Flexible Multimodal Camera Using a Light Field Architecture

Roarke Horstmeyer, Gary Euliss, Ravindra Athale, The MITRE Corporation
Marc Levoy, Stanford University

Abstract

We present a modified conventional camera that is able to collect multimodal images in a single exposure. Utilizing a light field architecture in conjunction with multiple filters placed in the pupil plane of a main lens, we are able to digitally reconstruct synthetic images containing specific spectral, polarimetric, and other optically filtered data. The ease with which these filters can be exchanged and reconfigured provides a high degree of flexibility in the type of information that can be collected with each image. This paper explores the various tradeoffs involved in implementing a pinhole array in parallel with a pupil-plane filter array to measure multi-dimensional optical data from a scene. It also examines the design space of a pupil-plane filter array layout. Images are shown from different multimodal filter layouts, and techniques to maximize resolution and minimize error in the synthetic images are proposed.

1. Introduction

Light field imaging has its origins in research completed over one hundred years ago. Lippmann [1] and Ives [2], early pioneers in the area of integral photography, sought to create cameras that could capture four dimensional information on a two dimensional plane. The advent of digital cameras and the theory behind a 4D light field introduced by Levoy and Hanrahan [3] and Gortler et al. [4] opened up new possibilities in the field. Adelson and Wang [5], Ng et al. [6], Georgiev and Intwala [7], and Veeraraghavan et al. [8] have each developed a unique modification to a conventional imaging system to record this 4D information in one image for the purpose of digital refocusing.

Recently, Raskar et al. [9] broadened the conceptual framework of a light field camera to include applications outside the realm of digital refocusing, by demonstrating reduced glare effects. Their approach used a pinhole array to achieve light field collection. In this paper we also describe a pinhole-based approach to imaging the light field, with the added utility provided by simultaneous placement of various optical filters directly in the camera's pupil plane – an idea recently proposed by Levoy et al. [10]. Using this approach, we show that it is possible in a single frame to collect not only multiple spectral components of a scene, but also polarization and images filtered for the purpose of extending the dynamic range of the imaging system. In the next section, we discuss pupil plane filtering with previous examples, followed by a brief tutorial on the basics of light field imaging. Design considerations associated with a pinhole-based light field approach are then discussed along with the specific implementation we have demonstrated incorporating pupil-plane filter arrays. Results are presented for arrays containing six, nine, and sixteen filters.

2. Background

We are not the first to implement a filter array in conjunction with a light field setup. Over a decade before Lippmann performed his seminal work in integral photography, numerous attempts were made to achieve color imaging by placing a line screen (essentially a pinhole array) over conventional film. In 1895, Lanchester [11] used a line screen in conjunction with a prism setup to achieve imaging with spectral diversity. Shortly thereafter, Liesegang [12], Branfill [13], and others combined the line screen with specific color filters in the pupil plane. Ten years later, Berthon [14] used color filters in the pupil plane along with a lenticular lens array in the focal plane to obtain a successful method for color imaging. This would eventually lead to the Kodak Kodacolor process, a color imaging technique used before the advent of color-emulsion film [15].

Various spatial and/or temporal approaches have previously been implemented to achieve filtering in an imaging system. For example, Smith [16] placed color filters in different areas of the pupil plane and sequentially shuttered them to achieve the first color motion picture process (Kinemacolor). Similar systems have been proposed using both polarization [17] as well as neutral density filters [18]. Bando et al. [19] also used color filters in the pupil plane along with a color sensor to extract a depth map and obtain refocusability. Schenchner and Nayar [20, 21] mixed spatial and temporal sampling

to extend dynamic range with a continuous filter placed in front of the pupil plane of a video camera. Mohan et al. [22] provide one more example of spatio-temporal sampling to obtain multidimensional information from a scene. Similar to the approach we describe in the following paper, they created spectral diversity at a conjugate of the pupil plane, and used a pinhole or small aperture at that plane to reduce confusion between the spectral and angular diversity of rays. However, they used a single pinhole, whereas we describe an approach based on an array of pinholes. This provides the added capability of measuring spectral information at each point in the field of view within a single frame.

A variety of systems have been designed to acquire a multidimensional image from a single snapshot. Plemmons et al. [23] implemented a camera array, with filter diversity being achieved through parallel channels imaging on a common sensor. Filters can also be placed directly on the focal plane (e.g., Bayer filters). Chun and Sadjadi [24] and Nayar and Narasimhan [25, 26] extended this concept to include polarization and density filters. Fife [27] has also proposed to segment the focal plane into filtered regions, while simultaneously integrating a light field system at the chip level. As demonstrated by the early analog designs, a light field architecture facilitates moving filter diversity into the pupil plane to create a more flexible system that can still capture multidimensional images in a single frame. In contrast with an approach that relies on a fixed pixel-scale array at the focal plan [24-26], a filter array can be easily reconfigured at the aperture. Furthermore, by preserving angular diversity in unfiltered aperture areas, our design retains the possibility of combining previously demonstrated [5-10] features such as refocusing and glare-reduction with multidimensional image capture. Moving away from a fixed filter array design is a step towards enabling a field-programmable imaging architecture.

3. Optical filtering in the light field

We begin with a brief discussion on the general concept of a light field camera and its terminology. A more detailed explanation of similar light field setups can be found in [5], [6], and [10]. The following analysis focuses on a pinhole-based light field system. A lenslet array could be used to achieve the desired results, but a pinhole array offers the advantage of a larger depth of field when imaging the pupil plane, thus relaxing design tolerances at the focal plane. Furthermore, it is relatively simple and inexpensive to produce multiple pinhole arrays of varying size and pitch that are optimized for different pupil plane filter configurations. This will be explained in Section 4. So although resolution and optical efficiency are generally going to be worse with a pinhole array than with a lenslet

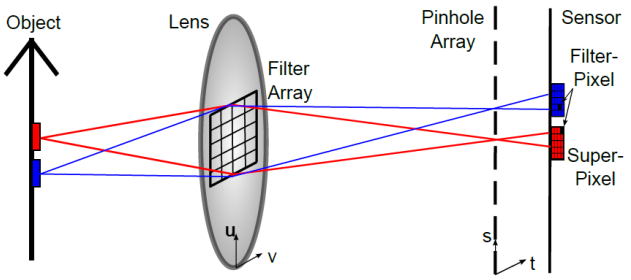

Figure 1: Diagram of a pinhole array light field with a pupil-plane filter array setup. Each filter is located at a specific (u,v) coordinate, and each pinhole is located at a specific (s,t) coordinate.

array, its advantages proved to be worthwhile for our experiments.

To begin, consider a conventional camera lens with a sensor located slightly behind its focal plane. At the focal plane, an array of pinholes is inserted so that a portion of the image from the main lens on the array will be sampled by each pinhole (Fig. 1). The pinholes spatially redistribute the impinging light depending upon the angle at which it passes through a given pinhole. A conventional camera loses this information when the sensor integrates all light incident at a certain spatial location.

One can think of a collected light ray as passing through two parallel planes: the main lens at coordinates (u,v) and a pinhole located at coordinates (s,t). The ability to characterize a ray of light with these four variables is a defining characteristic of a 4D light field camera. When a pinhole array is placed in front of the sensor, each pinhole with a given (s,t) coordinate images the pupil and maps it on the (u,v) plane. Filters placed in the pupil plane will be imaged by every pinhole in the array. Thus, under each pinhole with a particular (s,t) coordinate are spatially separated rays of light that have passed through different filters from unique (u,v) coordinates in the main lens.

In addition to thinking about a pinhole array as a means to record light in (u,v,s,t) space, it is helpful to think about the type of images the system records. Every pinhole creates a sub-image on the sensor array, which could be thought of as a "super-pixel." The super-pixels, when viewed collectively, roughly resemble the scene that is being imaged. Locally, each super-pixel is an image of the filter array. The smallest region of interest corresponds to a single filter in the super-pixel (or "filter-pixel"). Just like sub-pixels in other light field camera designs, these filter-pixels can be sorted and tiled together to create a synthetic image (Fig. 2). Sorting and combining filter-pixels from a particular filter allows an image of the entire scene to be constructed from a virtual aperture [3] containing the corresponding filter. Similarly, separate images can be reconstructed to correspond with the virtual aperture represented by each filter in the pupil-plane filter array.

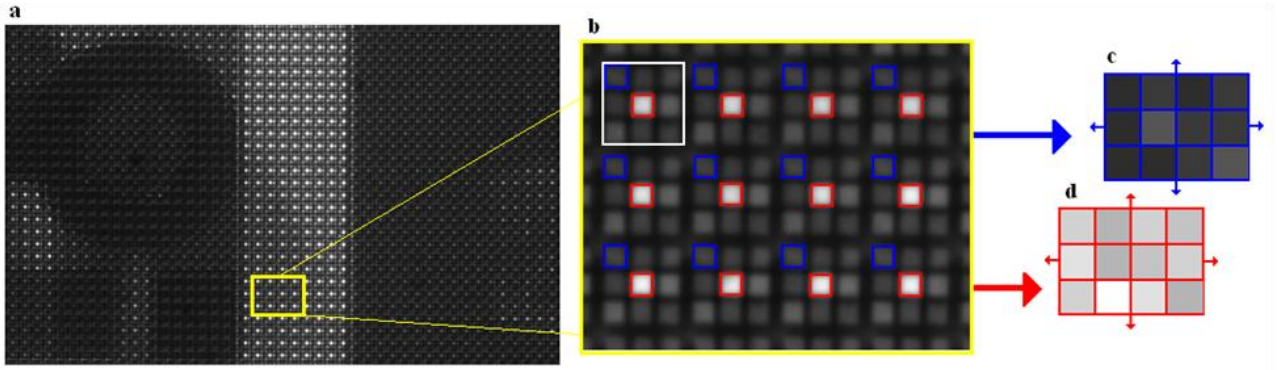

Figure 2: The process of parsing up a light field. (a) A magnified portion of a raw image captured with the pinhole-filter light field camera, in which each super-pixel is roughly distinguishable but a macroscopic image is still visible. (b) A magnified portion of (a), in which a super-pixel (white box) contains nine filter-pixels. All of the filter-pixels from the top left of each super-pixel (blue box) can be combined to create the synthetic image (c), while the filter-pixels from the center of each super-pixel (red) can be combined to create the different synthetic image (d).

4. Analysis of the design space

In this section, we address the design tradeoffs associated with constructing a pinhole array light field system with pupil-plane filters. A simple analysis will first be presented to determine the optimal size, pitch, and location of the pinhole array. Then, the effects of multiple filters in the pupil plane will be considered to determine how error in the light field system can be minimized with a proper filter array configuration.

Of interest is the tradeoff between the number of distinct filters that can be placed in the pupil-plane filter array and the resolution of the corresponding synthesized images. This is qualitatively analogous to a tradeoff between resolution in the (u,v)-plane vs. resolution in the (s,t)-plane. Resolution of the synthetic images increases with the number of sub-aperture images formed on the sensor, and hence the number of pinholes in the array. However, as will be shown, an increase in the number of pinholes reduces the upper limit on the number of filters that can be placed in the pupil plane, and therefore the associated degrees of freedom. The goal of the following analysis is to estimate a pinhole configuration – i.e., size, pitch, and location – that will enable a suitable balance to be achieved.

Much of the following analysis follows from Young [28]. Fig. 3 shows the geometrical layout of the pinhole array and filter setup. In a conventional camera lens, the aperture stop is located within a compound lens system, simplified to L_1 and L_2 in the figure. Each pinhole reimages a virtual image of the filter array formed by L_2 at a distance P. Referring to Fig. 3, we begin by *defining* the operating focal length f and magnification M of one pinhole in the array,

$$\frac{1}{f} = \frac{1}{P} + \frac{1}{Q} \tag{1}$$

$$M = \frac{Q}{P}. \tag{2}$$

In order to maximize the number of sub-aperture images on the sensor, the size of each sub-aperture image should be minimized. This can be achieved by placing the pinhole array as close to the sensor as possible. Theoretically there is a limit to how small Q can be while still forming an image. In practice, the sensor will normally have a cover glass limiting how close the pinhole array can be placed. We found that placing the pinhole array directly against the cover glass provided an adequate working distance, as others have previously observed as well [9]. In general, when considering pinhole sizes and image distances each two or three orders of magnitude larger than the operating wavelength, any analysis to optimize image formation is approximate [28].

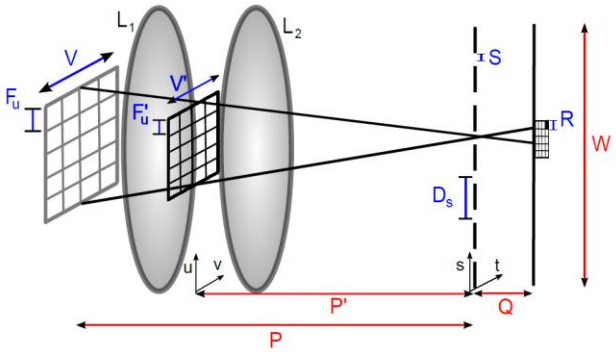

Figure 3: Camera diagram with fixed distances (red) and distances to be determined (blue) indicated for our setup. P' is the distance to the filter array and P is the distance to its virtual image. Here, the two variables we wish to maximize ($Filter_u$ = number of filters, $Resolution_s$ = light field image resolution) are five and six, respectively. Note that the pinhole is not imaging the filter itself, but instead its virtual image created by the lens group L_2 in between it and the pinhole array.

However, minor adjustments based on experimental optimization yield excellent results, and validate the theory as an effective design tool.

Once Q is set, the optimal pinhole size is estimated by

$$S = \sqrt{\lambda f} = \sqrt{\lambda\left(\frac{1}{P}+\frac{1}{Q}\right)^{-1}}, \quad (3)$$

where λ is the average wavelength of light being imaged. Eq. (3) is an estimate based on finding a compromise between the geometrical shadow model (valid for a large pinhole) and the Fraunhofer diffraction model (valid for a small pinhole). Estimates from (3) proved to be consistent with our experimental results.

Given an optimum pinhole size S, the pinhole resolution limit R can be conservatively estimated by

$$R \approx 1.5 \cdot S\left(1+\left|M\right|\right), \quad f \approx \frac{S^2}{\lambda}. \quad (4)$$

Using the approximation given by (4), the minimum resolvable spot size in the filter plane F will be given by

$$F = \left(\frac{1}{M}\right)\left(\frac{1}{M_2}\right)R = \left(\frac{1.5}{M_2}\right)\left[\lambda\left(\frac{1}{P}+\frac{1}{Q}\right)^{-1}\right]^{1/2}\left(\frac{P}{Q}+1\right), \quad (5)$$

which gives the minimum desired size of one pupil plane filter. Here, M_2 is the magnification between the virtual image and the actual size of the filter array determined by experiment. To maximize the number of filters in the pupil plane that can be imaged, all filters should be adjacent and correspond to this minimum size.

Consider for simplicity the total length of the pinhole array, determined by the length of the sensor (W), as well as the length of the pupil filter array (V), in just one direction along the parallel u and s axes. We can define the resolution of the light field images and the limit on the number of filters in the pupil by

$$Resolution_s = \frac{W}{D_s} \quad (6)$$

$$Filter_u = \frac{V}{F_u}. \quad (7)$$

Similar to the matching of F-numbers in lenslet-based light field camera setups, we want to maximize the size of the pinhole images without overlap, by allowing

$$D_s = Filter_u \cdot F \cdot \left(\frac{PM_2}{Q}\right) = Filter_u \cdot R. \quad (8)$$

This leads to the general relationship,

$$Resolution_s = \frac{W}{Filter_u \cdot R}, \quad (9)$$

in which a tradeoff becomes apparent. This tradeoff between synthetic image resolution and pupil plane diversity is almost identical to the tradeoff between spatial and angular resolution in other light field architectures [29]. A balance between maximizing information from (u,v) space and (s,t) space must be reached with all light field systems, but one distinction must be made regarding the pinhole array. Unlike other designs where adjacent *sensor* pixels are used to acquire angular information, the large blur (~50μ) created by a pinhole results in resolution limited by the optics (i.e, the pinhole lens). Therefore, it is the resolution at the pupil plane of the image produced by a pinhole that limits the size of each filter and therefore the number of filters in the array. In general, most of the above parameters are derived simply from P and Q, which will typically be set by the focal length of the main lens and the smallest achievable pinhole-sensor distance where image formation is still achieved.

Given an optimal pinhole array setup, a few qualitative observations will now be made regarding potential sources of and possible ways to minimize error. To begin, our system uses the (u,v) coordinates of the main lens to encode spectral, polarimetric, and optical density information. We assume that each pinhole is imaging the main lens such that these parameters do not vary over (u,v). Variation of filter information over (u,v) will introduce error when two or more filtered synthetic images are to be compared, like when a degree of polarization map is created. This type of error will increase for objects further from the idealized in-focus plane of the main lens. It can be minimized by grouping filters that will be compared, like orthogonal polarization

TABLE 1: Specific Optimal Setup Parameters

P	Q	S	R	F	M_2
51.23mm	1.14mm	25.32μ	38.83μ	1.51mm	1.15

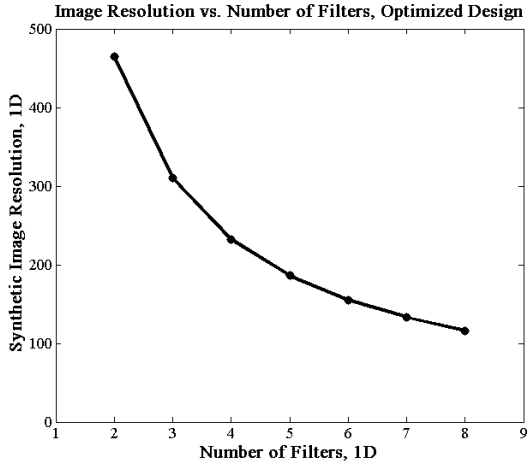

Figure 4: The resolution of filtered synthetic photographs is inversely proportional to the number of filters placed in the pupil plane in one dimension. The graph is based on using the parameters in Table 1 in (9), where W = 3.61 mm.

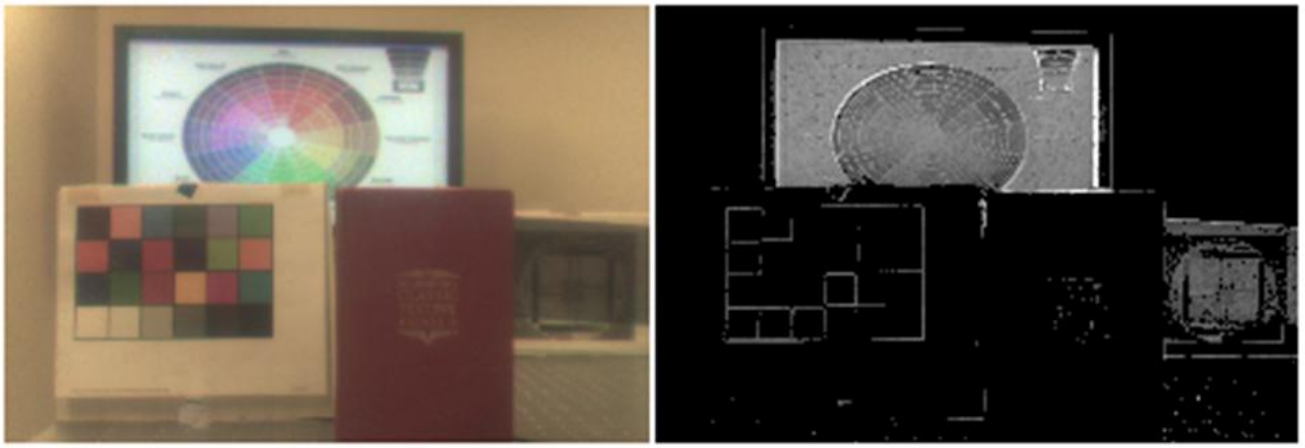

Figure 5: Synthetic images produced from a single frame using the six-filter configuration pictured in Fig. 7(b). An image prior of a Lambertian scene was collected in a separate frame and used for density comparison. The resolution of each image is 231 x 155. (a) Combining synthetic images from behind a red, green and blue filter on our monochrome CCD to create an RGB color image. (b) Combining synthetic images from behind polarizers oriented at 0º, 45º, and 90º to create a linear degree of polarization image.

filters, as close to one another as possible in the filter array. Doing so will minimize the (u,v) range over which we assume polarization, or any other parameter we are interested in measuring, remains angularly constant. The effects of this observation will be exemplified in the next section.

Second, the main lens (containing the filter array) and a pinhole will shape the wavefront incident on the sensor. The geometry of each will contribute to the net point spread function (PSF). We note that the influence of the pinhole at a short working distance will dominate the system PSF, reducing the significance of the filter array in predicting image quality. However, if a lenslet-based approach were used, the influence of the filter array on the overall image quality may be more significant. We observed that arranging lower optical density filters around the center of the filter array and higher density filters near the edges (or vice-versa) improved the PSF of the main lens.

5. Experiments

The light field imaging system used to demonstrate pupil-plane filtering was assembled with a Nikon 50mm f/1.8 lens and 4008 x 2672 board level Lumenera monochrome CCD sensor containing 9 µm pixels placed in the back focal plane. The pinhole array was printed on a transparency at 5080 dpi with a minimum resolution of 25 µm, resulting in pinholes that were roughly square. The pinhole array transparency was pressed against the sensor with a thin piece of glass. The thickness of the cover glass over the CCD is roughly 0.8 mm. Finally, the filters were placed in variously patterned 1 mm thick laser cut custom plastic filter holders. The filter holders were placed inside the lens directly in front of the aperture stop.

To process the images, a MATLAB program was written to find the center of each pinhole sub-image. Pixels at a given distance from this center are selected and combined to construct the filtered synthetic images. Each filter-pixel is typically 5x5 sensor pixels, from which a 2x2 or 3x3 block of pixels is sampled and averaged to create each synthetic photo pixel.

While not always necessary, comparing images to a prior image of a Lambertian scene can be useful. The Lambertian scene can be divided into an image to regain full resolution for objects in focus, as discussed in [8] and [9]. In addition, the printed pinhole masks have a number of slight defects, resulting in fixed pattern noise that can be removed with a previously acquired reference image. Finally, the different densities of each filter can be immediately accounted for with image prior knowledge. For example, an RGB image can be directly created when the densities of each filter are known from a prior. Otherwise, the gamma level of each color channel must be set in post processing. In general, a compromised exposure time must be chosen to account for varying filter densities. Saturation and under-exposure can be avoided when filters of relatively similar densities are used.

Fig. 5 shows results from a simple 2 x 3 filter bank arrangement, with red, green, and blue filters and three polarizers oriented at 0º, 45º and 90º. The filter bank was implemented in conjunction with an array of 50 µm pinholes on a 150 µm pitch. The three color-filtered synthetic images were combined, after being normalized to a prior image, to form a raw color image from the grayscale sensor in a standard format (Fig. 5a). Note that this reduced resolution RGB image is not intended to compete with conventional color photography techniques, but is instead included to highlight the ability of this design to simultaneously capture multiple spectral and

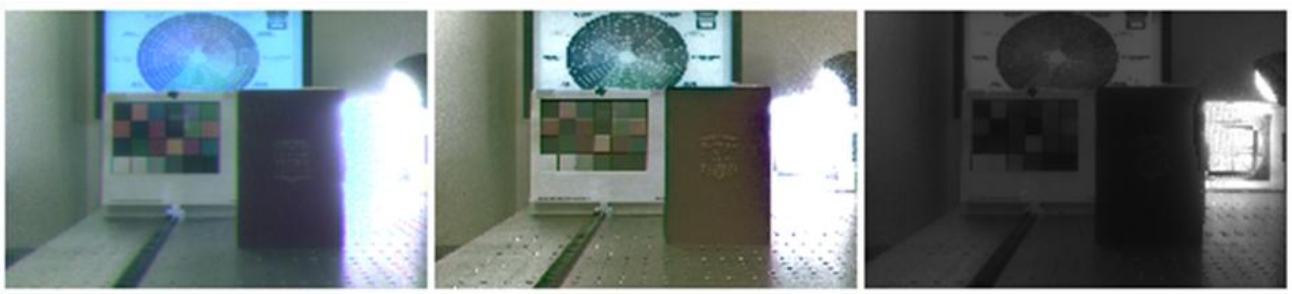

Figure 6: Three images produced from a single frame using the nine-filter configuration in Fig. 7(c). The resolution of each image is 177 x 117. (a) An RGB image, created using a Lambertian-prior for density correction but with no color correction, exhibits saturation over a large area. (b) A CMYK image, created with a K-value from a neutral density filter (0.6 optical density). Note that saturated areas under the lamp are reduced. (c) A grayscale image created from combining three synthetic images from neutral density filters (optical densities of 0.4, 0.6 and 1). A simple HDR image, saturated pixels were sequentially replaced with corresponding pixels from synthetic images of higher optical density. Note that the resolution chart is now visible.

polarization channels. The synthetic images from the three different polarizers were combined to form a linear degree of polarization (*ldop*) image (Fig. 5b) based on the expression,

$$ldop = \frac{\sqrt{(I_0 - I_{90})^2 + (2I_{45} - I_0 - I_{90})^2}}{I_0 + I_{90}}, \qquad (10)$$

where I_x is the intensity of light after passing through a linear polarizer oriented at x^o. As expected, the large LCD screen and the target on the right, which is obscured by a polarizing sheet, both exhibit a high *ldop*.

Erroneous *ldop* measurements, mostly visible along edges, can be attributed to combining filtered synthetic images from different angular perspectives (different *u*,*v* coordinates), and is more noticeable for objects far from the plane of focus of the main lens. Furthermore, the slight blurring of boundaries in the color image is also a result of this angular variation. Error associated with these effects can be minimized by following the first observation regarding pupil plane filter layouts made in the previous section. Linear image alignment techniques could also be used to minimize this disparity, or a more intensive angular displacement correction could be implemented as well.

The next group of images (Fig. 6) was created from a single light field image with a 6 mm x 6 mm array of nine filters in the pupil plane: red, green, blue, yellow, magenta and cyan, along with three different neutral density filters. These filters were used in a setup with an array of 50 µm pinholes on a 200 µm pitch (Fig. 7), and the images were formed after comparison to a prior. A variety of color and contrast comparisons can be made with this type of filter diversity. Spectral comparisons can be made between an RGB image (Fig. 6a) and a CMYK image (Fig. 6b). Additionally, saturation levels can be altered in the CMYK image, in which the K value can come from any one of the neutral density filter synthetic images. Fig. 6c demonstrates extended dynamic range imaging enabled by the three neutral density filters. As in the six-filter example, artifacts of the angular diversity between synthetic images from different areas of the pupil plane are also visible in this figure (e.g. the black rail along the table appears to tilt at different angles).

Fig. 8 contains all the raw synthetic images from a setup with an 8 mm x 8 mm array of 16 filters and an array of 50 µm pinholes on a 200 µm pitch. In this array, a clear aperture, an infrared filter and a 135º polarizer were combined with the filters used in the previous two designs. The top half of the figure contains all six spectral channels. Changes in the color wheel on the large LCD screen and the color calibration chart are apparent. Saturation levels vary in the neutral density filtered images in the lower right, where the visibility of the resolution chart fluctuates. Finally, polarization diversity is visible in the four images in the lower left, where the large LCD screen, the laptop computer screen, and the glare from the metallic optical table show varying

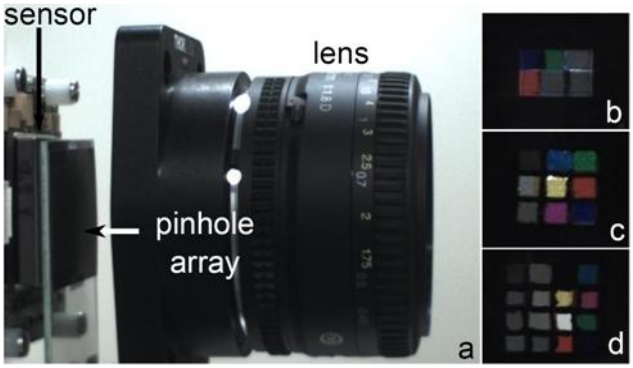

Figure 7: Photographs of the experimental camera. (a) Setup used to capture the image for Fig. 6, with pinhole array pressed against the sensor with a piece of glass. (b) The array of six filters used to create the images in Fig. 5. (c) The array of nine filters used to create the images in Fig. 6. (d) The array of sixteen filters used to create the images in Fig. 8 and Fig. 9.

Figure 8: Sixteen images created from a single frame using the 16-filter configuration in Fig. 7(d). The resolution of each image is 177 x 117. The filters corresponding to each image are (from left to right): first row: blue, green, magenta and cyan; second row: red, no filter, yellow and infrared-pass; third row: 0º polarizer, right circular polarizer, 0.4 OD filter, 0.6 OD filter; fourth row: 45º polarizer, 135º polarizer, 90º polarizer, 1 OD filter. Images are arranged in the same configuration as the filters in the pupil plane.

intensities. Neutral density and spectral filters are combined in Fig. 9, where six synthetic photographs are used to create a color image with an extended dynamic range.

6. Conclusions

An approach to parallel (i.e. single frame) multimodal image acquisition was described and demonstrated. Multimodal image diversity is achieved through the placement of filters in the field lens pupil of a light field imaging system. The advantage of the demonstrated approach is in its flexibility. The image diversity can be modified by simply reconfiguring the pupil plane filter array. The demonstrated imaging system is based on a pinhole array approach to collecting the light field, which provides the advantages of a large field of view and depth of field with respect to the pupil plane, but a lenslet array implementation is also possible. A maximum of 16 parallel filter channels was demonstrated, including spectral, polarization, and neutral density filters. A fundamental scaling analysis indicates a tradeoff between the number of parallel channels placed in the pupil (i.e., the degree of pupil-plane diversity) and synthetic image resolution. This represents an even more fundamental

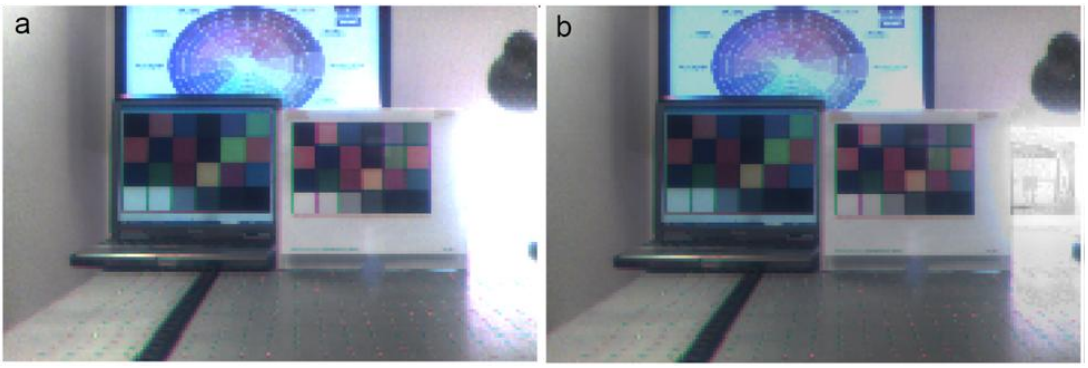

Figure 9: An example of combining six synthetic images in Fig. 8 from a sixteen-filter configuration. An image prior of a Lambertian scene was collected in a separate frame and used for density comparison. The resolution of each image is 177 x 117. (a) Combining synthetic images from behind a red, green and blue filter on our monochrome CCD to create an RGB color image, exhibiting saturation. (b) Combing the RGB image with three synthetic neutral density images to increase dynamic range, where the resolution chart behind the light becomes visible.

tradeoff between the system resources. The efficiency with which these resources are utilized presents an application specific optimization goal.

Additional applications for the described approach can be envisioned. For example, placement of several specifically tuned, narrow-band spectral filters (e.g., specially designed nanophotonic band-gap structures) in the pupil plane might enable parallel performance of multiple spectroscopic imaging tasks. The laboratory demonstrations described in this paper make use of discrete filter arrays. The narrow-band approach would represent another discrete filter implementation. However, it is worth noting that the architecture in general is well suited to implementation with continuously variable filters. For example, one could propose replacing the discrete color filters with a commercially available, continuously variable interference filter. This type of filter is available from sources including Schott and Ocean Optics, although not specifically configured for placement in the pupil plan of a conventional camera lens. In the case of a continuously variable filter, the spectral bands would be limited primarily by the spatial resolution (in the pupil plane) of the pinholes. Similarly, it is also possible to consider an implementation that would take advantage of a continuously variable neutral density filter in exchange for the discrete optical density elements demonstrated here. This would likewise be limited by the same resolution parameters, but would permit additional flexibility in the effective transfer function (charge collection vs. integration time) used in the synthesis of a high dynamic range image.

Continuous color filtering in the pupil plane using the described technique would enable a flexible approach to multispectral imaging in the visible/near IR band. We are aware of no other system that offers the potential to collect multispectral combined with polarimetric images, on a single focal plane, in a single frame, in a flexible, reconfigurable design.

References

[1] G. Lippmann. La photographie integrale. C.R. Acad, Sci. 146: 446-451, 1908

[2] H. E. Ives. Parallax Panoramagram. United States Patent 1918705, 1930

[3] M. Levoy and P. Hanrahan. Light field rendering. ACM Trans. Graphics (Proc. SIGGRAPH) 31-42, 1996

[4] S. J. Gortler, R. Grzeszczuk, R. Szeliski and M. F. Cohen. The lumigraph. ACM Trans. Graphics (Proc. SIGGRAPH) 43-54, 1996

[5] E. H. Adelson and J. Y. A. Wang. Single lens stereo with a plenoptic camera. IEEE Trans. Pattern Anal. Machine Intell. 14(2): 99-106, 1992

[6] R. Ng, M. Levoy, M. Bredif, G. Duval, M. Horowitz, and P Hanrahan. Light field photography with a hand-held plenoptic camera. Stanford Tech Report CTSR 2005-02, 2005

[7] T. Georgiev and C. Intwala. Light field camera design for integral view photography. Adobe Technical Report, 2006

[8] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan and J. Tumblin. Dappled photography: mask enhanced cameras for heterodyned light fields and coded aperture refocusing. ACM Trans. Graphics (Proc. SIGGRAPH) 26(3) 1-12, 2007

[9] R. Raskar, A. Agrawal, C.A. Wilson, and A. Veeraraghavan. Glare aware photography: 4D ray sampling for reducing glare effects of camera lenses. ACM Trans. Graphics (Proc. SIGGRAPH) 27(3), 56, 2008

[10] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz. Light Field Microscopy. ACM Trans. Graphics (Proc. SIGGRAPH) 25(3), 924-934, 2006

[11] F. N. Lanchester. English Patent No. 16548/95 (1895)

[12] R. E. Liesegang. British Journal of Photography Vol. 43: 569, 1896

[13] J. A. C. Branfill. British Journal of Photography Vol. 44: 142, 1897

[14] R. Berthon. English Patent No. 10611/09 (1909)

[15] J. S. Friedman. History of Color Photography. Read Books: 222-250, 2007

[16] G. A. Smith. Kinematograph apparatus for the production of color pictures. United States Patent 941960, 1909

[17] R. Horstmeyer, G. W. Euliss, R. A. Athale, R. A. Morrison, R. L. Stack, and J. Ford. Pupil plane multiplexing for multi-domain imaging sensors. Proc. SPIE 7096(45): 1-10 (2008)

[18] M. Aggarwal and N. Ahuja. Split aperture imaging for high dynamic range. Int. J. Comp. Vision 58(1): 7-17

[19] Y. Bando, B. Chen and T. Nishita. Extracting depth and matte using a color-filtered aperture. ACM Trans. Graphics (Proc. SIGGRAPH) 27(5), 134, 2008

[20] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing: high dynamic range in a wide field of view. Int. J. Comput. Vision 53(3): 245-267, 2003

[21] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing. Proc. ICCV, 2001

[22] A. Mohan, R. Raskar, and J. Tumblin. Agile spectral imaging: programmable wavelength modulation for cameras and projectors. Proc. Eurographics 27(2), 2008

[23] R. J. Plemmons, S. Prasad, S. Matthews, M. Mirotznik, R. Barnard, G. Gray, V. P. Pauca, T. C. Torgersen, J. van der Gracht, and G. Behrmann. PERIODIC: Integrated Computational Array Imaging Technology. Adaptive Optics OSA Technical Digest, CMA1, 2007

[24] C.S.L Chun and F. A. Sadjadi. Polarimetric imaging system for automatic target detection and recognition. DTIC No. ADA392865, 2000

[25] S. K. Nayar and S. G. Narasimhan. Assorted pixels: multi-sampled imaging with structural models. Proc ECCV, 2002

[26] S. G. Narasimhan and S. K. Nayar. Enhancing resolution along multiple imaging dimensions using assorted pixels. IEEE Trans. PAMI 27(4): 518-530, 2005

[27] K. Fife, A. El Gamal and H.-S. P. Wong. A 3MPixel Multi-Aperture Image Sensor with 0.7um Pixels in 0.11um CMOS. IEEE ISSCC Conference Slides, 2008

[28] M. Young. Pinhole Optics. Applied Optics 10(12): 2763-2767, 1971

[29] T. Georgiev, K.C. Zheng, B. Curless, D. Salesin, S. Nayar, and C. Intwala. Spatio-angular resolution tradeoff in integral photography. Eurographics Sym. Rendering, 2006