

Robustness of Planar Fourier Capture Arrays to Colour Changes and Lost Pixels

Patrick R. Gill^{ab*}, Changhyuk Lee^a, Sriram Sivaramakrishnan^a and Alyosha Molnar^a

*^aSchool of Electrical and Computer Engineering
Cornell University
223 Phillips Hall
Ithaca, NY 14853
USA*

*^bDepartment of Psychology
Sidney Smith Hall
University of Toronto
Toronto, ON M5S 3G3
Canada
E-mail: patrick.gill@utoronto.ca*

ABSTRACT: Planar Fourier capture arrays (PFCAs) are optical sensors built entirely in standard microchip manufacturing flows. PFCAs are composed of ensembles of angle sensitive pixels (ASPs) that each report a single coefficient of the Fourier transform of the far-away scene. Here we characterize the performance of PFCAs under the following three non-optimal conditions. First, we show that PFCAs can operate while sensing light of a wavelength other than the design point. Second, if only a randomly-selected subset of 10% of the ASPs are functional, we can nonetheless reconstruct the entire far-away scene using compressed sensing. Third, if the wavelength of the imaged light is unknown, it can be inferred by demanding self-consistency of the outputs.

KEYWORDS: Planar Fourier Capture Array; Compressive Sensing; Angle Sensitive Pixel.

*Corresponding author.

Contents

1. Introduction

1.1 Planar Fourier Capture Arrays

Planar Fourier capture arrays (PFCAs) [1] are ensembles of angle sensitive pixels (ASPs) [2, 3] that directly capture the Fourier transform of a far-away image. The light sensitivity of an individual ASP is a sinusoidal function of incident angle [2]. A carefully-chosen ensemble of ASPs can report the entire Fourier transform of the far-away scene up to the Nyquist limit set by the highest-frequency ASP in the ensemble [1]. As ASPs are extremely thin, light and cheap to manufacture, PFCAs have the potential to become a disruptive technology in the fields of sensing and robotics [4, 5, 6, 7] when volume, mass or cost constraints are paramount.

For PFCAs to reach their full potential, it would be useful to understand some of the limitations and freedoms afforded by this new class of imager. Our first PFCA was optimized for green incident light, and the reconstruction algorithms we used previously [1] assume all sensors are functional. In this publication we thoroughly report PFCA design considerations, then investigate the performance of PFCAs imaging light of a different colour than the design specification, both when that colour is known *a priori* and when it must be determined blindly. Further, we will show that by using a signal-processing technique called compressed sensing, we can reconstruct images even when the outputs of only a small fraction of the ASPs are available.

1.2 Angle Sensitive Pixels

Angle sensitive pixels (ASPs) [2, 3] are photosensitive pixels that, using diffraction gratings, couple to far-field light sources with an efficiency that is a sinusoidal function of incident angle. ASPs attain their angle sensitivity through the moiré effect using two metal gratings with identical spacing but at different heights above a photodiode below (see Figure 2). Due to the vertical displacement of the gratings, the angle subtended by one spatial period of the top grating viewed by the photodiode is slightly less than that subtended by the bottom grating. This angular disparity invokes the moiré effect, making the effective transmission aperture a sinusoidal function of incident angle. Both of the gratings can be manufactured using metal interconnect layers intrinsic to the CMOS process, and the photodiodes can be manufactured using intrinsic semiconductors, meaning ASPs can be manufactured entirely using existing unmodified CMOS process flows.

The moiré effect is essentially a ray optics effect, however at the small spatial scale of CMOS structures photons cannot be modeled as non-diffracting rays. Fortunately, the moiré effect can be rescued by exploiting the Talbot effect [8, 9]. The Talbot effect is the property that a regular periodic grating illuminated by monochromatic light creates self-images at integer multiples of the Talbot depth $Z_T = \frac{2p^2}{\lambda} - \frac{\lambda}{2}$ where p is the grating pitch (spatial period) and λ is the wavelength of

Talbot Effect




Figure 1. The Talbot Effect. Wave simulation of monochromatic light striking a grating at the normal. The grating periodicity is recapitulated at regularly-spaced depths below the grating; a phenomenon known as the Talbot effect. The Talbot depth Z_T is shown at right.

light (see Figures 1 and 2). By ensuring that the second grating is located at an integer multiple of half¹ the Talbot depth, the moiré effect is rescued and further diffraction effects can be neglected.

The transfer function of light incident on an ASP can be modeled as in Equation 1.1.

$$R = I_0(1 - m \cos(b\theta + \alpha))F(\theta)(1 + \eta), \quad (1.1)$$

where R is the readout of the ASP, I_0 is proportional to the light flux at the ASP, θ is the incident angle along the sensitive axis, b is the angular sensitivity (designed to range from 7 to 39 - see Equation 2.1), m is the modulation depth of the ASP (see Equation 2.2), α is a designable phase offset caused by a displacement between the top and bottom gratings, $F(\theta)$ is a slowly-varying aperture function and η is multiplicative noise.

2. Designing PFCAs

An ASP's angular sensitivity (b from Equation 1.1) can be found by considering the angular displacement required for a lobe of the Talbot diffraction pattern to traverse a full period of the analyzer grating, as in Equation 2.1.

$$b = 2\pi \frac{z_{\text{effective}}}{pn}, \quad (2.1)$$

where $z_{\text{effective}}$ is the effective vertical displacement between gratings (see Equation 2.3) and n is the refractive index of the medium, in this case SiO_2 . The depths z at which modulation depth m is maximal correspond to integer multiples of half the Talbot depth, as in Equation 2.2.

$$m \text{ maximal when } z = a \left(\frac{p^2}{\lambda} - \frac{\lambda}{4} \right); a \in \mathbb{I}. \quad (2.2)$$

¹At half-integer multiples of the Talbot depth the periodicity of the initial grating reemerges, but with its phase reversed; see Figures 1 and 2.




Figure 2. Interaction Between Incident Light and ASPs. Left: light incident at 10° produces spatial maxima that align with gaps in the second grating, resulting in a large photocurrent. Center: light incident at 25° produces spatial maxima that are largely blocked by the second grating, resulting in a small photocurrent. Right: photocurrent as a function of incident angle is a sinusoidal function.

Note that the $\frac{\lambda}{4}$ term is often omitted, but is derivable from a consideration of optical path lengths and provides a noticeable increase in accuracy when modeling and designing ASPs.

2.1 Effective Depth

Since the PFCA was built, detailed measurements have indicated that when using Equation 2.1, b is more accurately predicted by replacing the vertical separation between gratings z with an effective depth, $z_{\text{effective}}$, set by the nearest half-Talbot depth as follows.

$$a_{\text{closest Talbot}} = \text{round} \left(\frac{z}{\left(\frac{p^2}{\lambda} - \frac{\lambda}{4} \right)} \right)$$

$$z_{\text{effective}} = a_{\text{closest Talbot}} \left(\frac{p^2}{\lambda} - \frac{\lambda}{4} \right) \quad (2.3)$$

At optimal depths for a particular λ , $z_{\text{effective}} = z$. However, for non-optimal λ s, $z_{\text{effective}}$ and thus b vary as a function of λ , an effect not noted previously [1, 2, 3].

The phenomenon that the effective depth is the closest ideal Talbot pattern was demonstrated in a test structure accompanying our original PFCA prototype. A long, linear array of photodetectors was built with smoothly varying pitch. The observed angular sensitivity b as a function of pitch p is plotted in Figure 3, along with predictions based upon actual and effective depths applied to Equation 2.1. It can be seen that parameter b is better predicted by the effective depth from Equation 2.3 than the actual separation of the gratings.

2.2 Combining ASPs into a Fourier-complete array using CMOS

The metal interconnect layers of a CMOS process are manufactured at process-specific heights above the silicon substrate; to work within an established CMOS process it is necessary to manufacture all metals at one of the specified depths. There are therefore a discrete spectrum of inter-metal depths available. Using Equations 2.2 and 2.1, it is possible to determine the spectrum of manufacturable b values that have locally-maximal m (see filled circles of Figure 4).

For Fourier completeness, it is important to sample Fourier space with an adequate density to ensure coverage. The more densely-packed the Fourier frequency measurements, the larger




Figure 3. Closest Talbot Pattern Determines b . Observed $b(p)$ agrees better with the assumption that $z_{\text{effective}}$ equals the closest optimal Talbot depth, and not the manufactured depth. Blue circles: measured b as a function of pitch for incident green light on a series of devices spanning a range of p from 0.7 to 1.3 microns. Red xs: modeled b assuming Equations 2.1 and 2.3. Black line: modeled b assuming $z_{\text{effective}} = z$ and Equation 2.1. Grey bars: depth of closest Talbot pattern.




Figure 4. Selecting Devices for the PFCA. Filled black circles indicate manufacturable devices with maximal m for 520nm light; open circles indicate p and (λ -dependent) b for the suite of devices we manufactured.

the angular region that can be observed with Nyquist sampling. As reported previously [1], the relationship between the maximum allowable incident light angle h and the maximum difference in b between consecutive designs is

$$h = \frac{180^\circ}{\sqrt{2}\Delta b}. \quad (2.4)$$

The first prototype we built included devices with p tuned for a range of wavelengths other than the design point (520 nm) by manufacturing devices with p differing slightly from that of locally-optimal devices (see Equation 2.2). The ASPs included in the manufactured device are shown as empty circles in Figure 4; note each design of constant p and z has different b for different λ s. We arranged 1444 ASPs in two PFCAs with complementary α s (see Equation 1.1). All ASPs of a given design (combination of z and b) are found in one of 18 concentric rings arranged around four

low- b devices described elsewhere [10]. Rings with higher b are placed further from the center to allow for a greater variety of grating orientations, as is required for Fourier tiling.

Resolution limits are set by the maximum b of any device b_{max} . This Nyquist limit corresponds to two rows or columns of pixels per period of the highest- b device. For $h \approx 45^\circ$, the total number of effective pixels is $\left(\frac{b_{max}}{2}\right)^2$, or approximately 400 pixels for this prototype².

The prototype's bs span the space relatively well for blue and green light (largest Δb is 3.6 and 3.3 respectively; corresponding hs are 35° and 38°), but a gap in bs for red light in the $b = 13$ – 21 range reduces our expectations that red images will be recovered well for half-angles $h > 17^\circ$. Furthermore, as the pitch of a device approaches a single wavelength, its signal-to-noise ratios degrades, such that red light will intrinsically provide worse performance in fine-pitch ASPs. Both of these disadvantages working under red light contribute to poor reconstructions (see Section 3). The manufactured device is seen in Figure 5.




Figure 5. Manufactured PFCA. Concentric rings of ASPs with increasingly high b yield a complete Fourier description of the light intensity from the far field. Slowly varying orientation is evident from the enlarged section, where schematics show different metal layers in different colours.

2.3 Reconstructing Images

The prototype PFCA was presented with calibration (Figure 6A) and test images using a square CRT area 20cm on a side, 22.86cm from the PFCA³ for an h of 31.7° at the square's corners. This h is small enough to allow full Fourier coverage for blue and green, but not red light (see Equation 2.4).

²As ASPs' transfer functions are not pure Fourier components but contain some harmonics, a limited amount of information outside the expected Nyquist limit is available so long as the coverage of Fourier space is overcomplete and noise levels are low.

³As $23\text{cm} \gg 570\mu\text{m}$ (the PFCA's size), images presented are in the far field regime where the light field at each point in the PFCA is essentially identical.




Figure 6. PFCA Calibration. Transfer functions of each pixel are found by presenting 6710 random binary 50×50 calibration images (A) on a CRT screen (presentation time 16.7 ms each) to the array (B) and performing reverse correlation between the observed photocurrent of each sensor and the image presented. The kernels of three ASPs are shown (C); these resemble Fourier components.

To reconstruct images computationally, we performed the following operations. First, we computed the relationship between the photocurrent and the observed voltage drops at the photodiodes by fitting a quadratic function to $V(t)$ under constant illumination from an LED lit by a steady power source. Then we threw out all data from any ASP that was saturated for more than 30% of all image presentations (24.5% of all sensors). This was necessary because long range diffusion of photo-generated carriers tended to add a background photocurrent to many of the ASPs close to the edge of the array. Next, we performed reverse correlation [11] between inferred photocurrents and the calibration images with the following steps. We computed the pseudoinverse of the matrix⁴ of calibration stimuli $\mathbf{C}$, and multiplied it by the matrix of responses $\mathbf{R}$ to obtain an estimate of the individual transfer functions $\mathbf{H}$ as follows. We assume $\mathbf{R}$ is linearly related to $\mathbf{C}$ by $\mathbf{H}$, therefore

$$\mathbf{R} = \mathbf{CH}. \quad (2.5)$$

$\mathbf{C}$ is not square since there are 6710 patterns of 2500 free parameters each, so we can compute and multiply both sides of Equation 2.5 by $\mathbf{C}$'s pseudoinverse $(\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T$.

$$\begin{aligned} (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T \mathbf{R} &= (\mathbf{C}^T \mathbf{C})^{-1} (\mathbf{C}^T \mathbf{C}) \mathbf{H}. \\ (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T \mathbf{R} &= \mathbf{H}. \end{aligned} \quad (2.6)$$

Having computed $\mathbf{H}$ (sample rows of $\mathbf{H}$ are shown in Figure 6C), we need to find a way to invert it so we can reconstruct a new stimulus $\mathbf{s}_{\text{new}}$ given a new set of responses $\mathbf{r}_{\text{new}}$. We computed the eigenspectrum of $\mathbf{H}\mathbf{H}^T$ (see Figure 7) and found a regularized pseudoinverse of $\mathbf{H}$ using ridge regression [12] as follows. Ridge regression de-emphasizes eigenvectors of $\mathbf{H}\mathbf{H}^T$ that have little power. It is an acknowledgement that some stimulus eigenvalues are not represented strongly in $\mathbf{R}$ and thus should be attenuated to avoid letting noise in poorly-determined stimulus components overwhelm signal in well-determined components.⁵

$$\begin{aligned} \mathbf{r}_{\text{new}} &= \mathbf{s}_{\text{new}} \mathbf{H} \\ \mathbf{r}_{\text{new}} \mathbf{H}^T (\mathbf{H}\mathbf{H}^T + \Lambda \mathbb{I})^{-1} &\approx \mathbf{s}_{\text{new}} \end{aligned} \quad (2.7)$$

⁴Matrix quantities such as $\mathbf{C}$ shall be in capital bold, vector quantities in lowercase bold.

⁵We also attempted ridge regression while determining $\mathbf{H}$ in Equation 2.6, but it did not improve our results.




Figure 7. Eigenspectrum of $\mathbf{H}$. The strongest 36 eigenvectors of $\mathbf{H}$ for green light are shown in order from left to right, then top to bottom.

Here Λ was chosen by inspection of the eigenspectrum of $\mathbf{H}\mathbf{H}^T$, and values of 30, 20 and 20 were used for reconstructions with red, green and blue light, respectively (see Section 3) to reflect cutoffs below which noise dominates signal. These cutoffs preserved approximately 400 eigenvectors of $\mathbf{H}\mathbf{H}^T$ for green light, implying a 20×20 pixel resolution consistent with $\left(\frac{b_{max}}{2}\right)^2$. Using Equation 2.7, we reconstructed presented images up to this resolution limit as in Figure 8.




Figure 8. Image Reconstructions. Using the basis functions obtained in the calibration phase (Fig. 6), we reconstructed (B) the image presented (A) up to the Nyquist limit of our array. No off-chip optics were used; accumulation times were 16.7 ms.

3. Robustness to Changes in Imaged Wavelength

Although the prototype PFCA was designed for green light, we also calibrated and tested it with blue and red light from the CRT. For each wavelength of light, a separate $\mathbf{H}$ was calculated (see Equation 2.6) and images were reconstructed as for green light (see Equation 2.7). The reconstructed images for all three colours can be found in Figure 9.




Figure 9. Three Colour Image Reconstructions. Using the basis functions $\mathbf{H}$ obtained from red, green and blue calibration runs, the first prototype chip is able to resolve light images other than at the design wavelength (520nm). The center green image is identical to that of Figure 8B on the left.

As expected, blue light does not produce images of the same quality as green light, and red light yields much worse reconstructions. The poor reconstructions available from red light are due to the large gap in Fourier information available in the $b = 13\text{--}21$ range⁶ (see Figure 4) and the fact that red light cannot pass easily through the smallest gratings in use due to its larger wavelength resulting in lower signal to noise ratios.

4. Compressed Sensing Reconstructions

4.1 Exceedingly Brief Introduction to Compressed Sensing

Compressed sensing (CS) is a recently-developed signal processing tool useful for reconstructing signals with only partial observations [13]. In general, such underdetermined reconstructions are impossible, however if the signal is known to be sparse in some basis (having no more than k large-magnitude components), then incomplete measurements in some *other* basis (unrelated to the basis in which the signal is sparse) can be sufficient to determine the signal so long as comfortably⁷ more than k measurements are made. CS operates by finding a solution $\mathbf{x}$ that is sparse (i.e. containing mostly 0 entries) yet consistent with the signal $\mathbf{y}$. If the expected signal given a solution vector $\mathbf{x}$ is $\mathbf{Ax}$ and λ is a regularization parameter, sparse $\mathbf{x}$ consistent with $\mathbf{y}$ are found by solving Equation 4.1, known as the Basis Pursuit Denoising (BPDN) problem [14].

$$\operatorname{argmin}_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2 + \lambda \|\mathbf{x}\|_1 \quad (4.1)$$

⁶The PFCA’s blind spot to frequencies in this range is more evident in Figure 11, columns A1 and A2, rows G and B. These images are derived from Equation 2.7 using the $\mathbf{H}$ for red light but with green and blue light incident, leading to high noise. Noise in the $b = 13\text{--}21$ range is especially prominent, leading to artifactual “bull’s eyes” in the reconstruction.

⁷See [13] for a theoretical guarantee of what “comfortable” means in this context, although provable guarantees place much more stringent requirements than practically necessary.

Solving Equation 4.1 leads to a sparse explanation of the signal, which will coincide with the true signal if it is indeed sparse.

4.2 Compressed Sensing Permits Full Reconstructions from Incomplete PFCA Information

The basis in which natural images are sparse is close to a wavelet pyramid [15], which is different enough from the Fourier basis that a sparse collection of ASPs from a Fourier pixel should be sufficient to reconstruct a full image, provided it is sparse in the wavelet domain. We therefore projected $\mathbf{H}$ into an overcomplete (2500×8619) wavelet basis [16] $\mathbf{W}$ and declared $\mathbf{A}$ from Equation 4.1 to be $\mathbf{HW}$. The CS reconstruction problem for recovering the image given limited observations becomes

$$\mathbf{s}_{\text{CS}} = (\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T \left(\underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{r}_{\text{sub}} - \mathbf{H}_{\text{sub}} \mathbf{W} \mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1 \right) \quad (4.2)$$

where sub indicates only a randomly-chosen small subset (40%, 25%, 10% or 8%) of the responses $\mathbf{r}$ are kept to simulate incomplete measurements. CS reconstructions are of relatively high quality despite randomly subsampling Fourier space, as seen in Figure 10.




Figure 10. Compressed Sensing Results. We show compressed sensing reconstructions of the test image using a randomly-selected subset of ASP sensors. Proportion of sensors whose data are used shown in lower-left corner.

As part of related work solving underdetermined imaging problems [3], we wrote the in-crowd algorithm for solving BPDN [14], which is faster than all alternative methods on sparse, large-scale BPDN problems such as this one. On a modern desktop computer, computing the exact BPDN solution for our problems takes approximately one second with the in-crowd algorithm, whereas other popular alternatives such as GPSR may take longer than an hour for identical problems.

PFCAs are thus able to relate complex images to a certain degree even when over 90% of their constituent sensors malfunction. Alternatively, it would be possible to design a PFCA with a random collection of only 10% of the sensors traditionally considered essential for Fourier coverage and still be able to reconstruct images, provided they are sparse in some known, non-Fourier basis. This ability is derived from the fact that each ASP makes measurements in a basis in which natural images are not sparse (i.e. Fourier components), yet natural images have a known sparsity. There is only one sparse combination of wavelets that satisfies the incomplete Fourier measurements taken by the PFCA; this is the solution $\mathbf{s}_{\text{CS}}$ of Equation 4.2.

5. Determining Colour Blindly

Up until now, we have assumed the reconstruction process has access to information about the wavelength of the light imaged. Introducing uncertainty in the wavelength imaged makes reconstruction underdetermined by a factor equal to the number of possible wavelengths of light, since

the number of unknowns becomes the number of pixels times the number of λ s. However, recovering full colour information may still be possible. In Section 4.2 we showed that an undercomplete subset of ASP readings can be sufficient to reconstruct a full image. Moreover, as shown in Figure 4, the b of every device is wavelength-sensitive, meaning that PFCA signals should be sensitive enough to λ to discriminate colours. Given this wavelength sensitivity along with the robustness to undercomplete measurements as shown in Section 4.2, it should be possible to determine the colour of light incident on the PFCA merely by assuming the image is composed of a sparse combination of wavelets of monochromatic light.

To perform this blind chromatic reconstruction we first computed the matrix of the coupling of red, green and blue wavelets to the sensor ensemble: $[\mathbf{H}_{\text{red}} \mathbf{H}_{\text{green}} \mathbf{H}_{\text{blue}}] \mathbf{W}_3$ where $\mathbf{W}_3$ denotes a triple-tall matrix of wavelet coefficients. Next, we found the combination of these wavelets (of whatever mixture of colours) that is simultaneously sparse and satisfies the observed ASP outputs $\mathbf{r}$ by solving Equation 5.1.

$$\mathbf{s}_{\text{CS}} = (\mathbf{W}_3^T \mathbf{W}_3)^{-1} \mathbf{W}_3^T \left(\underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{r} - [\mathbf{H}_{\text{red}} \mathbf{H}_{\text{green}} \mathbf{H}_{\text{blue}}] \mathbf{W}_3 \mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1 \right) \quad (5.1)$$

The solution to Equation 5.1 is the triple-tall red, green and blue vector of pixel intensities composed of a combination of a small number of monochromatic wavelets. The solutions to Equation 5.1 are plotted in Figure 11. This technique does a remarkable job of discovering the correct colour. Note that there is no constraint that all wavelets found by Equation 5.1 must be of the same colour (and numerical inspection reveals that the images in D columns are not strictly monochromatic); merely enforcing sparsity of the solution and consistency with the observations selects wavelets that happen to be of the correct colour. Note also that using compressed sensing greatly improves the performance of the PFCA for red light: compare R, D1 to R, A1 of Figure 11. Unexpectedly, Fourier information from the missing band $b = 13\text{--}21$ has been filled in. CS gives no indication this completion should be possible since it prescribes random sampling (with no systematic, large blind spots), yet it performs admirably despite a lack of theoretical guarantee.

6. Conclusions

We have demonstrated techniques for reconstructing images acquired by PFCAs that provide robustness beyond what was initially imagined possible. First, we characterized wavelength robustness of our PFCA. Although designed for green light, our first prototype PFCA performs well under blue light using standard linear algebra techniques (Figure 9, right) and even under red light using basis pursuit denoising and compressed sensing (Figure 11 R, D1 and R, D2). Reconstruction with red incident light is particularly impressive given the large gap in spatial frequencies sampled under red light (Figure 4) and the poor performance of ASPs whose p approaches λ . Second, we demonstrated robustness against losing all signal from a randomly-chosen set of ASPs. Using compressed sensing, we formulated a basis pursuit denoising problem (Equation 4.2) that permits impressive reconstructions even with a randomly-chosen set of 90% of the signals destroyed (Figure 10). Last, we applied the same compressive sensing techniques to determine the colour of the images presented. The reconstructed wavelengths are nearly flawless (Figure 11), and in the case




Figure 11. Blind Colour Imaging. We show reconstructions of test images 1 & 2 without foreknowledge of the colour of incident light. R, G and B rows: presented images were red, green and blue. A columns: image reconstructed using Equation 2.7 based on $\mathbf{H}$ calibrated using red light. B & C columns: as for A, but using green- and blue-calibrated $\mathbf{H}$ matrices. D columns: reconstruction based on Equation 5.1 without any assumption about colour, or that the image is monochromatic.

of red illumination this technique yields an image of much higher quality than permitted by linear algebra reconstructions (Figure 9, left). In principle there is no obstacle to taking polychromatic images using a PFCA since Equation 5.1 does not require images to be monochromatic, and it should be possible as well to recover hyperspectral information (i.e. more than 3 colour channels) in a compressed sensing framework, especially using a PFCA designed with a greater number of device types. Each device type has an idiosyncratic $b(\lambda)$ with sharp discontinuities caused by abruptly rounding effective depth to the nearest optimal Talbot depth (see Figure 3 and Equation 2.3), providing much better λ sensitivity than initially thought.

In summary, PFCAs enjoy several unusual forms of robustness stemming from their operating principles, including insensitivity to lost ASP sensors, flexibility in terms of incident light wave-

length, and the ability to determine source image colour without employing any chromatic filters. These unexpected forms of redundancy enhance the range of applications appropriate for this new class of sensing device.

Acknowledgments

We would like to thank Eve De Rosa who helped support this work, Albert Wang for his input and ideas, and Igor Carron, whose suggestions and advice regarding compressed sensing and wavelet dictionaries proved invaluable. We would also like to acknowledge our funding sources, including DARPA, who supported this research via a YFA to Alyosha Molnar, and the NIH, who helped fund this work under R21 grant EB 009841-01.

References

- [1] P. R. Gill, C. Lee, D. Lee, A. Wang, and A. Molnar, *A microscale camera using direct Fourier-domain scene capture*, Optics Letters, vol. 36, no. 15, pp. 2949–2951, 2011.
- [2] A. Wang, P. Gill, and A. Molnar, *Light field image sensors based on the Talbot effect*, Applied optics, vol. 48, no. 31, pp. 5897–5905, 2009.
- [3] A. Wang, P. R. Gill, and A. Molnar, *Fluorescent imaging and localization with angle sensitive pixel arrays in standard cmos*, presented at the IEEE Sensors Conference, 2010.
- [4] B. Batchelor, D. Hill, and H. Hodgson, *Automated visual inspection*. Elsevier Science Pub., New York, NY, 1985.
- [5] T. Lillesand, R. Kiefer, and J. Chipman, *Remote sensing and image interpretation*. John Wiley & Sons Ltd, 2004, no. Ed. 5.
- [6] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer *et al.*, *Autonomous driving in urban environments: Boss and the urban challenge*, Journal of Field Robotics, vol. 25, no. 8, pp. 425–466, 2008.
- [7] N. Bergström, J. Bohg, and D. Kragic, *Integration of visual cues for robotic grasping*, Computer Vision Systems, pp. 245–254, 2009.
- [8] H. Talbot, *Lxxvi. facts relating to optical science. no. iv*, Philosophical Magazine Series 3, vol. 9, no. 56, pp. 401–407, 1836.
- [9] S. Teng, Y. Tan, and C. Cheng, *Quasi-talbot effect of the high-density grating in near field*, JOSA A, vol. 25, no. 12, pp. 2945–2951, 2008.
- [10] C. Koch, J. Oehm, J. Emde, and W. Budde, *Light source position measurement technique applicable in soi technology*, Solid-State Circuits, IEEE Journal of, vol. 43, no. 7, pp. 1588–1593, 2008.
- [11] F. E. Theunissen, K. Sen and A. J. Doupe, *Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds*, Journal of Neuroscience, vol. 20, no. 6, pp. 2315–2331, 2000.
- [12] A. N. Tychonoff, *On the stability of inverse problems*, Doklady Akademii SSSR, vol. 39, no. 5, pp. 195–198, 1943.
- [13] D. Donoho, *Compressed sensing*, Information Theory, IEEE Transactions on, vol. 52, no. 4, pp. 1289–1306, 2006.

- [14] P. R. Gill, A. Wang, and A. Molnar, *The in-crowd algorithm for fast basis pursuit denoising*, Signal Processing, IEEE Transactions on, vol. 59, no. 10, pp. 4595–4605, 2011. Code available at <http://molnargroup.ece.cornell.edu/files/InCrowdBeta1.zip>.
- [15] B. Olshausen and D. Field, *Sparse coding with an overcomplete basis set: A strategy employed by V1?* Vision research, vol. 37, no. 23, pp. 3311–3325, 1997.
- [16] E. P. Simoncelli and W. T. Freeman, *The steerable pyramid: a flexible architecture for multi-scale derivative computation*, IEEE Second Int’l Conf. on Image Processing, Washington DC, October 1995. Code available at <http://www.cns.nyu.edu/~eero/steerpyr/>.