

Self-Calibration of Optical Lenses

Michael Hirsch and Bernhard Schölkopf
 Department of Empirical Inference
 Max Planck Institute for Intelligent Systems
`firstname.lastname@tue.mpg.de`

Abstract

Even high-quality lenses suffer from optical aberrations, especially when used at full aperture. Furthermore, there are significant lens-to-lens deviations due to manufacturing tolerances, often rendering current software solutions like DxO, Lightroom, and PTLens insufficient as they don't adapt and only include generic lens blur models.

We propose a method that enables the self-calibration of lenses from a natural image, or a set of such images. To this end we develop a machine learning framework that is able to exploit several recorded images and distills the available information into an accurate model of the considered lens.

1. Introduction

High-quality lenses are typically the most expensive part in modern camera systems. Even high grade lenses exhibit severe lens aberrations at open aperture [38, 9]. In addition, lens-to-lens deviations due to manufacturing tolerances can be significant [5, 35]. Existing approaches aka DxO, Lightroom that use parametric models often neglect these individual errors and hence provide only limited image quality improvements.

Recently, Shih *et al.* [35] suggested an approach to calibrate individual lenses using a grid of light points and fitting parameters of the lens prescription to counter this issue. Schuler *et al.* [33] proposed a blind algorithm to correct for the individual errors of a lens. Our approach takes this one step further. It uses machine learning to estimate a regression model for the family of non-stationary lens blurs. This model represents point spread functions as real-valued functions of several parameters such as image position and orientation, allowing for continuous transfer across settings and multiple images. It makes use of major progress in the field of blind deconvolution (BD) during recent years [42], and it can exploit information present in sets of images in an elegant way, thus improving the blur model with the amount of photos taken over time. All these images are natural images, and none of them needs to be taken of a custom calibration pattern, which is why we refer to the process as *self-calibration*. Both high-quality lenses and simple lenses benefit from the approach that renders time and cost intensive physical calibration approaches unnecessary.

The optical aberrations of a lens are generally described by the point spread function (PSF). The overall PSF (or

blur) model is a high-dimensional function that for a given lens depends on a number of quantities, including position in the image plane, wavelength, aperture stop, focus length (for zooms), and focus setting. If we had a sufficiently large training set of accurate PSF observations¹ covering the space of input variability, we could simply apply machine learning to estimate the overall PSF. Each PSF measurement, however, would ideally require an optical measurement under controlled conditions, and it is not feasible to do this for each individual lens. The challenge addressed by the present paper thus is how to estimate this function from a relatively small set of imperfect PSF estimates extracted automatically from natural images, taking advantage of the symmetries of the problem. To this end, we introduce a non-parametric regression framework that learns a continuous PSF model as a function of lens settings. Our approach is generic and provides a framework that allows any local PSF estimation algorithm to be used and hence will benefit from future advances in the field of BD. At the same time any non-blind deconvolution (NBD) algorithm can be used for the final image correction, once the PSF family has been estimated. Therefore, our approach is not limited to a specific algorithm — it can be viewed as a novel generic tool for the continuous modeling of non-stationary point spread functions (PSFs).

Once we have a mathematical model of the blur at hand, we can use it to correct and enhance the captured image via NBD. We measure the improvement in image quality using test charts using the commercial image quality system iQ developed by Image Engineering [17].

Our contributions are three-fold:

- we present a benchmark for lens blur and make a comprehensive comparison of state-of-the-art PSF estimation algorithms
- we show how to get reliable PSF estimates from captured photos
- we introduce a new framework for spatially-varying PSF models depending on all relevant input variables that makes use of local PSF estimates and outputs a continuous PSF model tailored for each individual lens considered.

¹These local PSFs are often also referred to simply as PSFs, since they are also functions. We can think of them as functions obtained from the overall PSFs by restriction to a subset of the input arguments.

2. Problem formulation

The notion of a PSF is invaluable in the description of optical systems and directly owes itself to the linearity of the underlying physical equations. As mentioned above, the PSF of an imaging system describes its response to a point source or point object. Only in an ideal, i.e. aberration-free imaging system, each point of the object plane is mapped to a single point in the image plane, yielding a sharp image of the static scene being captured. However, in reality an ideal system does not exist,² and a point in the object plane will be spread or blurred according to the pattern determined by the *point spread function*.

In practice, capturing a picture with a digital image sensor yields a finite set of intensity values. A digital image can be represented by a matrix whose dimensions correspond to the resolution of the image. For simplicity, we consider x , y and f to be vectors, describing the underlying sharp image of the captured scene, the blurry possibly noise corrupted taken photo, and the PSF respectively. The case of matrices is analogous. Since the incoherent imaging equation [10] is linear in x , it can be written as

$$y = Fx. \quad (1)$$

where F is the convolution matrix of f . For incoherent light, the PSF fully characterizes an optical system and can be considered as its fingerprint. No two lenses have the same PSF. Blind deconvolution aims at inferring F from an image alone, i.e., given y , infer F without knowing x . Once F is known, one can remove the induced blur via non-blind deconvolution.

3. Related work

Blind image deconvolution. Blind image deconvolution (BD) has seen much progress in recent years. A good overview can be found in the review [42]. State-of-the-art methods include fast MAP based methods [4, 46] that employ heuristic non-linear filtering methods (e.g. shock filtering) for latent image reconstruction, the approach of [49] that use variational Bayesian inference for PSF and image estimation, and recently proposed methods [39, 26] that use powerful patch-based priors for image prediction. All of these methods have been developed for the removal of blur caused by unintended camera shake, which is a common image degradation that occurs independent of the lens quality.³ With the help of a new benchmark dataset, we test

²note that aperture diffraction, which invariably occurs, is also considered an aberration; for high quality and large f-stops, it can be the dominant aberration

³There is an interesting trade-off at work in that lens manufacturers try to counter this degradation using optical image stabilization. This, however, can lead to problems with lens centering (see e.g. http://www.photozone.de/fuji_x/879-fuji55200f3548?start=1), thus increasing the need for lens correction methods.

whether these methods are also able to estimate local PSFs stemming from optical aberrations and make a comprehensive comparison.

Other relevant work includes [46], arguing that many edges are actually not informative about the unknown PSF and may even misguide PSF estimation. A follow-up paper [15] trains a discriminative model and finds good regions for PSF estimation within an image. This idea has been further extended in [1], selecting several patches and building a synthetic mosaic comprising those parts of the image that are most informative about the PSF.

Non-blind image deconvolution. In contrast to BD, in non-blind image deconvolution (NBD) the PSF is assumed to be known. The task is to restore a sharp image from a blurry and possibly noisy one. Despite being a convex optimization problem, NBD is inherently ill-posed. This is due to its high sensitivity to noise, and the fact that blur can wipe out certain frequencies irretrievably.

One of the oldest methods, which still enjoys much popularity, is Richardson-Lucy deconvolution, originally proposed independently by [29] and [24]. It is an iterative method ensuring non-negativity of the estimated image. Only recently [44] proposed an extension that is able to deal with saturated pixels for which the linear model (1) does not hold true.

While non-negativity helps to constrain the solution space, it is often insufficient for the suppression of restoration artifacts such as ringing and the recovery of high-frequency detail. Other popular deconvolution routines make use of so-called natural image statistics [22, 21].

Recently proposed methods make use of learning approaches [30] and deep convolutional neural networks [31, 47], both of which achieve state-of-the-art results.

PSF modeling. While most works on BD and NBD assume an invariant PSF (i.e., a PSF model does not depend on the position in the image plane), in reality this is often not the case. [28] developed an extension of the convolutional model based on non-stationary combination [25]. In contrast, [14] proposed the so-called Efficient Filter Flow (EFF) model, based on non-stationary convolution, with a number of advantages when compared to Nagy's model [7]. Recently, [8] designed a model based on wavelets, which is also able to describe spatially-varying blur.

Optical aberration correction. Existing approaches to optical aberration correction can be categorized into those that measure the PSF in a calibration step, and those that estimate the PSF blindly.

[32] measure the PSF of a lens via point light sources and use the EFF framework to correct optical aberrations using subsequent non-blind deconvolution. Instead of point stimuli, [35] use a grid of illuminated pinholes for PSF measurement. Rather than using the PSF measurement directly, they fit a mechanical model (known as lens prescription model)

of the lens to the measured PSFs. The correction is again performed via NBD.

[19] estimate the PSF via a calibration pattern and fit a spatially-varying parametric model. [12] also use a calibration pattern for PSF estimation and present a non-blind deconvolution method that exploits cross-channel information to yield high-quality results. [18] propose a method that performs PSF estimation either with help of a calibration pattern or blindly through sharp edge prediction. [37] uses images of a point grid, of which they compute local moments, which are used to fit a polynomial description (Seidel polynomials) of lens aberrations.

Probably closest to our approach is the work of [33], using a blind deconvolution approach for PSF determination and a global lens PCA basis to constrain and connect local PSF estimates. The correction and image reconstruction is done via [21]. In contrast to our approach, this work only makes use of a single image and thus cannot propagate information about the PSF between images taken with the same lens; moreover, it cannot take into account all the input variables that our approach can handle.

[41] discusses the topic from a theoretical perspective and sheds light on the question how much information about the PSF can be obtained from a single image. [48] makes use of optical computing to correct for optical aberrations.

4. Overview

Our method consists of three essential ingredients:

- (1) **Local PSF estimation.** Each captured photo is assayed for informative regions suitable for reliable PSF estimation. If this is the case we employ a state-of-the-art PSF estimation technique to yield local PSF estimates. The PSF depends on a number of parameters including position within the image, color channel, and aperture setting.
- (2) **PSF modeling.** The local PSF estimates can be considered as samples of the continuous PSF that we seek to infer. To this end, we use efficient non-parametric regression methods from the ML literature to estimate the underlying physical PSF of the given lens.
- (3) **PSF correction.** With the PSF model at hand, we correct the images captured with it via non-blind non-stationary deconvolution.

EFF [14] provides an efficient way for the description and computation of non-stationary blur. Its drawback is the necessary increase in parameters for the modeling of large complex PSFs due to its limited interpolation capabilities. We propose a novel approach to handle this using non-parametric kernel regression. Previously, basis sets have been proposed for specific problems such as camera shake and optical aberrations. Presently, we describe a generic

learning framework and apply it to the problem of blind correction of non-stationary optical aberrations.

In the following sections we will discuss each of these steps in detail before we will present experimental results in Sec. 9, concluding in Sec. 10.

5. PSF estimation

5.1. Benchmark for optical aberrations

For building a model of a PSF that accurately describes optical lens aberrations, we need to sample it. Previously this has been accomplished by time-consuming measurement procedures as proposed in [32, 35, 12]. An alternative is to estimate local PSFs from images as proposed by [33].

In recent years single image blind deconvolution has seen much progress in dealing with blur arising from camera shake during exposure. However, it is not clear whether a state-of-the-art PSF estimation method developed for motion blur performs equally well for blur caused by optical aberrations.

Therefore, we created a benchmark and tested several state-of-the-art methods. The benchmark consists of 80 images \times 12 kernels, yielding 960 blurred images. These images are generated similarly to [39]: we use the same 80 ground truth images, convolve them with different PSFs (in a stationary way, i.e., using the same PSF for the whole image), and add 1% Gaussian noise. The PSFs were obtained by taking digital SLR images of a point light source in a distance of 150 cm, placed in one of the four image corners (in our case, the upper right), where the optical aberrations are most severe. For each PSF the corresponding lens was focused by first placing the above point light source into the image center. Then the camera was rotated to record the PSF⁴. For each PSF, three images were taken, which were then averaged and normalized. For simplicity, we use the green color channel and thus neglect chromatic aberrations. The table in Fig. 1 summarizes the lenses that were measured, providing a selection of Canon EOS lenses ranging from popular low-cost lenses (Canon EF 50mm f/1.4 USM) to professional lenses (Zeiss Otus 55mm f/1.4 APO-Distagon). Figure 1 shows the PSFs; note that they differ from PSFs typical for motion blur as measured and shown in [23, 20]. When comparing the PSFs, bear in mind that the scale of the PSF images varies, as specified in the table. We also want to stress, that through our particular focusing procedure no conclusions can be drawn from the size of the PSF to the imaging quality of the corresponding lens.

The result of a comparison of five state-of-the-art local PSF estimation methods is shown in the lower left panel of Fig. 1. All results were obtained using subsequent non-blind deconvolution using the method of [50]. The error

⁴Note that the focus distance was thus not optimal for the corner, leading to an error which depends on the field curvature of the lens.

ratio is defined and computed according to [23]. Images with values below 3 are visually of acceptable image quality with little or no artifacts. In particular, this means that all considered lenses mostly take pictures with acceptable image quality in the corners (dashed line). Interestingly, two of the evaluated methods [4, 46], which are based on heuristic filtering methods such as shock filtering, are hardly able to improve image quality on the present task. [39] is able to improve image quality in about 65% of the cases, [49] in about 90% of the cases. Only the recently proposed method of [26] dominates the dashed line and thus consistently improves image quality for all tested images.

5.2. PSF samples through patch-based estimation

All of the above tested BD methods assume that the blur is invariant across the image plane and return a single PSF for a given input image. While the assumption of a stationary PSF seems often acceptable for camera shake [20], it is unrealistic for blur stemming from optical aberrations [32].

Rather than adopting and modifying a state-of-the-art BD method for non-stationary lens correction as e.g. [33] did, we employ another strategy:

- (1) First we assay an image or a set of images (taken with the same lens and settings) to find image regions that are most informative about the unknown PSF. We consider two different approaches: 1) one based on the rmap idea of [46], which computes a mask for PSF estimation that neglects all edges that likely impede blur estimation; 2) the other uses the discriminative model of [15] that selects good regions to deblur. In practice we found both approaches to work equally well.
- (2) Second, we run [26], the best performing method in our lens blur benchmark, on the set of image patches selected by (1). For each analysed patch we get a single estimated blur kernel, which after centering (i.e. shifting the center of mass into the center) serves as a local sample of the global non-stationary PSF, to be used as input to the subsequent data-driven PSF modeling step.

The motivation for this approach is threefold: 1) it is completely generic and modular such that any future advances in one of these sub-tasks could be readily employed to improve the overall performance; 2) uninformative image regions do not hamper PSF estimation; 3) since only parts of the input image(s) are analysed, it is computationally more efficient and allows for massively parallel processing.

6. PSF modeling

6.1. Motivation

Although EFF provides an efficient way for space-variant filtering, it has some shortcomings if computational efficiency should be maintained:

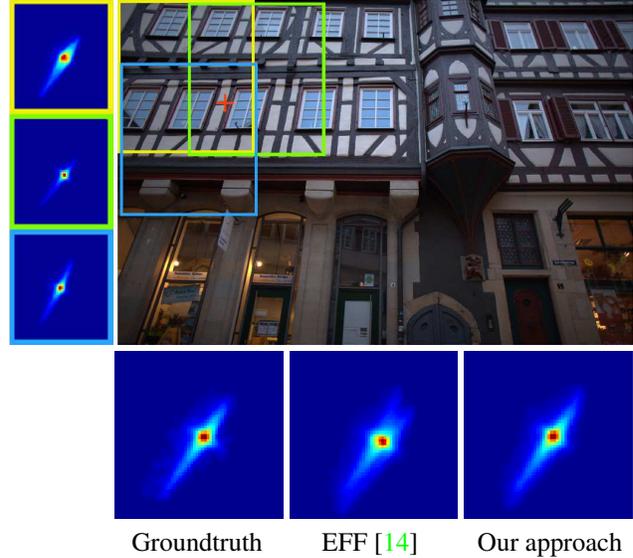


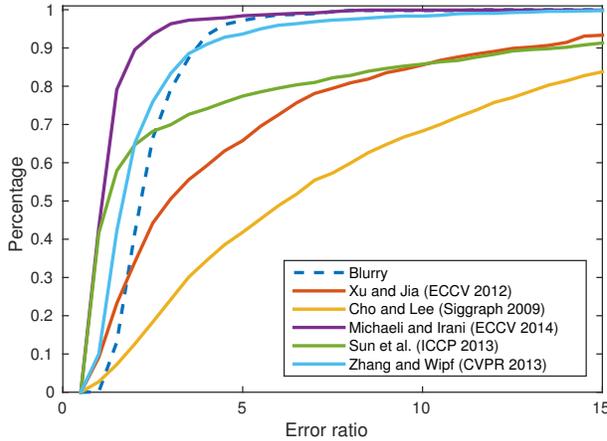
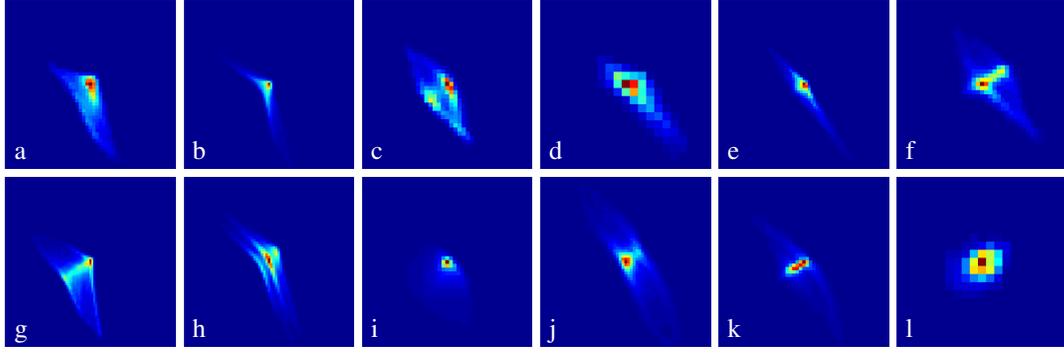
Figure 2. Comparison between EFF and our regression based approach: EFF computes the effective PSF at a certain pixel position (red cross) as a weighted linear interpolation between PSFs (left) from different image regions (colored boxes). For complex PSFs due to optical aberrations (Canon EF 24mm f/1.4L) such a description is insufficient and provides unsatisfactory results. In contrast to EFF, our approach models the overall PSF as a high-dimensional regression, continuously interpolating in all parameters and properly accounting for the location and orientation of the PSF samples. As the image on the bottom right illustrates, this truly interpolates, rather than taking a superposition (bottom middle), thus yielding a faithful PSF representation. For better visualization we applied a gamma of 0.5 to all PSF images.

P1 the patches must have the same size

P2 the local PSFs used to parametrise the model need to lie on a regular equi-spaced grid

P3 the weighting functions perform linear interpolation between the parametrising PSFs

In particular, property P3 has strong implications with regard to the accuracy of PSF modeling: non-trivial PSFs can only be represented faithfully if the number of parametrising PSF samples is high enough. Figure 2 illustrates this issue for a non-trivial PSF caused by optical aberrations of a Canon 24 mm f/1.4L lens. EFF performs a linear interpolation between the observed PSF samples to compute the local PSF applied at a certain pixel location. Note, that for each pixel the weighting is different according to the overlap and weighting function that have been chosen. However, irrespective of that choice the interpolation is always linear, i.e., the applied PSF is a weighted linear combination of the observed PSF samples. Figure 2 shows the resulting PSF for the popular choice of 50% overlap and a Bartlett-Hann window. The resulting PSF is a sub-



Lens	PSF size
a Canon EF 85mm f/1.2L II USM	51x51
b Canon EF 50mm f/1.4 USM	81x81
c Zeiss Otus 55mm f/1.4 APO-Distagon	41x41
d Walimex Pro 35mm f/1.4	21x21
e Canon EF 35mm f/1.4L USM	65x65
f Sigma 35mm f/1.4 DG HSM	41x41
g Sigma 50mm f/1.4 EX DG HSM	91x91
h Nikon AF-S Nikkor 50mm f/1.4G	71x71
i Sigma 85mm f/1.4 EX DG HSM	45x45
j Nikon AF Nikkor 20mm f/2.8D	61x61
k Voigtlander 20mm f/3.5 SL-II	51x51
l Nikon AF-S Nikkor 14-24mm f/2.8G ED (PSF taken at 14mm)	21x21

Figure 1. Benchmark for blind deconvolution of images that have been artificially blurred with real-world PSFs of popular photographic lenses (see table). The PSFs were recorded at open aperture with a light source placed in the upper right image corner. For the benchmark we used only the green color channel, comprising half of the pixels of the raw images. The plot on the left shows the percentage of images whose degradation is upper bounded by the given error rate, thus higher is better.

optimal approximation of the true PSF. To some extent, this shortcoming can be addressed by using more PSF samples, but this is not always possible: in the case of BD, where we do not measure but estimate the PSF, not all image regions yield good PSF estimates. For example, regions with few edges and little texture might not contain enough useful information for a reliable local PSF estimation.

This issue has been recognised before in the context of non-stationary motion blur [11]. There, the authors tried to mitigate this effect by enforcing similarity (as measured by Euclidean distance) between neighboring PSF estimates. In addition, they tried to identify bad estimates through an entropy measure, which once identified were then excluded and replaced.

A more sophisticated approach to constrain the estimated PSF during blind deconvolution has been proposed in [13] and [33]. Although tackling different problems (motion blur removal vs. correction of optical aberrations), the key idea is the same: in both cases the local PSF estimates are projected to a PSF basis which constrains the solution space to physically plausible PSFs only. The optimization is then per-

formed in the basis coefficients rather than the parametrising PSFs. Since the optimization problem size scales exponentially with the dimensionality of the basis, both the dimensionality as well as the number of discretization step has to be kept low. While in [33], this basis was orthogonal, [13] used a over-complete basis.

In this following section we present an alternative approach to this problem, which is not only computationally efficient but also offers a number of additional advantages.

6.2. Kernel regression for PSF modeling

Kernel regression denotes a class of non-parametric techniques in statistics used for the estimation of non-linear relations between random variables.

$$y_i = f(x_i) + \epsilon_i, \quad i = 1, 2, \dots, n. \quad (2)$$

Nadaraya [27] and Watson [43] proposed to estimate such a relation f as a locally weighted average of observations y_i using a kernel K as a weighting function:

$$\hat{f}(x) = \frac{\sum_{i=1}^n y_i K_h(x - x_i)}{\sum_{i=1}^n K_h(x - x_i)} \quad (3)$$

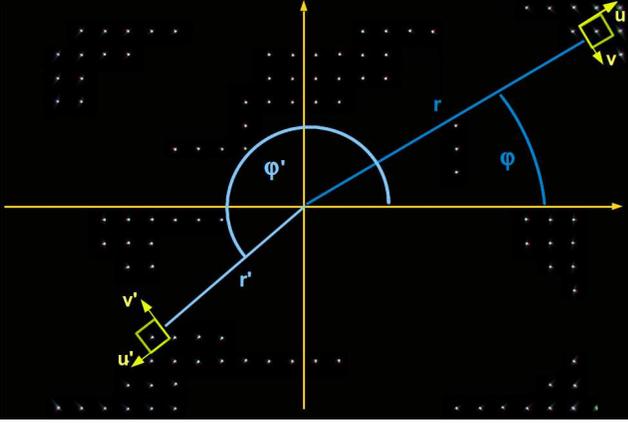


Figure 3. Illustration of coordinate parametrization used for multi-variate kernel regression in our PSF modeling step. The irregularly located bright “blobs” correspond to PSF estimates from selected image patches. Each PSF estimate within the image plane is parametrized by its distance to the image center r , its angle φ and a locally attached coordinate system (u, v) . Please note, that (u, v) and (u', v') denote the coordinates in a rotated local coordinate system and hence do not take integer but real values.

The Nadaraya-Watson estimator (NWE) is the simplest form of kernel regression and can be viewed as an adaptive filter. As pointed out in [40] many popular image processing filters such as bilateral filtering can be interpreted as NWE with a modified kernel definition. In the multivariate case the Nadaraya-Watson estimator (3) becomes

$$\hat{f}(\mathbf{x}) = \frac{\sum_{i=1}^n y_i K_{\mathbf{H}}(\mathbf{x} - \mathbf{x}_i)}{\sum_{i=1}^n K_{\mathbf{H}}(\mathbf{x} - \mathbf{x}_i)} \quad (4)$$

In our case, the function we want to estimate is the true underlying PSF, which is a real-valued continuous function of several variables including the positions of the local PSF sample within the image plane, given in polar coordinates (r, φ) , pixel position within the local PSF sample (u, v) , and possibly other parameters (Φ) such as wavelength or aperture setting, i.e.,

$$f = f(u, v, r, \varphi, \Phi) \quad (5)$$

Figure 3 illustrates the parametrization that we use in our experiments. Note that u and v denote the coordinates in the rotated local coordinate system and thus take real values rather than integers. Hence we have access to samples that lie on a non-regular grid and hence provide sub-pixel information about the unknown PSF. This is a serendipitous advantage of our choice to use rotated coordinate systems for the local PSF, a choice we originally made to be able to take into account the fact that PSFs tend to be close to rotationally symmetric with respect to the image center.

In summary, each local PSF sample thus provides us with a set of training points for the regression (one for each pixel), where each training point consists of (at least) four position variables as input, and one intensity value (the value of that pixel) as output.

6.3. Kernel choice

It is known [36] that for the case of classical kernel regression the choice of the kernel has only a small effect on the accuracy of estimation (in contrast to the choice of the bandwidth), and, therefore, preference is given to differentiable kernels with low computational complexity such as an RBF kernel. Furthermore, a reasonable assumption is that the kernel factorizes. We model the kernel as

$$K_{\mathbf{H}}(\mathbf{x}, \mathbf{x}') = \prod_{x \in \{u, v, r, \varphi, \Phi\}} K_{h_x}(x, x') \quad (6)$$

where for $x \in \{u, v, r, \Phi\}$ we assume a RBF kernel

$$K_{h_x}(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2h_x^2}\right). \quad (7)$$

For $x = \varphi$ we map the variable to the two-dimensional unit sphere and use a Gaussian RBF kernel in terms of distances in that space. This avoids discontinuous treatment of the circular variable ϕ , thus ensuring that any two PSF estimates that are close to each other in the image plane are encouraged to take similar values. The values of the bandwidth parameters h_x control the coupling and smoothing of the available PSF samples. For example to enforce rotational symmetry, i.e., no dependence of the PSF on φ , one would choose a very large bandwidth h_φ , such that PSF samples with a certain distance r to the image center are averaged during regression.

For automatic parameter determination we follow [3] and compute the bandwidth for variable x as

$$h_x = \sigma_x \left\{ \frac{4}{(d+2)n} \right\}^{1/(d+4)} \quad (8)$$

where d and n denote the dimensionality and number of input samples respectively; σ_x denotes the standard deviation of variable x and is computed via its sample estimator. To speed up computation we employ Kd-tree accelerated nearest neighbor search to approximate kernel computations, similar to [34]. This method has the advantage that the “training” only consists of building the Kd-tree, and it is possible to update the tree as additional training examples become available, hence allowing a lens model that improves over time as new images come in.

While we use kernel regression in our experiments, any non-parametric regression technique could be used. We also experimented with Gaussian process regression, however found it computationally prohibitive due to the large amount

of variables involved. A typical problem size in our experiments is of the order 10^6 , and this size could increase further for large image databases or bigger sized local PSFs.

7. PSF correction

With an accurate PSF model at hand, existing images as well as images to be captured subsequently can be corrected via non-blind deconvolution (NBD). In this work, we use the Richardson-Lucy NBD algorithm of [44], which is capable of dealing with partially saturated image regions. For efficient computation we use EFF [14] and accelerate the iterative Richardson-Lucy updates via [2].

8. Implementation

We run all our experiments with unoptimized Matlab code. For PSF estimation we use the Matlab code of [26], for patch selection the Matlab code of [15], and for PSF correction a modified version of the Matlab code of [44]. Due to the long processing time of [26], we distribute the workload to a compute cluster and process all patches in parallel. The size of the unknown blur kernel is kept fixed to 55×55 pixels in all our experiments. The runtime for each of the steps, i.e., PSF estimation, modeling and correction, is in the order of tens of minutes. For RAW development we use Dave Coffin’s `dcrw` [6]. For chromatic aberration correction we use `tca_correct` and `fulla` of the freely available panorama stitching software Hugin [16].

9. Experimental results

Due to limited space we show only a few selected results here. For more comparisons we refer the interested reader to the supplementary material⁵.

Comparison with state-of-the-art. Figures 5 and 6 show comparisons with state-of-the-art methods for both non-blind [32] and blind [33, 19] lens correction. In both cases our approach yields comparable if not superior results. Figure 4 depicts the estimated PSF for the example of Fig. 5 before and after applying kernel regression.

Comparison with commercial software. Commercial software packages such as DxO, Lightroom and PTLens that aim to correct optical aberrations work with pre-defined parametric lens models. While they manage to correct for geometric distortions and lateral chromatic aberrations rather well, they typically fall short in correcting other types of blur such as astigmatism, coma and longitudinal chromatic aberrations. In Figure 6 we compare with a result obtained with DxO Optics Pro 7.2, an image borrowed from [33]. The comparison demonstrates that our approach compares favourably in terms of recovered overall sharpness.

Quantitative evaluation. We measure the improvement in image quality based on a test chart (TE268) with an

MTF50	w/o corr.	$n=1$	$n=2$	$n=5$	non-blind
Corner	383	545	608	654	696
Center	613	923	1013	1076	1120

Table 1. Improvement in MTF50 for Canon 24mm f/1.4L lens measured in line pairs per image height (LP/PH) through test chart TE268 and iQ Analyzer of Image Engineering. Higher means better. n denotes the number of images used for PSF modeling.

off-the-shelf image quality system[17], developed by Image Engineering. A statistic that is commonly used for quantitative assessment of resolution and image quality is the MTF50, which denotes the value at which the Modulation Transfer Function drops down to half of its maximum value. In Tab. 1 we show how with an increasing number of analysed images the image quality improves. This experiment illustrates that our approach allows a self-calibration of lenses that improves with time.

10. Conclusion

Optical aberrations are inevitable and no lens is ever perfect. To mitigate the adverse effects of optical flaws we proposed a novel method for blind lens correction. It comprises three steps: local PSF estimation, overall PSF modeling, and PSF correction. Our approach is generic and allows simple integration of existing and future methods for any of these sub-tasks. For PSF modeling we use non-parametric kernel regression as a versatile tool that allows the incorporation of physical constraints such as symmetry conditions, and smoothly interpolates between local PSFs to build a continuous PSF family producing a meaningful estimate at any desired point. In a number of experiments, we demonstrate the robustness and effectiveness of our approach.

A drawback of our current system is the limited extrapolation capability of the proposed regression technique using localized kernels. While it performs well in interpolating between local PSF estimates and distributing information throughout the image plane, it does less so for extrapolating a non-trivial PSF pattern to image corners. In principle, this could be addressed by a localized criterion which adjusts the kernel bandwidth depending on the amount of data available in a neighborhood. Alternatively, recent advances in fast Gaussian process regression for multi-dimensional pattern extrapolation [45] might be able to help remedy this problem, suggesting another direction for future work.

References

- [1] H. Bae, C. C. Fowlkes, and P. H. Chou. Patch mosaic for fast motion deblurring. In *ACCV*, 2013. 2
- [2] D. Biggs. Acceleration of iterative image restoration algorithms. *Applied Optics*, 36(8):17661775, 1997. 7
- [3] A. W. Bowman and A. Azzalini. *Applied Smoothing Techniques for Data Analysis*. Oxford Univ. Press, 1997. 6
- [4] S. Cho and S. Lee. Fast Motion Deblurring. *SIGGRAPH*, 2009. 2, 4

⁵Available for download from the accompanying project webpage http://webdav.is.mpg.de/pixel/psf_regression

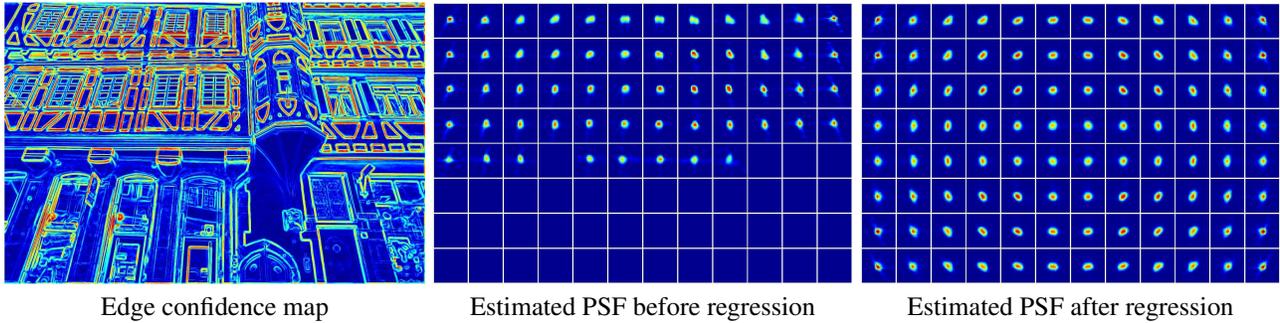


Figure 4. Our method first analyses an input image and determines those image regions that are most informative about the unknown PSF. To this end it computes an edge confidence map (left panel) that highlights useful edges for PSF estimation. For selected regions we compute local PSF estimates via [26] (middle panel), which serve as input to our PSF modeling step. Through regression we combine local information to a globally consistent model of the PSF (right panel). Only the green color channel of the example image of Fig. 5 is shown.

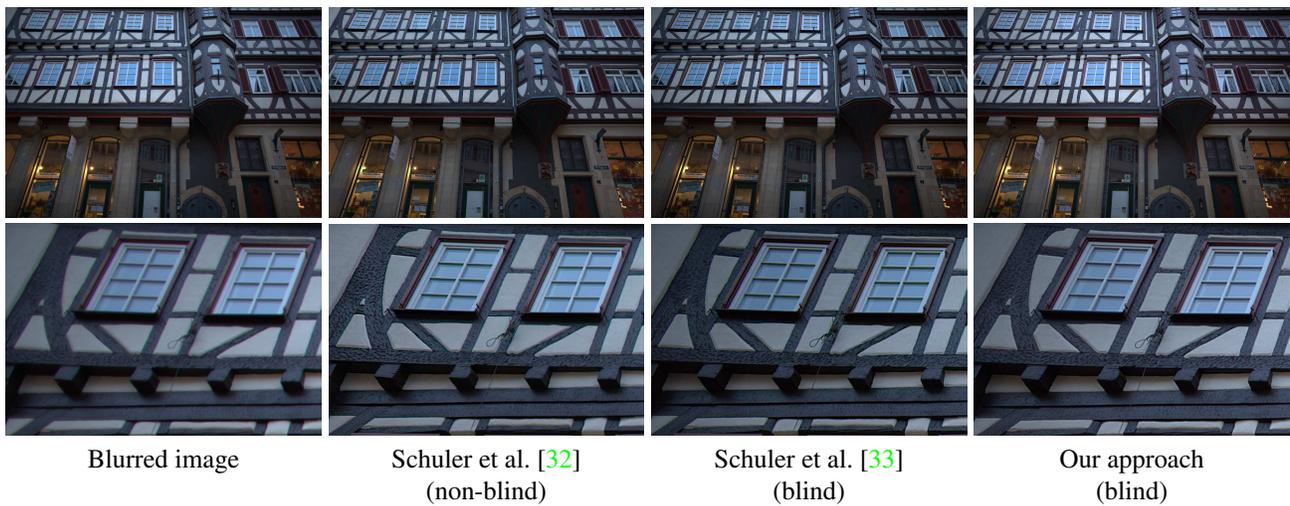


Figure 5. Image taken with a Canon 24mm f/1.4 lens: Comparison with the non-blind [32] and blind [33] approach of Schuler et al. Closeups show the upper left corner of the image, where optical aberrations are most severe. Note, that [33] was able to correct the image in only half of the original resolution due to its high memory demand. In contrast, [32] and our approach were applied at full image resolution, i.e. 5634×3753 pixels. In terms of level of detail and overall sharpness our result is on a par despite the larger effective blur when compared to [33]. Also, our two step approach to chromatic aberration correction seems to do a better job in suppressing color fringes.

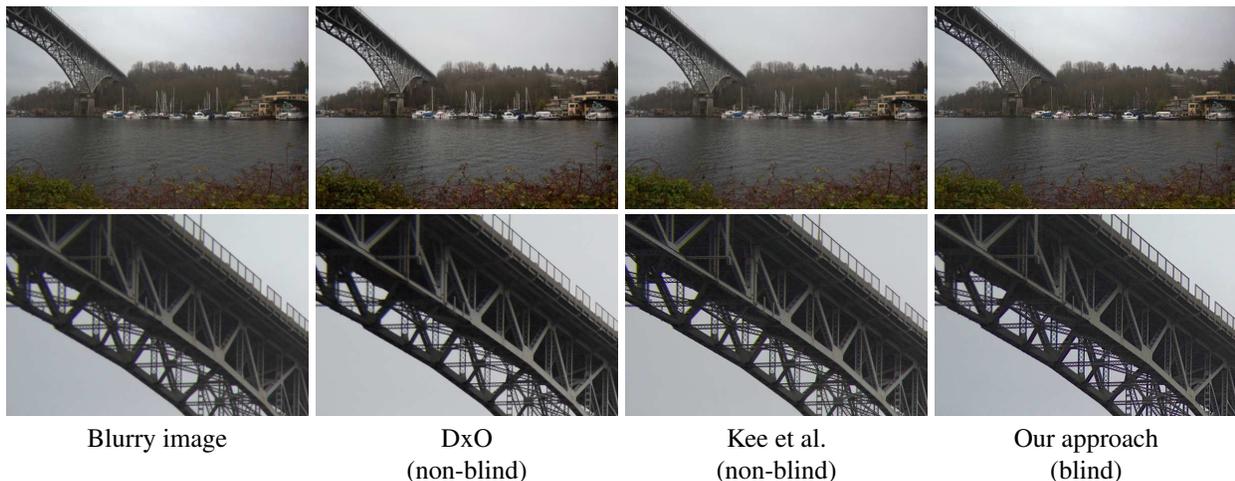


Figure 6. Comparison with a commercial software solution (DxO) to optical lens correction and the state-of-the-art non-blind approach of [19]. We stress the fact that our result was obtained blindly based on the input image only.

- [5] R. Cicala. <http://www.lensrentals.com/blog/2013/09/there-is-no-perfect-lens>, 2013. 1
- [6] DCRaw. <https://www.cybercom.net/~dcocoffin/dccraw>, Version: v9.26. 7
- [7] L. Denis, T. E., and F. Soulez. Fast Model of Space-Variant Blurring and Its Applications to Deconvolution in Astronomy. In *ICIP*, 2011. 2
- [8] P. Escande and P. Weiss. Numerical computation of spatially varying blur operators a review of existing approaches with a new one. *arXiv preprint arXiv:1404.1023*, 2014. 2
- [9] R. E. Fischer, B. Tadic-Galeb, and P. R. Yoder. *Optical System Design*. SPIE Press, 2008. 1
- [10] J. W. Goodman. *Introduction to Fourier optics*. McGraw-Hill, 2005. 2
- [11] S. Harmeling, M. Hirsch, and B. Schölkopf. Space-variant single-image blind deconvolution for removing camera shake. In *NIPS*, 2010. 5
- [12] F. Heide, M. Rouf, M. B. Hullin, B. Labitzke, W. Heidrich, and A. Kolb. High-quality computational imaging through simple lenses. In *SIGGRAPH*, 2013. 3
- [13] M. Hirsch, C. Schuler, S. Harmeling, and B. Schölkopf. Fast removal of non-uniform camera-shake. In *ICCV*, 2011. 5
- [14] M. Hirsch, S. Sra, B. Schölkopf, and S. Harmeling. Efficient filter flow for space-variant multiframe blind deconvolution. In *CVPR*, 2010. 2, 3, 4, 7
- [15] Z. Hu and M.-H. Yang. Good regions to deblur. In *ECCV*, 2012. 2, 4, 7
- [16] Hugin. <http://hugin.sourceforge.net>, Version: 2013.0.0.4692917e7a55. 7
- [17] Image Engineering - iQ-Analyzer. <http://www.image-engineering.de>, V6.0.33. 1, 7
- [18] N. Joshi, R. Szeliski, and D. J. Kriegman. PSF estimation using sharp edge prediction. In *CVPR*, 2008. 3
- [19] E. Kee, S. Paris, S. Chen, and J. Wang. Modeling and removing spatially-varying opt. blur. In *ICCP*, 2011. 3, 7, 8
- [20] R. Köhler, M. Hirsch, B. Mohler, B. Schölkopf, and S. Harmeling. Recording and playback of camera shake: benchmarking blind deconvolution with a real-world database. In *ECCV*, 2012. 3, 4
- [21] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-Laplacian priors. In *NIPS*, 2009. 2, 3
- [22] A. Levin. Blind motion deblurring using image statistics. In *NIPS*, 2006. 2
- [23] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding and evaluating blind deconvolution algorithms. In *CVPR*, 2009. 3, 4
- [24] L. Lucy. An iterative technique for the rectification of observed distributions. *The Astron. J.*, 79:745–754, 1974. 2
- [25] G. Margrave. Theory of nonstationary linear filtering in the fourier domain with application to time-variant filtering. *Geophysics*, 63(1):244–259, 1998. 2
- [26] T. Michaeli and M. Irani. Blind deblurring using internal patch recurrence. In *ECCV*, 2014. 2, 4, 7, 8
- [27] E. A. Nadaraya. On estimating regression. *Theory of Probability & Its Applications*, 9(1):141–142, 1964. 5
- [28] J. G. Nagy and D. P. O’Leary. Fast iterative image restoration with a space-varying psf. *Advanced Signal Processing Algorithms: Architectures, and Implementations IV*, 3162:388–399, 1997. 2
- [29] W. H. Richardson. Bayesian-based iterative method of image restoration. *J. Optical Soc. Amer.*, 62:55–59, 1972. 2
- [30] U. Schmidt, C. Rother, S. Nowozin, J. Jancsary, and S. Roth. Discriminative non-blind deblurring. In *CVPR*, 2013. 2
- [31] C. J. Schuler, H. C. Burger, S. Harmeling, and B. Schölkopf. A machine learning approach for non-blind image deconvolution. In *CVPR*, 2013. 2
- [32] C. J. Schuler, M. Hirsch, S. Harmeling, and B. Schölkopf. Non-stationary correction of optical aberrations. In *ICCV*, 2011. 2, 3, 4, 7, 8
- [33] C. J. Schuler, M. Hirsch, S. Harmeling, and B. Schölkopf. Blind correction of optical aberrations. In *ECCV*, 2012. 1, 3, 4, 5, 7, 8
- [34] Y. Shen, A. Ng, and M. Seeger. Fast gaussian process regression using kd-trees. In *NIPS*, 2006. 6
- [35] Y. Shih, B. Guenter, and N. Joshi. Image enhancement using calibrated lens simulations. In *ECCV*, 2012. 1, 2, 3
- [36] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, 1986. 6
- [37] J. D. Simpkins and R. L. Stevenson. A spatially varying psf model for seidel aberrations and defocus. In *IS&T/SPIE Electronic Imaging*, 2013. 3
- [38] W. J. Smith. *Modern Optical Engineering*. SPIE, 2007. 1
- [39] L. Sun, S. Cho, J. Wang, and J. Hays. Edge-based blur kernel estimation using patch priors. In *ICCP*, 2013. 2, 3, 4
- [40] H. Takeda, S. Farsiu, and P. Milanfar. Kernel regression for image processing and reconstruction. *PAMI*, 16(2):349–366, 2007. 6
- [41] H. Tang and K. N. Kutulakos. What does an aberrated photo tell us about the lens and the scene? In *ICCP*, 2013. 3
- [42] R. Wang and D. Tao. Recent progress in image deblurring. *arXiv preprint arXiv:1409.6838*, 2014. 1, 2
- [43] G. S. Watson. Smooth regression analysis. *Sankhyā: The Indian J. of Stat., Series A*, pages 359–372, 1964. 5
- [44] O. Whyte, J. Sivic, and A. Zisserman. Deblurring shaken and partially saturated images. *International Journal of Computer Vision*, 110(2):185–201, 2014. 2, 7
- [45] A. Wilson, E. Gilboa, J. P. Cunningham, and A. Nehorai. Fast kernel learning for multidimensional pattern extrapolation. In *NIPS*, 2014. 7
- [46] L. Xu and J. Jia. Two-phase kernel estimation for robust motion deblurring. In *ECCV*, 2010. 2, 4
- [47] L. Xu, J. S. Ren, C. Liu, and J. Jia. Deep convolutional neural network for image deconvolution. In *NIPS*, 2014. 2
- [48] T. Yue, J. Suo, Y. Xiao, L. Zhang, and Q. Dai. Image quality enhancement using original lens via optical computing. *Optics Express*, 22(24):29515–29530, 2014. 3
- [49] H. Zhang, D. Wipf, and Y. Zhang. Multi-image blind deblurring using a coupled adaptive sparse prior. In *CVPR*, 2013. 2, 4
- [50] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, 2011. 3