

Non-stationary Correction of Optical Aberrations

Christian J. Schuler, Michael Hirsch, Stefan Harmeling and Bernhard Schölkopf
Max Planck Institute for Intelligent Systems, Tübingen, Germany

{cschuler, mhirsch, harmeling, bs}@tuebingen.mpg.de

<http://webdav.is.mpg.de/pixel/lenscorrection>



Figure 1. Self-made photographic lens with one glass element only. Taken image without and with lens correction.

Abstract

Taking a sharp photo at several megapixel resolution traditionally relies on high grade lenses. In this paper, we present an approach to alleviate image degradations caused by imperfect optics. We rely on a calibration step to encode the optical aberrations in a space-variant point spread function and obtain a corrected image by non-stationary deconvolution. By including the Bayer array in our image formation model, we can perform demosaicing as part of the deconvolution.

1. Introduction

In an ideal optical system as described theoretically by paraxial optics, all light rays emitted by a point source converge to a single point in the focal plane, forming a clear and sharp image. Departures of an optical system from this behaviour are called aberrations, causing unwanted blurring of the image.

Manufacturers of photographic lenses attempt to minimize optical aberrations by combining several lenses. The design and complexity of a compound lens depends on various factors, *e.g.*, aperture size, focal length, and constraints on distortions. Optical aberrations are inevitable and the design of a lens is always a trade-off between various parameters, including price. To correct these errors in software is still an unresolved problem.

Rather than proposing new designs for complicated com-

pound lenses, we show that almost all optical aberrations can be corrected by digital image processing. For this, we note that optical aberrations of a linear optical system are fully described by their *point spread function* (PSF). We will show how PSFs encountered in real photographic lenses suffering from various optical aberrations can be approximated as non-stationary convolutions. For a given lens/camera combination, the parameters of the non-stationary convolution are estimated via an automated calibration procedure that measures the PSF at a grid covering the image. We also include demosaicing into our image reconstruction, because it fits naturally into our forward model. Our results surpass current state of the art.

Main contribution: We show how to reconstruct a full-color image, *i.e.*, all three color channels at full resolution, given a raw image that is corrupted by various monochromatic and chromatic aberrations, and Bayer filtered by a color filter array (CFA) of our off-the-shelf camera. This image reconstruction is even possible for heavily degraded images, taken with a self-constructed lens consisting of a *single lens element* attached to a standard camera, Fig. 1.

2. Related work

We are not aware of any work that tries to solve the demosaicing and the correction of lens errors simultaneously. There exist many different methods solely for demosaicing, for reviews see [16, 7, 1, 13]. However, none of them model

and exploit the aberration of the lens to facilitate demosaicing as our method does.

Chromatic aberrations arise because the refractive index of glass, and thus focal length and image scale, is dependent on the wave length. A common approach to correct for lateral chromatic aberrations is a non-rigid registration of the different color channels [2, 10, 14]. Such methods correspond to restricting our model to delta-peaked PSFs, and generally ignore other optical aberrations. The method of [4] measures chromatic aberration at edges through color differences and compensates locally, however without using a PSF model of the lens. The approach in [9] also relies on the estimation of sharp step edges and can be used in a non-blind fashion. Even though full PSF are estimated, they are only used to remove chromatic aberrations, where a rough knowledge of the PSF is sufficient. None of these approaches consider demosaicing.

A method that focuses on correcting coma has been proposed in [6], showing how to reduce coma by locally applying blind deconvolution methods to image patches. This method is designed for gray scale images and thus does neither consider chromatic aberration nor demosaicing.

Algorithmically related to our work is [5], considering sparsity regularization in the luminance channel, and Tikhonov regularization in the two chromaticity channels. However, [5] combines the image information from several images, while our method works with a single image. Also, [5] combines demosaicing with super-resolution, while we combine it with correction for chromatic aberrations.

The image reconstruction problem we are addressing can also be dealt with using the proprietary software “DxO Optics Pro 6” (DXO), which tries to correct for image aberrations. DXO is considered state of the art among professional photographers and presumably uses the same kind of information as our approach (it contains a custom database of lens/camera combinations). It has been developed over a number of years and is highly optimized. DXO states that it can correct for “lens softness”, which their website¹ defines as image blur that varies across the image and between color channels in strength and direction. It is not known to us whether DXO models the blur as space-variant defocus blur of different shapes or with more flexible PSFs as we do; neither do we know whether DXO demosaics and deblurs simultaneously as we do. In the experimental section we show that our results compare favorably against results obtained by DXO.

Using deconvolution to correct for lens aberrations is also discussed in [11]. This work focuses on removing lens blur across multiple aperture and zoom settings of a given lens. A calibration method similar to [9] is used. While this method can correct aberrations across many different

settings of a lens, the blur shape is modelled as a Gaussian. As can be seen in Fig. 6 this is not appropriate for strong aberrations considered in this paper. Also, the problem of demosaicing is not treated.

There exist several papers which suggest calibration procedures to measure the lens, *e.g.* [17, 19, 9]. However, they mainly focus on correcting geometric distortion or do not address monochromatic aberrations.

3. Aberrations as a non-stationary convolution

While the aberrations of an imaging system can be described as a simple matrix operator, the required matrix-vector multiplication would be computationally expensive. More efficient for describing blurs are convolutions, however, the usual *stationary* convolution applies the same blur kernel across the whole image and can thus only describe *space-invariant* PSFs, which are insufficient to model lens aberrations. As can be seen in Fig. 4 on the left, the PSF can vary in size, shape, orientation, position and intensity. How can we approximate such a *space-variant* PSF in an imaging model that allows efficient computation?

Hirsch *et al.* [8] presented the so-called *Efficient Filter Flow* (EFF) framework, which can model a PSF that smoothly varies across the image. The basic idea is to cover the image with overlapping patches each of which is assigned a blur kernel.

For notational simplicity, all images and blur kernels are column vectors. The generalization to two-dimensional matrices is straight-forward. Let x be some image, *i.e.*, a column vector of length n , and $f^{(r)}$ a blur kernel or filter, *i.e.*, a column vector of length k . The i th pixel value y_i in the blurred image y can be written as a linear combination of the p differently blurred patches,

$$y_i = \sum_{r=0}^{p-1} \sum_{j=0}^{k-1} f_j^{(r)} w_{i-j}^{(r)} x_{i-j} \text{ for } 0 \leq i < n \quad (1)$$

where $w^{(r)} \geq 0$ is a fixed weighting vector which is non-zero only on the r th patch. Since the patches are usually chosen to overlap, these weights smoothly interpolate between neighboring filters $f^{(r)}$. Note that the weighting vectors have to sum up to one, *i.e.*

$$\sum_{r=0}^{p-1} w_i^{(r)} = 1 \text{ for } 0 \leq i < n. \quad (2)$$

Let f be the column vector that we obtain by stacking all local filters $f^{(r)}$. Since the space-variant blur in Eq. (1) is linear in x and in f there exist matrices X and B such that

$$y = Xf = Bx. \quad (3)$$

Below, we call X and B the EFF matrices.

¹http://www.dxo.com/us/photo/dxo_optics_pro/optics_geometry_corrections/lens_softness

Note that if there is overlap between the patches, this can effectively imply a different PSF at each pixel. Basically, the method applies the overlap-add trick for convolution, so it can be computed as efficiently as the convolution, while being much more flexible. Next, we explain the lens aberrations we would like to correct.

Monochromatic aberrations. This class of aberrations include *spherical aberration* (in spherical lenses, the focal length is a function of the distance from the axis) as well as a number of off-axis aberrations: *coma* occurs in an oblique light bundle when the intersection of the rays is shifted w.r.t. its axis; *field curvature* occurs when the focal surface is non-planar; *astigmatism* denotes the case when the sagittal and tangential focal surfaces do not coincide (*i.e.*, the system is not rotationally symmetric for off axis light bundles); *distortion*, which is the only aberration we do not address, is related to a spatially varying image scale. All these monochromatic aberrations lead to blur that varies across the image. Any such blur can be expressed in the EFF framework by appropriately choosing the local blur filters $f^{(0)}, \dots, f^{(p-1)}$.

Chromatic aberration. The refraction index of most materials including glass is dependent on the wavelength of the transmitted light. Axially, this results in the focus of a lens being a function of the wavelength (*longitudinal chromatic aberration*); off-axis, we observe *lateral chromatic aberration* caused by the fact that the different focal lengths for different wavelengths directly imply that the image scale slightly varies with wavelength. By modeling the three color channels with separate space-variant PSFs, we are able to describe such chromatic aberration. This means on the color channels x_R , x_G , and x_B each acts a blur B_R , B_G and B_B , which we can also write as a blur B acting on the full color image x .

Vignetting. Because oblique light bundles do not reach the focal plane in their entirety, the intensity of the image falls off towards the image corners. This can be corrected by photographing a *flat field* frame, *i.e.*, an image of a homogeneous background, and dividing the image by it. While this is straightforward, the EFF framework can also include vignetting into our model by omitting the energy conservation constraint, in that case the filters $f^{(r)}$ in Eq. (1) do not have to sum up to one, *i.e.*, we only require $\sum_j f_j^{(r)} \leq 1$ and $f_j^{(r)} \geq 0$ for all j and r . By allowing dimmer filters we automatically correct for vignetting using our procedure. Note that Eq. (2) is unaffected by relaxing the energy conservation constraint.

4. Forward model including mosaicing

The image blurred by the blur B is the image that will enter the CFA, just before being mosaiced. The operation of the CFA can be described as a linear map represented by some matrix D , whose result will be the image that hits the photo-sensitive sensor *behind* the CFA. Note that D is a rectangular matrix with three times as many columns as rows.

The forward model combines the lens aberration and Bayer filtering into a single matrix A and adds noise n , *i.e.*,

$$y = DBx + n = Ax + n. \quad (4)$$

5. Recovering the corrected, full-color image

Assuming the weights in the Bayer matrix D to be fixed and known (we use a trivial Bayer matrix disregarding cross-talk between color channels) the linear transformation A , *i.e.*, the PSF, is parameterized by the set of filters that determine the EFF matrices B_R , B_G , and B_B for the three color channels. These filters depend on the lens and the camera used. In Sec. 6 we will detail the experimental setup and procedure how we measure these filters at regularly placed sites.

Assuming the noise in Eq. (4) to be Gaussian, we could recover the unknown full-color image x from a measured raw image y by solving a least-squares problem, *i.e.*, by minimizing $\|y - Ax\|_2^2$ w.r.t. x . However, the PSF parameterized by the EFF framework is only an approximation to the true PSF and is subject to errors. Using stochastic robust matrix approximation [3] and the assumption that each of the n elements of the PSF exhibits a standard deviation of σ with zero mean, we add a regularization term. Just for the EFF matrices this would result in $n\sigma^2\|x\|_2^2$, including the Bayer matrix the regularization can be approximated as $n\sigma^2(\|x_R\|_2^2/4 + \|x_G\|_2^2/2 + \|x_B\|_2^2/4)$.

One challenge of processing real photos is that pixels might be saturated, their true values may be clipped due to limited dynamic range. Thus the measured values of clipped pixels are not in agreement with the physical model of the blur. We exclude saturated pixels in the data-fidelity term $\|y - Ax\|_2^2$ by summing only over non-saturated pixels.

This term corresponds to the likelihood term (or data fit) of the implicitly underlying probabilistic model. However, because we are trying to estimate three color channels from a single raw image, which means there are three times as many unknowns as observations, our deblurring problem is ill-posed. To regularize it we include prior knowledge about natural images: it has been shown that the image gradients approximately follow a hyper-Laplacian distribution [12, 18]. This can be incorporated into the optimization problem by adding a regularization term of the form $\|\nabla x\|_\gamma^\gamma$ to the objective function. The effect of this regularization is

to penalize strong gradients and therefore to smooth the image. We follow Farsiu *et al.* [5] who transformed the RGB image to a luminance/chrominance color space (here we use YUV) before applying the regularization. This allows us to regularize more strongly in the chrominance channels, and less in luminance. Note that the human eye is more sensitive to differences in luminance than in chrominance, *i.e.*, a visually pleasing result has to be sharp in the luminance channel. The transformation from RGB to YUV is simply a matrix vector multiplication $[x_Y^T, x_U^T, x_V^T]^T = C[x_R^T, x_G^T, x_B^T]^T$ with appropriately chosen matrix C . With x_Y , x_U , and x_V we can write our combined objective function as

$$\|y - Ax\|_2^2 + \alpha \|\nabla x_Y\|_\gamma^\gamma + \beta \|\nabla x_U\|_\gamma^\gamma + \beta \|\nabla x_V\|_\gamma^\gamma \quad (5) \\ + n\sigma^2 (\|x_R\|_2^2/4 + \|x_G\|_2^2/2 + \|x_B\|_2^2/4).$$

We obtained good results by setting $\alpha = 10^{-4}$, $\beta = 10^{-3}$, $\gamma = 0.65$ and $\sigma = 10^{-3}$ in our simulated experiments. On real images, the optimal values for α and β were smaller by a factor of ten.

We minimize the objective function w.r.t. x adapting Krishnan and Fergus' [12] approach to our setup, alternating between a convex and a non-convex phase, with the non-convex phase being accelerated by a lookup table.

6. Estimating the non-stationary convolution

Leaving aside diffraction effects (*e.g.*, by ensuring the pixel size to be larger than the Airy disk), a point light source should influence just a single pixel on the imaging sensor of a digital camera. However, this would only happen if a digital camera was a perfect optical system. In practice, the various lens aberrations discussed above will spread out the point light source over a larger region of the imaging sensor. This local pattern characterizes the PSF, so by recording these patterns across the image plane we can set the filters of the non-stationary convolution described above.

To automate the measurements, we mounted a camera on a motor-driven platform with two rotational degrees of freedom. A lens measurement process is conducted in a completely dark room by remotely changing the angles of the camera towards a point light source (a gas lamp emitting light through an aperture of $100\mu\text{m}$ in 12 meters distance) such that in subsequent exposures the light point is captured at equidistant locations on the sensor.

In our experiments we use a 18 times 27 grid of supporting points for the EFF framework. The blur kernels were recorded by averaging three dark frame subtracted images of the point light source and thresholding noise. This simple setup gives sufficiently good measurements for the PSF, as can be seen in the deconvolution results in Sec. 8.

Image	(a) Deconv. then demosaic.	(b) Demosaic. then deconv.	(c) Joint approach
1	23.09	25.92	26.35
2	30.11	31.92	32.23
3	30.67	33.47	33.68
4	29.12	32.23	32.49
5	22.58	26.08	26.62
6	24.84	27.09	27.47
7	27.87	33.07	33.47
8	20.32	23.77	24.28
9	28.02	32.11	32.51
10	28.54	31.53	31.96
11	25.92	28.77	29.11
12	29.51	32.67	33.04
13	21.32	23.32	23.81
14	25.34	28.32	28.79
15	28.90	32.14	32.52
16	28.41	30.40	30.68
17	28.22	31.33	31.68
18	25.06	27.75	28.20
19	24.77	27.87	28.46
20	27.66	31.40	31.78
21	25.27	28.17	28.63
22	26.86	29.61	29.95
23	30.00	34.08	34.59
24	23.74	26.06	26.34
Average	26.51	29.54	29.94

Table 1. Comparison of peak signal-to-noise ratios (PSNR in dB) for Kodak image data set. Consistently, the joint approach outperforms the sequential demosaicing and deconvolution procedures (higher number means better reconstruction).



Figure 3. Point spread function used for our simulations on the Kodak image data set.

7. Results on simulated images

To test our method under controlled conditions we artificially blurred test images usually used for evaluating demosaicing algorithms from the Kodak PhotoCD. To simulate the lens aberrations, we created a 4×6 filter array containing measured blur kernels of a Canon 50mm f/1.4 lens at maximum aperture on a Canon 5D Mk II. This filter array are the parameters of a non-stationary convolution that represent our estimated model of the artificial lens aberrations. To account for the fact that the true PSF is not exactly known, we modify these filters with a low pass filter before convolving the ground truth images. In the image recon-

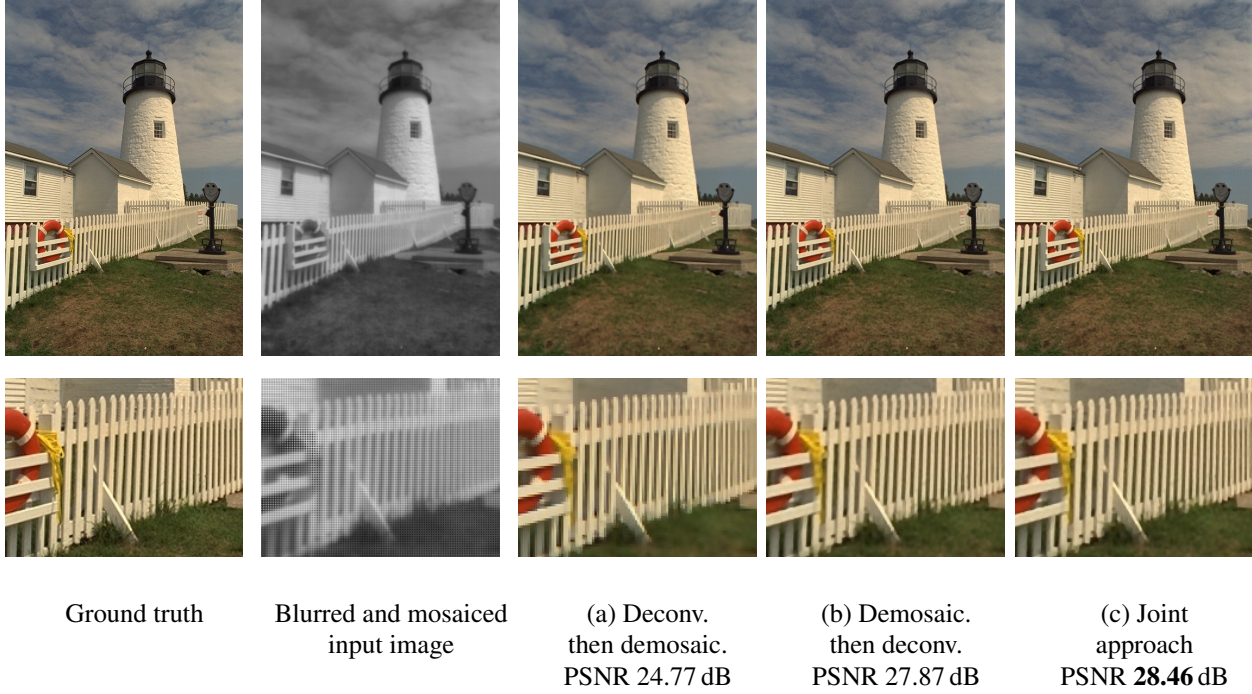


Figure 2. Comparison of our joint approach vs. sequential demosaicing and deconvolution procedures. The PSF used for the simulations are shown in Fig. 3. Gaussian noise with a SNR of 50 dB has been added.

struction process, the non-modified blur filters were used. We then added white noise with signal to noise ratio 50 dB and mosaiced the result with a Bayer filter array.

With the simulated experiments we want to investigate whether (a) we should apply the aberration correction separately on each color channel and subsequently demosaic with a state-of-the-art demosaicing algorithm [15], whether (b) our aberration correction should be better applied to images that have been already demosaiced by a standard demosaicing procedure, or whether (c) it is best to apply the forward model that includes the mosaicing (as described in Sec. 4), *i.e.*, to jointly correct the aberrations and the demosaicing.

Tab. 1 compares the peak-signal-to-noise ratios (PSNR) of the reconstructed images for the approaches (a), (b), and (c) on the image data set. For all 24 images the joint approach (c) leads to the best results, approach (b) being a close runner-up. This finding is also visually confirmed in Fig. 2 where approach (c) leads to the best reconstruction. Note that to suppress influence of the border region, a 15 pixel border on all edges has been excluded in the calculation of the PSNR.

We believe that our approach is able to compete with state-of-the-art demosaicing algorithm because separating demosaicing and deblurring has the disadvantage that it does not require the result to be consistent with the image formation model. Because of the blur, we gain knowledge




Figure 6. Interpolation of a mosaiced PSF at the example of a green PSF from the Canon 50mm f/1.4 lens.

about possible values for missing color information. For example, if we measure no light at a certain pixel, we can infer that in the deblurred image the surrounding region given by the size of the PSF also has to be dark. Furthermore, typical demosaicing algorithms do not take chromatic aberration into account, which lead to a spatial separation of edge information across different color channels.

8. Results on real images

Using the automated procedure from Sec. 6, we approximate the PSFs of three different lenses: (i) Canon 50mm f/1.4, (ii) Canon 24mm f/1.4 L, and (iii) a self-built lens consisting of a single glass element, see Fig. 1. For the Canon lenses, we took several pictures with a Canon 5D Mk II digital camera, for the self-built lens we used a Canon 5D Mk I. We applied our image reconstruction procedure described in Sec. 5 to these images and next describe the results.

In our PSF measurement we only obtain mosaiced versions. However, as can be seen in Fig. 6, the blur is sufficiently well behaved such that bilinear interpolation gives a good approximation to the true PSF.

Canon 50mm f/1.4. First, we use a Canon 50mm f/1.4 prime lens on a Canon 5D Mark II at maximum aperture. The comparison between original photo and the image corrected for lens errors is in Fig. 4. In Fig. 5, it is compared with the result of DXO (see Sec. 2), a software that is also able to correct for lens aberrations. Similar to our approach, it relies on previously recorded information about the error of a certain lens/camera-combination. In the comparison, all image improvements except the correction for “lens unsharpness”, chromatic aberration and vignetting were deactivated. While in the DXO result the edges are sharpened, the objects have a halo, e.g., around the wooden bars, which is not present in the original scene. This means the blur introduced by the lens is not completely removed.

Canon 24mm f/1.4. Furthermore, we correct the errors of a Canon EF 24mm f/1.4 at maximum aperture, which exhibits considerably visible errors in the border regions of the image at fully open aperture. The original and the corrected image can be seen in Fig. 4. In the recorded image strong chromatic aberration is visible as green and red lines near edges, which are reduced in the deconvolved result. This lens is not available in the DXO database for the Canon 5D Mk II, so DXO cannot be applied.

Self-built lens with a single lens element. The two lenses used above are high-end lenses with a complicated system of compound lenses that are built to minimize optical errors. Trying to make our algorithm fail, we constructed a simple photographic lens from a single convex-concave lens with focal length 120mm. Amazingly, the image can be well reconstructed as can be seen in Fig. 1 and 4. In Fig. 4, nearly no detail is recognizable in the grain of the wood in the original image. Also, the pegs on the right and upper edge of the image are hardly visible. The corrected image does not suffer from these problems.

Running time. For the 21.1 megapixel photos taken with the Canon lenses, the full-color non-convex optimization problem has more than 60M unknowns. It needs about 5 hours running time on a quad-core computer. For the self-built lens, we used a camera which produces 12.8 megapixel images and a blur size of 200x200. In the EFF framework with 27x18 supporting points, the processing takes about 7 hours using a MATLAB implementation of the algorithm.

This running time is impractical. However, we show how the EFF framework can be used to do *Direct Deconvolution*



Figure 7. Comparison of deconvolution with optimization (left) and direct method (right).

in Fourier space with a slightly modified version of our objective function. Since the demosaicing operator is not diagonal in Fourier space, we work on each already demosaiced color channel separately and solve the problem

$$\|y - Bx\|_2^2 + \alpha \|\nabla x\|_\gamma^\gamma + n\sigma^2 \|x\|_2^2. \quad (6)$$

This can be done with the approach of [12], however, the inversion of B is necessary. Using the expression from [8], the application of B in the EFF framework can be written as

$$y = \underbrace{\sum_r L_r^T F^H \text{Diag}(FPf^{(r)}) FK_r \text{Diag}(w^{(r)})}_B x. \quad (7)$$

In this summation over all patches, the matrix P zero-padds each patch, the matrices K_r and L_r are cropping matrices. F applies the discrete Fourier transform. This expression can be approximately inverted as

$$x \approx N \sum_r \text{Diag}(w^{(r)})^{1/2} \cdot K_r^T F^H \frac{\overline{FPf^{(r)}} \odot (FL_r \text{Diag}(w^{(r)})^{1/2} y)}{|FPf^{(r)}|^2 + |FR|^2}, \quad (8)$$

where $|z|$ and \bar{z} denote entry-wise absolute value and complex conjugate, respectively. The matrix R regularizes the result, e.g., a discrete Laplace operator. The weighting N is obtained by applying the inversion to a constant image and is necessary to remove artifacts stemming from inverting the windows. In Fig. 7 the results obtained by optimizing the more sophisticated objective function (6) are compared to the direct method. While losing a small amount of image quality, the running time is only 2 minutes for a 21.1 megapixel image.

9. Conclusion

We have proposed a method to correct the aberrations in optical imaging systems. A spatial-variant PSF is obtained in a calibration step, encoding the errors of the imaging system. These are then removed by non-stationary deconvolution. Furthermore, by requiring the corrected image to be

consistent with the image formation model, we are able to recover missing image information. We have shown this using the example of reconstructing color data lost in a mosaicing process.

Using controlled experiments on images artificially convolved with a non-stationary PSF, we have seen that our linear image formation model leads to better results than separately deblurring and demosaicing Bayer-filtered photos. More importantly, we were able to show that in a real imaging setup, we can correct the optical aberrations rather well both for commercial camera lenses and optically poor single element lenses. The results compare favorably to DXO, a commercially available software package considered state of the art in lens error correction among professional photographers.

9.1. Limitations

For the image taken with a one-element lens, we have seen that although a drastic improvement can be achieved, a perfect reconstruction was not possible. Moreover, our measurement procedure suffers from the fact that the PSF obtained are already subject to mosaicing, therefore the PSF used in the joint demosaicing/deblurring are only an approximation. A better PSF could, *e.g.*, be obtained with a monochromatic camera and color filters. The general quality of the PSF could for example be improved with wavefront measurement.

Also, the lens aberrations depend to a certain extent on the settings of the lens (aperture, focus, zoom), which can not be trivially modeled. In the case of lens blurs that can be approximated as a Gaussian with spatially varying parameters, it has been demonstrated how the change of the lens aberrations can be modeled [11]. In the case of non-Gaussian blurs, as treated in this publication, this problem has still to be solved. This would make our method feasible for lenses with a large number of possible settings, *e.g.*, zoom lenses.

9.2. Future Work

A further common error of imaging systems, distortions, can in principle also be encoded in a spatially varying PSF. However, in the case of strong distortions this would require PSFs as large as 500x500 pixels, say, and a large computational load. It would, however, be an elegant method for correcting all optical aberrations in one framework.

We believe that our work can have significant implications for the design of lenses, which today are probably the most expensive components of high-end camera systems.

References

- [1] D. Alleysson and B. de Lavarène. Frequency selection demosaicking: A review and a look ahead. In *Proc. SPIE Conf. Vis. Commun. and Image Proc.*, 2008. 1
- [2] T. Boulton and G. Wolberg. Correcting chromatic aberrations using image warping. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, pages 684–687, 1992. 2
- [3] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge Univ. Press, 2004. 3
- [4] S. Chung, B. Kim, and W. Song. Detecting and eliminating chromatic aberration in digital images. In *Proc. IEEE Int. Conf. Image Process.*, pages 3905–3908, 2009. 2
- [5] S. Farsiu, M. Elad, and P. Milanfar. Multiframe demosaicing and super-resolution of color images. *IEEE Trans. Image Process.*, 15(1):141–159, 2006. 2, 4
- [6] S. Gifford. Astronomical Coma Image Restoration Through The Use of Localized Deconvolution. In *Symp. Telescope Sci.*, page 141, 2008. 2
- [7] B. Gunturk, J. Glotzbach, Y. Altunbasak, R. Schafer, and R. Mersereau. Demosaicking: color filter array interpolation. *IEEE Signal Processing Mag.*, 22(1):44–54, 2005. 1
- [8] M. Hirsch, S. Sra, B. Schölkopf, and S. Harmeling. Efficient filter flow for space-variant multiframe blind deconvolution. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, pages 607–614, 2010. 2, 6
- [9] N. Joshi, R. Szeliski, and D. Kriegman. PSF estimation using sharp edge prediction. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, pages 1–8, June 2008. 2
- [10] V. Kaufmann and R. Ladstädter. Elimination of color fringes in digital photographs caused by lateral chromatic aberration. In *Proc. XXth International Int. CIPA Symp.*, volume 26, pages 403–408, 2005. 2
- [11] E. Kee, S. Paris, S. Chen, and J. Wang. Modeling and removing spatially-varying optical blur. In *Proc. IEEE Int. Conf. Computational Photography*, Pittsburgh, PA, 2011. 2, 7
- [12] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-Laplacian priors. In *Advances in Neural Inform. Process. Syst.*, 2009. 3, 4, 6
- [13] X. Lia, B. Gunturk, and L. Zhang. Image demosaicing: A systematic survey. In *Proc. SPIE Conf. Vis. Commun. and Image Proc.*, 2008. 1
- [14] J. Mallon and P. Whelan. Calibration and removal of lateral chromatic aberration in images. *Pattern Recognition Lett.*, 28(1):125–135, 2007. 2
- [15] D. Paliy, V. Katkovnik, R. Bilcu, S. Alenius, and K. Egiazarian. Spatially adaptive color filter array interpolation for noiseless and noisy data. *Int. J. Imaging Syst. and Technology*, 17(3):105–122, 2007. 5
- [16] R. Ramanath, W. Snyder, G. Bilbro, and W. Sander III. Demosaicking methods for Bayer color arrays. *J. Electron. Imaging*, 11(3):306–315, 2002. 1
- [17] S. Shah and J. Aggarwal. A simple calibration procedure for fish-eye (high distortion) lens camera. In *Proc. IEEE Conf. Robotics and Automation*, pages 3422–3427, 1994. 2
- [18] E. Simoncelli and E. Adelson. Noise removal via bayesian wavelet coring. In *Proc. Int. Conf. Image Process.*, volume 1, pages 379–382, Sept. 1996. 3
- [19] G. Stein. Lens distortion calibration using point correspondences. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, pages 602–608, 2002. 2

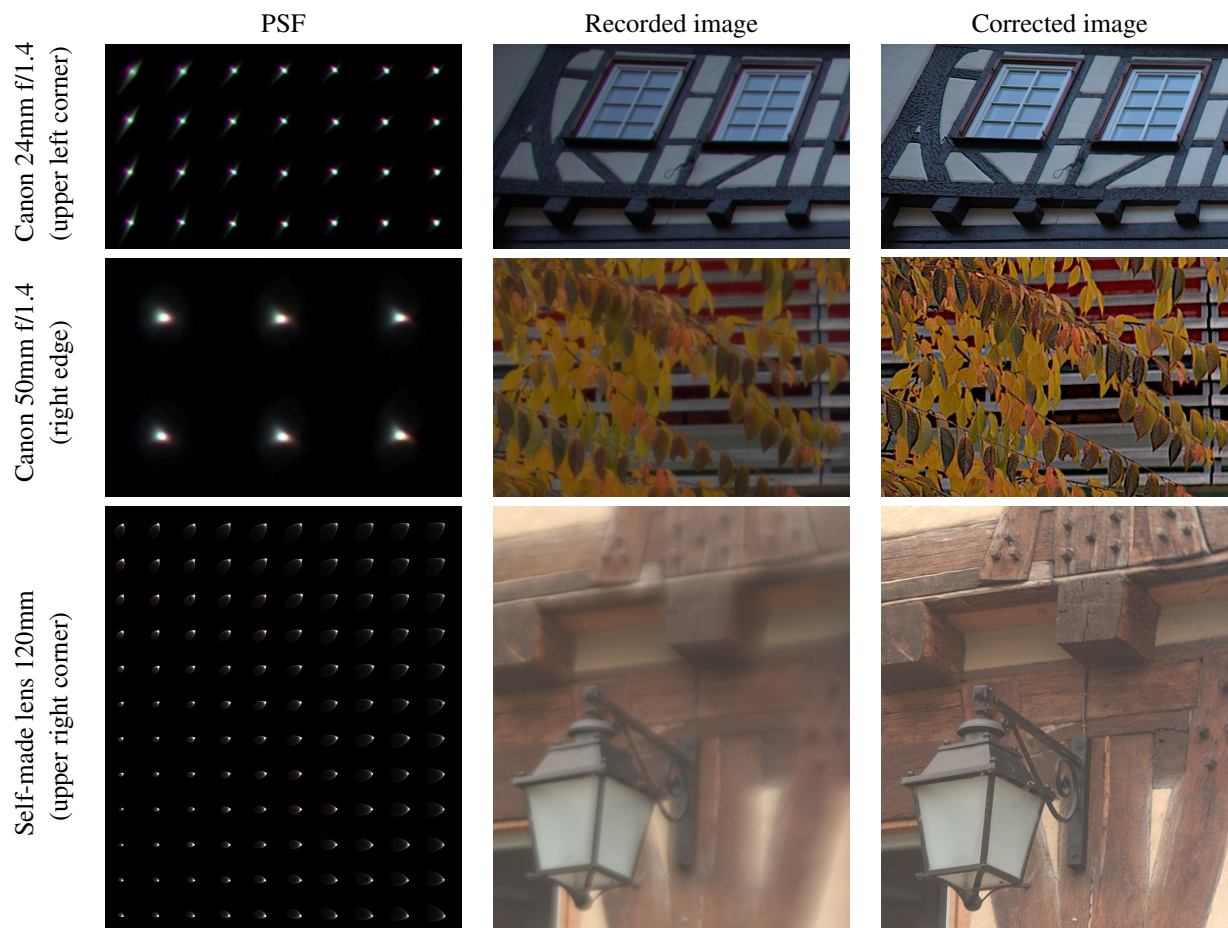


Figure 4. Comparison between original and corrected image and the respective PSFs.




Figure 5. Comparison with DXO for images taken with a Canon EF 50mm f/1.4 lens. Best viewed on screen.