Megapixel Adaptive Optics: Towards Correcting Large-scale Distortions in Computational Cameras

CONGLI WANG, QIANG FU, XIONG DUN, and WOLFGANG HEIDRICH,
King Abdullah University of Science and Technology, Saudi Arabia

Fig. 1. We propose to use a high resolution adaptive optics (AO) technique to correct for large distortions in photography and vision applications. AO relies on the combination of a wavefront sensor and a phase modulator to measure and optically correct for incoming phase distortions. Based on a previously proposed high resolution wavefront sensor, our proposed AO prototype achieves unprecedented AO sensing and correcting simultaneous high spatial resolution for large phase distortions. An example for vision deblurring is shown on the right where a defocus distortion is being well compensated.

Adaptive optics has become a valuable tool for correcting minor optical aberrations in applications such as astronomy and microscopy. However, due to the limited resolution of both the wavefront sensing and the wavefront correction hardware, it has so far not been feasible to use adaptive optics for correcting large-scale waveform deformations that occur naturally in regular photography and other imaging applications.

In this work, we demonstrate an adaptive optics system for regular cameras. We achieve a significant improvement in focus for large wavefront distortions by improving upon a recently developed high resolution coded wavefront sensor, and combining it with a spatial phase modulator to create a megapixel adaptive optics system with unprecedented capability to sense and correct large distortions.

CCS Concepts: • **Computing methodologies** → **Computational photography**;

Additional Key Words and Phrases: Adaptive Optics, Wavefront Sensing, Computational Cameras

ACM Reference Format:

Congli Wang, Qiang Fu, Xiong Dun, and Wolfgang Heidrich. 2018. Megapixel Adaptive Optics: Towards Correcting Large-scale Distortions in Computational Cameras. *ACM Trans. Graph.* 37, 4, Article 115 (August 2018), 12 pages. <https://doi.org/10.1145/3197517.3201299>

Authors' addresses: C. Wang, congli.wang@kaust.edu.sa; Q. Fu, qiang.fu@kaust.edu.sa; X. Dun, xiong.dun@kaust.edu.sa; W. Heidrich, wolfgang.heidrich@kaust.edu.sa, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia, 23955-6900.

This work is licensed under a Creative Commons
Attribution-NonCommercial International 4.0 License.

© 2018 Copyright held by the owner/author(s).
0730-0301/2018/8-ART115
<https://doi.org/10.1145/3197517.3201299>

1 INTRODUCTION

Adaptive Optics (AO) systems are highly effective in correcting dynamic aberrations in applications such as astronomy and microscopy [Beckers 1993; Marsh et al. 2003]. In these settings, the aberrations are typically small (e.g. atmospheric distortions in the case of telescopes), albeit significant in the case of otherwise diffraction limited optics.

Adaptive optics systems are comprised of two major components: a wavefront sensor, which measures the shape of the wavefront of light across the aperture of the optical system, and a phase modulator, which corrects the distorted wavefront to make it planar. In traditional AO systems, the wavefront sensor is usually a Shack-Hartmann [Shack and Platt 1971] or pyramid sensor [Ragazzoni 1996], while the phase modulation is achieved by tilting mirrors [Bifano et al. 1999; Perreault et al. 2002] or deforming a reflective membrane with pistons [Zhu et al. 1999]. Both the sensing and the modulation components in these systems suffer from low *spatial resolution* (e.g. at most hundreds to thousands of measurement points and actuators) as well as small *range* that limits the amplitude of the distortion. As a result, most AO systems can only correct for distortions corresponding to a few low-order Zernike polynomials and low amplitudes. Multi-conjugate AO systems [Beckers 1988; Diolaiti et al. 2001] are one possible solution to offer high resolution and large compensation amplitudes, but at the cost of additional wavefront sensors and correctors that have to be carefully aligned and calibrated to one another. This significantly increases the cost and overall system complexity compared to the single sensor / single corrector systems that we consider in our work.

Table 1. Performance characteristics of different wavefront sensing technologies.

	Resolution	Dynamic Range	Frame Rate	Pixel Usage	Calibration	Requires Coherent Light?	Light Efficiency
Shack-Hartmann sensor [Shack and Platt 1971]	Low/High	Large/Small	High	Low	Easy	No	High
Curvature sensor [Roddier 1988]	High	Large	High	High	Hard	Yes	High
Pyramid wavefront sensor [Ragazzoni 1996]	Medium	Medium	High	Medium	Medium	No	Low
Interferometry sensor	High	Small	High	High	Medium	Yes	High
Multi-lateral shearing interferometer sensor [Primot and Sogno 1995]	Medium	Medium	High	High	Easy	No	Medium
Coded Wavefront Sensor [Wang et al. 2017]	High	Large	Medium [†]	High	Easy	No	High [‡]

[†]Computationally bounded and inversely proportional to wavefront resolution.

[‡]Depends on the modulation mask. For phase mask modulation the light efficiency is almost 100 %.

In photography and most computer vision applications, the situation is different: normal camera optics are usually not diffraction limited, and less sensitive to small distortions. However, large distortions may occur regularly. Consider, for example, the case of a computer vision system of a self-driving car operating under rainy conditions. The camera of this vision system may have to re-focus through a dynamically changing water surface, as rain flows across the cover surface of the optical system. Current AO systems cannot deal with such large wavefront distortions.

To mitigate this situation, we propose a new AO setup, which is built around an improved version of the recently introduced coded wavefront sensor [Wang et al. 2017], and an LCoS spatial phase modulator (phase SLM). Both components have megapixel resolution, which allows us to detect and correct for wavefront distortions with complex shapes. Moreover, both the sensor and the modulator can also deal with distortions of significantly larger amplitude, so that much more severe defocus effects can be corrected. In particular, we can achieve Strehl ratios close to 1 even for large-scale deformation near the optical axis, although the off-axis performance is reduced. Specifically, our technical contributions are:

- An adaptive optics system design that is capable of sensing and correcting large wavefront distortions with megapixel wavefront sensing and phase modulation.
- An improved version of the coded wavefront sensor [Wang et al. 2017] with better light efficiency and improved algorithms.

Although this initial prototype system is limited in frame rate and light efficiency (see Section 6 for a full discussion), we believe this work is a major step towards utilizing adaptive optics in regular cameras for both photography and machine vision applications.

2 RELATED WORK

The purpose of adaptive optics systems is to re-focus optical systems through (often time-varying) distortions. In the following we provide a brief review of existing means to measure and correct defocus.

Coded Apertures and Image Deblurring. Coded apertures and image deblurring have been widely researched in the computational photography community. Researchers have delved into coded apertures for defocus blurring compensations in extended-depth-of-field applications [Cossairt and Nayar 2010; Dowski and Cathey 1995; Levin et al. 2007, 2009; Zhou and Nayar 2009], or in motion blur removal [Levin et al. 2008; Raskar et al. 2006]. Approaches on pure

software deconvolution have also been proposed, for example non-blind deconvolution algorithms [Cho et al. 2011; Hu et al. 2014; Ji and Wang 2012] given the blurring kernel, or the more challenging blind deconvolution case where algorithms design specifically for motion blurring [Cho and Lee 2009; Fergus et al. 2006; Xu and Jia 2010], or with utilization of natural priors [Dong et al. 2017; Krishnan et al. 2011; Pan et al. 2016; Xu et al. 2013]. These techniques either rely on specific assumptions on the scene and blurring kernels, or consume heavy computations, thus not applicable to general distortions for real time applications. Moreover, purely software-based solutions cannot deal with large aberrations as we demonstrate in Section 5.

Light Field Cameras. Light field cameras are closely related to wavefront sensors. In particular, the Shack-Hartmann sensor [Shack and Platt 1971], discussed in detail below, is essentially a lenslet-based light field camera [Ng et al. 2005]. One of the core benefits of light field cameras is the ability to refocus images post-capture [Ng 2005]. This feature has so far been primarily applied for artistic purposes, although refocusing through dynamic distortions should at least in principle be possible as well.

Recent work on light field cameras replaces the microlens array with specifically designed modulators, for example using amplitude masks [Marwah et al. 2013; Veeraraghavan et al. 2007; Xu et al. 2012], a single diffuser [Antipa et al. 2016], or from a partially wetted window with water drops [Iseringhausen et al. 2017]. These designs are closely related to the coded wavefront sensor [Wang et al. 2017], discussed below.

Classical Wavefront Sensors. Classical wavefront sensors are key components of adaptive optics systems, where they are responsible for measuring the incoming wavefront distortion, and for providing enough information for the control system to update the wavefront correction component. Table 1 summarizes some key performance characteristics of different wavefront sensor designs, which we will discuss in the following.

The classical Shack-Hartmann wavefront sensor [Shack and Platt 1971] tracks the 2D motion of focus spots generated by a microlens array to recover the unknown wavefront slopes. It offers high frame rates, but the spatial wavefront resolution is limited to the number of lenslets. High spatial resolution (e.g. 2×2 pixels per lenslet in Altair [Saddlemeyer et al. 1998]) can be achieved by increasing the number of lenslets, but proportionally at the cost of decreased ability to measure large wavefront slopes. Similar tradeoff also exists for other slope-tracing wavefront sensors, for example the pyramid wavefront sensors [Ragazzoni 1996] and the quadriwave lateral

Fig. 2. (a) Hardware overview. A shortpass dichroic mirror reflects the infrared (IR) light from the guide star, while transmitting the visible light from the object, merging the two into a single beam that undergoes a distortion before entering the AO system. Inside the AO system, a linear polarizer ensures the SLM operates in phase modulation mode. After being modulated by a phase-only SLM, the joint light cone is split back into IR and visible components, with the visible light being directed to the camera, while the IR light is directed to the wavefront sensor. The dichroic mirrors have the same cutoff wavelength so that the visible and IR light paths do not interfere with each other. Relay lenses ensure the SLM and the Coded Wavefront Sensor are in conjugate. For easy visualization only the IR light is drawn for the broadband white lamp. (b) At k^{th} AO iteration, the GPU workstation takes in the captured image i_k from the wavefront sensor, computes the reconstructed phase ϕ_k , and renders the SLM with a new correction phase u_k .

shearing interferometric wavefront sensors [Primot and Guérineau 2000; Primot and Sogno 1995].

On the other hand, curvature wavefront sensors [Rodier 1988], based on the Transport of Intensity Equation [Teague 1983], offer full sensor spatial wavefront resolution at high frame rates. Typical curvature sensors require coherent illumination and precise mechanical scanning to obtain multiple images for later computational phase reconstruction [Waller et al. 2010b]. Recently, there has been work on curvature wavefront sensing using a single color sensor [Waller et al. 2010a], exploiting the lens chromatic aberration, and thus avoiding the mechanical scanning. However, because of its comparatively complicated setup and high demands on sensor sensitivity, no large phase distortions have been experimentally demonstrated. Similar limitations also hold for interferometric wavefront sensors, which are accurate for small phase distortions, but require phase unwrapping for large ones. Interferometric wavefront sensors also require coherent illumination and a complicated setup with highly sensitive alignment.

Coded Wavefront Sensors. Coded wavefront sensors replace the microlens array of the Shack-Hartmann designs with a binary amplitude mask in close proximity to a bare image sensor [Wang et al. 2017], and are thus related to mask-based light field cameras, as discussed above. In the Coded Wavefront Sensor, the slope of the wavefront is tracked using numerical methods related to optical flow [Wang et al. 2017], and as such they combine a full spatial resolution with the ability to measure large distortions. Other works with similar ideas have been presented, for example using a diffuser for visible light [Berto et al. 2017], or high frequency phase objects for X-ray [Bérubon et al. 2012; Morgan et al. 2012]. In this work, we improve the coded wavefront sensor both by introducing a phase mask instead of an amplitude mask (thereby improving light sensitivity by a factor of 2), and by improving the software algorithms

for tracking the slope of the wavefront. Finally the Coded Wavefront Sensor also bears similarity to Background Oriented Schlieren imaging of transparent phenomena [Atcheson et al. 2008; Richard and Raffel 2001], except that the patterned background has been moved into the camera.

Adaptive Optics. Adaptive optics techniques were originally developed for military and astronomical telescopes to sharpen stellar observations by measuring and compensating atmospheric turbulence. Since then, AO has found its other applications in ophthalmology [Fernández et al. 2001; Liang et al. 1997; Roorda et al. 2002], microscopy [Booth 2014; Ji et al. 2010], and optical coherent tomography [Hermann et al. 2004; Zhang et al. 2006]. Classical AO systems employ a Shack-Hartmann wavefront sensor to observe the wavefront of a single point-shaped source, also known as the *guide star*. Any measured atmospheric distortion equally affects the full field of view of the main camera that observes the object of interest. However, this distortion can be corrected using a deformable mirror. In *open loop* systems, only the view of the main camera is corrected, while in *closed loop* systems the wavefront sensor also images the guide star through the deformable mirror, so that in each time step only a smaller differential deformation needs to be measured.

Due to the low resolution nature of deformable mirrors, spatial light modulators have been utilized to improve correction resolution, e.g. [Hu et al. 2004; Li et al. 2009; Love 1997; Vargas-Martin et al. 1998]. However, existing AO systems have not been able to demonstrate the combination of high spatial resolution and the ability to measure and correct for large distortions, which is required in many photography and computer vision applications.

The need for a guide star is a potential impediment for the use of AO approaches in regular imaging. Possible solutions include a laser-generated dot, or special application-specific setups. For example, in the above-mentioned example of a car vision system,

the camera could be behind the windshield, and the guide star could be integrated in the hood of the car to compensate for distortions caused by raindrops on the wind shield.

To our knowledge, this paper is the first AO system that uses a phase-only SLM and a correspondingly high resolution wavefront sensor, to optically sharpen heavily blurred vision images.

3 MEGAPIXEL ADAPTIVE OPTICS

Our AO system makes innovations on both the hardware and the software aspects. Figure 2(a) shows a diagram of the hardware setup. A Coded Wavefront Sensor observes an infrared guide star through the same optical system and distortions through which a regular machine vision camera observes a visible light scene. Dichroic mirrors are used to separate the IR and visible light paths while maximizing light efficiency. A phase-only SLM is used to correct for distortions. We utilize a closed loop design, i.e. the SLM is in the optical path of both the camera and the wavefront sensor.

In the following, we first discuss the software aspects of our system, before returning to the hardware prototype in Section 4.

3.1 Closed Loop Adaptive Optics

The software part of the system is the closed control loop for the adaptive optics system, which consists of a wavefront sensing component and an update pattern for the phase SLM. Figure 2(b) depicts a control block diagram for the closed loop system. At time step k , we denote the distorted wavefront from the guide star as $\hat{\phi}_k$. This wavefront is partially corrected by the correction phase $\mathbf{u}_{k-1}$ computed for the phase SLM in the previous time step, resulting in a raw sensor image $\mathbf{i}_k$ being observed by the wavefront sensor. From this image, the Coded Wavefront Sensor measures an estimate of the residual distortion:

$$\phi_k = \hat{\phi}_k - \mathbf{D}\mathbf{u}_{k-1}, \quad (1)$$

where $\mathbf{D}$ is a geometric alignment transformation that needs to be calibrated. If SLM and wavefront sensor are perfectly aligned, then $\mathbf{D}$ is the identity. In the first iteration, the SLM is initialized with a flat phase, i.e. $\mathbf{u}_0 = \mathbf{0}$.

One goal in closed loop AO is to stabilize the loop, i.e. to design a control strategy to update $\mathbf{u}_k$ such that $\|\phi_k\|_2 \rightarrow 0$ with $k \rightarrow \infty$. For sufficiently fast-response AO systems, or slowly time-varying distortion wavefronts, stabilization can be achieved using a digital integrator as a controller:

$$\mathbf{u}_k = \mathbf{u}_{k-1} + \alpha \mathbf{D}^{-1} \phi_k, \quad (2)$$

where α is a loop gain parameter. Ideally for $\alpha = 1$ the loop converges in one iteration and $\mathbf{u}_k = \mathbf{D}^{-1} \hat{\phi}_k$.

3.2 Wavefront Solver

We now describe the Coded Wavefront Sensor and its associated numerical solver for computing the observed wavefront ϕ_k from a captured image $\mathbf{i}_k$, at the k^{th} AO iteration.

Sensing Model. Figure 3 depicts the working principle, which is briefly reviewed in the following. The Coded Wavefront Sensor consists of a bare image sensor with a diffraction mask placed in closed proximity. The original Coded Wavefront Sensor [Wang

Fig. 3. Principle of the Coded Wavefront Sensor.

et al. 2017], used a binary amplitude mask, which absorbs 50% of the incident light, while we utilize a binary phase mask, where the absorption is negligible. Irrespective of the type of mask used, an incident plane wave creates a characteristic diffraction pattern on the image sensor, which is measured and saved as image $\mathbf{i}_0(\mathbf{x})$, $\mathbf{x} = (x, y)^T$ in a calibration phase.

Wang et al. [2017] were able to show that, if a distorted wavefront ϕ_k is incident on the same configuration, it results in an image $\mathbf{i}_k$ that is a locally warped version of $\mathbf{i}_0$ according to the following relationship:

$$\mathbf{i}_k(\mathbf{x}) \approx \mathbf{i}_0\left(\mathbf{x} - \frac{\lambda z}{2\pi} \nabla \phi_k\right), \quad (3)$$

where z is the spacing between the mask and the image sensor, and λ is the wavelength of light. That is, the local shift of the pattern is proportional to the gradient of the wavefront. Intuitively, since the pattern has a high spatial frequency content, the 2D distortion can be tracked with optical flow-style methods, and the shape of the wavefront ϕ_k can be recovered from the pair of images $\mathbf{i}_0, \mathbf{i}_k$.

We note that, although Eq. (3) contains the wavelength λ , this principle actually does not assume coherent light and works well for broadband illumination.

Optimization. To simplify the notation, we will in the following absorb the constant $\lambda z/2\pi$ into the ϕ_k variable. The process of recovering ϕ_k can be expressed as a least-squares optimization problem with an additional smoothness regularizer:

$$\underset{\phi_k}{\text{minimize}} \quad \|\mathbf{i}_k(\mathbf{x}) - \mathbf{i}_0(\mathbf{x} - \nabla \phi_k)\|_2^2 + \beta \|\nabla \phi_k\|_2^2, \quad (4)$$

where $\beta > 0$ is a weighting factor.

To solve this problem, we take inspiration from the optical flow literature [Brox et al. 2004; Horn and Schunck 1981], and linearize $\mathbf{i}_0(\mathbf{x} - \nabla \phi_k)$ around $\mathbf{x}$ in Eq. (3). This yields the following approximation:

$$\nabla \mathbf{i}_0(\mathbf{x}) \cdot \nabla \phi_k + \mathbf{i}_k(\mathbf{x}) - \mathbf{i}_0(\mathbf{x}) \approx 0, \quad (5)$$

where $\cdot$ denotes the inner product. We apply Eq. (5) to Eq. (4), and denote image gradient fields as $\mathbf{g}_x$ and $\mathbf{g}_y$ along x and y directions, and $\mathbf{g}_t$ as the difference between the two images, and finally define $\mathbf{G} = [\text{diag}(\mathbf{g}_x) \text{ diag}(\mathbf{g}_y)]$ where $\text{diag}(\cdot)$ denotes a diagonal matrix formed by the corresponding vector. With these definitions, Eq. (4) is reduced to a linear equation:

$$\underset{\phi_k}{\text{minimize}} \quad \|\mathbf{G}\mathbf{M}\nabla \phi_k + \mathbf{g}_t\|_2^2 + \beta \|\nabla \phi_k\|_2^2. \quad (6)$$

We solve this optimization problem with Neumann boundary conditions. To avoid boundary artifacts, we have introduced a spatial

Fig. 4. Our ADMM solver efficiently converges to the minimum with a good visual plausibility, whereas conjugate gradient method takes much more iterations even when the energy decreasing is barely perceptible.

selection matrix $\mathbf{M}$ to include the boundary values of ϕ_k as additional unknowns, to be determined by the optimization as that in [Almeida and Figueiredo 2013].

This optimization problem bears a strong similarity to the classical Horn-Schunck optical flow problem [Horn and Schunck 1981]. However, in our case the flow vectors $\nabla\phi_k$ have a physical meaning as the gradients of the phase function ϕ_k , so we seek to directly solve for ϕ_k , instead of solving for the optical flow.

We solve this joint optimization problem by introducing a slack variable $\mathbf{w}$ that physically represents the wavefront gradient $\nabla\phi_k$, and apply the Alternating Direction Method of Multipliers (ADMM) [Boyd et al. 2011], yielding Algorithm 1. Here η is the dual variable, μ is a proximal parameter, $\mathcal{F}_{\text{DCT}}$ and $\mathcal{F}_{\text{DCT}}^{-1}$ respectively denote forward and inverse Discrete Cosine Transforms (DCT). To suppress noise, median filtering is applied to the gradient estimation before a final integration to get the output wavefront solution ϕ_{estimate} , as suggested in [Sun et al. 2010]. More details regarding the solver can be found in the Supplementary Material.

The superiority of our ADMM solver over the usual conjugate gradient method is illustrated in Figure 4, where a 10λ Gaussian wavefront is simulated. The conjugate gradient method is initialized with a rough wavefront estimation ϕ^0 , which is the Poisson integration from a one-step flow estimation $\nabla\phi_k$ in Eq. (6) (see the Supplementary Material). Our ADMM solver converges efficiently with low error, whereas the conjugate gradient method requires a large number of iterations. In real experiments, the ADMM solver only runs for 10 iterations, which provides a good compromise between accuracy and speed.

Nonlinear Warping Scheme. This linear solver can be improved by including non-linear warping [Brox et al. 2004]. We first use a

Algorithm 1 ADMM linear solver for Eq. (6).

```

1: function RECONSTRUCT WAVEFRONT( $\mathbf{g}_x, \mathbf{g}_y, \mathbf{g}_t$ )
2:   Initialize  $\phi^0, \mathbf{w}^0$  and  $\eta^0$  from previous frame, set  $\mu > 0$ ;
3:   while not converge do
4:      $\phi^{k+1} \leftarrow \frac{\mu}{\beta+\mu} \mathcal{F}_{\text{DCT}}^{-1} \left( \mathcal{F}_{\text{DCT}}(\nabla^T(\mathbf{w}^k - \eta^k)) / \mathcal{F}_{\text{DCT}}(\nabla^2) \right)$ ;
5:      $\mathbf{w}^{k+1} \leftarrow (\mu \mathbf{I} + \mathbf{M}^T \mathbf{G}^T \mathbf{G} \mathbf{M})^{-1} \left( \mu(\nabla\phi^{k+1} + \eta^k) - \mathbf{M}^T \mathbf{G}^T \mathbf{g}_t \right)$ ;
6:      $\eta^{k+1} \leftarrow \eta^k + \nabla\phi^{k+1} - \mathbf{w}^{k+1}$ ;
7:   end while
8:    $\hat{\mathbf{w}} = \text{MEDIAN FILTER}(\mathbf{w}^k - \eta^k)$ ;
9:   return  $\phi_{\text{estimate}} = \frac{\mu}{\beta+\mu} \mathcal{F}_{\text{DCT}}^{-1} \left( \mathcal{F}_{\text{DCT}}(\nabla^T \hat{\mathbf{w}}) / \mathcal{F}_{\text{DCT}}(\nabla^2) \right)$ ;
10: end function

```

Algorithm 2 A warping scheme for Algorithm 1.

```

1: Initialize  $\phi^0$  with zeros;
2: for pyramid level  $s$  do                                      $\triangleright$  Pyramid-level warping
3:    $i_0^s(\mathbf{x}) \leftarrow i_0(\mathbf{x}) \downarrow$ ;                                $\triangleright \downarrow$ : Down-sampling
4:    $i_k^s(\mathbf{x}) \leftarrow i_k(\mathbf{x}) \downarrow$ ;
5:   while  $\|\Delta\phi^j\|_2^2 > \epsilon$  do                                    $\triangleright$  In-level warping,  $\phi^j \leftarrow \phi^s$ 
6:     Compute  $\mathbf{g}_x^s, \mathbf{g}_y^s, \mathbf{g}_t^s$  from  $i_0^s(\mathbf{x} - \nabla\phi^j)$  and  $i_k^s(\mathbf{x})$ ;
7:      $\Delta\phi^j \leftarrow \text{RECONSTRUCT WAVEFRONT}(\mathbf{g}_x^s, \mathbf{g}_y^s, \mathbf{g}_t^s)$ ;
8:      $\phi^{j+1} \leftarrow \phi^j + \Delta\phi^j$ ;
9:   end while                                                    $\triangleright \phi^s \leftarrow \phi^j$ 
10:  if  $s$  is not the final level then
11:     $\phi^{s+1} \leftarrow \phi^s \uparrow$ ;                                      $\triangleright \uparrow$ : Up-sampling
12:  end if
13: end for

```

linear approximation according to Eq. (5) to obtain a preliminary wavefront estimation, then warp the reference image $i_0(\mathbf{x})$ towards the gradients of the preliminary estimation, and then re-linearize the system to obtain an improved estimate in the next round.

Algorithm 2 shows this non-linear variant of Algorithm 1. This scheme contains a coarse-to-fine strategy and an in-level nonlinear warping at each pyramid level. The pyramid level s increases from 0 to a given pyramid level number, namely from the smallest down-sampled image size to the original size. At each warping step j , the pyramid level image $i_0^s(\mathbf{x})$ is warped to a new image $i_0^s(\mathbf{x} - \nabla\phi^j)$ according to current wavefront estimation ϕ^j , then the linearized problem Eq. (6) is solved with the new gradient estimations $\mathbf{g}_x^s, \mathbf{g}_y^s$ and $\mathbf{g}_t^s$. The unknown wavefront ϕ_k is iteratively updated in this way until convergence.

The improvement of the warping strategy is illustrated in Figure 5, where a turbulence phase screen is simulated using the sub-harmonic method [Lane et al. 1992]. Algorithms 2 and 1 (with and without warping, respectively) are compared for different pyramid levels. The warping strategy greatly improves large wavefront sensing to a fine accuracy.

4 PROTOTYPE

In the following, we describe our prototype in terms of both hardware components and software implementation details.

Fig. 5. For large phase distortions, the proposed warping scheme wavefront solver enables an improved estimation accuracy. Simulated wavelength $\lambda = 550$ nm.

4.1 Hardware

Mask Fabrication and Assembly. In our experiments, the Coded Wavefront Sensor is built around a monochromatic $2/3''$ CCD camera (PointGrey GS3-U3-15S5M-C), with a sensor resolution of 1384×1032 and a pixel pitch of $6.45 \mu\text{m}$. This camera was operated without a lens, and the IR filter was removed. Instead, the mask was mounted onto the sensor at a distance of ca. 1.5 mm from the light sensitive surface.

The binary phase mask itself is a random, binary height field fabricated on a 0.5 mm thick $4''$ fused silica wafer via photolithography followed by Reactive Ion Etching (RIE). Each mask pixel is $12.9 \mu\text{m} \times 12.9 \mu\text{m}$, and the etching depth is chosen such that the corresponding phase delay of the pixels is either 0 or π . The dimension of the mask is adjusted to $20 \text{ mm} \times 17 \text{ mm}$ to match the cover glass on the sensor. Refer to the Supplementary Material for more details.

Optical Setup. Figure 6 shows the experimental setup. Our prototype employs largely off-the-shelf optical and mechanical components from Thorlabs, with a few customized 3D-printed components. We use a HOLOEYE PLUTO phase-only SLM (PLUTO-2-VIS-014), with a pixel pitch of $8.0 \mu\text{m}$ and a maximum 3π phase retardation for wavelength $\lambda = 532.8$ nm, and the refresh rate is 60 Hz (same as V-Sync). The frame rate of the wavefront sensor is set to be its maximum of 45 Hz, with streaming mode enabled. To match sizes between SLM and the wavefront sensor, a relay lens system scales down the SLM plane by a factor of $4/3$ onto the wavefront sensor plane. Due to the limited wavefront sensor size, not all SLM pixels are covered (see also Figure 9), so a 3D-printed square aperture is designed to prevent extra light from impinging onto the unobservable areas of the SLM. Black flocked paper is used for blocking stray light. All lenses are achromatic doublets.

To prevent the guide star from being seen by the vision camera, the two dichroic mirrors have been selected to operate at the same cutoff wavelength of 650 nm. However in practice, it is still needed to have a longpass filter in front of the broadband white light to further suppress the visible spectrum from the lamp. Figure 7 illustrates the effect.

The USB output of the PointGrey sensor and the HDMI input for the phase-only SLM, are both connected to a host workstation running Ubuntu 16.04.3 LTS, with a NVIDIA GTX Titan X (Pascal) graphics card, 2.70 GHz Intel Xeon E5-2680 processors ($\times 32$) and 64 GB RAM.

Fig. 6. Our high resolution AO prototype setup. (a) shows the setup overview. Closeup (b) and (c) respectively shows how we obtain collimated, polarized infrared light, and how iris (d) is attached close to the SLM. Some distortion candidates are in (e). Note the phase-only SLM and wavefront sensor are mirror dual to each other.

Fig. 7. Adding a longpass filter in front of the guide star to suppress it from being visible to the vision camera.

4.2 Software

All control code and processing algorithms are implemented using C++ and CUDA. To avoid unnecessary CPU-GPU data transfer, the output phase image is rendered directly to the SLM through CUDA-OpenGL interoperability. All textures are operated with replicate boundary conditions. To synchronize the camera and SLM so the camera shutter will not start to integrate when the SLM is updating, we use the workstation output V-sync signal (60 Hz) as a hardware trigger. For maximum speed, all unknown sizes are set to be power of two ($1024, 512, \dots$) to make use of the radix-2 FFT algorithm.

Interpolation. Interpolations are needed when doing warps. Empirically for best performance, we use cubic spline interpolation for in-level warping, and bilinear interpolation (with anti-aliasing pre-filtering) for pyramid-level warping. To avoid expensive linear system solvers for cubic coefficient calculations each time, all image warping operations are performed on the reference image $i_0(\mathbf{x})$. For

Fig. 8. Bilateral filtering of detected wavefront ϕ_k and the accumulated SLM phase u_k . Smoother phase reduces undesired higher-order SLM diffraction and hence stabilizing the closed loop. Bilateral filter spatial and intensity window sizes are both 15×15 . Phases are visualized by 2π wrapping.

optimum performance, before running the GPU solver Algorithm 2 online, the reference image $i_0(\mathbf{x})$ is first decomposed in-place with its cubic coefficients, thus the subsequent online warping operations become trivially fast convolutions.

Denoising. To reduce wavefront sensing noise and the unavoidable phase wrapping artifacts produced by the phase-only SLM, bilateral filtering is applied to both the reconstructed wavefront ϕ_k before updating u_k , and the accumulated phase image u_k before final rendering onto the SLM. Figure 8 shows an experimentally obtained performance comparison between an original and a denoised version, where the latter one produces smoother phase profile with more continuous wrapping rings, and stably converges to a much smaller residual in terms of wavefront Root-Mean-Square (RMS).

Calibration. AO systems are designed to detect and correct for very small wavefront distortions, and as such they are extremely sensitive to mis-alignment and mis-calibration. For a successful AO system, one has to perform careful calibration to ensure the wavefront sensor and the corrector are in good alignment. To measure the wanted inverse alignment transform D^{-1} in Eq. (2), a microlens array phase image is shown on the SLM, and is then measured by the Coded Wavefront Sensor. By comparing the original and detected center points, assuming a 2D homography, D^{-1} can be fitted by an over-determined system. Figure 9 depicts this calibration pipeline, as an improved variant to that in [Jang et al. 2014]. To reduce noise and system uncertainty, one thousand repeatedly captured measurement images are averaged as one single measurement image for the wavefront solver. Before starting the AO correction, under collimated illumination, a reference image $i_0(\mathbf{x})$ is captured by showing a “black” screen (zero phase) on the SLM.

Algorithm Parameters. For pyramid up-sampling and down-sampling, a factor of 2 along each coordinate direction is used. In practice for

Fig. 9. Calibration of inverse misalignment transform D^{-1} . To obtain D^{-1} , a 2D perspective matrix is calculated by fitting from the detected center points (Coded Wavefront Sensor) to the original ones (phase-only SLM). Detected center points are generated by circular Hough transforms.

Table 2. Software performance of our algorithm (1024×1024 unknowns).

Forward/Inverse DCTs	21.24 ms	Bilateral Filters	2.75 ms
Convolutions	2.61 ms	Elementwise	6.19 ms
CPU-GPU Data Transfer	2.88 ms	Warping	0.12 ms
Others	5.01 ms	Overall	40.80 ms

the linear solver $\beta = 5$, $\mu = 100$ or $\mu = 10$ depends on the specific pyramid level, with an ADMM iteration of 10. For the nonlinear warping scheme 2 pyramid levels and 2 in-level warping operations are used. The AO loop gain is set as $\alpha = 1$. The bilateral filter window size is 15×15 .

Performance. To avoid tearing artifacts on the phase SLM, the phase-only SLM is set to be V-Sync double buffered. Our controller writes and renders directly to the back buffer, which swaps itself with the front one at V-Sync rate. The wavefront sensor is externally triggered by V-Sync as well, and the wavefront solver runs freely. The SLM, however, suffers from a relatively slow response time and in practice one can see the phase lagging in the measurements made by the wavefront sensor. In other words, phase measurement in our prototype is much faster than the phase correction. To avoid wrong iterations, our total frame rate has to drop down and meet with SLM’s update speed, and hence the overall AO performance is limited. Please refer to the Supplementary Material for more graphical details on our prototype timing. Despite hardware latency, Table 2 shows the overall benchmark performance for software solely.

5 RESULTS

In the following, we present experiments with both simulated as well as real-world distortions. In addition to visual presentation of results, we also use the *Strehl ratio* as a quantitative metric. The Strehl ratio is commonly used to evaluate AO systems, and is defined as the ratio of the peak brightness of the *actual* PSF of the system to

the peak brightness of the *ideal* PSF. The value is therefore between 0 and 1, with 1 being the best. In diffraction limited systems, the ideal PSF is usually the diffraction limited Airy disk. However, since our main camera is not diffraction limited, we instead use the measured PSF of the camera without optical distortions as the reference “ideal” PSF.

5.1 Simulation

To have a performance comparison among our proposed high resolution AO system and others, we present here a numerical simulation where large phase distortions are introduced, hence only the Shack-Hartmann and curvature sensors are simulated and compared. For a fair comparison, the free parameters of each wavefront sensor are tuned to best match the incoming phase distortion. For example, we maximize the lenslet number of the Shack-Hartmann wavefront sensor while keeping the maximum target phase distortion slope resolvable. Note that no such parameter tuning is necessary for our own system. All simulated AO systems are equipped with the same high resolution phase-only SLM, the only difference is the wavefront sensors in use. For full implementation parameters details please consult the Supplementary Material.

Figure 10 shows a comparison for two different wavefronts, a smooth wavefront (cubic phase) and a turbulent wavefront. The Shack-Hartmann AO system corrects well for smooth wavefronts, but is slow in convergence because it mainly corrects for low frequency terms due to its low spatial wavefront resolution nature, and requires a large number of iterations for final convergence. The Curvature AO system, is not capable of correcting for offset wavefronts, is noise-sensitive and hence not able to maintain stable performance. Our proposed high-resolution AO system outperforms the other two existing AO opponents in terms of both correction (by Strehl ratio) and speed (by number of AO iterations). We also show experiments with lower resolution wavefront correction. As can be expected, the low-resolution actuators are not capable of correcting for highly detailed aberrations, resulting in a large residual error. This demonstrated that both the sensor and the actuator should have a high spatial resolution in order to apply AO to regular photography and machine vision applications.

5.2 Real Experiments

Experimental results are presented, with a nominal wavelength 532.8 nm used to quantify wavefront errors.

Point Spread Function Evaluation. To test the performance of our prototype AO system and evaluate the Strehl ratio, a point light source is placed at the object plane. Artificial wrapped phase images are generated on the SLM as initial phase distortions for the AO system to correct with. Figure 11 shows the evolution of the AO system, quantified by wavefront RMS and image quality Strehl ratio. Our AO system performs fast convergence to the null state. More animated examples are demonstrated in our Video.

Static Deblurring. In order to demonstrate the deblurring capability of our AO system, we introduce phase distortions in the optical path, as shown in Figure 6. The arbitrary irregular phase distortions were generated by warping transparent polycarbonate plates

when being heated up. After cooling down to room temperature, local phase distortions are accumulated owing to the thermally introduced curvatures (see Figure 6 (e)). This allows us to repeat measurements with and without AO correction. Results are shown in Figure 12 where the target objects show significant improvement in focus when the AO system is switched on.

Comparison with Software Deblurring. We also compare our AO deblurring results with software-only deconvolution algorithms (blind and non-blind) for real data. For blind deconvolution (Figure 13, top), the methods are directly applied to the blurred, uncorrected image seen by the camera. For non-blind deconvolution (Figure 13, bottom), we assume that a wavefront sensor is available to measure the distortion, but that no actuator is available to correct it. From the measured distortion, and calibrated propagation distances and aperture sizes, it is possible to analytically derive the PSF for the non-blind deconvolution. Due to the prototype nature of our setup, instead of calibrating propagation distance and aperture size, we sampled this parameter space to produce a family of PSFs, and show results for the PSF that produces the best deconvolved result. The results show that being able to measure the PSF with a wavefront sensor does result in sharper reconstructions than fully blind deconvolution. However, the blur is so strong that even the non-blind deconvolution problem is severely ill-posed. As a result, AO as a hardware solution outperforms both blind and non-blind deconvolution by a large margin.

Dynamic Deblurring. We also show dynamic deblurring by translating the distorted polycarbonate plate horizontally, such that at each instant of time the AO system sees different phase distortions. Our system can detect phase distortions at a maximum of 30 Hz (depends on specific algorithm parameters), but is only able to do correction at around 12 Hz. This is limited by the liquid crystal response speed of our current phase SLM (see Section 6 for more discussions), and could be potentially improved to real time by employing SLMs with higher modulation rates.

6 DISCUSSION AND CONCLUSIONS

In this paper, we have demonstrated an adaptive optics system that can sense and correct optical distortions with megapixel spatial resolution and a large magnitude. This system is based on the combination of an improved custom Coded Wavefront Sensor and a readily available phase SLM. With this combination we have for the first time demonstrated the ability to correct for large-scale distortions relevant to machine vision and photography applications using adaptive optics. However, our prototype system does suffer from several limitations, which would be interesting to address in the future.

One limitation of our prototype is that the field-of-view (FoV) over which our system is able to compensate is limited. This is known as anisoplanatism in astronomy AO research. The wavefront we sense is only correct for points near the optical axis, and as such the AO performance will degrade with distance from the optical axis. Figure 15 demonstrates this effect, as two different phase distortions are imposed in front of the system. If the distortion has a simple structure, for example by introducing a spherical lens or

Fig. 10. Synthetic comparison of large wavefront distortions corrected by AO systems with different wavefront sensors and correctors.

defocusing the main lens, the measured on-axis distortion is a good approximation for a relatively wide field of view. However, for a more complex distortion such as warped glass, the wavefront shape and thus the PSF is different for different points on the image plane. While the AO system generally still improves the PSF, it is then not

capable of achieving perfect focus. One remedy is multi-conjugate AO [Marchetti et al. 2007], where multiple optically conjugated wavefront sensor-corrector pairs are separately grouped and then

Fig. 11. PSF experiment with test phase generated on the SLM. Our high resolution AO system is capable to correct distortion and converges in a few iterations. For more visual demonstrations please refer to the Video.

Fig. 12. Static deblurring for large distortions.

recombined to correct for volumetric phase distortions. This approach, however, has higher computational cost, and we leave a study of its feasibility in our setting to future work.

Other limitations of the prototype are related to the specific phase SLM that we use. The first limitation is the size – at only 8.6 mm × 8.6 mm, which limits the total aperture size of the system, and leads to a long and bulky prototype. We chose the Holoeye phase SLM in

Fig. 13. Comparison of our AO approach with software-only methods.

Fig. 14. Dynamic deblurring of translated polycarbonate plates. Our AO system is able to compensate moving dynamic phase distortions. See the Video for full frame performance.

part for its large range of phase modulation – up to 3π . However, this large range comes at the expense of reduced phase stability and switching speed; the frame rate of the SLM is the limiting factor for the frame rate of our prototype. However, other phase SLMs with quite different performance tradeoffs are commercially available, and we intend to explore these options in the future.

For the future, we believe it will be exciting to bring adaptive optics systems to photography and machine vision. These are applications where cameras are typically not being operated in a diffraction limited setting, but where optical distortions can be much more severe than in traditional applications of AO. Our work shows that AO is promising in these scenarios, and believe that it presents a major step towards a more widespread adaptation of AO.

ACKNOWLEDGMENTS

This work was supported by King Abdullah University of Science and Technology under Individual Baseline Funding.

REFERENCES

- Mariana SC Almeida and Mario Figueiredo. 2013. Deconvolving images with unknown boundaries using the alternating direction method of multipliers. *IEEE Trans. Image Process* 22, 8 (2013), 3074–3086.
- Nicholas Antipa, Sylvia Necula, Ren Ng, and Laura Waller. 2016. Single-shot diffuser-encoded light field imaging. In *Proc. ICCP*. IEEE, 1–11.

(a) Deblurring with anisoplanatism

(b) Measured Strehl ratios with anisoplanatism

Fig. 15. Anisoplanatism. (a) A limitation of AO systems like ours is that they only estimate the distorted waveform for one viewing direction. Depending on the complexity of the distortion, this measurement may not accurately represent off-axis distortions, which can therefore not be compensated exactly. In this experiment, center region (the wing) is well-compensated whereas surrounding regions (the head) suffer from incomplete compensation. (b) Strehl ratios decay with increased FoV, and the deblurring performance decreases. The Strehl ratios for the uncorrected wavefronts were so small that they could not be measured accurately with the method from Section 5.

- Bradley Atcheson, Ivo Ihrke, Wolfgang Heidrich, Art Tevs, Derek Bradley, Marcus Magnor, and Hans-Peter Seidel. 2008. Time-resolved 3d capture of non-stationary gas flows. *ACM Trans. Graph.* 27, 5 (2008), 132.
- Jacques M Beckers. 1988. Increasing the size of the isoplanatic patch with multiconjugate adaptive optics. In *Proc. ESOCW*, Vol. 30. 693.
- Jacques M Beckers. 1993. Adaptive optics for astronomy: principles, performance, and applications. *Annu. Rev. Astron. Astrophys.* 31, 1 (1993), 13–62.
- Pascal Berto, Hervé Rigneault, and Marc Guillon. 2017. Wavefront sensing with a thin diffuser. *Opt. Lett.* 42, 24 (2017), 5117–5120.

- Sébastien Bérubon, Eric Ziegler, Roberto Cerbino, and Luca Peverini. 2012. Two-dimensional x-ray beam phase sensing. *Phys. Rev. Lett.* 108, 15 (2012), 158102.
- Thomas G Bifano, Julie Perreault, R Krishnamoorthy Mali, and Mark N Horenstein. 1999. Microelectromechanical deformable mirrors. *IEEE J. Select. Topics Quantum Electron* 5, 1 (1999), 83–89.
- Martin J Booth. 2014. Adaptive optical microscopy: the ongoing quest for a perfect image. *Light Sci. Appl.* 3, 4 (2014), e165.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.
- Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. 2004. High accuracy optical flow estimation based on a theory for warping. In *Proc. ECCV*. Springer, 25–36.
- Sunghyun Cho and Seungyong Lee. 2009. Fast motion deblurring. *ACM Trans. Graph.* 28, 5 (2009), 145.
- Sunghyun Cho, Jue Wang, and Seungyong Lee. 2011. Handling outliers in non-blind image deconvolution. In *Proc. ICCV*. IEEE, 495–502.
- Oliver Cossairt and Shree Nayar. 2010. Spectral focal sweep: Extended depth of field from chromatic aberrations. In *Proc. ICCP*. IEEE, 1–8.
- E Diolaiti, R Ragazzoni, and M Tordi. 2001. Closed loop performance of a layer-oriented multi-conjugate adaptive optics system. *Astron. Astrophys* 372, 2 (2001), 710–718.
- Jiangxin Dong, Jinshan Pan, Zhixun Su, and Ming-Hsuan Yang. 2017. Blind Image Deblurring with Outlier Handling. In *Proc. CVPR*. 2478–2486.
- Edward R Dowski and W Thomas Cathey. 1995. Extended depth of field through wave-front coding. *Appl. Opt.* 34, 11 (1995), 1859–1866.
- Rob Fergus, Aaron Hertzmann, Sam T Roweis, and William T Freeman. 2006. Removing camera shake from a single photograph. *ACM Trans. Graph.* 25, 3 (2006), 787–794.
- Enrique J Fernández, Ignacio Iglesias, and Pablo Artal. 2001. Closed-loop adaptive optics in the human eye. *Opt. Lett.* 26, 10 (2001), 746–748.
- B Hermann, EJ Fernández, A Unterhuber, H Sattmann, AF Fercher, W Drexler, PM Prieto, and P Artal. 2004. Adaptive-optics ultrahigh-resolution optical coherence tomography. *Opt. Lett.* 29, 18 (2004), 2142–2144.
- Berthold KP Horn and Brian G Schunck. 1981. Determining optical flow. *Artificial intelligence* 17, 1-3 (1981), 185–203.
- Lifa Hu, Li Xuan, Yongjun Liu, Zhaoliang Cao, Dayu Li, and QuanQuan Mu. 2004. Phase-only liquid-crystal spatial light modulator for wave-front correction with high precision. *Opt. Express* 12, 26 (2004), 6403–6409.
- Zhe Hu, Sunghyun Cho, Jue Wang, and Ming-Hsuan Yang. 2014. Deblurring low-light images with light streaks. In *Proc. CVPR*. 3382–3389.
- Julian Iseringhausen, Bastian Goldlücke, Nina Pesheva, Stanimir Iliev, Alexander Wender, Martin Fuchs, and Matthias B Hullin. 2017. 4D imaging through spray-on optics. *ACM Trans. Graph.* 36, 4 (2017), 35.
- Mooseok Jang, Haowen Ruan, Haojiang Zhou, Benjamin Judkewitz, and Changhui Yang. 2014. Method for auto-alignment of digital optical phase conjugation systems based on digital propagation. *Opt. Express* 22, 12 (2014), 14054–14071.
- Hui Ji and Kang Wang. 2012. Robust image deblurring with an inaccurate blur kernel. *IEEE Trans. Image Process.* 21, 4 (2012), 1624–1634.
- Na Ji, Daniel E Milkie, and Eric Betzig. 2010. Adaptive optics via pupil segmentation for high-resolution imaging in biological tissues. *Nat. Methods* 7, 2 (2010), 141–147.
- Dilip Krishnan, Terence Tay, and Rob Fergus. 2011. Blind deconvolution using a normalized sparsity measure. In *Proc. CVPR*. IEEE, 233–240.
- RG Lane, A Glindemann, JC Dainty, et al. 1992. Simulation of a Kolmogorov phase screen. *Waves in random media* 2, 3 (1992), 209–224.
- Anat Levin, Rob Fergus, Frédo Durand, and William T Freeman. 2007. Image and depth from a conventional camera with a coded aperture. *ACM Trans. Graph.* 26, 3 (2007), 70.
- Anat Levin, Samuel W Hasinoff, Paul Green, Frédo Durand, and William T Freeman. 2009. 4D frequency analysis of computational cameras for depth of field extension. *ACM Trans. Graph.* 28, 3 (2009), 97.
- Anat Levin, Peter Sand, Taeg Sang Cho, Frédo Durand, and William T Freeman. 2008. Motion-invariant photography. *ACM Trans. Graph.* 27, 3 (2008), 71.
- Chao Li, Mingliang Xia, Quanquan Mu, Baoguang Jiang, Li Xuan, and Zhaoliang Cao. 2009. High-precision open-loop adaptive optics system based on LC-SLM. *Opt. Express* 17, 13 (2009), 10774–10781.
- Junzhong Liang, David R Williams, and Donald T Miller. 1997. Supernormal vision and high-resolution retinal imaging through adaptive optics. *JOSA A* 14, 11 (1997), 2884–2892.
- Gordon D Love. 1997. Wave-front correction and production of Zernike modes with a liquid-crystal spatial light modulator. *Appl. Opt.* 36, 7 (1997), 1517–1524.
- Enrico Marchetti, Roland Brast, Bernard Delabre, Robert Donaldson, Enrico Fedrigo, Christoph Frank, Norbert Hubin, Johann Kolb, Jean-Louis Lizon, Massimiliano Marchesi, et al. 2007. On-sky testing of the multi-conjugate adaptive optics demonstrator. *The Messenger* 129, 8 (2007).
- PN Marsh, D Burns, and JM Girkin. 2003. Practical implementation of adaptive optics in multiphoton microscopy. *Opt. Express* 11, 10 (2003), 1123–1130.
- Kshitij Marwah, Gordon Wetzstein, Yosuke Bando, and Ramesh Raskar. 2013. Compressive light field photography using overcomplete dictionaries and optimized projections. *ACM Trans. Graph.* 32, 4 (2013), 46.
- Tomer Michaeli and Michal Irani. 2014. Blind deblurring using internal patch recurrence. In *Proc. ECCV*. Springer, 783–798.
- Kaye S Morgan, David M Paganin, and Karen KW Siu. 2012. X-ray phase imaging with a paper analyzer. *Appl. Phys. Lett.* 100, 12 (2012), 124102.
- Ren Ng. 2005. Fourier slice photography. *ACM Trans. Graph.* 24, 3 (2005), 735–744.
- Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz, and Pat Hanrahan. 2005. Light field photography with a hand-held plenoptic camera. *CSTR* 2, 11 (2005), 1–11.
- Jinshan Pan, Deqing Sun, Hanspeter Pfister, and Ming-Hsuan Yang. 2016. Blind image deblurring using dark channel prior. In *Proc. CVPR*. 1628–1636.
- Julie A Perreault, Thomas G Bifano, B Martin Levine, and Mark N Horenstein. 2002. Adaptive optic correction using microelectromechanical deformable mirrors. *Opt. Eng.* 41, 3 (2002), 561–566.
- Jérôme Primot and Nicolas Guérineau. 2000. Extended Hartmann test based on the pseudoguiding property of a Hartmann mask completed by a phase chessboard. *Appl. Opt.* 39, 31 (2000), 5715–5720.
- JetLSOGNO Primot and L Sogno. 1995. Achromatic three-wave (or more) lateral shearing interferometer. *JOSA A* 12, 12 (1995), 2679–2685.
- Roberto Ragazzoni. 1996. Pupil plane wavefront sensing with an oscillating prism. *J MOD OPTIC* 43, 2 (1996), 289–293.
- Ramesh Raskar, Amit Agrawal, and Jack Tumblin. 2006. Coded exposure photography: motion deblurring using fluttered shutter. *ACM Trans. Graph.* 25, 3 (2006), 795–804.
- Hughes Richard and Markus Raffel. 2001. Principle and Applications of the Background Oriented Schlieren (BOS) Method. *Meas. Sci. and Techn.* (2001), 1576–1585.
- Francois Roddier. 1988. Curvature sensing and compensation: a new concept in adaptive optics. *Appl. Opt.* 27, 7 (1988), 1223–1225.
- Austin Roorda, Fernando Romero-Borja, William J Donnelly III, Hope Queener, Thomas J Hebert, and Melanie CW Campbell. 2002. Adaptive optics scanning laser ophthalmoscopy. *Opt. Express* 10, 9 (2002), 405–412.
- Leslie K Saddlemyer, Glen Herriot, Jean-Pierre Véran, and Murray Fletcher. 1998. Design aspects of the reconstructor for the Gemini Adaptive Optics System (Altair). 3353 (1998), 150–159.
- Roland V Shack and BC Platt. 1971. Production and use of a lenticular Hartmann screen. *JOSA A* 61, 5 (1971), 656.
- Deqing Sun, Stefan Roth, and Michael J Black. 2010. Secrets of optical flow estimation and their principles. In *Proc. CVPR*. IEEE, 2432–2439.
- Libin Sun, Sunghyun Cho, Jue Wang, and James Hays. 2013. Edge-based blur kernel estimation using patch priors. In *Proc. ICCP*. IEEE, 1–8.
- Michael Reed Teague. 1983. Deterministic phase retrieval: a Green’s function solution. *JOSA A* 73, 11 (1983), 1434–1441.
- Fernando Vargas-Martin, Pedro M Prieto, and Pablo Artal. 1998. Correction of the aberrations in the human eye with a liquid-crystal spatial light modulator: limits to performance. *JOSA A* 15, 9 (1998), 2552–2562.
- Ashok Veeraraghavan, Ramesh Raskar, Amit Agrawal, Ankit Mohan, and Jack Tumblin. 2007. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph.* 26, 3 (2007), 69.
- Laura Waller, Shan Shan Kou, Colin JR Sheppard, and George Barbastathis. 2010a. Phase from chromatic aberrations. *Opt. Express* 18, 22 (2010), 22817–22825.
- Laura Waller, Lei Tian, and George Barbastathis. 2010b. Transport of intensity phase-amplitude imaging with higher order intensity derivatives. *Opt. Express* 18, 12 (2010), 12552–12561.
- Congli Wang, Xiong Dun, Qiang Fu, and Wolfgang Heidrich. 2017. Ultra-high resolution coded wavefront sensor. *Opt. Express* 25, 12 (2017), 13736–13746.
- Li Xu and Jiaya Jia. 2010. Two-phase kernel estimation for robust motion deblurring. In *Proc. ECCV*. Springer, 157–170.
- Li Xu, Shicheng Zheng, and Jiaya Jia. 2013. Unnatural l0 sparse representation for natural image deblurring. In *Proc. CVPR*. IEEE, 1107–1114.
- Zhimin Xu, Jun Ke, and Edmund Y Lam. 2012. High-resolution lightfield photography using two masks. *Opt. Express* 20, 10 (2012), 10971–10983.
- Yan Zhang, Barry Cense, Jungtae Rha, Ravi S Jonnal, Weihua Gao, Robert J Zawadzki, John S Werner, Steve Jones, Scot Olivier, and Donald T Miller. 2006. High-speed volumetric imaging of cone photoreceptors with adaptive optics spectral-domain optical coherence tomography. *Opt. Express* 14, 10 (2006), 4380–4394.
- Changyin Zhou and Shree Nayar. 2009. What are good apertures for defocus deblurring?. In *Proc. ICCP*. IEEE, 1–8.
- Lijun Zhu, Pang-Chen Sun, Dirk-Uwe Bartsch, William R Freeman, and Yesiaiahu Fainman. 1999. Adaptive control of a micromachined continuous-membrane deformable mirror for aberration compensation. *Appl. Opt.* 38, 1 (1999), 168–176.