*. [BSN*09] and Du *et al*. [DTCL09] propose a similar concept by using an array of pinholes that limits the rays that can reach the sensor from the scene. The pinholes are arranged such that a prism following in the optical path disperses the spectrum and fills the pixels with spectral information. A different approach is taken by Fletcher-Holmes *et al*. [FHH05]. They are interested in only providing a small ‘foveal region’ in the centre of the image with multi-spectral information. For this, the centre of the image is probed with fibre optic cables which are fed into a standard spectrometer. Mathews *et al*. [Mat08] and Horstmeyer *et al*. [HEAL09] describe light field cameras with spectrally filtered sub-images. An issue with this design is the problem of motion parallax induced by the different view points when registering the images (see Section 4). In general, this registration problem is difficult and requires knowledge of scene geometry and reflectance which cannot easily be estimated.

Multi-device set-ups are similar in spirit to spectral scanning spectrometers, replacing the scanning process by additional hardware. A straightforward solution recording five spectral bands is presented by Lau *et al*. [LY05]. They use a standard multi-video array where different spectral filters are mounted on each camera. The motion-parallax problem mentioned previously is even worse in this case. McGuire *et al*. [MMP*07] discuss optical splitting trees where the individual sensors are aligned such that they share a single optical axis. The design of beam-splitter/filter trees is non-trivial and the authors propose an automatic solution based on optimization. A hybrid camera was proposed by Cao *et al*. [CTDL11], where the output of a high resolution RGB camera was combined with that of a prism-based low spatial-resolution, high spectral-resolution camera.

Computational spectral imaging aims at trading computational complexity for simplified optical designs. *Computed* *tomography image spectrometry* (CTIS) was developed by Okamoto and Yamaguchi [OY91]. They observed that by placing a diffraction grating in the optical path, several spectral copies overlay on the image sensor. Every pixel is measuring a line integral along the spectral axis. Knowing the imaging geometry enables a tomographic reconstruction of the spectra. A drawback to this technique is that not all data can be measured and thus an ill-conditioned problem, similar to limited angle tomography, is encountered. The technique was extended to single shot imaging by Descour *et al*. [DD95, DC*97].

A relatively novel technique is referred to as *Coded Aperture Snapshot Spectral Imaging* (CASSI) [GJB*07, WPSB08, WPSB09]. In a series of papers the authors show how to construct different devices to exploit the compressive sensing paradigm [CRT06] which promises to enable higher resolution computational reconstructions with less samples than predicted by the Shannon-Nyquist sampling theorem.

The results presented for both CTIS and CASSI have only been demonstrated for relatively low-resolution, low-quality spectral images. Therefore, these approaches are not yet suitable for high-quality photographic applications.

3.4. Applications

There is a huge amount of applications for multi-spectral imaging and we are just beginning to explore the possibilities in computer graphics and vision. Traditional users of multi-spectral imaging technology are in the fields of astronomy and remote sensing where, for instance, the mapping of vegetation, minerals, water surfaces, and hazardous waste monitoring are of interest. In addition, multi-spectral imaging is used for material discrimination [DTCL09], ophthalmology [LFHHM02], the study of combustion dynamics [HP01], cellular dynamics [KYN*00], surveillance [HBG*00], for deciphering ancient scrolls [Man05], flower photography [Ror08], medicine, agriculture, manufacturing, forensics and microscopy. It should not be forgotten that the military is an interested party [Vag07].

4. Light Field Acquisition

Light fields are sets of 2D images, each depicting the same scene from a slightly different viewpoint [LH96, GGSC96]. In a region free of occluders, such an image-based scene representation fully describes a slice of constant time and wavelength of the full plenoptic function. A light field includes all global illumination effects for objects of any geometric and radiometric complexity in a fixed illumination environment. As discussed in more detail in Section 4.4, applications include novel viewpoint rendering, synthetic aperture photography, post-capture image refocus as well as geometry and material reconstruction.

The concept of a light field predates its introduction in computer graphics. The term itself dates to the work of Gershun [Ger36], who derived closed-form expressions for illumination patterns projected by area light sources. Ashdown [Ash93] continued this line of research. Moon and Spencer [MS81] introduced the equivalent concept of a *photic field* and applied it to topics spanning lighting design, photography, and solar heating. The concept of a light field is similar to epipolar volumes in computer vision [BBM87]. As demonstrated by Halle [Hal94], both epipolar volumes and holographic stereograms can be captured by uniform camera translations. The concept of capturing a 4D light field, for example by translating a single camera [LH96, GGSC96] or by using an array of cameras [WSLH02], is predated by integral photography [Lip08], parallax panoramagrams [Ive03] and holography [Gab48].

This section catalogues existing devices and methods for light field capture, as well as applications enabled by such data sets. Note that a sensor pixel in a conventional camera averages the radiance of light rays impinging over the full hemisphere of incidence angles, producing a 2D projection of the 4D light field. In contrast, light field cameras prevent such averaging by introducing spatio-angular selectivity. Such cameras can be classified into those that primarily rely on multiple sensors or a single sensor augmented by temporal, spatial or frequency-domain multiplexing.

4.1. Multi-sensor capture

As described by Levoy and Hanrahan [LH96], a light field can be measured by capturing a set of photographs taken by an array of cameras distributed on a planar surface. Each camera measures the radiance of light rays incident on a single point, defined in the plane of the cameras, for a set of angles determined by the field of view of each camera. Thus, each camera records a 2D slice of the 4D light field. Concatenating these slices yields an estimate of the light field. Wilburn *et al*. [WSLH02, WJV*05] achieve dynamic light field capture using an array of up to 125 digital video cameras (see Figure 4, left). Yang *et al*. [YEBM02] propose a similar system using 64 cameras. Nomura *et al*. [NZN07] create scene collages using up to 20 cameras attached to a flexible plastic sheet, combining the benefits of both multiple sensors and temporal multiplexing. Custom hardware allows accurate calibration and synchronization of the camera arrays. Such designs have several unique properties. Foremost, as demonstrated by Vaish *et al*. [VSZ*06], the captured light field can be considered as if it were captured using a single camera with a main lens aperture extending over the region occupied by the cameras. Such large-format cameras cannot be practically constructed using refractive optics. Vaish *et al*. exploit this configuration by applying methods of synthetic aperture imaging to obtain sharp images of objects obscured by thick foliage.

Figure 4: *Light field cameras can be categorized by how a 4D light field is encoded in a set of 2D images. Methods include using multiple sensors or a single sensor with temporal, spatial or frequency-domain multiplexing. (Top, Left) Wilburn et al. [WSLH02] describe a camera array. (Top, Middle) Liang et al. [LLW*08] achieve temporal multiplexing with a programmable aperture. (Top, Right) Georgiev et al. [GIBL08] capture spatially multiplexed light fields using an array of lenses and prisms. (Bottom) Raskar et al. [RAWV08] capture frequency-multiplexed light fields by placing a heterodyne mask [VRA*07, VRA*08, LRAT08] close to the sensor. (Figures reproduced from [WSLH02], [LLW*08], [GIBL08] and [RAWV08].)*

4.2. Time-sequential imaging

Camera arrays have several significant limitations; foremost, a sparse array of cameras may not provide sufficient light field resolution for certain applications. In addition, the cost and engineering complexity of such systems prohibit their use for many consumer applications. As an alternative, methods using a single image sensor have been developed. For example, Levoy and Hanrahan [LH96] propose a direct solution; using a mechanical gantry, a single camera is translated over a spherical or planar surface, constantly reoriented to point towards the object of interest. Alternatively, the object can be mechanically rotated on a computer-controlled turntable. Ihrke *et al.* [ISG*08] substitute mechanical translation of a camera with rotation of a planar mirror, effectively creating a time-multiplexed series of virtual cameras. Thus, by distributing the measurements over time, single-sensor light field capture is achieved. Taguchi *et al.* [TARV10] show how capturing multiple images of rotationally symmetric mirrors from different camera positions allow wide field of view light fields to be captured. Gortler *et al.* [GGSC96] propose a similar solution; the camera is manually translated and computer vision algorithms are used to estimate the light field from such uncontrolled translations. These approaches trace their origins to the method introduced by Chen and Williams [CW93], which is implemented by QuickTime VR.

The preceding systems capture the light field impinging on surfaces enveloping large regions (e.g. a sphere encompassing the convex hull of a sculpture). In contrast, hand-held light field photography considers capturing the light field passing through the main lens aperture of a conventional camera. Adelson and Wang [AW92], Okano *et al.* [OAHY99] and Ng *et al.* [NLB*05] extend integral photography to spatially multiplex a 4D light field onto a 2D image sensor, as discussed in the following subsection. However, time-sequential capture can also achieve this goal.

Liang *et al.* [LLW*08] propose programmable aperture photography to achieve time-multiplexed light field capture. While Ives [Ive03] uses static parallax barriers placed close to the image sensor, Liang *et al.* use dynamic aperture masks (see Figure 4, middle). For example, consider capturing a sequence of conventional photographs. Between each exposure a pinhole aperture is translated in raster scan order. Each photograph records a pencil of rays passing through a pinhole located at a fixed position in the aperture plane for a range of sensor pixels. Similar to multiple sensor acquisition schemes, each image is a 2D slice of the 4D light field and the sequence can be concatenated to estimate the radiance for an arbitrary light ray passing through the aperture plane. To reduce the necessary exposure time, Liang *et al*. further apply Hadamard aperture patterns, originally proposed by Schechner and Nayar [SNB07], that are 50% transparent.

The preceding methods all consider conventional cameras with refractive lens elements. Zhang and Chen [ZC05] propose a lensless light field camera. In their design, a bare sensor is mechanically translated perpendicular to the scene. The values measured by each sensor pixel are recorded for each translation. By the Fourier projection-slice theorem [Ng05], the 2D Fourier transform of a given image is equivalent to a 2D slice of the 4D Fourier transform of the light field; the angle of this slice is dependent on the sensor translation. Thus, tomographic reconstruction yields an estimate of the light field using a bare sensor, mechanical translation, and computational reconstruction methods.

4.3. Single-shot multiplexing

Time-sequential acquisition reduces the cost and complexity of multiple sensor systems, however it has one significant limitation: dynamic scenes cannot be readily captured. Thus, either a high-speed camera is necessary or alternative means of multiplexing the 4D light field into a 2D image are required. Ives [Ive03] and Lippmann [Lip08] provide two early examples of spatial multiplexing with the introduction of parallax barriers and integral photography, respectively. Such spatial multiplexing allows light field capture of dynamic scenes, but requires a trade-off between the spatial and angular sampling rates. Okano *et al*. [OAHY99] and Ng *et al*. [NLB*05] describe modern, digital implementations of integral photography, however numerous other spatial multiplexing schemes have emerged.

Instead of affixing an array of microlenses directly to an image sensor, Georgiev *et al*. [GZN*06] add an external lens attachment with an array of lenses and prisms (see Figure 4, right). Ueda *et al*. [UKTN08, ULK*08] consider similar external lens arrays; however, in these works, an array of variable focus lenses, implemented using liquid lenses controlled by electrowetting, allow the spatial and angular resolution to be optimized depending on the observed scene.

Rather than using absorbing masks or refractive lens arrays, Unger *et al*. [UWH*03], Levoy *et al*. [LCV*04], Lanman *et al*. [LWTC06] and Taguchi *et al*. [TAV*10] demonstrate that a single photograph of an array of tilted, planar mirrors or mirrored spheres produces a spatially multiplexed estimate of the incident light field. Yang *et al*. [YLIM00] demonstrate a large-format, lenslet-based architecture by combining an array of lenses and a flatbed scanner. Related compound imaging systems, producing a spatially multiplexed light field using arrays of lenses and a single sensor, were proposed by Ogata *et al*. [OIS94], Tanida *et al*. [TKY*01, TSK*03] and Hiura *et al*. [HMR09].

Spatial multiplexing produces an interlaced array of elemental images within the image formed on the sensor. Veeraraghavan *et al*. [VRA*07] introduce frequency multiplexing as an alternative method for achieving single-sensor light field capture. The optical heterodyning method proposed by Veeraraghavan *et al*. encodes the 4D Fourier transform of the light field into different spatio-angular bands of the Fourier transform of the 2D sensor image. Similar in concept to spatial multiplexing, the sensor spectrum contains a uniform array of 2D spectral slices of the 4D light field spectrum. Such frequency-domain multiplexing is achieved by placing non-refractive, light-attenuating masks slightly in front of a conventional sensor (see Figure 4, bottom).

As described by Veeraraghavan *et al*., masks allowing frequency-domain multiplexing (i.e. heterodyne detection) must have a Fourier transform consisting of an array of impulses (i.e. a 2D Dirac comb). In [VRA*07], a *Sum-of-Sinusoids (SoS) pattern*, consisting of a weighted harmonic series of equal-phase sinusoids, is proposed. As shown in Figure 5, such codes transmit significantly more light than traditional pinhole arrays [Ive03]; however, as shown by Lanman *et al*. [LRAT08], these patterns are equivalent to a truncated Fourier series approximation of a pinhole array for high angular sampling rates. In [LRAT08], Lanman *et al*. propose *tiled-broadband patterns*, corresponding to periodic masks with individual tiles exhibiting a broadband Fourier transform. This family includes pinhole arrays, SoS patterns and the tiled-MURA patterns proposed in that work (see Figure 5). Such patterns produce masks with 50% transmission, enabling shorter exposures than existing methods.

In subsequent work, Veeraraghavan *et al*. [VRA*08] propose adaptive mask patterns, consisting of aharmonic sinusoids, optimized for the spectral bandwidth of the observed scene. Georgiev *et al*. [GIBL08] analyse such heterodyne cameras and further propose masks placed external to the camera body. Rather than using a global, frequency-domain decoding scheme, Ihrke *et al*. [IWH10] demonstrate how spatial-domain decoding methods can be extended to frequency-multiplexed light fields.

4.4. Applications

Given the wide variety of light field capture devices, a similarly diverse set of applications is enabled by such high-dimensional representations of light transport. While Kanolt [Kan18] considers the related concept of a parallax panoramagram to achieve 3D display, light fields have also proven useful for applications spanning computer graphics, digital photography and 3D reconstruction.

In the field of computer graphics, light fields were introduced to facilitate image-based rendering [LH96, GGSC96].

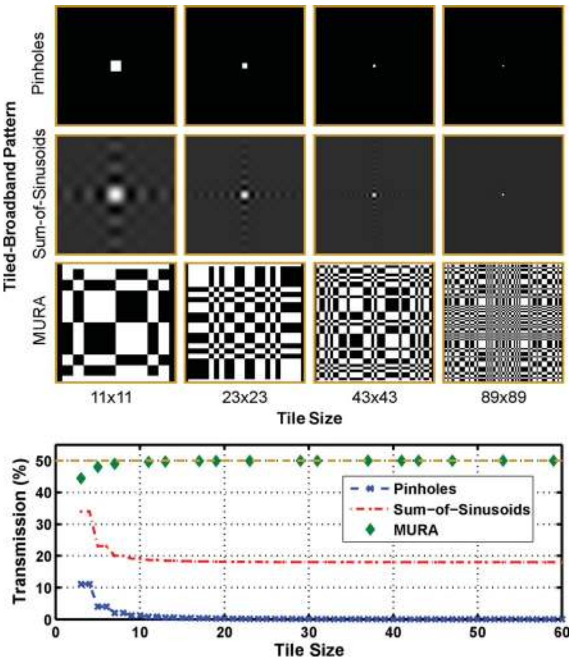

Figure 5: *Lanman et al. [LRAT08] introduce tiled-broadband patterns for mask-based, frequency-multiplexed light field capture. (Top) Each row, from left to right, shows broadband tiles of increasing spatial dimensions, including: pinholes [Ive28], Sum-of-Sinusoids (SoS) [VRA*07] and MURA [GF89, LRAT08]. (Bottom) The SoS tile converges to 18% transmission, whereas the MURA tile remains near 50%. Note that frequency multiplexing with either SoS or MURA tiles significantly outperforms conventional pinhole arrays in terms of total light transmission and exposure time. (Figures reproduced from [Lan10].)*

In contrast to the conventional computer graphics pipeline, novel 2D images are synthesized by resampling the 4D light field. With sufficient light field resolution, views are synthesized without knowledge of the underlying scene geometry. Subsequent to these works, researchers continued to enhance the fidelity of image-based rendering. For example, a significant limitation of early methods is that illumination cannot be adjusted in synthesized images. This is in stark contrast to the conventional computer graphics pipeline, wherein arbitrary light sources can be supported using ray tracing together with a model of material reflectance properties. Debevec *et al.* [DHT*00] address this limitation by capturing an 8D *reflectance field*. In their system, the 4D light field reflected by an object is measured as a function of the 4D light field incident on the object. Thus, an 8D reflectance field maps variations in the input radiance to variations in the output radiance, allowing image-based rendering to support variation of both viewpoint and illumination.

Light fields parametrize every possible photograph that can be taken outside the convex hull of an object; as a result, they have found widespread application in 3D television, also known as free-viewpoint video. Carranza *et al.* [CTMS03] describe a system with an array of cameras surrounding one or more actors. Similar systems have been developed by Matusik *et al.* [MBR*00] and Starck *et al.* [SH08]. Image-based rendering allows arbitrary adjustment of the viewpoint in real-time. Vlasic *et al.* [VPB*09] further demonstrate 3D reconstruction of human actors from multiple-camera sequences captured under varying illumination conditions.

Light fields, given their similarity to conventional parallax panoramagrams [Ive28], have also found application in the design and analysis of 3D displays. Okano *et al.* [OAHY99] adapt integral photography to create a 3D television system supporting both multi-view capture and display. Similarly, Matusik and Pfister [MP04] achieve light field capture using an array of 16 cameras and implement light field display using an array of 16 projectors and lenticular screens. Zwicker *et al.* [ZMDP06] develop antialiasing filters for automultiscopic 3D display using a signal processing analysis. Hirsch *et al.* [HLHR09] develop a *BiDirectional (BiDi) screen*, supporting both conventional 2D image display and real-time 4D light field capture, facilitating mixed multitouch and gesture-based interaction; the device uses a lensless light field capture method, consisting of a tiled-MURA pattern [LRAT08] displayed on an LCD panel and a large-format sensor. Recently, Lanman *et al.* [LHKR10] use an algebraic analysis of light fields to characterize the rank constraints of all dual-layer, attenuation-based light field displays; through this analysis they propose a generalization of conventional parallax barriers, using content-adaptive, time-multiplexed mask pairs to synthesize high-rank light fields with increased brightness and spatial resolution. Wetzstein *et al.* [WLHR11] demonstrate how a stack of attenuating layers, such as spaced transparencies or LCD panels, can be used in combination with computerized tomographic reconstruction to display natural light fields of 3D scenes.

Post-processing of captured light fields can resolve long-standing problems in conventional photography. Ng [Ng05] describes efficient algorithms for digital image refocusing, allowing the plane of focus to be adjusted after a photograph has been taken. Recorded with a camera array, light fields allow photographs to be simulated that exhibit an aperture size corresponding to the size of the array rather than any individual camera aperture (see Figure 6) . In addition, Talvala *et al.* [TAHL07] and Raskar *et al.* [RAWV08] demonstrate that high-frequency masks can be combined with light field photography to eliminate artefacts due to glare and multiple scattering of light within camera lenses. Similarly, light field capture can be extended to microscopy and confocal imaging, enabling similar benefits in extended depth of field (DOF) and reduced scattering [LCV*04, LNA*06]. Smith *et al.* [SZJA09] improve conventional image stabilization algorithms using light fields captured with an array of 25 cameras. As described, most single-sensor acquisition schemes trade increased angular resolution for

Figure 6: *The images of a light field can be processed to reveal parts of the scene (right) that are occluded in any of the individual views (left) of a camera array. (Figures reproduced from [VWJL04].)*

decreased spatial resolution [GZN*06]; Bishop *et al.* [BZF09] and Lumsdaine and Georgiev [LG09] apply priors regarding the statistics of natural images and modified imaging hardware, respectively, to achieve super-resolution light field capture that, in certain conditions, mitigates this resolution loss.

As characterized throughout this report, the plenoptic function of a given scene contains a large degree of redundancy; both the spatial and angular dimensions of light fields of natural scenes are highly correlated. Recent work is exploring the benefits of compressive sensing for light field acquisition. Fergus *et al.* [FTF06] introduce *random lens imaging*, wherein a conventional camera lens is replaced with a random arrangement of planar mirrored surfaces, allowing super-resolution and 3D imaging applications. Babacan *et al.* [BAL*09] propose a compressive sensing scheme for light field capture utilizing randomly coded, non-refractive masks placed in the aperture plane. Ashok and Neifeld [AN10] propose compressive sensing schemes, again using non-refractive masks, allowing either spatial or angular compressive light field imaging. As observed in that work, future capture methods will likely benefit from joint spatio-angular compressive sensing; however, as discussed later in this report, further redundancies exist among all the plenoptic dimensions, not just the directional variations characterized by light fields.

5. Multiplexing Space and Focal Surfaces

The ability to resolve spatial light variation is an integral part of any imaging system. For the purpose of this report we differentiate between spatial variation on a plane perpendicular to the optical axis and variation along the optical axis inside a camera behind the main lens. The former quantity, transverse light variation, is what all 2D sensors measure. In this section, we discuss approaches for very high-resolution imaging (Section 5.1), panoramic and gigapixel imaging (Section 5.2), focal surface curvature correction techniques of the light field inside a camera (Section 5.3), and extended DOF photography (Section 5.4).

5.1. Super-resolution imaging

Although exotic camera systems can resolve structures in the order of 100 nm [vPAB*11], the resolution of standard photographs is usually limited by the physical layout and size of the photosensitive elements, the optical resolution of employed optical elements and the diffraction limit. Attempts to break these limits, which are significantly larger than 100 nm, are referred to as super-resolution imaging. Such techniques have been of particular interest to the vision community for many years. In most cases a sequence of slightly shifted low-resolution photos is captured and fused into a single high-resolution image. The shifts are usually smaller than the pixel size; an extensive review of such techniques can be found in [BK02, BS98]. Sub-pixel precise shifts of low-resolution images can be achieved by mechanical vibrations [LMK01, BEZN05], by coding the camera's aperture using phase [AN07] and attenuation [MHRT08] masks, by exploiting object motion in combination with temporally coded apertures [AR07], or by analysing multiple frames of a video [LS11]. For an increased resolution in space and time, successive frames in videos of a single camera [SFI11] or multiple devices [SCI02, SCI05] (see Section 6.2) can be analysed instead. All super-resolution approaches require an optimization problem to be solved for the unknown super-resolved image given multiple low-resolution measurements. This is computationally expensive for higher resolutions and is usually an ill-posed problem requiring additional image priors [BK02].

5.2. Panoramic and gigapixel photography

Recording photographs that exceed the field of view of a camera's lens are referred to as panoramic images. Usually, a series of images, each covering a small part of the desired field of view, is recorded with a rotating camera. The individual images are then stitched together to form a single, wide field of view photograph. Computerized solutions have been under investigation for decades (e.g. [Mil75]); robust, fully automatic algorithms are available today [BL07].

Gigapixel imaging is a relatively new field that, similar to super-resolution, aims at capturing very high-resolution imagery. The main difference is that gigapixel imaging approaches generally do not try to beat the limits of sensor resolution, but rather stitch a gigapixel panoramic image together from a set of megapixel images. These can be photographed by mounting a camera on a computer-controlled rotation stage [KUDC07, Nou10], or a high-resolution small-scale sensor that is automatically moved in the image plane of a large-format camera [BE11]. Both of these techniques implement the concept of capturing a sequence of images with a single device that are composited into a high-quality photograph. In this case, the parameter that is varied for each image in the sequence is the camera pose. Alternatively, the optics of a camera can be modified, for instance with custom

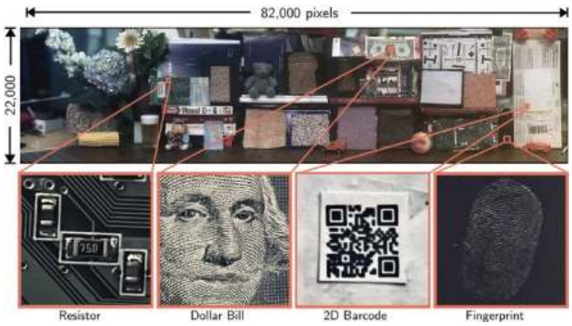

Figure 7: *A wide field of view 1.7 gigapixel image captured by Cossairt et al. [CMN11].*

spherical lens elements, to allow a single very high-resolution image to be captured instantaneously with multiple sensors [CMN11]. An example scene captured with this technique is shown in Figure 7.

5.3. Optical field correction

Not only is the actual resolution of digital photographs limited by the pixel count and the diffraction limit, but also by the applied optical elements. Standard spherical lenses have a focal surface that is, unlike most sensors, not actually planar but curved. Significant engineering effort is put into the commercial development of complex lens systems, especially in variable-focus camera objectives, that correct for the resulting image blur at sensor locations away from the optical axis. Several approaches have been proposed to correct for what is usually called field curvature, or more simply lens aberrations, and also phase aberrations caused by atmospheric turbulences in astronomical imaging [Tys91]. These usually integrate secondary optical assemblies into the system, such as fibre optics [KH57], prisms [Ind73], lenslet arrays [HN06, BH09], coded attenuation masks [PKR10], or spatial light modulators [Tys91] and oftentimes require computational processing of the measured data.

5.4. Extended depth of field photography

DOF, that is a depth-dependent (de)focus of a pictured scene, is a photographic effect that is caused by the finite aperture size of cameras (see standard literature, e.g. [Hec02]). Whereas objects in the photographed scene that are located around the focal plane are imaged sharply onto the sensor image, objects farther away from that plane are out of focus and, therefore, blurred. In this section, we discuss approaches for extending the DOF of an optical system; all such techniques aim at joint optical light modulation and computational processing to either capture all-focused images or allow the focal plane to be adjusted after an image was taken. Light fields, as discussed in Section 4, can practically enhance DOF but require an increased amount of visual information to be captured. The approaches presented in this section, on the other hand, usually manipulate the camera optics so that a single recorded image contains all the necessary information to refocus it or remove all defocus, without requiring the full 4D light field to be reconstructed.

Although a photographer can interactively adjust the focal plane before a standard photograph is captured, removing the DOF blur in such an image is a very difficult problem. DOF blur inversion requires a deconvolution of the image with a spatially varying point spread function (PSF). The PSF shape corresponds to that of the camera aperture, which can often be well approximated with a Gaussian distribution; unfortunately, a deconvolution with a Gaussian is an ill-posed inverse problem, because high frequencies are irreversibly lost in the image capture. Applying natural image priors can improve reconstructions (see e.g. [LFDF07b]). The spatially varying PSF size is directly proportional to the depth of the scene, which is in most cases unknown. A common approach to alleviate this problem is to mechanically or optically modify the depth-dependent PSF of the imaging system so that it becomes depth-invariant resulting in a reconstruction that only requires a spatially invariant deconvolution, which is much easier and does not require knowledge of the scene depth.

One family of techniques that only requires a single shot to capture a scene with a depth-invariant PSF is called *Focal Sweep*. Here, the PSF modulation is achieved by moving the object [Hä72] or the sensor [NKZN08] during the exposure time, or by exploiting the wavelength-dependency of the PSF to multiplex multiple focal planes in the scene onto a single sensor image [CN10].

Alternatively, the apertures of the imaging system can be coded with cubic phase plates [DC95] or other phase masks [OCLE05, BEMKZ05, CG01], diffusers [GGML07, CZN10], attenuation patterns [LFDF07a], polarization filters [CCG06] or multi-focal elements [LHG*09].

All of the above listed focal sweep and coded aperture approaches optically modify the PSF of the optical system for an extended DOF. The captured images usually need to be post-processed, for instance by applying a deconvolution. An analysis of quality criteria of attenuation-based aperture masks for defocus deblurring was presented by Zhou and Nayar [ZN09]; this analysis was extended to also consider PSF invertibility [Bae10].

Focal Stacks are series of images of the same scene, where the focal plane differs for each photograph in the sequence. A single, focused image can be composited by selecting the best-focused match in the stack for each image region [PK83]. The optimal choice of parameters, including focus and aperture, for the images in a focal stack are well established [HKDF09, HK08]. Capturing a focal stack with a large-scale high-resolution camera was implemented by Ben-Ezra [BE10]. Kutulakos and Hasinoff [KH09] proposed

Figure 8: *Multiple frames of a flying bird multiplexed into a single photograph (left). These kinds of photographs were shot with a photographic gun (right) by Etienne-Jules Marey as early as 1882.*

to multiplex a focal stack into a single sensor image in a similar fashion as CFAs multiplex different colour channels into a RAW camera image. However, to the knowledge of the authors, this camera has not yet been built.

Green *et al.* [GSMD07] split the aperture of a camera using circular mirrors and multiplex the result into different regions of a single photograph. In principle, this approach captures multiple frames with varying aperture settings at a reduced spatial resolution in a single snapshot.

Other applications for flexible focus imaging include 3D shape reconstruction with shape from (de)focus (see e.g. [NN94, ZLN09]) or *Confocal Stereo* [HK06, HK09], video matting [MMP*05] and extended DOF projection [GWGB10].

6. Multiplexing Time

Capturing motion and other forms of movement in photographs has been pursued since the invention of the daguerreotype. Early pioneers in this field include Eadweard Muybridge (e.g. [Muy57]) and Etienne-Jules Marey (e.g. [Bra92]). As illustrated in Figure 8, much of the early work on picturing time focused on the study of anatomy and locomotion of animals and humans; photographic apparatuses were usually custom built at that time (Figure 8, right). In this section, we discuss two classes of techniques for picturing motion: image capture at temporal resolutions that are significantly lower (Section 6.1) or higher (Section 6.2) than the resolving capabilities of the human visual system and approaches for joint optical and computational motion deblurring (Section 6.3).

6.1. Time lapse photography

Photographing scenes at very low temporal sampling rates is usually referred to as time lapse photography. Technically, time lapses can simply be acquired by taking multiple photographs from the same or a very close camera position at larger time intervals and assembling them in a video. In order to avoid temporal aliasing, or in simpler terms provide naturally looking motion, the exposure times should ideally be as long as the interval between successive shots. Timothy Allen, photographer for the BBC, provides a very informative tutorial on time lapse photography on his website [All10]. The BBC has produced a number of astounding time lapse videos, including many scenes in their Planet Earth and Life series.

6.2. High-speed imaging

Analogue high-speed film cameras have been developed throughout the last century. A variety of technologies exist that expose film at very high speeds including mechanical movement through temporally registered pins and rotating prisms or mirrors. For a detailed discussion of the history of high-speed photography, applications, and the state of the art about 9 years ago, the reader is referred to the book by Ray [Ray02].

Single sensor approaches

Today, high-speed digital cameras are commercially available. Examples are the Phantom Flex by Vision Research [Res10], which can capture up to 2570 frames per second (fps) at HD resolution, and the FASTCAM SA5 by Photron, which captures 7500 fps at megapixel resolution or up to one million frames per second at a reduced resolution (64×16 pixels) [Pho10]; both cameras employ CMOS sensor technology. A modified CCD is used in the HyperVision HPV-2 by Shimadzu [Shi10], which operates at one million fps for an image resolution of 312×260 pixels. The Dynamic Photomechanics Laboratory at the University of Rhode Island (mcise.uri.edu/dpml/facilities.html) houses an IMACON 468-MkII digital camera operating at 200 million fps, but exact specifications of that camera are unknown to the authors. With the introduction of Casio's Exilim camera series (exilim.casio.com), which records low resolution videos at up to 1000 fps, high-speed cameras have entered the consumer market. Temporal resampling and filtering of high-speed footage is explored by Fuchs *et al.* [FCW*10].

An alternative to high-speed sensors is provided by *Assorted Pixels* [NN05], where spatial resolution is traded for temporal resolution by measuring spatially interleaved, temporally staggered exposures on a sensor. This approach is very similar to what standard CFAs do to acquire colour information (see Section 3.2). While this concept was initially only theoretical, it has recently been implemented by aligning a DMD with a CCD sensor [BTH*10]. Alternatively, the sensor readout could be controlled on a per-pixel basis, as for instance provided by non-destructive sensor readout (e.g. [Sem10]). Reddy *et al.* [RVC11] built an LCOS-based camera prototype that modulates the exposure of each pixel randomly throughout the exposure time. In combination with a non-linear sparse reconstruction algorithm, the 25 fps prototype has been shown to capture imagery with up to 200 frames per second without loss of spatial resolution by exploiting sparsity in the spatio-temporal volume. Coded rolling shutters [GHMN10] have the potential to implement this concept on a per-scanline basis.

Agrawal *et al*. [AVR10] demonstrated how a pinhole in the aperture plane of a camera, which moves throughout the exposure time, allows the captured data to be adaptively re-interpreted. For this purpose, temporal light variation is directly encoded in the different views of the light field that is simultaneously acquired with a Sum-of-Sinusoids (SoS) attenuation-mask (see Section 4.3) in a single shot. Temporal variation and different viewpoints cannot be separated in this approach.

Multiple devices

Rather than photographing a scene with a single high-speed camera, multiple synchronized devices can be used. One of the most popular movie scenes that shows high-speed motion captured by a camera array is the bullet time effect in *The Matrix*. Here, a rig of digital SLR cameras, arranged along a virtual camera path, captures a scene at precisely controlled time steps so that a virtual, high-speed camera can be simulated that moves along the predefined path.

The direct capture of high-speed events with camera arrays was scientifically discussed by Wilburn *et al*. [WJV*04, WJV*05]. In this approach, the exposure windows of the cameras are slightly staggered so that a high-speed video can be composed by merging the data of the individual cameras. Shechtman *et al*. [SCI02, SCI05] proposed to combine the output of multiple low-resolution video cameras for space-time super-resolution. Coded exposures have been shown to optimize temporal super-resolution from multi-camera arrays [AGVN10] by alleviating the ill-posedness of the reconstruction. As required for spatial super-resolution (see Section 5.1), temporal super-resolution requires computationally expensive post-processing of the measured data.

High-speed illumination

High-speed imagery can also be acquired by utilizing high-speed illumination. Harold ‘Doc’ Edgerton [Pro09] created this field by inventing electronic strobes and using them to depict very fast motions in a similar fashion as Eadweard Muybridge and Etienne-Jules Marey had done with more primitive, mechanical technologies decades before him. Today, high-speed illumination, in an attosecond time scale, is more conveniently achieved with lasers rather than stroboscopes [BRH*06, Ray02].

Stroboscopic illumination can be used to compensate for rolling shutter effects and synchronize an array of consumer cameras [BAIH09]. Narasimhan *et al*. [NKY08] exploited the high-speed temporal dithering patterns of DLP-based illumination for a variety of vision problems, including photometric stereo and range imaging.

Compressive sensing of dynamic scenes

Coded strobing, by either illumination or controlled sensor readout, in combination with reconstructions developed in the compressive sensing community, allows high-speed periodic events to be acquired [VRR11]. Another high-speed imaging approach that is inspired by compressive sensing was proposed by Gupta *et al*. [GAVN10]. Here, a 3D spatio-temporal volume is adaptively encoded with a fixed voxel budget. This approach encodes fast motions with a high temporal, but lower spatial resolution, while the spatial resolution in static parts of the scene is maximized. A co-axial projector-camera pair was used to simulate controllable per-pixel exposures.

Exotic ultra high-speed imaging

Other imaging devices that capture ultra high-speed events are streak cameras. Rather than recording standard 2D photographs, these devices capture 2D images that encode spatial information in one dimension and temporal light variation in the other. These systems are usually combined with pulsed laser illumination and operate at temporal resolutions of about one hundred femtoseconds [Ham10], corresponding to a framerate of ten trillion fps. Another exotic ultra high-speed imager is the STREAM camera [GTJ09], which optically converts a 2D image into a serial time-domain waveform that is recorded with a single high-speed photodiode at 6.1 million fps.

6.3. Motion deblurring

Motion deblurring has been an active area of research over the last few decades. It is well known that deblurring is an ill-posed problem, which is why many algorithms apply regularizers [Ric72, Luc74] or natural image statistics (e.g. [FSH*06, LFDF07b]) to solve the problem robustly. Usually, high-frequency spatio-temporal image information is irreversibly lost in the image formation because a standard shutter along with the sensor integration time create a temporal rect-filter which has many zero-crossings in the Fourier domain. In this section, we review coded image acquisition techniques that optically modify the motion PSF so that the reconstruction becomes a well-posed problem.

Coded single capture approaches

One of the earliest approaches of coded temporal sampling was introduced by Raskar *et al*. [RAT06] as the *Fluttered*

Shutter. The motion PSF in a single sensor image was modified by mounting a programmable liquid crystal element in front of the camera lens and modulating its transmission over the exposure time with optimized binary codes. These codes were designed to preserve high temporal frequencies, so that the required image deconvolution becomes well-posed. Optimized codes and an algorithm for the problem of combined motion deblurring and spatial super-resolution of moving objects with coded exposures were analysed by Agrawal and Raskar [AR07]. Both approaches use programmable, attenuation-based shutters to apply the codes, thereby sacrificing light transmission, and require a manual rectification of object motion in the captured images. Optimality criteria for motion PSF invertibility were extended to also allow high-quality PSF estimation [AX09]; this automates the motion rectification step.

Inspired by approaches that create depth-invariant point spread functions (see Section 5.4), Levin *et al*. [LSC*08] showed how one-dimensional parabolic sensor motion during the exposure time can achieve a motion-invariant PSF along the line of sensor motion. Compared to attenuation-coded temporal exposures, this method does not sacrifice light transmission but requires prior knowledge of object motion and only works along one spatial direction. The optimal tradeoffs for single image deblurring from either attenuation-coded exposures or sensor motion, in terms of signal-to-noise ratio of the reconstructions, were analysed by Agrawal and Raskar [AR09].

Image sequences or multiple cameras

Synthetic Shutter Speed Imaging [TSY*07] combines multiple sharp but noisy images captured with short exposure times. The resulting image has a lower noise level; motion blur is reduced by aligning all images before they are fused.

A *Hybrid Camera* for motion deblurring, consisting of a rig of two cameras, was introduced by Ben-Ezra and Nayar [BEN03, BEN04]. One of the cameras captures the scene at a high temporal, but low spatial resolution; the output of this camera is used to estimate the motion PSF, which in turn is used to deblur the high-quality image captured by the other camera. Improvements of reconstructions for hybrid cameras have recently been presented [THBL10]. Hybrid camera architectures also provide the opportunity to simultaneously deblur captured images and reconstruct a high-resolution depth map of the photographed scene [LYC08].

Motion blur in a video can be synthetically removed by applying super-resolution techniques to multiple successive frames [BBZ96]. Agrawal *et al*. [AXR09] showed that improved results can be achieved by modulating the exposure times for successive frames in video sequences so that a reconstruction from multiple images becomes a well-posed problem. An example for this is shown in Figure 9.

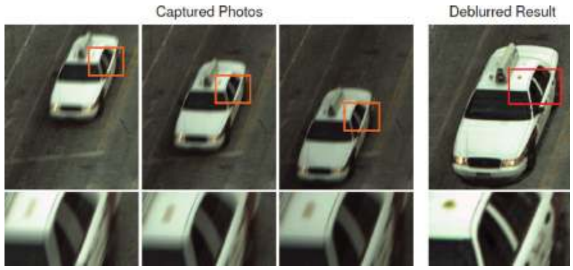

Figure 9: *By varying the exposure time for successive frames in a video (left), multi-image deblurring (right) can be made invertible [AXR09].*

7. Acquisition of Further Light Properties

In this section, we review acquisition approaches for light properties that are not considered dimensions of the plenoptic function, but are closely related in terms of capture or application. These properties include polarization, phase and time-of-flight, which are all attributes that can be associated with individual light rays in addition to the plenoptic dimensions.

7.1. Polarization

Polarization is an inherent property of the wave nature of light [Col05]. Generally, polarization describes the oscillation of a wave travelling through space in the transverse plane, perpendicular to the direction of propagation. Linear polarization refers to transverse oscillation along a line, whereas spherical or elliptical polarization describe corresponding oscillation trajectories.

Although some animals, including mantis shrimp [MO99], cephalopods (squid, octopus, cuttlefish) [MSH09] and insects [Weh76], are reported to have photoreceptors that are sensitive to polarization, standard solid state sensors are not. The most straightforward way of capturing this information is by taking multiple photographs of a scene with different polarizing filters mounted in front of the camera lens. These filters are standard practice in photography to reduce specular reflections, increase the contrast of outdoor images and improve the appearance of vegetation. Alternatively, this kind of information can be captured using polarization filter arrays [SN03b] which, similar to generalized mosaics [SN05], require multiple photographs to be captured. Recently, polarized illumination [GCP*10] has been shown to have the potential to acquire all Stokes parameters necessary to describe polarization.

Applications for the acquisition of polarized light include shape [MTHI03, MKI04, AH05, WRHR11] and BRDF [AH08] estimation, image dehazing [SNN01, SNN03, NS05], improved underwater vision [SK04, SK05], specular highlight removal [WB91, NFB93, Mül96, UG04], material

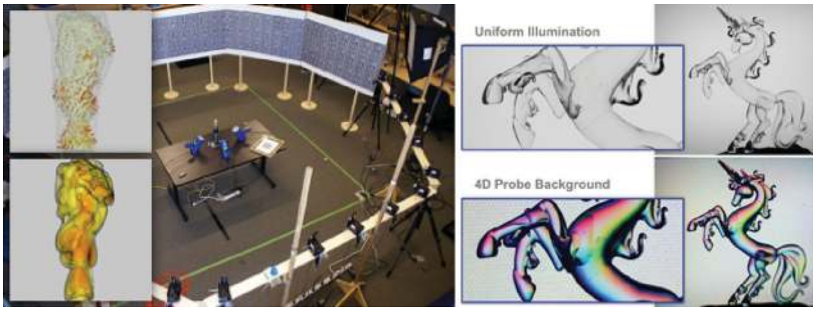

Figure 10: *Schlieren imaging for tomographic gas reconstruction (left) and capture of refractive events using light field probes (right). (Figures reproduced from [AIH*08] and [WRH11].)*

classification [CW98], light source separation [CDPW07], surface normal acquisition [MHP*07], surface normal and refractive index estimation [Sad07, GCP*10], separation of transparent layers [SSK99] and optical communication [SZ91, Sch03, Yao08].

7.2. Phase imaging

A variety of techniques has been proposed to visualize and quantify phase retardation in transparent microscopic organisms [Mur01]. Many of these phase-contrast imaging approaches, such as Zernike phase contrast and differential interference contrast, require coherent illumination and are qualitative rather than quantitative. This implies that changes in phase or refractive events are encoded as intensity variations in captured images, but remain indistinguishable from the intensity variations caused by absorption in the medium. Quantitative approaches exist [BNBN02], but require multiple images, are subject to a paraxial approximation, and are limited to orthographic cameras.

Schlieren and shadowgraph photography are alternative, non-intrusive imaging methods for dynamically changing refractive index fields. These techniques have been developed in the fluid imaging community over the past century, with substantial improvements in the 1940s. An extensive overview of different optical set-ups and the historic evolution of Schlieren and Shadowgraph imaging can be found in the book by Settles [Set01]. As illustrated in Figure 10, recently proposed applications of Schlieren imaging include the tomographic reconstruction of transparent gas flows using a camera array [AIH*08] and the capture of refractive events with 4D light field probes [WRH11, WRHR11].

7.3. LIDAR and time-of-flight

LIDAR (LIght Detection and Ranging) [Wan05] is a technology that measures the time of a laser pulse from transmission to detection of the reflected signal. It is similar to radar, but uses different wavelengths of the electromagnetic spectrum, typically in the infrared range. Combining such a pulsed laser with optical scanners and a positioning system such as GPS allows very precise depth or range maps to be captured, even from airplanes. Overlaying range data with standard photographs provides a powerful tool for aerial surveying, forestry, oceanography, agriculture, and geology. Flash LIDAR [LS01] or time-of-flight cameras [KBKL10] capture a photograph and a range map simultaneously for all pixels. Although spatial resolution of the range data is often poor, these cameras usually capture at video rates.

Streak cameras operate in the picosecond [CS83] or even attosecond [IQY*02] range and usually capture 2D images, where one dimension is spatial light variation and the other dimension is time-of-flight. These cameras have recently been used to reveal scene information outside the line of sight of a camera, literally behind corners [KHDR09, PVB*11].

8. Discussion and Conclusions

In summary, we have presented a review of approaches to plenoptic image acquisition. We have used an intuitive categorization based on plenoptic dimensions and hardware set-ups for the acquisition. Alternative categorizations may be convenient for the discussion of the more general field of computational photography [RT09]. The increasingly growing number of publications in this field is one of the main motivations for this state-of-the-art report, which focuses specifically on joint optical encoding and computational reconstruction approaches for the acquisition of the plenoptic function.

Based on the literature reviewed in this report, we make the following observations:

- most of the discussed approaches either assume that some plenoptic dimensions are constant, such as time in sequential image capture, or otherwise restricted, for instance spatially band-limited in single sensor interleaved capture; these assumptions result in fixed plenoptic resolution tradeoffs,
- however, there are strong correlations between the dimensions of the plenoptic function; these are, for instance,

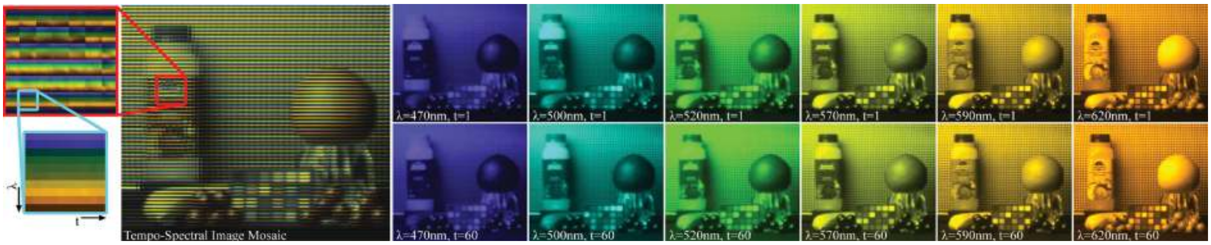

Figure 11: *Dataset containing multi-spectral variation and controlled object movement in the scene. Left: image mosaic illustrating the correlation between the spectral and temporal plenoptic dimension of natural scenes. Right: six spectral, colour-coded slices of the dataset with two temporal snapshots each. We recorded these datasets using a custom multi-spectral camera that consists of collimating optics, a liquid crystal tunable filter and a USB machine vision camera.*

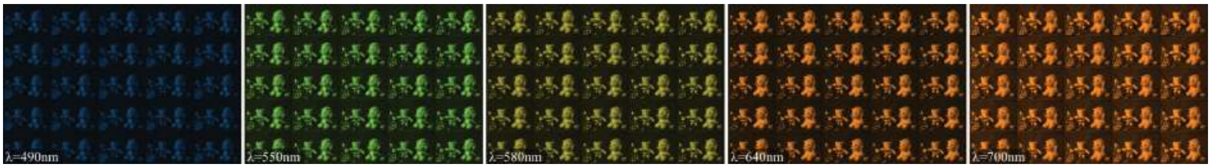

Figure 12: *Five colour channels of a multi-spectral light field with 5 × 5 viewpoints and 10 colour channels for each viewpoint. This dataset was captured by mounting our multi-spectral camera on an X-Y translation stage and sequentially capturing the spectral bands for each viewpoint.*

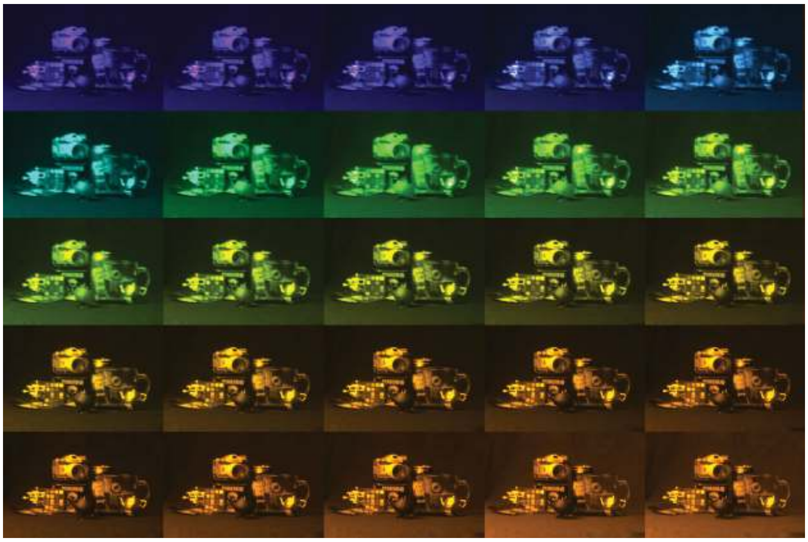

Figure 13: *Another multi-spectral light field dataset with 15 × 15 viewpoints and 23 narrow-band colour channels for each viewpoint. The spectral channels range from 460 nm to 680 nm in 10 nm increments. Only 5 × 5 viewpoints are shown in this mosaic and each of those is colour-coded. The photographed scene includes a variety of illumination effects including diffraction, refraction, inter-reflections and specularities.*

exploited in colour demosaicing, optical flow estimation, image and video compression, and stereo reconstruction,

- therefore, natural image statistics that can be used as priors in computational image acquisition, incorporating all plenoptic dimensions with their correlations, are desirable; so are sophisticated reconstruction techniques employing these.

It has recently been shown that all approaches for interleaved plenoptic sampling on a single sensor, including spatial [NN05] and Fourier multiplexing [VRA*07, LRAT08] methods, can be cast into a common reconstruction framework [IWH10]. While the exploitation of correlations between plenoptic dimensions, for example spatial and spectral light variation, is common practice for imaging with CFAs and subsequent demosaicing, there is significant potential to develop similar techniques for demosaicing other multiplexed plenoptic information, for instance light fields [LD10].

Priors for the correlations between plenoptic dimensions can be very useful for plenoptic super-resolution or generally more sophisticated reconstructions. These could, for instance, be derived from plenoptic image databases [WIGH11]; we show examples of such data in Figures 11, 12 and 13.

Another promising avenue of future research is adaptive imaging. Precise control of the sampled plenoptic information is the key for flexible and adaptive reconstruction. An intuitive next step for sophisticated imaging with respect to temporal light variation and dynamic range is pixel-precise, non-destructive sensor readout. In the future, however, it is desirable to being able to control the optical modulation of all plenoptic dimensions.

While most of the reviewed approaches make fixed plenoptic resolution tradeoffs, some already show a glimpse of the potential of adaptive re-interpretation of captured data [AVR10]. Ideas from the compressive sensing community (see e.g. [CRT06]) have also started to play an important role in adaptive plenoptic imaging [GAVN10, VRR11]. In these approaches optical coding is combined with content adaptive reconstructions that can dynamically trade higher-dimensional resolution in post-processing to best represent the recorded data.

Picturing space, colour, time, directions and other light properties has been of great interest to science and art alike for centuries. With the emergence of digital light sensing technology and computational processing power, many new and exciting ways to acquire some of the visual richness surrounding us have been presented. We have, however, only begun to realize how new technologies allow us to transcend the way evolution has shaped visual perception for different creatures on this planet.

References

[AA04] Aggarwal M., Ahuja N.: Split aperture imaging for high dynamic range. *International Journal Computer Vision 58*, 1 (2004), 7–17.

[AB91] Adelson E. H., Bergen J. R.: The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing* (1991), MIT Press, Cambridge, MA, pp. 3–20.

[AGVN10] Agrawal A., Gupta M., Veeraraghavan A., Narasimhan S.: Optimal coded sampling for temporal super-resolution. In *Proceedings of the IEEE CVPR* (2010), pp. 374–380.

[AH05] Atkinson G., Hancock E.: Multi-view surface reconstruction using polarization. In *Proceedings of ICCV* (2005), vol. 1, pp. 309–316.

[AH08] Atkinson G. A., Hancock E. R.: Two-dimensional BRDF estimation from polarisation. *Computer Vision and Image Understanding 111*, 2 (2008), 126–141.

[AIH*08] Atcheson B., Ihrke I., Heidrich W., Tevs A., Bradley D., Magnor M., Seidel H.: Time-resolved 3D capture of non-stationary gas flows. *ACM Transactions on Graphics (SIGGRAPH Asia) 27*, 5 (2008), 132, pp. 1–9.

[All10] Allen T.: Time Lapse Tutorial, 2010. `http://timothyallen.blogs.bbcearth.com/2009/02/24/time-lapse-photography/`. Accessed 9 October 2011.

[AN07] Ashok A., Neifeld M. A.: Pseudorandom phase masks for superresolution imaging from subpixel shifting. *Applied Optics 46*, 12 (2007), 2256–2268.

[AN10] Ashok A., Neifeld M. A.: Compressive light field imaging. In *Proceedings of SPIE 7690* (2010), p. 76900Q.

[AR07] Agrawal A., Raskar R.: Resolving objects at higher resolution from a single motion-blurred image. In *Proceedings of the IEEE CVPR* (2007), pp. 1–8.

[AR09] Agrawal A., Raskar R.: Optimal single image capture for motion deblurring. In *Proceedings of the IEEE CVPR* (2009), pp. 1–8.

[Ash93] Ashdown I.: Near-field photometry: A new approach. *Journal of the Illuminating Engineering Society 22*, 1 (1993), 163–180.

[ASH05] Alleyson D., süsstrunk S., hérault J.: Linear demosaicing inspired by the human visual system. *IEEE Transactions on Image Processing 14*, 4 (2005), 439–449.

[ATP*10] Adams A., Talvala E.-V., Park S. H., Jacobs D. E., Ajdin B., Gelfand N., Dolson J., Vaquero D., Baek J., Tico M., Lensch H. P. A., Matusik W., Pulli K., Horowitz M., Levoy M.: The Frankencamera: An experimental platform for computational photography. *ACM Transactions on Graphics (SIGGRAPH) 29* (2010), 29:1–29:12.

[AVR10] Agrawal A., Veeraraghavan A., Raskar R.: Reinterpretable imager: Towards variable post-capture space, angle and time resolution in photography. In *Proceedings of the Eurographics* (2010), pp. 1–10.

[AW92] Adelson E., Wang J.: Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence 14*, 2 (1992), 99–106.

[AX09] Agrawal A., Xu Y.: Coded exposure deblurring: Optimized codes for PSF estimation and invertibility. In *Procedings of the IEEE CVPR* (2009), pp. 1–8.

[AXR09] Agrawal A., Xu Y., Raskar R.: Invertible motion blur in video. *ACM Transactions on Graphics (Siggraph) 28*, 3 (2009), 95.

[Bae10] Baek J.: Transfer efficiency and depth invariance in computational cameras. In *Proceedings of the ICCP* (2010), pp. 1–8.

[BAIH09] Bradley D., Atcheson B., Ihrke I., Heidrich W.: Synchronization and rolling shutter compensation for consumer video camera arrays. In *Proceedings of ProCams* (2009), pp. 1–8.

[BAL*09] Babacan S. D., Ansorge R., Luessi M., Molina R., Katsaggelos A. K.: Compressive sensing of light fields. In *Proceedings of ICIP* (2009), pp. 2313–2316.

[Bay76] Bayer B. E.: Color imaging array, 1976. US Patent 3,971,065.

[BBM87] Bolles R. C., Baker H. H., Marimont D. H.: Epipolar-plane image analysis: An approach to determining structure from motion. *International Journal of Computer Vision 1*, 1 (1987), 7–55.

[BBZ96] Bascle B., Blake A., Zisserman A.: Motion deblurring and super-resolution from an image sequence. In *Proceedings of ECCV* (1996), pp. 573–582.

[BE10] Ben-Ezra M.: High resolution large format tile-scan camera. In *Proceedings of IEEE ICCP* (2010), pp. 1–8.

[BE11] Ben-Ezra M.: A digital gigapixel large-format tile-scan camera. *IEEE Computer Graphics and Applications 31* (2011), 49–61.

[BEMKZ05] Ben-Eliezer E., Marom E., Konforti N., Zalevsky Z.: Experimental realization of an imaging system with an extended depth of field. *Applied Optics 44*, 11 (2005), 2792–2798.

[BEN03] Ben-Ezra M., Nayar S.: Motion deblurring using hybrid imaging. In *Proceedings of IEEE CVPR* (2003), pp. 657–664.

[BEN04] Ben-Ezra M., Nayar S.: Motion-based motion deblurring. *IEEE Transactions on Pattern Analysis and Machine Intelligence 26*, 6 (2004), 689–698.

[BEZN05] Ben-Ezra M., Zomet A., Nayar S.: Video super-resolution using controlled subpixel detector shifts. *IEEE Transactions on Pattern Analysis and Machine Intelligence 27*, 6 (2005), 977–987.

[BH09] Brady D. J., Hagen N.: Multiscale lens design. *Optics Express 17*, 13 (2009), 10659–10674.

[BK02] Baker S., Kanade T.: Limits on super-resolution and how to break them. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24* (2002), 1167–1183.

[BL07] Brown M., Lowe D.: Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision 74*, 1 (2007), 59–73.

[BNBN02] Barone-Nugent E. D., Barty A., Nugent K. A.: Quantitative phase-amplitude microscopy I: Optical Microscopy. *Journal of Microscopy 206*, 3 (2002), 194–203.

[Bra92] Braun M.: *Picturing Time: The Work of Etienne-Jules Marey (1830-1904)*. The University of Chicago Press, 1992.

[BRH*06] Baker S., Robinson J. S., Haworth C. A., Teng H., Smith R. A., Chirila C. C., Lein M., Tisch J. W. G., Marangos J. P.: Probing proton dynamics in molecules on an attosecond time scale. *Science 312*, 5772 (2006), 424–427.

[BS98] Borman S., Stevenson R.: Super-resolution from image sequences - A review. In *Proceedings of the Symposium on Circuits and Systems* (1998), pp. 374–378.

[BSN*09] Bodkin A., Sheinis A., Norton A., Daly J., Beaven S., Weinheimer J.: Snapshot hyperspectral imaging - the hyperspectral array camera. In *Proceedings of SPIE 7334* (2009), pp. 1–11.

[BTH*10] Bub G., Tecza M., Helmes M., Lee P., Kohl P.: Temporal pixel multiplexing for simultaneous high-speed, high-resolution imaging. *Nature Methods 7* (2010), 209–211.

[BZF09] Bishop T., Zanetti S., Favaro P.: Light-field super-resolution. In *Proceedings of ICCP* (2009), pp. 1–9.

[CCG06] Chi W., Chu K., George N.: Polarization coded aperture. *Optics Express 14*, 15 (2006), 6634–6642.

[CDPW07] Cula O. G., Dana K. J., Pai D. K., Wang D.: Polarization multiplexing and demultiplexing for appearance-based modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*, 2 (2007), 362–367.

[CG01] Chi W., George N.: Electronic imaging using a logarithmic asphere. *Optics Letters 26*, 12 (2001), 875–877.

[ci09] Cri Inc: VariSpec Liquid Crystal Tunable Filters. www.cri-inc.com/varispec, 2009. Accessed 9 October 2011.

[CMN11] Cossairt O., Miau D., Nayar S. K.: Gigapixel computational imaging. In *Proceedings of ICCP* (2011).

[CN10] Cossairt O., Nayar S. K.: Spectral focal sweep: Extended depth of field from chromatic aberrations. In *Proceedings of ICCP* (2010), pp. 1–8.

[Col05] Collett E.: *Field Guide to Polarization*. SPIE Press, Bellingham, WA, 2005.

[CRT06] Candès E., Romberg J., Tao T.: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory 52*, 2 (2006), 489–509.

[CS83] Campillo A., Shapiro S.: Picosecond streak camera fluorometry-A review. *Journal of Quantum Electronics 19*, 4 (1983), 585–603.

[CTDL11] Cao X., Tong X., Dai Q., Lin S.: High resolution multispectral video capture with a hybrid camera system. In *Proceedings of the IEEE CVPR* (2011), pp. 1–8.

[CTMS03] Carranza J., Theobalt C., Magnor M. A., Seidel H.-P.: Free-viewpoint video of human actors. *ACM Transactions on Graphics 22*, 3 (2003), 569–577.

[CW93] Chen S. E., Williams L.: View interpolation for image synthesis. In *Proceedings of the ACM SIGGRAPH* (1993), pp. 279–288.

[CW98] Chen H., Wolff L. B.: Polarization phase-based method for material classification in computer vision. *International Journal of Computer Vision 28*, 1 (1998), 73–83.

[CZN10] Cossairt O., Zhou C., Nayar S. K.: Diffusion coded photography for extended depth of field. *ACM Transactions on Graphics (Siggraph) 29*, 3 (2010), 31.

[DC95] Dowski E., Cathey T.: Extended depth of field through wave-front coding. *Applied Optics 34*, 11 (1995), 1859–1866.

[DC*97] Descour M. E., Volin C. E., Dereniak E., Thome K. J.: Demonstration of a high-speed nonscanning imaging spectrometer. *Optics Letters 22*, 16 (1997), 1271–1273.

[DD95] Descour M., Dereniak E.: Computed-tomography imaging spectrometer: Experimental calibration and reconstruction results. *Applied Optics 34*, 22 (1995), 4817–4826.

[Deb02] Debevec P.: Image-based lighting. *IEEE Computer Graphics and Applications* (2002), 26–34.

[DHT*00] Debevec P., Hawkins T., Tchou C., Duiker H.-P., Sarokin W., Sagar M.: Acquiring the reflectance field of a human face. In *Proceedings of the ACM SIGGRAPH* (2000), pp. 145–156.

[DM97] Debevec P. E., Malik J.: Recovering high dynamic range radiance maps from photographs. In *Proceedings of the ACM Siggraph* (1997), pp. 369–378.

[DTCL09] Du H., Tong X., Cao X., Lin S.: A prism-based system for multispectral video acquisition. In *Proceedings of the IEEE ICCV* (2009), pp. 175–182.

[Eas] Eastman Kodak Company: PhotoCD PCD0992. http://r0k.us/graphics/kodak. Accessed 9 October 2011.

[Fai05] Fairchild M. D.: *Color Appearance Models*. Addison-Wesley, Reading, MA, 2005.

[FCW*10] Fuchs M., Chen T., Wang O., Raskar R., Seidel H.-P., Lensch H. P.: Real-time temporal shaping of high-speed video streams. *Computers & Graphics 34* (2010), 575–584.

[FHH05] Fletcher-Holmes D. W., Harvey A. R.: Real-time imaging with a hyperspectral fovea. *Journal of Optics A: Pure and Applied Optics 7* (2005), S298–S302.

[Fov10] Foveon: X3 Technology, 2010. www.foveon.com. Accessed 9 October 2011.

[FSH*06] Fergus R., Singh B., Hertzmann A., Roweis S. T., Freeman W. T.: Removing camera shake from a single photograph. *ACM Transactions on Graphics 25* (2006), 787–794.

[FTF06] Fergus R., Torralba A., Freeman W. T.: *Random Lens Imaging*. Tech. Rep. MIT-CSAIL-TR-2006-058, National Bureau of Standards, 2006.

[Gab48] Gabor D.: A new microscopic principle. *Nature* (1948), 777–778.

[Gat00] Gat N.: Imaging spectroscopy using tunable filters: A review. In *Proceedings of SPIE 4056* (2000), pp. 50–64.

[GAVN10] Gupta M., Agrawal A., Veeraraghavan A., Narasimhan S. G.: Flexible voxels for motion-aware videography. In *Proceedings of the ECCV* (2010), pp. 100–114.

[GAW*10] Granados M., Ajdin B., Wand M., Theobalt C., Seidel H.-P., Lensch H. P. A.: Optimal HDR reconstruction with linear digital cameras. In *Proceedings of the CVPR* (2010), pp. 215–222.

[GCP*10] GHOSH A., CHEN T., PEERS P., WILSON C. A., DEBEVEC P.: Circularly polarized spherical illumination reflectometry. *ACM Transactions on Graphics (Siggraph Asia) 27*, 5 (2010), 134.

[Ger36] GERSHUN A.: The light field. *Journal of Mathematics and Physics XVIII* (1936), 51–151. Translated by P. Moon and G. Timoshenko.

[GF89] GOTTESMAN S. R., FENIMORE E. E.: New family of binary arrays for coded aperture imaging. *Applied Optics 28*, 20 (1989), 4344–4352.

[GFHH10] GORMAN A., FLETCHER-HOLMES D. W., HARVEY A. R.: Generalization of the Lyot filter and its application to snapshot spectral imaging. *Optics Express 18*, 6 (2010), 5602–5608.

[GGA*05] GUNTURK B., GLOTZBACH J., ALTUNBASAK Y., SCHAFER R., MERSEREAU R.: Demosaicking: Color filter array interpolation in single-chip digital cameras. *IEEE Signal Processing 22*, 1 (2005), 44–54.

[GGML07] GARCIA-GUERRERO E. E., MENDEZ E. R., LESKOVA H. M.: Design and fabrication of random phase diffusers for extending the depth of focus. *Optics Express 15*, 3 (2007), 910–923.

[GGSC96] GORTLER S., GRZESZCZUK R., SZELINSKI R., COHEN M.: The lumigraph. In *Proceedings of the ACM Siggraph* (1996), pp. 43–54.

[GHMN10] GU J., HITOMI Y., MITSUNAGA T., NAYAR S. K.: Coded rolling shutter photography: Flexible space-time sampling. In *Proceedings of the IEEE ICCP* (2010), pp. 1–8.

[GIBL08] GEORGIEV T., INTWALA C., BABACAN S., LUMSDAINE A.: Unified frequency domain analysis of lightfield cameras. In *Proceedings of the ECCV* (2008), pp. 224–237.

[GJB*07] GEHM M. E., JOHN R., BRADY D. J., WILLETT R. M., SCHULZ T. J.: Single-shot compressive spectral imaging with a dual-disperser architecture. *Optics Express 15*, 21 (2007), 14013–14027.

[GKT09] GAO L., KESTER R. T., TKACZYK T. S.: Compact image slicing spectrometer (ISS) for hyperspectral fluorescence microscopy. *Optics Express 17*, 15 (2009), 12293–12308.

[GN03] GROSSBERG M. D., NAYAR S. K.: High dynamic range from multiple images: Which exposures to combine. In *Proceedings of the ICCV Workshop CPMCV* (2003).

[GSMD07] GREEN P., SUN W., MATUSIK W., DURAND F.: Multi-aperture photography. In *Proceedings of the ACM Siggraph* (2007), p. 68.

[GTJ09] GODA K., TSIA K. K., JALALI B.: Serial time-encoded amplified imaging for real-time observation of fast dynamic phenomena. *Nature*, *458* (2009), 1145–1149.

[GWGB10] GROSSE M., WETZSTEIN G., GRUNDHÖFER A., BIMBER O.: Coded aperture projection. *ACM Transactions on Graphics 29* (2010), 22:1–22:12.

[GWHD09] GAIGALAS A. K., WANG L., HE H.-J., DEROSE P.: Procedures for wavelength calibration and spectral response correction of CCD array spectrometers. *Journal of Research of the National Institute of Standards and Technology 114*, 4 (2009), 215–228.

[GZN*06] GEORGIEV T., ZHENG C., NAYAR S., CURLESS B., SALESIN D., INTWALA C.: Spatio-angular resolution trade offs in integral photography. In *Proceedings of the EGSR* (2006), pp. 263–272.

[Hal94] HALLE M. W.: Holographic stereograms as discrete imaging systems. In *SPIE Practical Holography* (1994), pp. 73–84.

[Ham10] HAMAMATSU: Streak Systems, 2010. `http://sales.hamamatsu.com/en/products/system-division/ultra-fast/streak-systems.php`. Accessed 9 October 2011.

[Häu72] HÄUSLER G.: A method to increase the depth of focus by two step image processing. *Optics Communications 6*, 1 (1972), 38–42.

[HBG*00] HARVEY A. R., BEALE J., GREENAWAY A. H., HANLON T. J., WILLIAMS J.: Technology options for imaging spectrometry imaging spectrometry. In *Proceedings of SPIE 4132* (2000), pp. 13–24.

[HDF10] HASINOFF S. W., DURAND F., FREEMAN W. T.: Noise-optimal capture for high dynamic range photography. In *Proceedings of the IEEE CVPR* (2010), pp. 1–8.

[HEAL09] HORSTMEYER R., EULISS G., ATHALE R., LEVOY M.: Flexible multimodal camera using a light field architecture. In *Proceedings of the ICCP* (2009), pp. 1–8.

[Hec02] HECHT E.: *Optics* (4th edition). Addison Wesley, 2002.

[HFHG05] HARVEY A. R., FLETCHER-HOLMES D. W., GORMAN A.: Spectral imaging in a snapshot. In *Proceedings of SPIE 5694* (2005), pp. 1–10.

[HK06] HASINOFF S. W., KUTULAKOS K. N.: Confocal stereo. In *Proceedings of the ECCV* (2006), pp. 620–634.

[HK08] Hasinoff S. W., Kutulakos K. N.: Light-efficient photography. In *Proceedings of the ECCV* (2008), pp. 45–59.

[HK09] Hasinoff S. W., Kutulakos K. N.: Confocal stereo. *International Journal of Computer Vision 81*, 1 (2009), 82–104.

[HKDF09] Hasinoff S. W., Kutulakos K. N., Durand F., Freeman W. T.: Time-constrained photography. In *Proceedings of the ICCV* (2009), pp. 333–340.

[HLHR09] Hirsch M., Lanman D., Holtzman H., Raskar R.: BiDi screen: A thin, depth-sensing LCD for 3D interaction using light fields. In *ACM Trans. Graph. (SIGGRAPH Asia)* (2009), pp. 1–9.

[HMR09] Hiura S., Mohan A., Raskar R.: Krill-eye: Superposition compound eye for wide angle imaging via GRIN lenses. In *Proceedings of the OMNIVIS* (2009), pp. 1–8.

[HN06] Hanrahan P., Ng R.: Digital correction of lens aberrations in light field photography. In *International Optical Design Conference* (2006), pp. 1–3.

[HP01] Hunicz J., Piernikarski D.: Investigation of combustion in a gasoline engine using spectrophotometric methods. In *Proceedings of SPIE 4516* (2001), pp. 307–314.

[HP06] Hirakawa K., Parks T. W.: Joint demosaicing and denoising. *IEEE Transactions on Image Processing 15*, 8 (2006), 2146–2157.

[Hun91] Hunt R. W. G.: *Measuring Color* (3rd edition). Fountain Press, 1991.

[Hun04] Hunt R. W. G.: *The Reproduction of Color* (6th edition). John Wiley and Sons, Hoboken, NJ, 2004.

[HW07] Hirakawa K., Wolfe P.: Spatio-spectral color filter array design for enhanced image fidelity. In *Proceedings of the ICIP* (2007), pp. II-81–II-84.

[HW08] Hirakawa K., Wolfe P.: Spatio-spectral color filter array design for optimal image recovery. *IEEE Transactions on Image Processing 17*, 10 (2008), 1876–1890.

[IKL*10] Ihrke I., Kutulakos K. N., Lensch H. P. A., Magnor M., Heidrich W.: Transparent and specular object reconstruction. *Computer Graphics Forum 29*, 8 (2010), 2400–2426.

[Ind73] Inderhees J.: Optical field curvature corrector. US patent 3,720,454, 1973.

[IQY*02] Itatani J., Quéré F., Yudin G. L., Ivanov M. Y., Krausz F., Corkum P. B.: Attosecond streak camera. *Physical Review Letters 88*, 17 (2002), 173903.

[ISG*08] Ihrke I., Stich T., Gottschlich H., Magnor M., Seidel H.: Fast incident light field acquisition and rendering. In *Proceedings of WSCG* (2008), pp. 177–184.

[Ive03] Ives H.: Parallax Stereogram and Process of Making Same. US patent 725,567, 1903.

[Ive28] Ives H.: Camera for making parallax panoramagrams. *Journal of the Optical Society of America 17* (1928), 435–439.

[IWH10] Ihrke I., Wetzstein G., Heidrich W.: A theory of plenoptic multiplexing. In *Proceedings of the IEEE CVPR* (2010), pp. 1–8.

[Kan18] Kanolt C. W.: Parallax Panoramagrams. US patent 1,260,682, 1918.

[KBKL10] Kolb A., Barth E., Koch R., Larsen R.: Time-of-flight cameras in computer graphics. *Computer Graphics Forum 29*, 1 (2010), 141–159.

[KH57] Kapany N. S., Hopkins R. E.: Fiber optics. Part III. Field flatteners. *Journal of the Optical Society of America 47*, 7 (1957), 594–595.

[KH09] Kutulakos K. N., Hasinoff S. W.: Focal stack photography: High-performance photography with a conventional camera. In *Proceedings of the IAPR Conference on Machine Vision Applications* (2009), pp. 332–337.

[KHDR09] Kirmani A., Hutchison T., Davis J., Raskar R.: Looking around the corner using transient imaging. In *Proceedings of the ICCV* (2009), pp. 1–8.

[KUDC07] Kopf J., Uyttendaele M., Deussen O., Cohen M. F.: Capturing and viewing gigapixel images. *ACM Transactions on Graphics (SIGGRAPH) 26*, 3 (2007).

[KUWS03] Kang S. B., Uyttendaele M., Winder S., Szeliski R.: High dynamic range video. In *Proceedings of the ACM Siggraph* (2003), pp. 319–325.

[KYN*00] Kindzelskii A. L., Yang Z. Y., Nabel G. J., Todd R. F., Petty H. R.: Ebola virus secretory glycoprotein (sGP) diminishes Fc gamma RIIIB-to-CR3 proximity on neutrophils. *Journal of Immunology 164* (2000), 953–958.

[Lan10] Lanman D.: Mask-based Light Field Capture and Display. Brown University, School of Engineering, 2010.

[LCV*04] Levoy M., Chen B., Vaish V., Horowitz M., McDowall I., Bolas M.: Synthetic aperture confocal imaging. *ACM Transactions on Graphics (SIGGRAPH) 23*, 3 (2004), 825–834.

[LD10] Levin A., Durand F.: Linear view synthesis using a dimensionality gap light field prior. In *Proceedings of the IEEE CVPR* (2010), pp. 1–8.

[Lev10] LEVOY M.: Computational Photography and the Stanford Frankencamera. Technical Talk, 2010. www.graphics.stanford.edu/talks/. Accessed 9 October 2011.

[LFDF07a] LEVIN A., FERGUS R., DURAND F., FREEMAN W.: Image and depth from a conventional camera with a coded aperture. *ACM Transactions on Graphics (Siggraph) 26*, 3 (2007), 70.

[LFDF07b] LEVIN A., FERGUS R., DURAND F., FREEMAN W. T.: Deconvolution using Natural Image Priors, 2007. groups.csail.mit.edu/graphics/CodedAperture/SparseDeconv-LevinEtAl07.pdf. Accessed 9 October 2011.

[LFHHM02] LAWLOR J., FLETCHER-HOLMES D. W., HARVEY A. R., MCNAUGHT A. I.: In vivo hyperspectral imaging of human retina and optic disc. *Investigative Ophthalmology and Visual Science 43* (2002), 4350.

[LG09] LUMSDAINE A., GEORGIEV T.: The focused plenoptic camera. In *Proceedings of the ICCP* (2009), pp. 1–8.

[LGZ08] LI X., GUNTURK B., ZHANG L.: Image demosaicing: A systematic survey. In *SPIE Conference on Visual Communication and Image Processing* (2008), pp. 68221J–68221J–15.

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proceedings of the ACM Siggraph* (1996), pp. 31–42.

[LHG*09] LEVIN A., HASINOFF S. W., GREEN P., DURAND F., FREEMAN W. T.: 4D frequency analysis of computational cameras for depth of field extension. *ACM Transactions on Graphics (Siggraph) 28*, 3 (2009), 97.

[LHKR10] LANMAN D., HIRSCH M., KIM Y., RASKAR R.: Content-adaptive parallax barriers: Optimizing dual-layer 3D displays using low-rank light field factorization. *ACM Transactions on Graphics (Siggraph Asia) 28*, 5 (2010), 1–10.

[Lip08] LIPPMANN G.: La Photographie Intégrale. *Academie des Sciences 146* (1908), 446–451.

[LLW*08] LIANG C.-K., LIN T.-H., WONG B.-Y., LIU C., CHEN H. H.: Programmable aperture photography: Multiplexed light field acquisition. *ACM Transactions on Graphics (Siggraph) 27*, 3 (2008), 1–10.

[LMK01] LANDOLT O., MITROS A., KOCH C.: Visual sensor with resolution enhancement by mechanical vibrations. In *Proceedings of the Advanced Research in VLSI* (2001), pp. 249–264.

[LNA*06] LEVOY M., NG R., ADAMS A., FOOTER M., HOROWITZ M.: Light field microscopy. *ACM Transactions on Graphics (Siggraph) 25*, 3 (2006), 924–934.

[LRAT08] LANMAN D., RASKAR R., AGRAWAL A., TAUBIN G.: Shield fields: Modeling and capturing 3D occluders. *ACM Transactions on Graphics (Siggraph Asia) 27*, 5 (2008), 131.

[LS01] LANGE R., SEITZ P.: Solid-state time-of-flight range camera. *Journal of Quantum Electronics 37*, 3 (2001), 390–397.

[LS11] LIU C., SUN D.: A Bayesian approach to adaptive video super resolution. In *Proceedings of the IEEE CVPR* (2011), pp. 1–8.

[LSC*08] LEVIN A., SAND P., CHO T. S., DURAND F., FREEMAN W. T.: Motion-invariant photography. *ACM Transactions on Graphics (Siggraph) 27*, 3 (2008), 71.

[Luc74] LUCY L. B.: An iterative technique for the rectification of observed distributions. *The Astronomical Journal 79* (1974), 745–754.

[LV09] LU Y. M., VETTERLI M.: Optimal color filter array design: Quantitative conditions and an efficient search procedure. In *Proceedings of the SPIE 7250* (2009), pp. 1–8.

[LWTC06] LANMAN D., WACHS M., TAUBIN G., CUKIERMAN F.: Spherical catadioptric arrays: Construction, multi-view geometry, and calibration. In *Proceedings of the 3DPVT* (2006), pp. 81–88.

[LY05] LAU D. L., YANG R.: Real-time multispectral color video synthesis using an array of commodity cameras. *Real-Time Imaging 11*, 2 (2005), 109–116.

[LYC08] LI F., YU J., CHAI J.: A hybrid camera for motion deblurring and depth map super-resolution. In *Proceedings of the IEEE CVPR* (2008), pp. 1–8.

[Lyo44] LYOT B.: Le filtre monochromatique polarisant et ses applications en physique solaire. *Annales d'Astrophysique 7* (1944), 31.

[Man05] MANSFIELD C. L.: Seeing into the past, 2005. www.nasa.gov/vision/earth/technologies/scrolls.html. Accessed 9 October 2011.

[Mat08] MATHEWS S. A.: Design and fabrication of a low-cost, multispectral imaging system. *Applied Optics 47*, 28 (2008), 71–76.

[Max60] MAXWELL J. C.: On the theory of compound colours, and the relations of the colours of the spectrum. *Philosophical Transactions of the Royal Society of London 150* (1860), 57–84.

[MBR*00] Matusik W., Buehler C., Raskar R., Gortler S. J., McMillan L.: Image-based visual hulls. In *ACM SIGGRAPH* (2000), pp. 369–374.

[MDMS05] Mantiuk R., Daly S., Myszkowski K., Seidel H.-P.: Predicting visible differences in high dynamic range images - model and its calibration. In *Proceedings Human Vision and Electronic Imaging X*. B. E. Rogowitz, T. N. Pappas, and S. J. Daly (Eds.). (San Jose, CA, USA, 2005), *vol. 5666*, pp. 204–214.

[MHP*07] Ma W.-C., Hawkins T., Peers P., Chabert C.-F., Weiss M., Debevec P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of the EGSR* (2007), pp. 183–194.

[MHRT08] Mohan A., Huang X., Raskar R., Tumblin J.: Sensing increased image resolution using aperture masks. In *Proceedings of the IEEE CVPR* (2008), pp. 1–8.

[Mil75] Milgram D. L.: Computer methods for creating photomosaics. *IEEE Transactions on Computers 24*, 11 (1975), 1113–1119.

[MKI04] Miyazaki D., Kagesawa M., Ikeuchi K.: Transparent surface modelling from a pair of polarization images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 26*, 1 (2004), 73–82.

[MKRH11] Mantiuk R., Kim K. J., Rempel A., Heidrich W.: HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics (Siggraph) 30*, 3 (2011), 1–12.

[MMP*05] McGuire M., Matusik W., Pfister H., Hughes J. F., Durand F.: Defocus video matting. *ACM Transactions on Graphics (SIGGRAPH) 24*, 3 (2005), 567–576.

[MMP*07] McGuire M., Matusik W., Pfister H., Chen B., Hughes J. F., Nayar S. K.: Optical splitting trees for high-precision monocular imaging. *IEEE Computer Graphics and Applications 27*, 2 (2007), 32–42.

[MN99] Mitsunaga T., Nayar S. K.: Radiometric self calibration. In *Proceedings of the IEEE CVPR* (1999), pp. 374–380.

[MO99] Marshall J., Oberwinkler J.: Ultraviolet vision: The colourful world of the mantis shrimp. *Nature 401*, 6756 (1999), 873–874.

[MP95] Mann S., Picard R. W.: Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proceedings of the IS&T* (1995), pp. 442–448.

[MP04] Matusik W., Pfister H.: 3d tv: A scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Transactions on Graphics 23* (2004), 814–824.

[MRT08] Mohan A., Raskar R., Tumblin J.: Agile spectrum imaging: Programmable wavelength modulation for cameras and projectors. *Computer Graphics Forum (Eurographics) 27*, 2 (2008), 709–717.

[MS81] Moon P., Spencer D. E.: *The Photic Field*. MIT Press, Cambridge, MA, 1981.

[MSH09] Mäthger L. M., Shashar N., Hanlon R. T.: Do cephalopods communicate using polarized light reflections from their skin?. *Journal of Experimental Biology 212* (2009), 2133–2140.

[MTHI03] Miyazaki D., Tan R. T., Hara K., Ikeuchi K.: Polarization-based inverse rendering from a single view. In *Proceedings of the ICCV* (2003), pp. 982–998.

[Mül96] Müller V.: Elimination of specular surface-reflectance using polarized and unpolarized light. In *Proceedings of the IEEE ECCV* (1996), pp. 625–635.

[Mur01] Murphy D. B.: *Fundamentals of Light Microscopy and Electronic Imaging*. Wiley-Liss, Hoboken, NJ, 2001.

[Muy57] Muybridge E.: *Animals in Motion* (Ist edition). Dover Publications, Chapman and Hall 1899, 1957.

[NB03] Nayar S., Branzoi V.: Adaptive dynamic range imaging: Optical control of pixel exposures over space and time. In *Proceedings of the IEEE ICCV* (2003), vol. 2, pp. 1168–1175.

[NBB04] Nayar S., Branzoi V., Boult T.: Programmable imaging using a digital micromirror array. In *Proceedings of the IEEE CVPR* (2004), vol. *I*, pp. 436–443.

[NBB06] Nayar S. K., Branzoi V., Boult T. E.: Programmable imaging: Towards a flexible camera. *International Journal of Computer Vision 70*, 1 (2006), 7–22.

[NFB93] Nayar S., Fang X.-S., Boult T.: Removal of specularities using color and polarization. In *Proceedings of the IEEE CVPR* (1993), pp. 583–590.

[Ng05] Ng R.: Fourier slice photography. *ACM Transactions on Graphics (Siggraph) 24*, 3 (2005), 735–744.

[NKY08] Narasimhan S. G., Koppal S. J., Yamazaki S.: Temporal dithering of illumination for fast active vision. In *Proceedings of the ECCV* (2008), pp. 830–844.

[NKZN08] Nagahara H., Kuthirummal S., Zhou C., Nayar S.: Flexible depth of field photography. In *Proceedings of the ECCV* (2008), pp. 60–73.

[NLB*05] Ng R., Levoy M., Brédif M., Duval G., Horowitz M., Hanrahan P.: *Light Field Photography with a Hand-Held Plenoptic Camera*. Tech. Rep. Computer Science CSTR 2005-02, Stanford University, 2005.

[NM00] Nayar S., Mitsunaga T.: High dynamic range imaging: Spatially varying pixel exposures. In *Proceedings of the IEEE CVPR* (2000), vol. 1, pp. 472–479.

[NN94] Nayar S. K., Nakagawa Y.: Shape from focus. *IEEE Transactions on Pattern Analysis and Machine Intelligence 16*, 8 (1994), 824–831.

[NN05] Narasimhan S., Nayar S.: Enhancing resolution along multiple imaging dimensions using assorted pixels. *IEEE Transactions on Pattern Analysis and Machine Intelligence 27*, 4 (2005), 518–530.

[Nou10] Nourbakhsh I.: GigaPan. `http://gigapan.org/`, 2010. Accessed 9 October 2011.

[NS05] Namer E., Schechner Y. Y.: Advanced visibility improvement based on polarization filtered images. In *Proceedings of SPIE 5888* (2005), pp. 36–45.

[NZN07] Nomura Y., Zhang L., Nayar S.: Scene collages and flexible camera arrays. In *Proceedings of the EGSR* (2007), pp. 1–12.

[OAHY99] Okano F., Arai J., Hoshino H., Yuyama I.: Three-dimensional video system based on integral photography. *Optical Engineering 38*, 6 (1999), 1072–1077.

[OCLE05] Ojeda Castaneda J., Landgrave J. E. A., Escamilla H. M.: Annular phase-only mask for high focal depth. *Optics Letters 30*, 13 (2005), 1647–1649.

[OIS94] Ogata S., Ishida J., Sasano T.: Optical sensor array in an artificial compound eye. *Optical Engineering 33*, 11 (1994), 3649–3655.

[Opt11] Optec: Separation prism technical data, Jan 2011. `www.alt-vision.com/color_prisms_tech_data.htm`. Accessed 9 October 2011.

[OY91] Okamoto T., Yamaguchi I.: Simultaneous acquisition of spectral image information. *Optics Letters 16*, 16 (1991), 1277–1279.

[Pan10] Panavision: Genesis, 2010. `www.panavision.com`. Accessed 9 October 2011.

[PG12] Prokudin-Gorskii S. M.: The Prokudin-Gorskii Photographic Records Recreated. `www.loc.gov/exhibits/empire/`, 1912. Accessed 9 October 2011.

[Pho10] Photron: FASTCAM SA5, 2010. `www.photron.com/datasheet/FASTCAM_SA5.pdf`. Accessed 9 October 2011.

[Pix10] Pixim: Digital Pixel System, 2010. `www.pixim.com`. Accessed 9 October 2011.

[PK83] Pieper R. J., Korpel A.: Image processing for extended depth of field. *Applied Optics 22*, 10 (1983), 1449–1453.

[PKR10] Pandharkar R., Kirmani A., Raskar R.: Lens aberration correction using locally optimal mask based low cost light field cameras. In *Proceedings of the OSA Imaging Systems* (2010), pp. 1–3.

[PLGN07] Park J. I., Lee M.-H., Grossberg M. D., Nayar S. K.: Multispectral imaging using multiplexed illumination. In *Proceedings of the IEEE ICCV* (2007), pp. 1–8.

[PR06] Parmar M., Reeves S. J.: Selection of optimal spectral sensitivity functions for color filter arrays. In *Proceedings of the of ICIP* (2006), pp. 1005–1008.

[PR10] Parmar M., Reeves S. J.: Selection of optimal spectral sensitivity functions for color filter arrays. *IEEE Transactions on Image Processing 19*, 12 (Dec 2010), 3190–3203.

[Pro09] Project E. D. C.: Harold 'Doc' Edgerton, 2009. `www.edgerton-digital-collections.org/techniques/high-speed-photography`. Accessed 9 October 2011.

[PVB*11] Pandharkar R., Velten A., Bardagjy A., Raskar R., Bawendi M., Kirmani A., Lawson E.: Estimating motion and size of moving non-line-of-sight objects in cluttered environments. In *Proceedings of the ICCC CVPR* (2011), pp. 1–8.

[RAT06] Raskar R., Agrawal A., Tumblin J.: Coded exposure photography: Motion deblurring using fluttered shutter. *ACM Transactions on Graphics (Siggraph) 25*, 3 (2006), 795–804.

[RAWV08] Raskar R., Agrawal A., Wilson C. A., Veeraraghavan A.: Glare aware photography: 4D ray sampling for reducing glare effects of camera lenses. *ACM Transactions on Graphics (Siggraph) 27*, 3 (2008), 56.

[Ray02] Ray S. F.: *High Speed Photography and Photonics*. SPIE Press, Bellingham, WA, 2002.

[RBS99] Robertson M. A., Borman S., Stevenson R. L.: Estimation-theoretic approach to dynamic range enhancement using multiple exposures. *Journal of Electronic Imaging 12* (1999), 2003.

[Res10] RESEARCH V.: Phantom Flex, 2010. www.visionresearch.com/Products/High-Speed-Cameras/Phantom-Flex/. Accessed 9 October 2011.

[RFMM06] RI S., FUJIGAKI M., MATUI T., MORIMOTO Y.: Accurate pixel-to-pixel correspondence adjustment in a digital micromirror device camera by using the phase-shifting moiré method. *Applied Optics 45*, 27 (2006), 6940–6946.

[Ric72] RICHARDSON H. W.: Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America 62*, 1 (1972), 55–59.

[RKAJ08] REINHARD E., KHAN E. A., AKYÜZ A. O., JOHNSON G. M.: *Color Imaging*. A K Peters Ltd., 2008.

[RMH*11] ROUF M., MANTIUK R., HEIDRICH W., TRENTACOSTE M., LAU C.: Glare encoding of high dynamic range images. In *Proceedings of the IEEE CVPR* (2011), pp. 1–8.

[Ror08] RORSLETT B.: Uv flower photographs, 2008. www.naturfotograf.com/index2.html. Accessed 9 October 2011.

[RSBS02] RAMANATH R., SNYDER W., BILBRO G., SANDER W.: Demosaicking methods for Bayer color arrays. *Journal of Electronic Imaging 11*, 3 (2002), 306–315.

[RT09] RASKAR R., TUMBLIN J.: *Computational Photography: Mastering New Techniques for Lenses, Lighting, and Sensors*. A. K. Peters, 2009.

[Ruc61] RUCK P.: Photoreceptor Cell Response and flicker fusion frequency in the compound eye of the fly, Lucilia Sericata (Meigen). *Biological Bulletin 120* (1961), 375–383.

[RVC11] REDDY D., VEERARAGHAVAN A., CHELLAPPA R.: P2C2: Programmable pixel compressive camera for high speed imaging. In *Proceedings of the IEEE CVPR* (2011), pp. 1–8.

[RWD*10] REINHARD E., WARD G., DEBEVEC P., PATTANAIK S., HEIDRICH W., MYSZKOWSKI K.: *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann Publishers, 2010.

[Sad07] SADJADI F.: Extraction of surface normal and index of refraction using a pair of passive infrared polarimetric sensors. In *Proceedings of the IEEE CVPR* (2007), pp. 1–5.

[Sch03] SCHÖNFELDER T.: Polarization Division Multiplexing in Optical Data Transmission Systems, 2003. US Patent 6,580,535.

[SCI02] SHECHTMAN E., CASPI Y., IRANI M.: Increasing space-time resolution in video. In *Proceedings of the ECCV* (2002), pp. 753–768.

[SCI05] SHECHTMAN E., CASPI Y., IRANI M.: Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence 27*, 4 (2005), 531–545.

[Sem10] SEMICONDUCTOR C.: LUPA Image Sensors, 2010. www.cypress.com/?id=206. Accessed 1 February 2011.

[Set01] SETTLES G.: *Schlieren & Shadowgraph Techniques*. Springer, 2001.

[SFI11] SHAHAR O., FAKTOR A., IRANI M.: Space-time super-resolution from a single video. In *Proceedings of the IEEE CVPR* (2011), pp. 1–8.

[SH08] STARCK J., HILTON A.: Model-based human shape reconstruction from multiple views. *Computer Vision and Image Understanding (CVIU) 111*, 2 (2008), 179–194.

[Shi10] SHIMADZU: HyperVision HPV-2, 2010. http://www.shimadzu.com/an/test/hpv-2/hpv1.html. Accessed 9 October 2011.

[SHS*04] SEETZEN H., HEIDRICH W., STUERZLINGER W., WARD G., WHITEHEAD L., TRENTACOSTE M., GHOSH A., VOROZCOVS A.: High dynamic range display systems. *ACM Transactions on Graphics (SIGGRAPH 2004) 23*, 3 (2004), 760–768.

[SK04] SCHECHNER Y. Y., KARPEL N.: Clear underwater vision. In *Proceedings of the IEEE CVPR* (2004), pp. 536–543.

[SK05] SCHECHNER Y. Y., KARPEL N.: Recovery of underwater visibility and structure by polarization analysis. *IEEE Journal of Oceanic Engineering 30*, 3 (2005), 570–587.

[SMH*11] SAJADI B., MAJUMDER A., HIWADA K., MAKI A., RASKAR R.: Switchable primaries using shiftable layers of color filter arrays. *ACM Transactions on Graphics (Siggraph) 30*, 3 (2011), 1–10.

[SN01] SCHECHNER Y., NAYAR S.: Generalized mosaicing. In *Proceedings of the IEEE ICCV* (2001), vol. *1*, pp. 17–24.

[SN02] SCHECHNER Y., NAYAR S.: Generalized mosaicing: Wide field of view multispectral imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24*, 10 (2002), 1334–1348.

[SN03a] SCHECHNER Y., NAYAR S.: Generalized mosaicing: High dynamic range in a wide field of view. *Interna-*

tional Journal of Computer Vision 53, 3 (2003), 245–267.

[SN03b] SCHECHNER Y., NAYAR S. K.: Polarization mosaicking: High dynamic range and polarization imaging in a wide field of view. In *Proceedings of SPIE 5158* (2003), pp. 93–102.

[SN04] SCHECHNER Y., NAYAR S.: Uncontrolled modulation imaging. In *Proceedings of the IEEE CVPR* (2004), pp. 197–204.

[SN05] SCHECHNER Y., NAYAR S.: Generalized mosaicing: Polarization panorama. *IEEE Transactions on Pattern Analysis and Machine Intelligence 27*, 4 (2005), 631–636.

[SNB07] SCHECHNER Y., NAYAR S., BELHUMEUR P.: Multiplexing for optimal lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*, 8 (2007), 1339–1354.

[SNN01] SCHECHNER Y., NARASIMHAN S. G., NAYAR S. K.: Instant dehazing of images using polarization. In *Proceedings of the IEEE CVPR* (2001), pp. 325–332.

[SNN03] SCHECHNER Y., NARASIMHAN S. G., NAYAR S. K.: Polarization-based vision through haze. *Applied Optics 42*, 3 (2003), 511–525.

[Sph10] SpheronVR: SpheroCam HDR, 2010. `www.spheron.com`. Accessed 9 October 2011.

[SSK99] SCHECHNER Y., SHAMIR J., KIRYATI N.: Polarization-based decorrelation of transparent layers: The inclination angle of an invisible surface. In *Proceedings of the ICCV* (1999), pp. 814–819.

[SZ91] SIDDIQUI A. S., ZHOU J.: Two-channel optical fiber transmission using polarization division multiplexing. *Journal of Optical Communications 12*, 2 (1991), 47–49.

[SZJA09] SMITH B. M., ZHANG L., JIN H., AGARWALA A.: Light field video stabilization. In *Proceedings of the of ICCV* (2009), pp. 1–8.

[TAHL07] TALVALA E.-V., ADAMS A., HOROWITZ M., LEVOY M.: Veiling glare in high dynamic range imaging. *ACM Transactions on Graphics (Siggraph) 26*, 3 (2007), 37.

[TAR05] TUMBLIN J., AGRAWAL A., RASKAR R.: Why I want a gradient camera. In *Proceedings of the IEEE CVPR* (2005), pp. 103–110.

[TARV10] TAGUCHI Y., AGRAWAL A., RAMALINGAM S., VEERARAGHAVAN A.: Axial light fields for curved mirrors: Reflect your perspective, widen your view. In *Proceedings of the IEEE CVPR* (2010), pp. 1–8.

[TAV*10] TAGUCHI Y., AGRAWAL A., VEERARAGHAVAN A., RAMALINGAM S., RASKAR R.: Axial-cones: Modeling spherical catadioptric cameras for wide angle light field rendering. *ACM Transactions on Graphics 29* (2010), 172:1–172:8.

[TH97] TOYOOKA S., HAYASAKA N.: Two-dimensional spectral analysis using broad-band filters. *Optical Communications 137* (Apr 1997), 22–26.

[THBL10] TAI Y., HAO D., BROWN M. S., LIN S.: Correction of spatially varying image and video motion blur using a hybrid camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence 32*, 6 (2010), 1012–1028.

[TKTS11] TOCCI M. D., KISER C., TOCCI N., SEN P.: A versatile HDR video production system. *ACM Transactions on Graphics (Siggraph) 30*, 4 (2011).

[TKY*01] TANIDA J., KUMAGAI T., YAMADA K., MIYATAKE S., ISHIDA K., MORIMOTO T., KONDOU N., MIYAZAKI D., ICHIOKA Y.: Thin observation module by bound optics (TOMBO): Concept and experimental verification. *Applied Optics 40*, 11 (2001), 1806–1813.

[TSK*03] TANIDA J., SHOGENJI R., KITAMURA Y., YAMADA K., MIYAMOTO M., MIYATAKE S.: Color imaging with an integrated compound imaging system. *Optics Express 11*, 18 (2003), 2109–2117.

[TSY*07] TELLEEN J., SULLIVAN A., YEE J., WANG O., GUNAWARDANE P., COLLINS I., DAVIS J.: Synthetic shutter speed imaging. *Computer Graphics Forum (Eurographics) 26*, 3 (2007), 591–598.

[Tys91] TYSON R. K.: *Principles of Adaptive Optics*. Academic Press, Maryland Heights, MO, 1991.

[UG04] UMEYAMA S., GODIN G.: Separation of diffuse and specular components of surface reflection by use of polarization and statistical analysis of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 26*, 5 (2004), 639–647.

[UKTN08] UEDA K., KOIKE T., TAKAHASHI K., NAEMURA T.: Adaptive integral photography imaging with variable-focus lens array. In *Proc SPIE: Stereoscopic Displays and Applications XIX* (2008), pp. 68031A–9.

[ULK*08] UEDA K., LEE D., KOIKE T., TAKAHASHI K., NAEMURA T.: Multi-focal compound eye: Liquid lens array for computational photography. ACM SIGGRAPH New Tech Demo, 2008.

[UWH*03] UNGER J., WENGER A., HAWKINS T., GARDNER A., DEBEVEC P.: Capturing and rendering with incident light fields. In *Proceedings of the EGSR* (2003), pp. 141–149.

[Vag07] VAGNI F.: *Survey of Hyperspectral and Multispectral Imaging Technologies*. Tech. Rep. TR-SET-065-P3, NATO Research and Technology, 2007.

[Val10] VALLEY G.: Viper FilmStream Camera, 2010. www.grassvalley.com. Accessed 9 October 2011.

[vPAB*11] VAN Putten E., AKBULUT D., BERTOLOTTI J., VOS W., LAGENDIJK A., MOSK A.: Scattering lens resolves sub-100 nm structures with visible light. *Physical Review Letters 106*, 19 (2011), 1–4.

[VPB*09] VLASIC D., PEERS P., BARAN I., DEBEVEC P., POPOVIĆ J., RUSINKIEWICZ S., MATUSIK W.: Dynamic shape capture using multi-view photometric stereo. In *ACM Trans. Graph. (SIGGRAPH Asia)* (2009), pp. 1–11.

[VRA*07] VEERARAGHAVAN A., RASKAR R., AGRAWAL A., MOHAN A., TUMBLIN J.: Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocussing. *ACM Transactions on Graphics (Siggraph) 26*, 3 (2007), 69.

[VRA*08] VEERARAGHAVAN A., RASKAR R., AGRAWAL A., CHELLAPPA R., MOHAN A., TUMBLIN J.: Non-refractive modulators for encoding and capturing scene appearance and depth. In *Proceedings of the IEEE CVPR* (2008), pp. 1–8.

[VRR11] VEERARAGHAVAN A., REDDY D., RASKAR R.: Coded strobing photography: Compressive sensing of high speed periodic videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence 33*, 4 (2011), 671–686.

[VSZ*06] VAISH V., SZELISKI R., ZITNICK C. L., KANG S. B., LEVOY M.: Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures. In *Proceedings of the IEEE CVPR* (2006), pp. 23–31.

[VWJL04] VAISH V., WILBURN B., JOSHI N., LEVOY M.: Using plane + parallax for calibrating dense camera arrays. In *Proceedings of the IEEE CVPR* (2004), pp. 1–8.

[Wan05] WANDINGER U.: *Lidar: Range-Resolved Optical Remote Sensing of the Atmosphere*. Springer, 2005.

[WB91] WOLFF L. B., BOULT T. E.: Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence 13*, 7 (1991), 635–657.

[Weh76] WEHNER R.: Polarized-light navigation by insects. *Scientific American 235* (1976), 106–115.

[WH04] WANG S., HEIDRICH W.: The design of an inexpensive very high resolution scan camera system. *Computer Graphics Forum (Eurographics) 23*, 10 (2004), 441–450.

[WIGH11] WETZSTEIN G., IHRKE I., GUKOV A., HEIDRICH W.: Towards a database of high-dimensional plenoptic images. In *Proceedings of the ICCP (Poster)* (2011).

[WIH10] WETZSTEIN G., IHRKE I., HEIDRICH W.: Sensor saturation in fourier multiplexed imaging. In *Proceedings of the IEEE CVPR* (2010), pp. 1–8.

[WJV*04] WILBURN B., JOSHI N., VAISH V., LEVOY M., HOROWITZ M.: High speed video using a dense array of cameras. In *Proceedings of the IEEE CVPR* (2004), pp. 1–8.

[WJV*05] WILBURN B., JOSHI N., VAISH V., TALVALA E.-V., ANTUNEZ E., BARTH A., ADAMS A., HOROWITZ M., LEVOY M.: High performance imaging using large camera arrays. *ACM Transactions on Graphics (Siggraph) 24*, 3 (2005), 765–776.

[WLHR11] WETZSTEIN G., LANMAN D., HEIDRICH W., RASKAR R.: Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Transactions on Graphics (Siggraph)* (2011).

[WPSB08] WAGADARIKAR A., PITSIANIS N., SUN X., BRADY D.: Spectral image estimation for coded aperture snapshot spectral imagers. In *Proceedings of SPIE 7076* (2008), p. 707602.

[WPSB09] WAGADARIKAR A., PITSIANIS N., SUN X., BRADY D.: Video rate spectral imaging using a coded aperture snapshot spectral imager. *Optics Express 17*, 8 (2009), 6368–6388.

[WRH11] WETZSTEIN G., RASKAR R., HEIDRICH W.: Handheld schlieren photography with light field probes. In *Proceedings of the ICCP* (2011), pp. 1–8.

[WRHR11] WETZSTEIN G., ROODNICK D., HEIDRICH W., RASKAR R.: Refractive shape from light field distortion. In *Proceedings of the ICCV* (2011), pp. 1–7.

[WS82] WYSZECKI G., STILES W.: *Color Science*. John Wiley and Sons, Inc., Hoboken, NJ, 1982.

[WSLH02] WILBURN B., SMULSKI M., LEE K., HOROWITZ M. A.: The Light field video camera. In *SPIE Electronic Imaging* (2002), pp. 29–36.

[Yao08] YAO S.: Optical Communications Based on Optical Polarization Multiplexing and Demultiplexing, 2008. US Patent 7,343,100.

[YEBM02] YANG J. C., EVERETT M., BUEHLER C., MCMILLAN L.: A real-time distributed light field camera. In *Proceedings of the EGSR* (2002), pp. 77–86.

[YLIM00] Yang J., Lee C., Isaksen A., McMillan L.: A Low-Cost Portable Light Field Capture Device. ACM SIGGRAPH Technical Sketch, 2000.

[YMIN10] Yasuma F., Mitsunaga T., Iso D., Nayar S. K.: Generalized assorted pixel camera: Post-capture control of resolution, dynamic range and spectrum. *IEEE Transactions on Image Processing 99*, *19*, 9 (2010), 2241–2253.

[ZC05] Zhang C., Chen T.: Light field capturing with lensless cameras. In *Proceedings of the ICIP* (2005), pp. III–792–5.

[ZLN09] Zhou C., Lin S., Nayar S. K.: Coded aperture pairs for depth from defocus. In *Proceedings of the ICCV* (2009), pp. 1–8.

[ZMDP06] Zwicker M., Matusik W., Durand F., Pfister H.: Antialiasing for automultiscopic 3D displays. In *Proceedings of Eurographics Symposium on Rendering* (Nicosia, Cyprus, 2006).

[ZN09] Zhou C., Nayar S.: What are good apertures for defocus deblurring?. In *Proceedings of the ICCP* (2009), pp. 1–8.