

A Theory of Plenoptic Multiplexing - Supplemental Material

Ivo Ihrke
ivoihrke@cs.ubc.ca

Gordon Wetzstein
wetzstel@cs.ubc.ca
The University of British Columbia

Wolfgang Heidrich
heidrich@cs.ubc.ca

In this document we would like to give additional information and detailed derivations that have been omitted in the paper due to space restrictions. It is structured as follows:

- Section 1 contains proofs of the PSM and PFM theorems.
- Section 2 discusses the applicability of our theory to optical systems with refractive elements behind the main lens and shows additional results, in particular, reconstruction with non-linear spatial filters.
- Section 3 shows additional results for noise analysis in the case of different ratios between signal-dependent and signal-independent noise.

1. Additional Details for Plenoptic Multiplexing

In this section we would like to give detailed proofs of Theorems 1 and 2. Corollary 1 then directly follows from Theorem 1. For the purpose of this discussion please keep in mind that we are working with continuously defined sensor images $\mathbf{i}(\vec{x})$. For ease of reference we repeat Equation 3 from the paper:

$$\begin{aligned}\mathbf{i}(\vec{x}) &= \sum_{j=1}^N \sigma_j(\vec{x}) \rho_j(\vec{x}) \\ &= \sum_{j=1}^N \sigma_j(\vec{x}) \int_{\mathcal{P}} \pi_j(\vec{p}) l_\lambda(\vec{x}, \vec{p}) d\vec{p}.\end{aligned}\quad (1)$$

1.1. Proof of Theorem 1

Assuming spatial $(\Delta\vec{x}_s^x, \Delta\vec{x}_s^y)$, i.e. super-pixel, periodicity of all spatial basis functions $\sigma_j(x, y) = \sigma_j(x + u\Delta\vec{x}_s^x, y + v\Delta\vec{x}_s^y)$, $(u, v) \in Z$, and defining a sampling operator with the same sampling period as

$$\mathbb{I}_k(\vec{x}) = \sum_{u, v \in Z} \delta(\vec{x} - \left(\begin{array}{c} u\Delta\vec{x}_s^x \\ v\Delta\vec{x}_s^y \end{array} \right) + \Delta\vec{x}_k), \quad (2)$$

where $\Delta\vec{x}_k$ is the sampling offset of individual measurement channels within each super-pixel, we see that sampling a measurement channel $\tilde{\mathbf{c}}_k$ from sensor image $\mathbf{i}(\vec{x})$, results in constant values σ_j^k for the spatial basis functions σ_j :

$$\begin{aligned}\tilde{\mathbf{c}}_k(\vec{x}) &= \mathbb{I}_k(\vec{x}) \mathbf{i}(\vec{x}) \\ &= \mathbb{I}_k(\vec{x}) \left(\sum_{j=1}^N \sigma_j(\vec{x}) \rho_j(\vec{x}) \right) \\ &= \sum_{j=1}^N \sigma_j^k \mathbb{I}_k(\vec{x}) \rho_j(\vec{x}),\end{aligned}\quad (3)$$

i.e. the spatial basis functions become independent of the spatial position \vec{x} in the sampled measurement channels. Each sampled measurement channel \tilde{c}_k is thus associated with N spatial basis constants σ_j^k . Reconstruction of a measurement channel c_k from its sampled representation \tilde{c}_k is performed by convolution with an appropriate filter kernel $f(\vec{x})$:

$$c_k(\vec{x}) = \tilde{c}_k(\vec{x}) \otimes f(\vec{x}). \quad (4)$$

According to the sampling theorem the original signal must be spatially band-limited for this reconstruction to be a faithful representation.

Since the spatial basis is a set of constants for every measurement channel c_k , and the plenoptic basis is independent of \vec{x} , convolution with the filter kernel can be moved inside of sum and integral yielding

$$\begin{aligned} c_k(\vec{x}) &= \sum_{j=1}^N \sigma_j^k \left(\text{III}_k(\vec{x}) \rho_j(\vec{x}) \otimes f(\vec{x}) \right) \\ &= \sum_{j=1}^N \sigma_j^k \int_{\mathcal{D}} \pi_j(\vec{p}) \left(\text{III}_k(\vec{x}) l_\lambda(\vec{x}, \vec{p}) \otimes f(\vec{x}) \right) d\vec{p} \\ &\approx \sum_{j=1}^N \sigma_j^k \sum_{l=1}^M \pi_j^l \left(\text{III}_k(\vec{x}) l_\lambda(\vec{x}, \vec{p}^l) \otimes f(\vec{x}) \right), \end{aligned} \quad (5)$$

where the last row approximates the continuous plenoptic basis by discretization. Equation 5 shows that all measurement channels c_k are locally related to the plenoptic function via a linear combination. It also shows that sampling and reconstruction on the measurement channels, i.e. before de-correlation, is equivalent to sampling and reconstructing the plenoptic function itself, i.e. after de-correlation. This follows from spatial constancy of the spatial basis and spatial independence of the plenoptic basis for individual reconstructed measurement channels. \square

For convenience, we give the explicit form of Eq. 5 in matrix notation:

$$\begin{pmatrix} c_1(\vec{x}) \\ \vdots \\ c_K(\vec{x}) \end{pmatrix} = \begin{pmatrix} \sigma_1^1 & \dots & \sigma_1^N \\ \vdots & \ddots & \vdots \\ \sigma_1^K & \dots & \sigma_1^N \end{pmatrix} \begin{pmatrix} \pi_1^1 & \dots & \pi_1^M \\ \vdots & \ddots & \vdots \\ \pi_1^N & \dots & \pi_1^M \end{pmatrix} \begin{pmatrix} l_\lambda(\vec{x}, \vec{p}^1) \\ \vdots \\ l_\lambda(\vec{x}, \vec{p}^M) \end{pmatrix}, \quad (6)$$

which we write in short as $\vec{c}(\vec{x}) = \Sigma \Pi^T \vec{l}_\lambda(\vec{x})$, the form given in Theorem 1. In general, the number of measurement channels K and the number of spatial bases N does not need to match. In case of $K > N$ the system is over-determined. This case occurs when the sensor over-samples the modulator. If $K < N$, the system becomes ill-conditioned and suitable priors have to be found to enable a reconstruction. An interesting research direction here would be compressive sampling [2] where typically $K \ll N$. In this case, the linear system is typically inverted using a less-than-one norm.

Note that even though the proof has been presented for a rectangular sampling lattice, it remains valid for more general sampling patterns. The only requirement for the validity of the proof are that the periodicity of the spatial basis functions and the sampling period match, i.e. that sampling a measurement channel results in a set of constants σ_j^k . Theorem 1 thus holds for more exotic sampling patterns like e.g. hexagonal pixels [4] as well. In other cases, e.g. for the recently proposed Penrose pixels [1], novel spatial bases that are periodic in the aperiodic Penrose sampling pattern would have to be constructed.

1.2. Proof of Theorem 2

The proof of Theorem 2 follows Equations 5-7 from the paper. To provide additional detail, we start with Equation 7

$$\mathcal{F}_x \mathbf{i} = \sum_{k=1}^N \left(\delta(\vec{\omega}_x - \Delta \vec{\omega}_x^k) \otimes \sum_{j=1}^N \hat{\sigma}_j^k \hat{\rho}_j(\vec{\omega}_x) \right), \quad (7)$$

where $k\Delta \vec{\omega}_x$ has been replaced by $\Delta \vec{\omega}_x^k$, a vector to the center frequency of a Fourier channel, Figure 1 (right), to allow for generalized sampling patterns. If the signal is properly band-limited the image's Fourier transform separates into disjoint sets,

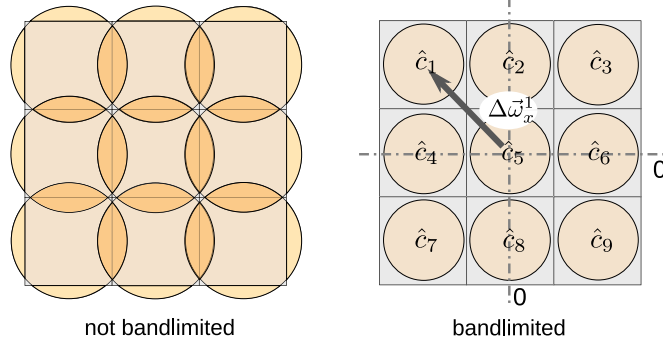


Figure 1. Information overlaps in the Fourier domain if the signal is not suitably band-limited (left). With the appropriate band-limitation, the Fourier representation decomposes into distinct correlated Fourier channels \hat{c}_k . The arrow in the right figure indicates the position $\Delta\vec{\omega}_x^k$ of the Dirac peaks, i.e. the center frequencies, of a Fourier channel.

each encoding one Fourier channel \hat{c}_k , Figure 1 (right). In the following we again consider a rectangular sampling lattice. In this case a Fourier channel can be obtained by masking the region around the channel's Dirac peak using a rectangle function¹:

$$\begin{aligned}
\hat{c}_k(\vec{\omega}'_x) &= \text{rect}^k(\vec{\omega}'_x) \delta(\vec{\omega}'_x) \otimes \sum_{j=1}^N \hat{\sigma}_j^k \hat{\rho}_j(\vec{\omega}'_x). \\
&= \text{rect}^k(\vec{\omega}'_x) \sum_{j=1}^N \hat{\sigma}_j^k \int_{\mathcal{D}} \pi(\vec{p}) \hat{l}_\lambda(\vec{\omega}'_x, \vec{p}) d\vec{p}. \\
&\approx \text{rect}^k(\vec{\omega}'_x) \sum_{j=1}^N \hat{\sigma}_j^k \sum_{l=1}^M \pi_j^l \hat{l}_\lambda(\vec{\omega}'_x, \vec{p}^l).
\end{aligned} \tag{8}$$

To separate sensor image and Fourier channel spatial frequencies, we have used the substitution $\vec{\omega}'_x = \vec{\omega}_x - \Delta\vec{\omega}_x^k$. These are the Fourier channel spatial frequencies introduced in the paper. Since the $\vec{\omega}'_x$ describe corresponding channel spatial frequencies for the different Fourier channels \hat{c}_k we use a single symbol to describe them. The convolution with the comb function in Eq. 7 reduces to a convolution with a single Dirac peak $\delta(\vec{\omega}'_x)$ because of band-limitation. This is a unit operation and is thus removed from the equations. In the last row, again an approximation is introduced by discretizing the plenoptic basis. Equation 8 proves Theorem 2. \square

In addition, by inverse Fourier transforming and sampling the Fourier channel \hat{c}_k we see that the reconstruction filter is indeed a sinc as claimed in the paper:

$$\hat{c}_k(\vec{x}) = \sum_{j=1}^N \sigma_j^k \sum_{l=1}^M \pi_j^l \left(\text{III}_k(\vec{x}) \mathbf{l}_\lambda(\vec{x}, \vec{p}^l) \otimes \text{sinc}(\vec{x}) \right). \tag{9}$$

Please note that the inverse transformed Fourier channels \hat{c}_k , Eq. 9, and the measurement channels c_k , Eq. 5 are not direct Fourier equivalents but are related by the Fourier transform of the spatial basis in a single super-pixel.

The explicit form of Eq. 8 is very similar to the spatial case, Eq. 6. The short form is $\vec{c}(\vec{\omega}'_x) = \hat{\Sigma} \mathbf{\Pi}^T \vec{l}_\lambda(\vec{\omega}'_x)$ as provided in the paper.

2. Additional Details for Case Studies

2.1. Imaging through Color Filter Arrays

In this subsection, we would like to give mathematical expressions for the spatial basis functions of a Bayer pattern, as space restrictions allow us to only show this in Figure 3 of the manuscript otherwise. Section 4.1 in the manuscript discusses the plenoptic bases and also points out that the spatial basis functions are differently shifted Dirac peaks, thus

¹For other sampling patterns corresponding, e.g. hexagonal, masking functions would be used.

$$\Sigma_{Bayer} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (10)$$

which is defined up to a permutation of the rows, corresponding to different arrangements of the R,G,B filter elements. The Fourier correlation matrix $\hat{\Sigma}_{Bayer}$ is then given as

$$\hat{\Sigma}_{Bayer} = \begin{pmatrix} 1 & 2 & 1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & -2 & 1 \end{pmatrix}. \quad (11)$$

2.2. Light Field Acquisition

2.2.1 Non-Refractive Modulators

We would like to show additional results for light fields that were captured with attenuation masks and previously reconstructed in the Fourier domain. Figure 2 shows an additional viewpoint of the reconstructed Mannequin dataset (Lanman et al. [5]) and, more importantly, a spatial reconstruction using the PSM theorem with a non-linear filter, specifically a joint bilateral filter. The non-linear filter is just one of many possible choices for sophisticated spatial reconstruction filters that are facilitated by the PSM theorem. The joint bilateral filter can, in this case, reconstruct a slightly sharper image than a bicubic filter.

Although Corollary 1 is theoretically only valid for linear filters (as proven in Section 1), we demonstrate in Figure 2 that it can practically also work well with non-linear filters.




Figure 2. Additional reconstructions for Mannequin data set [5]. The results are three-times upsampled and show two different views of the reconstruction. The right-most column shows an additional example (cf. Figure 4 in the manuscript) that was computed with a non-linear spatial filter - the joint bilateral filter.

2.2.2 Lens-Based Modulators

For completeness, we would like to show how our theory applies to refractive optical elements. Instead of modeling refractions as a general geometric mapping from sensor space to world space defined in the camera, we can model it as a convolution of the plenoptic function before modulation (including refraction) and a plenoptic modulation kernel $\mathbf{k}(\vec{x}, \vec{x}_c, \vec{v}_c)$ (see Levin et al. [6]). The convolution disregards area integration over pixel surfaces and photometric terms. A sensor image is then defined as

$$\mathbf{i}(\vec{x}) = \int_{\vec{x}_c} \int_{\vec{v}_c} \mathbf{k}(\vec{x}, \vec{x}_c, \vec{v}_c) \mathbf{l}_\lambda(\vec{x}_c, \vec{v}_c) d\vec{x}_c d\vec{v}_c, \quad (12)$$

where \vec{x} is the spatial coordinate on the sensor surface and \vec{x}_c is the spatial coordinate defined on the plane of the refractive element, which is parallel to the sensor plane at a distance z . The directional coordinates \vec{v}_c describe the directional variation of the plenoptic function before modulation by the kernel inside the camera behind the main lens. Plenoptic modulation

kernels for different optical elements are well known. Under paraxial approximations and disregarding the element's aperture and wavelength of light, the kernel for most refractive elements is of the form

$$\mathbf{k}(\vec{x}, \vec{x}_c, \vec{v}_c) = \delta(\vec{x} - z\vec{v}_c - \phi(\vec{x}_c)). \quad (13)$$

The specific kernel for a lens at focal distance from the sensor is given by $\phi_{z=f}(\vec{x}_c) = 0$. The term becomes non-zero when the lens is moved away from the focal length, as proposed by Lumsdaine and Georgiev [7], $\phi_{z \neq f}(\vec{x}_c) = s\vec{x}_c$, where $s = 1 - z/f$ is the slope of the integration surface. The kernel for a cubic phase plate (Dowski and Cathey [3]) is a parabola, thus $\phi_{cubic}(\vec{x}_c) = a\vec{x}_c^2$, where a is proportional to the curvature of the parabola. The optical setups and integration surfaces for all of these cases are illustrated in Figure 3.

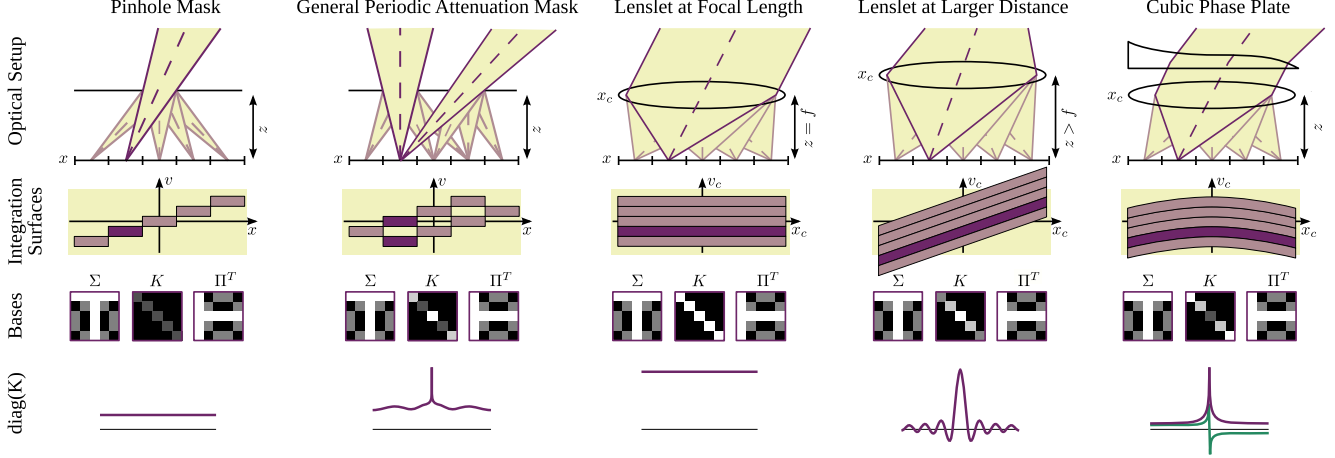


Figure 3. In addition to the standard plenoptic cameras already shown in Figure 3 in the manuscript, we also illustrate the optical setups, integration surfaces, spatial and plenoptic bases as well as the Fourier weighting factors κ for the focused plenoptic camera (center right column) and a lens at focal distance combined with a cubic phase plate (right column) here. Please note that the camera and sensor coordinates for non-refractive elements are identical. Furthermore, the integration surfaces for the refractive elements are illustrated in camera coordinates.

In the following, we demonstrate how the separation of the plenoptic modulator into a spatial and a plenoptic basis can be performed. This is shown for a single optical element from an array as employed by all plenoptic cameras. We assume that the plenoptic function is spatially band-limited, which in this case requires it not to spatially vary over a *single optical element*, i.e. $\mathbf{l}_\lambda(\vec{x}_c, \vec{v}_c) = \mathbf{l}_\lambda(\vec{v}_c)$.

Combining Equations 12 and 13 yields

$$\begin{aligned} \mathbf{i}(\vec{x}) &= \int_{\vec{x}_c} \int_{\vec{v}_c} \mathbf{l}_\lambda(\vec{v}_c) \delta(\vec{x} - z\vec{v}_c - \phi(\vec{x}_c)) d\vec{x}_c d\vec{v}_c \\ &= \int_{\vec{x}_c} \int_{\vec{v}_c} \mathbf{l}_\lambda(\vec{v}_c) \int_{\vec{\omega}_x} e^{2\pi i(\vec{x} - z\vec{v}_c - \phi(\vec{x}_c)) \cdot \vec{\omega}_x} d\vec{\omega}_x d\vec{x}_c d\vec{v}_c \\ &= \int_{\vec{\omega}_x} e^{2\pi i\vec{x} \cdot \vec{\omega}_x} \int_{\vec{v}_c} \mathbf{l}_\lambda(\vec{v}_c) e^{-2\pi i z\vec{v}_c \cdot \vec{\omega}_x} d\vec{v}_c \int_{\vec{x}_c} e^{-2\pi i\phi(\vec{x}_c) \cdot \vec{\omega}_x} d\vec{x}_c d\vec{\omega}_x. \end{aligned} \quad (14)$$

We see that just like in the case of general attenuation masks, refractive optical elements can be separated into a spatial and a plenoptic basis. Again, the former is the inverse Fourier basis, while the latter is the Fourier basis. The extra term $\kappa_j = \int_{\vec{x}_c} e^{-2\pi i\phi(\vec{x}_c) \cdot \vec{\omega}_x} d\vec{x}_c$ varies for different optical elements and represents the optical transfer function of that particular setup (see Figure 3). Lenses at a distance to the sensor that is equal to the focal length are not affected by the extra term, i.e. $\kappa_j = 1$, j . Therefore, the combined spatial and plenoptic basis is the identity basis, which we illustrated in Figure 3 in the manuscript.

Based on the band-limited assumption, the scalars κ_j can be separated from the plenoptic basis, which allows us to generalize this separation for arrays of optical elements. Just like in the case of non-refractive modulation masks, the PSM theorem can be applied by upsampling the sensor measurements before a local decorrelation is performed.

Let us show a comparison of spatial and Fourier reconstructions of spatially multiplexed light fields with a lenslet array mounted at the focal distance to the sensor. A decorrelation of the plenoptic channels in the spatial domain is in this case redundant, as every sensor pixel measures uncorrelated directional samples of the plenoptic function. The re-weighting step in the Fourier reconstruction by the factors κ_j becomes obsolete as all $\kappa_j = 1$. In this example, the main difference between the spatial reconstruction (PSF theorem) and the Fourier reconstruction (PFM theorem) is the cropping operation of the Fourier tiles followed by an inverse 4D Fourier transform. As discussed in the manuscript, the operation is equivalent to a sinc reconstruction filter in the spatial domain and results in ringing artifacts.

Figure 4 shows a comparison between different upsampled reconstructions. The dataset was captured with a lenslet-based light field microscope (lightfield.stanford.edu). As we can see in the magnification and contrast enhanced difference images, the Fourier reconstructions result in ringing artifacts. Although the artifacts are clearly visible, they are more subtle than those in the mask-based light field reconstruction of Figure 4 in the manuscript. This can be attributed to the pixels under each microlens integrating spatial light variation over the entire area of the corresponding lenslet, which provides a proper spatial band-limitation and therefore minimizes aliasing.




Figure 4. Reconstruction of spatially multiplexed light fields. The images show a three-times upsampled spatial reconstruction with a bicubic filter (left), a Fourier reconstruction in the original resolution followed by a bicubic upsampling (center), and a Fourier reconstruction with padding in the Fourier domain (right). One of the views and a contrast enhanced difference image to the spatial reconstruction are shown in the magnifications.

Employing lenslets at a distance to the sensor that is different to their focal length, i.e. $z = f$, was explored by Lumsdaine and Georgiev [7]. It was shown that this allows light fields with a higher spatial, but lower angular resolution to be captured. In this optical setup each sensor pixel integrates over a range of incident light directions, which is equivalent to a convolution with a rect kernel. A deconvolution by the corresponding optical transfer function, that is a sinc as seen in Figure 3, is ill-posed because of its zero crossings. It was shown that a custom resorting algorithm [7] can perform the proposed spatio-angular resolution tradeoff. Applying our theory would aim at reconstructing the full spatial and directional resolution of the light field. This particular optical setup is, due to the zero-crossings of the optical transfer function, ill-conditioned in our reconstruction framework unless statistical priors are incorporated.

Cubic phase plates have been well-studied for extended depth-of-field photography [3]. Employing these refractive elements in combination with a lenslet array in plenoptic cameras, as illustrated in Figure 3, could also extend the depth-of-field of light fields and make a reconstruction at full spatial and directional resolution with our theory feasible. However, this optical setup is to be practically explored in the future.

Due to the space restrictions of the manuscript we could not include the analysis for refractive optical elements in our submission.

3. Additional Details for Noise Analysis

In this section we show additional results and a more detailed analysis of the noise in light fields that are reconstructed from simulated sensor images of different multiplexing schemes. Specifically, we focus on a comparison of various attenuation mask types (pinholes, sum-of-sinusoids, and MURAs) and the standard plenoptic camera that employs a lenslet array.

Assuming only additive, signal independent noise in the captured sensor images allows us to quantify the noise of the reconstructed light field with the covariance matrix C :

$$C = \zeta^2 \left((\Sigma \Pi^T)^T (\Sigma \Pi^T) \right)^{-1}, \quad (15)$$

where ζ^2 is the variance of an additive Gaussian noise distribution with zero-mean in the sensor image $i(\vec{x})$. Figure 5 shows the magnitudes of C assuming $\zeta^2 = 1$. Values larger than 1 amplify noise in the camera image and off-diagonal entries accumulate noise from different regions of the captured images. The covariance matrix of the sum-of-sinusoids (SoS) mask has many large-valued off-diagonal entries, which indicates noise amplification in the reconstruction, whereas the MURA mask does have off-diagonal entries, but with much smaller magnitudes than both pinhole and SoS masks. The same can be inferred from the plots of the singular values of the multiplexing matrix $\Sigma\Pi^T$ in Figure 5 (right).




Figure 5. Covariance matrices and eigenvalues for different light field multiplexing schemes. Large values and especially off-diagonal entries in the covariance matrices indicate amplification of additive noise in the sensor images. The sum-of-sinusoids mask is thus expected to perform worse than a pinhole mask for dominating dark current noise in the sensor image, which can also be inferred from the plots of the multiplexing matrix’s singular values (right).

Figure 6 shows more results of noise amplification for light field acquisition with different multiplexing schemes. In this case both signal-independent additive noise and signal-dependent photon noise are considered. The parameter χ^2 is the ratio of photon noise variance and additive noise variance. As expected from the noise performance plot in our manuscript and the covariance matrices and singular values in Figure 5, lenslet arrays and MURA masks have a better noise performance than pinhole masks, whereas SoS masks perform worse (see Figure 6, row 2). The value $\log(\chi^2) = -0.93$ corresponds to a PointGrey Dragon y camera given by Schechner et al. [8], where the additive noise term dominates. The lower two rows in Figure 6 show how the reconstruction noise increases for an increasingly dominating photon noise term up to a point where a pinhole mask performs better than even the MURA mask. Please note that for our analysis the exposure times of the simulated sensor were equal for each method (thus the intensity difference of the sensor images in Fig. 6, upper row). Furthermore, we disregard quantization and other sensor non-linearities for our analysis.

References

- [1] M. Ben-Ezra, Z. Lin, and B. Wilburn. Penrose Pixels Super-Resolution in the Detector Layout Domain. In *Proc. of ICCV*, pages 1–8, 2007. 2
- [2] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Information Theory*, 52(2):489–509, Feb. 2006. 2
- [3] E. Dowski and T. Cathey. Extended Depth of Field through Wave-Front Coding. *Applied Optics*, 34(11):1859–1866, 1995. 5, 6
- [4] X. He and W. Jia. Hexagonal Structure for Intelligent Vision. In *Proc. of International Conference on Information and Communication Technologies (ICICT)*, pages 52–64, 2005. 2
- [5] D. Lanman, R. Raskar, A. Agrawal, and G. Taubin. Shield Fields: Modeling and Capturing 3D Occluders. *ACM Trans. Graph. (Siggraph Asia)*, 27(5):1–10, 2008. 4
- [6] A. Levin, S. W. Hasinoff, P. Green, F. Durand, and W. T. Freeman. 4D Frequency Analysis of Computational Cameras for Depth of Field Extension. *ACM Trans. Graph. (Siggraph)*, 28(3):1–14, 2009. 4
- [7] A. Lumsdaine and T. Georgiev. The Focused Plenoptic Camera. In *Proc. IEEE ICCP*, 2009. 5, 6
- [8] Y. Schechner, S. Nayar, and P. Belhumeur. Multiplexing for Optimal Lighting. *IEEE Trans. PAMI*, 29(8):1339–1354, 2007. 7




Figure 6. Comparison of noise amplification for different light field acquisition schemes on the golgi stained neuron dataset (light-field.stanford.edu). The upper row shows simulated sensor images with contrast enhanced close-ups. The other rows show a single view of the reconstructed light field from a noisy sensor image. The ratio χ^2 of signal-dependent photon noise and signal independent dark current noise varies for the different reconstructions. Row 2 simulates a reconstruction with a dominating additive noise term, while rows 3 and 4 show the effect of an increasingly dominating photon noise term in the sensor images.