

The Diffractive Achromat

Full Spectrum Computational Imaging with Diffractive Optics

Yifan Peng^{2,1*} Qiang Fu^{1†} Felix Heide^{2,1‡} Wolfgang Heidrich^{1,2§}

¹ King Abdullah University of Science and Technology ² The University of British Columbia

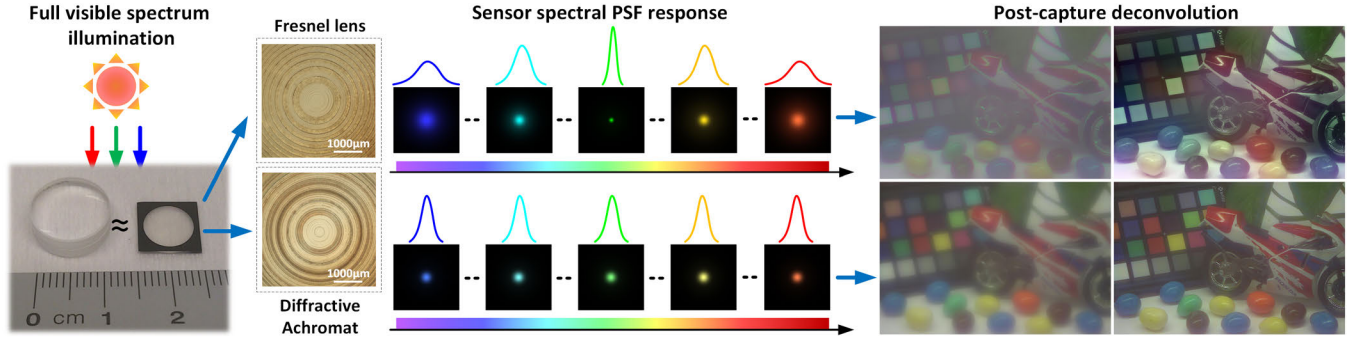


Figure 1: The diffractive achromat is a computationally optimized diffractive lens for full visible spectrum imaging, which is used jointly with a computational image reconstruction algorithm. The microscope images show a traditional Fresnel diffraction grating (top) and our diffractive achromat (bottom). In full visible spectrum illumination, the former can only be focused only at one specific wavelength (e.g. green here) while all other wavelengths are out of focus. This results in highly nonuniform spatial and spectral response (color PSFs) on the image plane coupled with Bayer filters (top middle). In particular, metamerism introduces a data dependency in the PSF shape for any kind of broadband image sensor. Our diffractive achromat is optimized to equalize the spectral focusing performance within the whole visible spectrum. Consequently, the PSFs for all wavelengths are nearly identical to each other (bottom middle). The captured blurry image shows much higher color fidelity than the conventional diffractive lens (right). Our diffractive achromat is much thinner and lighter than an refractive achromatic lens with the same optical power (left-most bottom).

Abstract

Diffractive optical elements (DOEs) have recently drawn great attention in computational imaging because they can drastically reduce the size and weight of imaging devices compared to their refractive counterparts. However, the inherent strong dispersion is a tremendous obstacle that limits the use of DOEs in full spectrum imaging, causing unacceptable loss of color fidelity in the images. In particular, metamerism introduces a data dependency in the image blur, which has been neglected in computational imaging methods so far. We introduce both a *diffractive achromat* based on computational optimization, as well as a corresponding algorithm for correction of residual aberrations. Using this approach, we demonstrate high fidelity color diffractive-only imaging over the full visible spectrum. In the optical design, the height profile of a diffractive lens is optimized to balance the focusing contributions of different wavelengths for a specific focal length. The spectral point spread functions (PSFs) become nearly identical to each other, creating approximately spectrally invariant blur kernels. This property guarantees good color preservation in the captured image and facilitates the correction of residual aberrations in our fast two-step deconvolution without additional color priors. We demonstrate our design of diffractive achromat on a 0.5mm ultrathin substrate by photolithography techniques. Experimental results show that our achromatic diffractive lens produces high color fidelity and better image quality in the full visible spectrum.

Keywords: achromatic, ultrathin, DOE, computational imaging

*e-mail: evanpeng@cs.ubc.ca

†e-mail: vorlahm@gmail.com

‡e-mail: fheide@cs.ubc.ca

§e-mail: wolfgang.heidrich@kaust.edu.sa

Concepts: •Computing methodologies → Computational photography;

1 Introduction

High quality imaging with reduced optical complexity has for a long time been the target of investigation in both academic and industrial research and development. In conventional imaging systems, ever increasing optical complexity is inevitable because higher and higher sensor resolutions require ever improved correction of aberrations of all kinds. Recent advances in computational imaging have introduced computation as a virtual component that can shift the burden from optics to algorithms. This allows for significantly reduced optical complexity while maintaining high image fidelity at full sensor resolution and realistic apertures (e.g. [Heide et al. 2013; Schuler et al. 2011]). In particular, diffractive optical elements (DOEs) have drawn great attention because of their ultrathin and lightweight physical structure, a large, flexible design space, availability of mature fabrication techniques, as well as better off-axis imaging behavior. Integrating diffractive imaging elements and computational methods

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 ACM. SIGGRAPH '16 Technical Paper, July 24-28, 2016, Anaheim, CA, ISBN: 978-1-4503-4279-7/16/07 DOI: <http://dx.doi.org/10.1145/2897824.2925941>

in a single imaging system has resulted in several new computational imaging devices with ultra compactness in the past few years (e.g. [Gill and Stork 2013; Stork and Gill 2014; Peng et al. 2015]).

Unfortunately, two main problems still exist for diffractive imaging applied in the full visible spectrum — first, the wavelength-dependency of diffraction leads to strong chromatic aberrations that degrade the image quality with blurs of a very large diameter; second, the strong chromatic aberrations in addition cause a significant wavelength dependency of the point spread functions (PSFs) even within a single color channel. In particular, this wavelength dependency means that objects with the same RGB color are blurred differently if the underlying spectral distributions differ. This *metamerism* problem means that image restoration algorithms in practice use approximate PSFs based on some fixed spectral distribution that is implicitly derived from the lighting conditions during calibration. This approximation can result in a significant loss of color fidelity even after image reconstruction. While metamerism affects all imaging systems exhibiting chromatic aberration, the problem is particularly pronounced for the strong dispersion present in diffractive optics.

In this paper, we aim to overcome the above limitations of diffractive imaging in the full visible spectrum by introducing not only improvements in the deconvolution, but more importantly, optimizing the diffractive optics itself. We find that the color fidelity loss in conventional diffractive imaging is caused by the inherent non-uniformity of spectral PSFs, which inspires us to design a *diffractive achromat* by optimizing the surface profile to produce nearly identical spectral PSF distribution for a series of wavelengths (see Figure 1). The benefit of this strategy is twofold. On the one hand, chromatic aberrations are reduced because of the balance among spectral PSFs. On the other hand, the quasi-uniform spectral PSFs significantly improve the color reproduction of the captured image. Effectively, we sacrifice sharpness in a single channel for spectral uniformity of PSF, which we can exploit to facilitate robust and efficient deconvolution. In addition to validating the usability of diffractive imaging under state-of-the-art deconvolution schemes, we explore a two-step cross-scale deconvolution scheme to recover sharp as well as color fidelity lossless images.

Particularly, our technical contributions are as follows:

- We introduce the diffractive achromat, a diffractive imaging system for full-spectrum visible light imaging that combines optimization in both diffractive optical design and post-capture image reconstruction.
- We employ an effective optimization method for designing achromatic DOEs subject to diverse achromatic requirements, which rearranges the spatial and spectral distributions of PSFs so that chromatic aberrations and color corruption are largely eliminated in hardware.
- We propose a cross-scale prior in the deconvolution to further mitigate the aberrations introduced by the diffractive optics. Benefiting from the optimized uniform PSFs, our method is more robust and efficient than state of the art.
- We build a prototype achromatic diffractive lens on a 0.5mm ultrathin glass plate to validate the practicality of our ultrathin and lightweight diffractive imaging in full visible spectrum for different indoor and outdoor scenarios.

2 Related Work

Color fidelity in computational imaging For conventional image acquisition, color fidelity is mostly preserved since most refractive lenses are carefully designed with spectral invariant focus-

ing power [Smith 2005; Yamaguchi et al. 2008]. Then state-of-the-art deconvolution formations can be directly applied to computationally recover color images. However, the spectral response of a diffractive lens affects color fidelity drastically because its highly spectral variant focusing power goes against the blur kernel convolution model in sRGB color space (see Section 3.1 for theoretical analysis). Consequently, the ill-posed inverse problem may lead to unacceptable color artifacts, even if the perceptual blur has been mostly removed. Preserving color fidelity is crucial for any consumer imaging device [Su et al. 2014], which is also the goal we seek in our diffractive achromat design.

Broadband diffractive imaging Despite advantages such as thin structure and design flexibility, the severe chromatic aberration in diffractive optics has limited their applications in imaging under broadband illumination. A limited amount of work has considered applying DOEs in consumer imaging devices, but only in collaboration with refractive lenses. A multilayer DOE has been used to correct chromatic aberrations in refractive lens systems [Nakai and Ogawa 2002], although it still relies on a multitude of refractive lens elements to provide most of the focal power. A recent report on multi-wavelength achromatic metasurfaces [Aieta et al. 2015] has revealed the potential for use in lightweight collimators. As such, a chromatic-aberration-corrected diffractive lens for broadband focusing has been designed [Wang et al. 2016]. Note that these works don’t address the reconstruction problem as we do and they are only designed at three wavelengths.

In computational imaging, two types of imaging devices based on DOEs have recently been investigated: lensless computational sensors [Gill and Stork 2013; Monjur et al. 2015] and Fresnel lens imaging with post processing [Nikonorov et al. 2015; Peng et al. 2015]. The former two integrate DOEs into the sensor, resulting in an ultra-miniature structure with medium image quality. The latter two share a similar idea with ours to apply Fresnel lens to replace bulky and heavy refractive lenses in a camera. Although the chromatic aberrations can be partially mitigated by optimizing for three discrete wavelengths or digitally removed in the post-capture step, the color fidelity is significantly reduced in the final image due to metamerism. In our work we present a combination of optical design and computational reconstruction that allows us to perform color imaging at realistic image resolutions (full resolution on a >5Mpixel image sensor), to our knowledge for the first time.

Digital correction of aberrations Optical aberrations can be corrected by utilizing state-of-the-art image deblurring methods. The principle is to formulate the image formation as a convolution process and apply statistical priors [Chan et al. 2011] to obtain an optimal solution with reasonable complexity [Shan et al. 2008]. Usually, a denoising step is added to improve image quality [Schuler et al. 2013]. Existing image deblurring techniques either assume the aberration-induced blur kernel is known [Joshi et al. 2008] or use an expected maximization-type approach to blindly estimating the blur kernel [Krishnan et al. 2011]. Both techniques involve a convolution based image blurring model without considering the spectral variance in PSFs.

For the correction of chromatic aberrations, cross-channel optimization has proven to be effective [Heide et al. 2013]. This method models the three color channels separately and exploits the correlation between gradients in different channels, which provides better localization. Since existing deblurring techniques for removing chromatic aberrations all rely on additional color priors [Yue et al. 2015], the computational complexity is considerable.

Although the correction of chromatic aberrations for diffractive imaging is also plausible [Nikonorov et al. 2015; Peng et al. 2015] using existing techniques, the residual color artifacts are inevitable. We will show that this is due to the inherent nonuniform distribution of spectral PSFs, which leads to a metamerism problem, where the PSF becomes scene dependent. Although this problem affects all broadband optical systems with chromatic aberration, it is especially pronounced in diffractive optics due to the large wavelength dependency of diffraction. Special attention has to be paid to the optimization of the diffractive lens in order to solve this problem.

Computational DOE design A huge amount of research has been done in designing DOEs used for multiple wavelengths, either by using harmonic diffraction and multiple layers to create high efficiencies in multiple wavebands [Singh et al. 2014], or by introducing computation to the design footprint to redefine the light transmission function. The latter case is of significance not only to mitigate dispersion, but also to encode the engineered phase distribution on the DOE profile for the tasks of image recovery or special focusing expectation [Quirin and Piestun 2013]. The principle is to design Fractal diffraction elements with variable transmittance [Muzychenko et al. 2011]. Besides, iterative methods based on greedy algorithms, such as Gerchberg-Saxton, genetic algorithms, simulated annealing, and direct binary search, have been extensively applied for optimizing both monochromatic and broadband DOEs [Kim et al. 2012; Zhou et al. 1999; Jiang et al. 2013]. Researchers in graphics have proposed alternative optimization methods in simulating and creating wave optics imaging for visualization purposes [Ye et al. 2014; Schwartzburg et al. 2014]. However, these approaches fail in our case for two reasons: first, they are not intended for broadband imaging DOE designs; second, they rely on either matured random search algorithms or their extended versions, lowering the computational efficiency and optimization robustness in our diffractive achromat designs.

3 Diffractive Imaging Model

3.1 Image Formation

In an imaging system, the recorded image in channel c is an integration of spectral images over the wavelength range Λ , weighted by the spectral response $Q_c(\lambda)$ of the sensors for that channel. Each spectral image reflects the joint modulation of illumination, surface reflectance and sensor response. This process can be written as

$$b_c(x, y) = \int_{\Lambda} Q_c(\lambda) \cdot \mathcal{A}(i(x, y; \lambda)) d\lambda, \quad (1)$$

where $i(x, y; \lambda)$ is the latent *spectral* image, and $\mathcal{A}(\cdot)$ denotes an operator describing the aberrations of the lens.

In Fourier optics, the incoherent imaging process is modeled as a convolution of the latent image and the system (intensity) PSF, so the aberration operator is defined as [Goodman 2008]

$$\mathcal{A}(i(x, y; \lambda)) = i(x, y; \lambda) \otimes |g(x, y; \lambda)|^2, \quad (2)$$

where \otimes denotes 2D convolution. $g(x, y; \lambda)$ is spectral *amplitude* PSF, from which the spectral *intensity* PSF can be derived as $k(x, y; \lambda) = |g(x, y; \lambda)|^2$. The amplitude PSF can be further derived from scalar diffraction theory [Goodman 2008] as

$$g(x, y; \lambda) = \frac{A}{\lambda z_i} \iint \mathcal{P}(u, v; \lambda) \exp\left(-j \frac{2\pi}{\lambda z_i} (ux + vy)\right) dudv, \quad (3)$$

where A is a constant amplitude, z_i is the distance from the lens to the image plane and (u, v) are coordinates on the lens plane. A generalized pupil function $\mathcal{P}(u, v; \lambda)$ accounts for the lens:

$$\mathcal{P}(u, v; \lambda) = P(u, v) \exp(j\Phi(u, v)), \quad (4)$$

where the aperture function $P(u, v)$ is usually a circ function. The phase term $\Phi(u, v)$ describes the phase retardation of light for each point on the aperture, which in a general imaging system could be caused by either refractive or diffractive optics, or a combination of the two. In our case, $\Phi(u, v)$ is the function that will be optimized to achieve a desired lensing effect and PSF. Using Eq. (4), we can rewrite Eq. (1) as

$$\begin{aligned} b_c(x, y) &= \int_{\Lambda} Q_c(\lambda) i(x, y; \lambda) \otimes k(x, y; \lambda) d\lambda \\ &= \int_{-\infty}^{\infty} \int_{\Lambda} Q_c(\lambda) i(\xi, \eta; \lambda) k(x - \xi, y - \eta; \lambda) d\lambda d\xi d\eta. \end{aligned} \quad (5)$$

For a conventional diffractive lens, the spectral PSFs $k(x, y; \lambda)$ are highly wavelength dependent, which has already been shown in Fig. 1. Therefore, the PSF is not separable from the inner integration over wavelength. This effect is usually neglected in state-of-the-art image formation models for deblurring, where the blur kernel is assumed to be convolved with the latent color image. However, this approximation does not hold for large chromatic aberrations, and therefore current deconvolution algorithms fail to recover the latent image with high color fidelity.

To guarantee the deconvolution algorithms work in RGB space, we design the spectral PSFs to be nearly wavelength independent, at least over the spectral support of each color channel. On only then can we use the approximation $k(x, y; \lambda) \approx k_c(x, y)$, resulting in

$$\begin{aligned} b_c(x, y) &\approx \int_{-\infty}^{\infty} k_c(x - \xi, y - \eta) \underbrace{\int_{\sigma} Q_c(\lambda) i(\xi, \eta; \lambda) d\lambda}_{i_c(x, y)} d\xi d\eta \\ &= k_c(x, y) \otimes i_c(x, y), \end{aligned} \quad (6)$$

where $i_c(x, y)$ is the latent color image in channel c .

For an RGB image, the vector form of Eq. (6) can then be written as

$$\mathbf{b}_c = \mathbf{K}_c \mathbf{i}_c, \quad c = 1, 2, 3. \quad (7)$$

As long as we can design a diffractive lens with nearly constant spectral PSFs (i.e. a diffractive achromat), the convolutional image formation model holds again.

3.2 Imaging Approach Overview

Our achromatic diffractive imaging approach consists of two main parts — achromatic lens design and post-capture processing.

In the lens design, we devise a method that optimizes a DOE to focus a range of wavelengths at a certain target focal length while maintaining both a compact PSF as well as spectral uniformity. This optimization is performed for a dense set of discrete wavelengths that are uniformly distributed over the target range.

Once we have a diffractive lens that exhibits spectral invariant kernel behavior, the post-capture processing is modeled as an optimization problem to solve the inverse problem of Eq. 7. A two-step, cross-scale deconvolution is presented to reconstruct the image.

4 Optical Design Optimization

4.1 Optimization Model

We have seen from Eqs. (3) and (4) that the spectral PSFs are determined by the phase profile of the diffractive lens. Further, the phase profile is a function of the height map $h(u, v)$ of a transmissive substrate

$$\Phi(u, v) = \frac{2\pi}{\lambda} (n_\lambda - 1) h(u, v), \quad (8)$$

where n_λ is the refractive index of the substrate. We can control the height map of the diffractive lens so that the PSFs can fit our target PSFs in the ℓ_1 sense. We choose the ℓ_1 norm instead of least squared error term because penalizing the absolute value is more robust to outliers, which in our case are the sparse high frequency components (e.g. glitch intensity) of PSFs. These high spatial frequencies are caused by optimizing discrete wavelengths individually, which leads to coherent interference patterns. Note that we are targeting imaging with incoherent light, such that the real PSFs are always smooth in practice. We simulate incoherence by low pass filtering and outlier removal.

The minimization problem is written as

$$h_{opt} = \operatorname{argmin}_h \sum_{\lambda_i \in \Lambda} w_i \|p_i(h) - t\|_1, \quad (9)$$

where we have omitted the spatial coordinates for brevity. Here $p_i(h)$ are the optimized PSFs, and t is the wavelength-independent target PSF. The weights w_i are assigned to balance relative diffraction efficiencies among wavelengths (see below). Note that the optimization uses a discrete set of design wavelengths λ_i which densely sample the target spectral range. This is possible because PSFs vary smoothly with wavelength. The implementation of $p_i(h)$ follows directly from Eq. (3): we first calculate the amplitude PSFs by Fresnel diffraction propagation and then take the magnitude squared.

In this work, we consider only rotationally symmetric patterns. As a result, we can reduce the optimization to a 1D problem. Operating directly on height profile is beneficial as fabrication constraints and other proximity effects can be incorporated into the model. The proposed method is summarized in Alg. 1. We discuss the algorithm in detail below.

Target PSFs In optical design, blur kernels are usually represented as Gaussian distributions with different variances. However, in our design we will end up sacrificing resolution in the central wavelengths for improved resolution at both longer and shorter wavelengths, so that the final PSF is achromatic but not as sharp. We expect this process to introduce longer tails that are not represented well with a single Gaussian.

To seek an optimal distribution that is feasible with the current physical profile, we adaptively tune the target function. Specifically, after a few iterations we average the PSF distributions of all wavelengths and fit this average to a mixture model of three (centered) Gaussians to represent the PSF

$$t = \sum_{j=1,2,3} a_j \mathcal{G}(\mu, \sigma_j), \quad (10)$$

where a_j are the weights for each component function and $\sum_{j=1}^3 a_j = 1$. The coefficients and parameters of the fitting model are then tuned (e.g. σ_j are shrunk by a few pixels) to generate new target PSFs with a sharper distribution. Via this update strategy, we

Algorithm 1 Optimization on diffractive achromat

```

1:  $k = 0, h_m^0 = h_{init}, t = t_{init}, w_i^0 = w_{init}, v_m^0 = v_{init}$ 
2: for iterations do ▷ Repeat till convergence/termination
3:   for all seeds  $1, \dots, m$  do ▷ Implement  $m$  seeds in-parallel
4:      $h_m^{k+1} = h_m^k + v_m^k$  ▷ Update height profile
5:      $h_m^{k+1} = \text{OPT}(h_m^k, h_m^{k+1})$  ▷ Update local optimal
6:      $\bar{h}^{k+1} = \text{OPT}(h_m^{k+1})$  ▷ Update global optimal
7:      $v_m^{k+1} = v_m^k + c_1(\bar{h}_m^{k+1} - h_m^{k+1}) + c_2(\bar{h}^{k+1} - h_m^{k+1})$  ▷
      Update velocity vector for each seed
8:   end for
9: end for
10: function OPT( $h$ ) ▷ Store updated height profile
11:   for  $i = 1$  to  $N$  do
12:      $p_i = p_i(h) \cdot f$  ▷ Update PSFs
13:   end for
14:    $h_{opt} = \min_h \sum_i w_i^k \|p_i - t\|_1$  ▷ Evaluate objective
15:    $w_i^{k+1} = \frac{\|p_i - t\|_1}{\sum_i \|p_i - t\|_1}$  ▷ Update weights
16:   return  $h_{opt}$ 
17: end function

```




Figure 2: The initial height profile for the optimization is a mixture of subregions screened from Fresnel lenses at the same focal length for different wavelengths. The whole area is divided into N subregions for N different wavelengths. The resulting initial height profile is the superposition of such subregions.

only need to initialize the target PSFs once based on preliminary simulation. Repeatedly fitting the averaged distribution assures that we have considered the focusing contribution of all wavelengths.

Despite its achromatic focusing constraint, this PSF representation helps maintain a relatively sharp peak in the PSFs so that high-frequency features can be preserved. This is particularly beneficial for deconvolution under the challenge of large kernels.

Initialization The optimization begins by calculating an initial guess for the starting height profile. A purely random height profile leaves our input too far away from the solution. We could start from a Fresnel phase plate designed for a single, central wavelength (e.g. 550nm for the visible spectrum). However, we found that a better choice is to use a composite of multiple zone plates. When optimizing for N discrete wavelengths, we divide the aperture into N rings of equal area. Within each ring, we initialize the height field to a Fresnel phase plate for a specific wavelength λ_i at the target focal length (see Fig.2).

Adaptive weights The final goal of our optimization is to uniformly assign the focal contribution of each wavelength. However, during the optimization we still adaptively tune the weights w_i in Eq.(9) for selected wavelengths according to their deviations of fitting errors. Specially, if the current design suffers from a weaker

optical power at one wavelength, which is explicitly reflected by larger fitting error, the weight on this wavelength in the cost function will be adjusted in the next iteration following a simple rule as $w_i^{k+1} = \frac{\|p_i - t\|_1}{\sum_i \|p_i - t\|_1}$. We see that all weights w_i will be updated toward nearly identical value if the optimization approaches the optimal.

Low-pass filtering The resulting PSFs after diffraction propagation contain narrow spectral peaks and valleys (Fig. 3). These artifacts are due to two types of discretization. First, both the DOE plane and the image plane are represented as point samples, which introduces high spatial frequencies that in reality get averaged out by integrating over finite pixel areas. The second, and maybe more important effect is that our simulation is also discretized along the wavelength direction, which effectively treats the light as coherent and introduces artificial interference patterns that are not present in real-world broad-band imaging scenarios. As a result of this analysis, we treat the spectral peaks and valleys as outliers that we filter out by applying a blur along the spectral dimension. Our experience indicates this filtering benefits the robustness and convergence speed of the optimization.

4.2 Stochastic Optimization Algorithm

We choose Particle Swarm Optimization (PSO) algorithm to solve our optimization [Eberhart and Shi 2001]. The advantage of PSO algorithm is its high computing efficiency compared with other stochastic algorithms, e.g. Genetic Algorithm etc. By implementing a series of seeds (i.e. m seeds in Alg. 1) in parallel at each iteration, the height profile update is more robust and faster to converge.

At each iteration, the seeds update their height profiles and velocities of current design by tracking the optimal solution of their own h and that of the group \bar{h} following a strategy $h^{k+1} = h^k + v^k$, where v^k is a velocity vector indicated in line 7 of Alg. 1. The two weights c_1, c_2 are randomly assigned between (0, 1). We further set a constraint in the implementation that $|v_m| \leq 0.25h_{max}$.

The idea is to assume individual seeds can evolve according to the information gathered from their own experience and that of the group, so that the focal power change for each individual wavelength is not be drastic and purely random. This update strategy is beneficial to avoid falling into a local minimum, as well as leveraging parallelism.

We borrow the *from-coarse-to-fine* strategy from multi-scale optimization to divide the N wavelengths to be optimized into several scales. For instance, we start by optimizing $N_1 = 9$ sampled wavelengths in the spectrum. Once the optimization for this scale has converged, we increase the number of wavelengths to a second level $N_2 = 15$, and $N_3 = 29$ eventually. We can accept an $10nm$ interval in the wavelength sampling to already approximate achromaticity among wavelengths in the full visible spectrum.

Figure 3 shows the comparison of simulated PSF cross-sections for three selected wavelengths $\lambda_1 = 650nm$, $\lambda_2 = 550nm$ and $\lambda_3 = 450nm$ in the full spectrum. We compare a single wavelength Fresnel phase plate (left) with the multi-ring initialization (center) and the final, optimized result (right). Our algorithm sacrifices the performance of the central wavelength to compromise the spectral focal contributions. It's worth noting that the shown PSFs correspond to single wavelengths, which explains the high spatial frequencies, which are interference patterns due to coherence. These patterns average out for incoherent illumination (see Section 6 and Fig. 11).




Figure 3: Cross sections of normalized PSFs for a regular Fresnel lens (left), the initial guess (center), and the optimized diffractive achromat (right) at three selected typical wavelengths $\lambda_1 = 650nm$, $\lambda_2 = 550nm$ and $\lambda_3 = 450nm$. Note that here we optimize for 29 wavelengths from $410nm$ to $690nm$ with an interval of $10nm$. We sacrifice the performance at the central wavelength to equalize PSF distributions.

5 Image Reconstruction

In this section we introduce the algorithms for solving the inverse problem to recover sharp and high-fidelity color images using a cross-scale prior in a two-step fashion.

5.1 Optimization Method

In order to recover latent images, our approach seeks the solution to Eq. 7 by solving the following minimization problem

$$\mathbf{i}_c = \underset{\mathbf{i}_c}{\operatorname{argmin}} \frac{\mu_c}{2} \|\mathbf{b}_c - \mathbf{K}_c \mathbf{i}_c\|_2^2 + \Gamma(\mathbf{i}_c), \quad (11)$$

where the first term is a standard least-square data fitting term, μ_c is the weight of data fitting term for channel $c = 1, 2, 3$. The second term $\Gamma(\mathbf{i}_c)$ is a regularization term that enforces natural image priors in the solution, as explained in detail below.

We investigate an efficient non-blind deconvolution accounting for the following properties of our diffractive achromat design. First, despite the large PSF size of our lens, they show a preserved central intensity peak as well as quasi-uniform intensity distribution for all the color channels, therefore we do not need additional color priors in our problem. Second, inspired by the image-pyramid strategy, we solve our problem with at two scales, enforcing similar gradient distributions between scales. Intuitively, the edges in natural images always exist at the same locations and are barely affected when the image is downsampled to a lower scale. However, down-sampling an image would lead to an improved signal-to-noise ratio, which helps improving the conditioning of the problem. The latter point can be particularly beneficial in our case.

Fast deconvolution and denoising at downsampled scale

We propose to implement the first step deconvolution on a downsampled image, for instance, half the size of the original image, to deblur large edges and remove strong color corrupted noise. By defining the regularization term $\Gamma(\mathbf{i}_c)$ in this scale, the cost function for a single channel in Eq. (11) is reformatted as

$$\begin{aligned} \mathbf{i}_c^d &= \underset{\mathbf{i}_c^d}{\operatorname{argmin}} \mu \|\mathbf{b}_c^d - \mathbf{K} \mathbf{i}_c^d\|_2^2 + \beta \|\mathbf{D} \mathbf{i}_c^d\|_2^2 \\ \Leftrightarrow \mathbf{i}_{\text{opt}}^d &= \mathcal{F}^{-1} \left(\frac{\mu \mathcal{F}(\mathbf{K})^* \mathcal{F}(\mathbf{b}_c^d)}{\mu \mathcal{F}(\mathbf{K})^* \mathcal{F}(\mathbf{K}) + \beta \mathcal{F}(\mathbf{D})^* \mathcal{F}(\mathbf{D})} \right), \end{aligned} \quad (12)$$

where $\mathcal{F}(\cdot)$ represents the Fourier transform, $\mathcal{F}^{-1}(\cdot)$ is its inverse. The superscript $*$ indicates the complex conjugate operation, \mathbf{D} is the first-order derivative filter matrix, and μ, β are the respective weights for each term. The superscript d denotes that all images are in the downsampled scale. This quadratic problem leads to a closed-form solution in frequency domain, such that we can directly use fast inversion to recover the sharp image \mathbf{i}_c^d at downsampled scale. In practice, we suggest to apply an additional denoising solver here at this scale if the captured image suffers from strong noise.

Cross-scale deconvolution at full scale In the second step, we apply a cross-scale prior in our regularization term, which borrows the relatively sharp and denoised edge information from the upsampled image of the 1st step’s result to benefit the deconvolution at full scale. Our cross-scale prior is inspired by cross-channel prior [Heide et al. 2013] and the multi-scale deconvolution scheme [Yuan et al. 2008]. It is reasonable to assume large edges and shapes shall be located where they are at both the upsampled version of the downsampled scale image and the original full scale image, i.e.

$$\nabla \mathbf{i}_c \approx \nabla \mathbf{i}_c^s \Leftrightarrow \mathbf{D} \mathbf{i}_c \approx \mathbf{D} \mathbf{i}_c^s, \quad (13)$$

where \mathbf{i}_c is the latent image at full scale, and \mathbf{i}_c^s is the upsampled version of the deconvolved image \mathbf{i}_c^d in the 1st step using a bicubic sampling scheme.

Note that strong noise in the original blurry image has been smoothed and mostly removed during the processing at downsampled scale, then the second step can be run with relatively weak regularization. Our approach differs from the multi-scale deconvolution in [Yuan et al. 2008], which progressively refines details in multiple scales such that at each scale iterative residual deconvolution is necessary. In our approach, we instead directly add the upsampled deconvolved image \mathbf{i}_c^s as an additional prior term to transfer their edge information between two scales. The remainder of our deconvolution problem still follows the same fast deconvolution scheme described above. By tuning the weights of two prior terms, we can flexibly compromise on sharpness and smoothness of the recovered image in a simple scheme. Then, by rewriting $\Gamma(\mathbf{i}_c)$ that includes the gradient prior as well as our cross-scale prior, the cost function in Eq. 11 is reformulated as

$$\mathbf{i}_c = \arg\min_{\mathbf{i}_c} \frac{\mu}{2} \|\mathbf{b}_c - \mathbf{K} \mathbf{i}_c\|_2^2 + \beta \|\mathbf{D} \mathbf{i}_c\|_1 + \gamma \|\mathbf{D} \mathbf{i}_c - \mathbf{D} \mathbf{i}_c^s\|_1. \quad (14)$$

Adding the cross-scale prior results in a non-linear optimization problem that can be solved by introducing slack variables for the ℓ_1 term. Specially, we form the proximal operators [Boyd et al. 2011] for the subproblems, thus turning the ℓ_1 terms into shrinkage operators. We define $\mathbf{p} = \mathbf{D} \mathbf{i}_c$ as a slack variable for

$$\begin{aligned} \text{prox}_{\theta \|\cdot\|_1}(\mathbf{p}) &= \max \left(1 - \frac{\theta}{|\mathbf{p}|}, 0 \right) \odot \mathbf{p} \\ \text{prox}_{\theta \|\cdot\|_1 - \alpha \|\cdot\|_1}(\mathbf{p}) &= \max \left(1 - \frac{\theta}{|\mathbf{p} - \alpha|}, 0 \right) \odot \mathbf{p} + \alpha, \end{aligned} \quad (15)$$

where $\alpha = \mathbf{D} \mathbf{i}_c^s$. The proximal operators for their convex conjugates can be derived from [Boyd et al. 2011]. We then use a similar half-quadratic penalty scheme as in [Krishnan and Fergus 2009] to solve Eq. (14). (see Supplementary document for details).

5.2 Efficiency and Robustness Analysis

In the first step, the direct division in frequency domain is very fast, but may result in the recovered image being strongly corrupted by noise. One can further apply a multi-layer perceptron (MLP)

approach in [Schuler et al. 2013] or similar fast method to denoise the lowest scale before upsampling. Here we directly use the MLP-based denoiser with off-line learned data in [Schuler et al. 2013] at the downsampled scale. Although the kernel of our lens differs from the Gaussian kernels with which the system has been trained the results are still perceptually pleasing for our purpose.

In the second step, we can recover the latent image efficiently at the full scale due to the introduction of the cross-scale prior. We have compared our implementation with recently reported non-blind deconvolution methods in Tab. 1. See supplementary document for PSNR results of the full dataset [Chakrabarti and Zickler 2011]. The last column in Tab. 1 indicates the case of a regular Fresnel lens with cross-channel deconvolution. We find that all the results reconstructed from the diffractive achromat show much higher PSNR than those from a regular Fresnel lens. This validates our design motivation that diffractive achromat preserves higher color fidelity than conventional diffractive lens. Further, by running our two-step deconvolution, denoising and cross-scale edge preservation, our results outperform existing methods.

Table 1: Averaged PSNR comparisons of recovered images from different deconvolution schemes. The first 5 columns indicate the results of our diffractive achromat using respectively the deconvolution by Krishnan, Schuler, multi-scale Krishnan, and multi-scale Schuler+Krishnan (see Fig. 4), while the last column shows the result for a standard Fresnel lens with a cross-channel prior.

Method	1	2	3	4	Ours	Fresnel
PSNR/dB	25.86	26.05	26.51	26.73	27.10	22.54

At either scale, we do not rely on a color regularizer in the cost function, which lowers significant computing burden. Additionally, the high frequency components in the compromised PSF is usually of very low intensity. They are actually mixed with additive noise, and are smoothed in the downsampling process. Although the theoretical size of PSFs can be considerably large, in practice, it doesn’t degrade the image so much. In our experiment, for a 20 megapixel RGB image and 200-pixel PSFs, the running time of our two-step algorithm in Matlab IDE is around 250 seconds on a PC with Intel Xeon i7@2.70Hz CPU. Further performance improvements could be achieved through GPU optimization and parallelization.

6 Implementation

We show our implementation of the proposed achromatic diffractive imaging method, including the prototype design, fabrication and a number of experimental results for synthetic data as well as real indoor and outdoor scenes at full sensor resolution. For figures some figures we present only cropped regions to highlight detail structure. The full resolution images can be found in the supplementary materials.

Prototype We designed two types of diffractive lenses, our diffractive achromat and a conventional diffractive lens with the same optical parameters for comparison. The focal length is designed at $f = 100\text{mm}$ with an aperture diameter 8mm for both cases. The conventional diffractive lens is designed at central wavelength 550nm . Our diffractive achromat is optimized for wavelengths from 410nm to 690nm with a 10nm sampling interval. Both lenses are attached on a Canon 70D camera body that has 5740×3648 pixels with the pixel pitch of $4.1\mu\text{m}$.

Fabrication We fabricate the designed diffractive achromat using multi-level photolithography techniques. In the lithography step,




Figure 4: Comparisons for different combinations of deblurring step and denoising step on synthetic dataset (top row) and real capture (bottom row), each with input blurry noised image, fast LUT deconvolution [Krishnan and Fergus 2009], direct deconvolution + MLP [Schuler et al. 2013], their multi-scale mix-implementations, and our approach with cross-scale prior. For the top row, the blurry image is synthesized using hyperspectral images with 29 wavelengths, blurred by the kernel in Fig. 10, and added with $\sigma = 0.005$ Gaussian white noise. The inset numbers indicate the PSNR and runtime on the synthetic image with 1.45 megapixels size in Matlab IDE on a commercial PC. Ours is proved to be robust and efficient.

an auxiliary Cr layer and a photoresist layer are first deposited and coated on the Fused Silica wafer. Patterns on the mask are transferred to the photoresist through exposure to the UV light. After development and Cr etching, a patterned area on the wafer becomes exposed to the ion beam in the following reactive ion etching (RIE) step. By controlling the etching duration, a certain depth on the wafer is obtained. A mixture of SF_6 and Ar are used as the ion beam in RIE.

The substrate in our implementation is a 0.5mm-thick 4 inch Fused Silica wafer. Each lens is fabricated by repeatedly applying the photolithography and RIE techniques. We choose 16-level micro-structures to approximate the continuous surface profile. Diffractive lenses with 2D micro-structures approximated by 16-level achieve a theoretical diffraction efficiency of up to 95%, and increasing the number of levels to 32 yields almost no improvement [Fischer et al. 2008]. The 16 levels can be achieved by repeating four iterations of the basic fabrication cycle with different amounts of etching. The total height for 2π phase modulation corresponds to 1195nm etching depth on the wafer. See supplementary document for more detail on the fabrication process.

Experimental results We show in Fig.4 the comparison results using different deblurring and denoising methods, including single scale fast LUT deconvolution [Krishnan and Fergus 2009], single scale MLP-based deconvolution [Schuler et al. 2013], two-scale fast LUT deconvolution, one scale MLP-based deconvolution followed by one scale fast LUT deconvolution, and our two-step cross-scale deconvolution. Our algorithm produces sharp, high color fidelity images in the full visible spectrum due to the achromatic design. We run our experiments for a hyperspectral image database with 50 images and the results yield an averaged PSNR 26.2dB with our proposed algorithm, even the worst result still stays above 25.0dB (See Tab. 1). Figure 5 shows some synthetic results of natural scenes from the hyperspectral datasets in [Chakrabarti and Zickler 2011] and [Skauil and Farrell 2013].

Figures 6 and 7 show the experimental results captured using our diffractive achromat. We have presented diverse natural scenes, including indoor, outdoor, rich color, high reflection, etc.. The results show that our method is close to spatial and depth invariant achromatic aberration within a long range. Refer to the caption of each figure for scene details.

Compared to synthetic results, the captured results suffer from




Figure 8: Blurred (left) and deblurred (right) results of capturing a standard resolution chart image which is projected by a projector onto a white plane. The capture distance is around 1.8m. The same PSF estimation used in Fig. 7 is applied here.

an additional haze effect, which degrades the image quality. We identify several sources for these deviations: First, the discrete height profile and limited spatial resolution of our photolithography process reduces the diffraction efficiency of the prototype. This could be alleviated by moving to an electron-beam lithography process. Second, the engineering errors derived from the custom optical mounts, e.g. custom holder and aperture result in some light leaks in the camera. Finally, our prototype is designed to be optimal for a spectral range from 410nm to 690nm, while the sensor may have a wider spectral response.

7 Evaluation and Discussion

We have demonstrated our diffractive achromat is able to image natural scenes with competitive resolution and color fidelity. In this section, we analyze the imaging performance from the perspectives of spatial resolution, off-axis behavior and color performance. The potential applications and limitations are discussed as well.

Resolution measurement We evaluate the resolution of our achromatic lens by taking an image of the ISO 12233 resolution chart, shown in Fig. 8. The captured image is very blurry due to the large blur kernel from our achromatic lens. High-frequency features, such as small edge patterns shown in the close-ups, are indistinguishable. After the deconvolution step, the recovered image preserves most of the low-frequency and mid-frequency features. Image contrast is also improved. The resolution comparison with a pure refractive lens is provided in the following.



Figure 5: Illustration of synthetic results with blurred (top row) and deblurred (bottom) results of selected scenes, with $\sigma = 0.005$ Gaussian white noise added, the left 3 pairs are synthesized with 29 wavelengths' hyperspectral images, while the right 2 pairs 71 wavelengths. The quantitative evaluations on full dataset are also provided, refer to Supplementary document for details.




Figure 6: Blurred (top) and deblurred (bottom) results of real captured scenes. All four scenes are captured at different depths with single exposure and no gain, under indoor artificial and mixed illumination (left two pairs), and natural sunlight illumination (right two pairs), using a single 0.5 mm ultrathin diffractive achromat we fabricate. Note that we roughly use a white light source attached with a pinhole to calibrate the PSF only at one depth (2 m), and we use it for all deconvolutions.

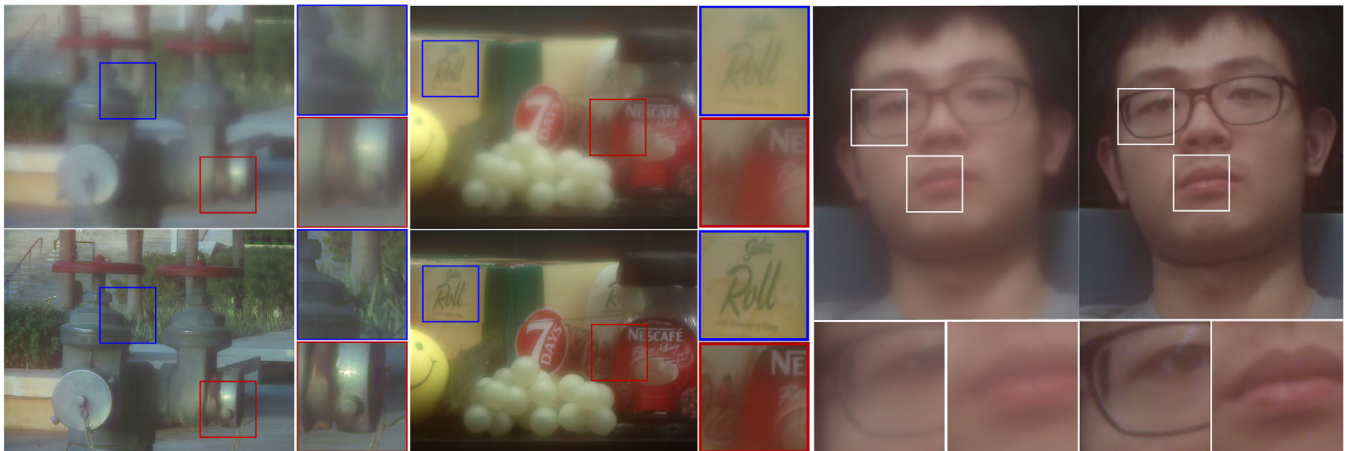


Figure 7: Blurred and deblurred results of outdoor scene with a large depth variance and reflection feature (left pair), desktop stuff with rich colors (center pair) and natural human face (right pair). The same PSF calibrated for Fig. 6 is used here.

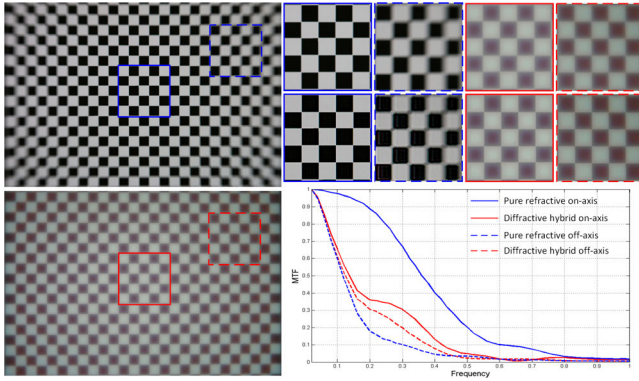


Figure 9: On-axis and off-axis behavior comparison of an achromatic refractive lens (top-left) and a hybrid diffractive-refractive lens (bottom-left). From the blurred and deblurred patch pairs presented in the right hand side, we observe that embedding our DOE design in a lens exhibits better spatial uniformity, despite the residual aberration. Here we assume within each selected patch the PSF is locally invariant. Accordingly, the MTFs estimated from gray-scale slant edges inside each deblurred patch are provided (bottom-right). The auxiliary refractive optics used are Thorlabs achromatic lenses, with focal length 50mm, 100mm, and thickness 8.7mm, 4.7mm, respectively. The equivalent focal lengths are all 50mm and full field of view is around 30° .

Off-axis behavior In addition to the benefit of ultrathin and lightweight structure, our diffractive achromat exhibits lower off-axis distortion than that of a refractive lens with the same focal length, as illustrated in Fig. 9. Compared to the simple refractive lens that has highly spatially variant PSFs across the field of view, our diffractive achromat exhibits smaller field curvature, thereby the resulting PSFs are almost uniform across the image plane. This property simplifies PSF calibration that otherwise is time-consuming or even impractical for refractive lenses. Specifically, after the deconvolution step, the off-axis image patch of a hybrid lens exhibits sharper edges than that of a pure refractive lens (see the zoom-in patches). The MTFs estimated from the slant edge method [Samei et al. 1998] are presented in the bottom-right of Fig. 9. We see that the diffractive achromat results in a good compromise between on-axis and off-axis performance. The computational burden for splitting image patches to account for the spatial variant blur kernels is also eliminated and is further beneficial to the deconvolution step.

Owing to the spatial invariant PSFs of our diffractive achromat, we can also introduce a hybrid refractive-diffractive lens design as in [Peng et al. 2015]. Our hybrid design concept differs from conventional hybrid designs in that we do not leverage the negative chromatic aberration of DOEs to compensate for the positive chromatic aberration of refractive lenses [Meyers 1998]. One can combine our diffractive achromat with any off-the-shelf refractive lens to assemble a hybrid lens that has improved spatial uniformity compared to purely refractive designs. See supplementary document for the extensional experimental results of an achromatic diffractive-refractive hybrid lens.

Color performance The advantage of our diffractive achromat is to generate spectrally invariant PSFs, which is usually neglected by conventional diffractive imaging methods. An ideal diffractive imaging device should preserve high spatial frequencies across all color channels, similar to a refractive optical system. Our proposed optimization method achieves this for all the three color channels. We show the Modulation Transfer Functions (MTFs) for each chan-

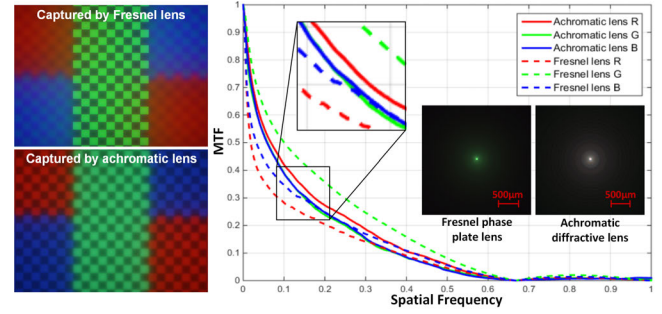


Figure 10: Color performance comparison of conventional diffractive lens and our achromatic diffractive lens. Conventional diffractive lens has a sharp green channel but severely blurred blue and red channels. Our achromatic lens balances three channels to show averaged performance for all. The MTF plots for our design (solid) are closer to each other, compared to those of the conventional diffractive lens (dashed). The insets show the respective color PSFs of the two lenses.

nel to illustrate the benefit of our design in Fig. 10. Compared with a conventional diffractive lens, our design shows balanced performance in three channels indicated by the MTF curves getting close to each other (solid plots), while the MTFs for the conventional diffractive lens are separated far away from each other (dashed plots). This can be more clearly seen from the captured color PSFs for both lenses. The conventional diffractive lens has a peak for the green channel and the other channels are of very low intensity, so the PSF looks green. For our diffractive achromat, since we optimize for 29 wavelengths to equalize the intensity distribution, the color PSF looks much more natural and closer to that of the refractive lens. The measured results of PSFs are presented in Fig. 11, from which we see that the quasi-uniform spectral PSF behavior has been established.

Limitations Our prototype suffers from a few shortcomings, which can be traced back to limitations of the current manufacturing process. Like other DOEs with discrete levels of surface relief, our achromatic prototype cannot achieve 100% diffraction efficiency for all the wavelengths. The spread of the lost energy results in slightly foggy appearance in the captured image, which is still difficult to eliminate especially in high contrast scenes.

Our prototype is also limited by the resolution of the photolithography process that we employ for fabrication. This results in a minimum feature size of $1\mu\text{m}$, and places a limit on the aperture sizes and focal lengths that are possible with this process. We note that the resolution could be improved by two orders of magnitude switching to an electron-beam lithography process, which, however, was beyond the scope of this work.

Finally, we also note that there is a small residual wavelength dependency left in the optimized design of the diffractive achromat. For imaging in natural environments with broad-band, incoherent light, we have demonstrated that this is not an issue. However, for partially coherent or fully coherent light, the PSF will contain interference patterns with high spatial frequencies that are sensitive to small shifts in wavelength. The deconvolution method will not be able to restore a high-quality image in this scenario.

Potential applications and future work We have shown in this paper that diffractive achromats are viable for full spectrum imaging in the visible domain. However, our method is not limited to applications in the visible band. Diffractive lenses are very prom-

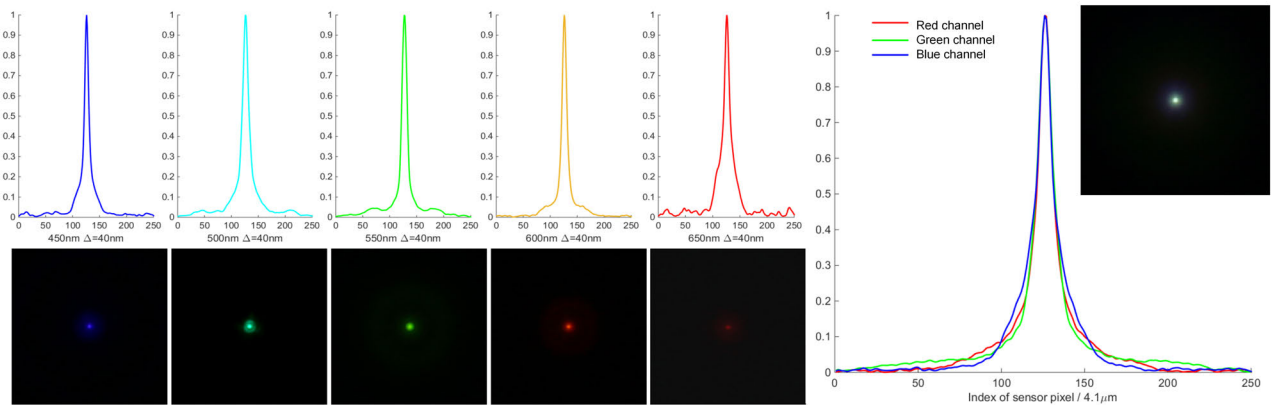


Figure 11: The measured PSFs on a sRGB color sensor for 5 selected spectrum bands (left), and the PSF we have calibrated for all above deconvolutions (right-most). Note that for individual measurement, a bandpass filter with a 40nm FWHM is attached in front of our lens.

using for imaging in the ultraviolet or far infrared spectrum, where refractive lenses are not able to transmit the desired wavelengths with high efficiency [Wang et al. 2003; Kang et al. 2010]. The ultrathin structure of diffractive optics drastically improves transmission for these wavelengths, while simultaneously reducing weight.

Moreover, it is possible to use our approach to design custom optics for specialized purposes that require imaging of multiple discrete wavelengths. By optimizing only for these wavelengths while neglecting others, better PSFs can be achieved than with a broad band optimization. Due to its flat field property, our diffractive achromat design is an alternative option to be combined with off-the-shelf lenses to correct off-axis aberrations flexibly such that the number of lenses in an imaging system can be significantly reduced without decreases in image quality [Aieta et al. 2012]. Additional work on how to apply the known kernel distributions to benefit the deconvolution process can also be interesting.

8 Conclusion

In this paper, we have proposed a novel achromatic diffractive imaging method that bridges diffractive optical elements and computational algorithms to build lightweight and thin optics for the full visible spectrum. By introducing optimization to the design of diffractive optics, we develop a diffractive achromat that trades off color fidelity and overall spatial resolution. The residual aberrations resulted from this compromise are tackled by employing a two-step image deconvolution with our proposed cross-scale prior. The algorithm includes a downsampled scale fast deconvolution and denoising step, and a full scale cross-scale deconvolution. Both steps are implemented without additional color priors to rapidly recover high quality color. We envision our method offers an opportunity for the joint design of ultrathin diffractive achromats and computational algorithms, further to boost the potential application of compact imaging devices in many broadband illumination scenarios.

Acknowledgments

This work was in part supported by King Abdullah University of Science and Technology (KAUST) baseline funding, the KAUST Advanced Nanofabrication Imaging and Characterization Core Lab, as well as a Doctoral 4 Year Fellowship from The University of British Columbia (UBC). The authors thank Robin Swanson and Jinhui Xiong for volunteering in the video audio recording and the capture experiment.

For the author contribution, Y.P. and Q.F. conceived the idea and the proofs. Y.P. proposed the optical design method, implemented the algorithms and the simulations. Q.F. fabricated the lenses. Y.P. conducted the experiments and the reconstructions. F.H. helped with the reconstructions. W.H. coordinated and instructed the whole project. All authors took part in writing the paper.

References

- AIETA, F., GENEVET, P., KATS, M. A., YU, N., BLANCHARD, R., GABURRO, Z., AND CAPASSO, F. 2012. Aberration-free ultrathin flat lenses and axicons at telecom wavelengths based on plasmonic metasurfaces. *Nano letters* 12, 9, 4932–4936.
- AIETA, F., KATS, M. A., GENEVET, P., AND CAPASSO, F. 2015. Multiwavelength achromatic metasurfaces by dispersive phase compensation. *Science* 347, 6228, 1342–1345.
- BOYD, S., PARIKH, N., CHU, E., PELEATO, B., AND ECKSTEIN, J. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* 3, 1, 1–122.
- CHAKRABARTI, A., AND ZICKLER, T. 2011. Statistics of real-world hyperspectral images. In *Proc. CVPR*, IEEE, 193–200.
- CHAN, S. H., KHOSHABEH, R., GIBSON, K. B., GILL, P. E., AND NGUYEN, T. Q. 2011. An augmented lagrangian method for total variation video restoration. *IEEE Trans. Image Process.* 20, 11, 3097–3111.
- EBERHART, R. C., AND SHI, Y. 2001. Particle swarm optimization: developments, applications and resources. In *Proc. of 2001 Congress on Evolutionary Computation*, IEEE, 81–86.
- FISCHER, R. E., TADIC-GALEB, B., YODER, P. R., AND GALEB, R. 2008. *Optical system design*. McGraw Hill.
- GILL, P. R., AND STORK, D. G. 2013. Lensless ultra-miniature imagers using odd-symmetry spiral phase gratings. In *Computational Optical Sensing and Imaging*, OSA, CW4C–3.
- GOODMAN, J. 2008. *Introduction to Fourier optics*. McGraw-hill.
- HEIDE, F., ROUF, M., HULLIN, M. B., LABITZKE, B., HEIDRICH, W., AND KOLB, A. 2013. High-quality computational imaging through simple lenses. *ACM Trans. Graph.* 32, 5, 149.

- JIANG, W., WANG, J., AND DONG, X. 2013. A novel hybrid algorithm for the design of the phase diffractive optical elements for beam shaping. *Opt. Laser Tech.* 45, 37–41.
- JOSHI, N., SZELISKI, R., AND KRIEGMAN, D. 2008. Psf estimation using sharp edge prediction. In *Proc. CVPR*, IEEE, 1–8.
- KANG, G., TAN, Q., WANG, X., AND JIN, G. 2010. Achromatic phase retarder applied to mwir & lwir dual-band. *Opt. Express* 18, 2, 1695–1703.
- KIM, G., DOMÍNGUEZ-CABALLERO, J. A., AND MENON, R. 2012. Design and analysis of multi-wavelength diffractive optics. *Opt. Express* 20, 3, 2814–2823.
- KRISHNAN, D., AND FERGUS, R. 2009. Fast image deconvolution using hyper-laplacian priors. In *Advances in Neural Information Processing Systems*, NIPS, 1033–1041.
- KRISHNAN, D., TAY, T., AND FERGUS, R. 2011. Blind deconvolution using a normalized sparsity measure. In *Proc. CVPR*, IEEE, 233–240.
- MEYERS, M. M., 1998. Hybrid refractive/diffractive achromatic camera lens, Feb. 3. US Patent 5,715,091.
- MONJUR, M., SPINOULAS, L., GILL, P. R., AND STORK, D. G. 2015. Ultra-miniature, computationally efficient diffractive visual-bar-position sensor. In *Proc. SensorComm*, IEIFSA.
- MUZYCHENKO, Y. B., ZINCHIK, A., STAFEEV, S., AND TOMILIN, M. 2011. Fractal diffraction elements with variable transmittance and phase shift. In *International Commission for Optics (ICO 22)*, International Society for Optics and Photonics, 80112F–80112F.
- NAKAI, T., AND OGAWA, H. 2002. Research on multi-layer diffractive optical elements and their application to camera lenses. In *Diffractive Optics and Micro-Optics*, OSA, DMA2.
- NIKONOROV, A., SKIDANOV, R., FURSOV, V., PETROV, M., BIBIKOV, S., AND YUZIFOVICH, Y. 2015. Fresnel lens imaging with post-capture image processing. In *Proc. CVPR Workshops*, IEEE, 33–41.
- PENG, Y., FU, Q., AMATA, H., SU, S., HEIDE, F., AND HEIDRICH, W. 2015. Computational imaging using lightweight diffractive-refractive optics. *Opt. Express* 23, 24, 31393–31407.
- QUIRIN, S., AND PIESTUN, R. 2013. Depth estimation and image recovery using broadband, incoherent illumination with engineered point spread functions [invited]. *Applied optics* 52, 1, A367–A376.
- SAMEI, E., FLYNN, M. J., AND REIMANN, D. A. 1998. A method for measuring the presampled mtf of digital radiographic systems using an edge test device. *Medical physics* 25, 1, 102–113.
- SCHULER, C. J., HIRSCH, M., HARMELING, S., AND SCHOLKOPF, B. 2011. Non-stationary correction of optical aberrations. In *Proc. ICCV*, IEEE, 659–666.
- SCHULER, C. J., BURGER, H. C., HARMELING, S., AND SCHOLKOPF, B. 2013. A machine learning approach for non-blind image deconvolution. In *Proc. CVPR*, IEEE, 1067–1074.
- SCHWARTZBURG, Y., TESTUZ, R., TAGLIASACCHI, A., AND PAULY, M. 2014. High-contrast computational caustic design. *ACM Trans. Graph.* 33, 4, 74.
- SHAN, Q., JIA, J., AND AGARWALA, A. 2008. High-quality motion deblurring from a single image. *ACM Trans. Graph.* 27, 3, 73.
- SINGH, M., TERVO, J., AND TURUNEN, J. 2014. Broadband beam shaping with harmonic diffractive optics. *Opt. Express* 22, 19, 22680–22688.
- SKAULI, T., AND FARRELL, J. 2013. A collection of hyperspectral images for imaging systems research. In *Proc. IS&T/SPIE Electronic Imaging*, SPIE, 86600C–86600C.
- SMITH, W. J. 2005. *Modern lens design*. McGraw-Hill.
- STORK, D. G., AND GILL, P. R. 2014. Optical, mathematical, and computational foundations of lensless ultra-miniature diffractive imagers and sensors. *International Journal on Advances in Systems and Measurements* 7, 3, 4.
- SU, Z., ZENG, K., LIU, L., LI, B., AND LUO, X. 2014. Corruptive artifacts suppression for example-based color transfer. *IEEE Trans. Multimedia* 16, 4, 988–999.
- WANG, Y., YUN, W., AND JACOBSEN, C. 2003. Achromatic fresnel optics for wideband extreme-ultraviolet and x-ray imaging. *Nature* 424, 6944, 50–53.
- WANG, P., MOHAMMAD, N., AND MENON, R. 2016. Chromatic-aberration-corrected diffractive lenses for ultra-broadband focusing. *Scientific reports* 6.
- YAMAGUCHI, M., HANEISHI, H., AND OHYAMA, N. 2008. Beyond red–green–blue (rgb): Spectrum-based color imaging technology. *JIST* 52, 1, 10201–1.
- YE, G., JOLLY, S., BOVE JR, V. M., DAI, Q., RASKAR, R., AND WETZSTEIN, G. 2014. Toward bxd display using multilayer diffraction. *ACM Trans. Graph.* 33, 6, 191.
- YUAN, L., SUN, J., QUAN, L., AND SHUM, H.-Y. 2008. Progressive inter-scale and intra-scale non-blind image deconvolution. *ACM Trans. Graph.* 27, 3, 74.
- YUE, T., SUO, J., WANG, J., CAO, X., AND DAI, Q. 2015. Blind optical aberration correction by exploring geometric and visual priors. In *Proc. CVPR*, IEEE, 1684–1692.
- ZHOU, G., CHEN, Y., WANG, Z., AND SONG, H. 1999. Genetic local search algorithm for optimization design of diffractive optical elements. *Appl. Opt.* 38, 20, 4281–4290.