Blind 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/blind_lenscorrection/

Abstract. Camera lenses are a critical component of optical imaging systems, and lens imperfections compromise image quality. While traditionally, sophisticated lens design and quality control aim at limiting optical aberrations, recent works [1,2,3] promote the correction of optical flaws by computational means. These approaches rely on elaborate measurement procedures to characterize an optical system, and perform image correction by non-blind deconvolution.

In this paper, we present a method that utilizes physically plausible assumptions to estimate non-stationary lens aberrations *blindly*, and thus can correct images without knowledge of specifics of camera and lens. The blur estimation features a novel preconditioning step that enables fast deconvolution. We obtain results that are competitive with state-of-theart non-blind approaches.

1 Introduction

Optical lenses image scenes by refracting light onto photosensitive surfaces. The lens of the vertebrate eye creates images on the retina, the lens of a photographic camera creates images on digital sensors. This transformation should ideally satisfy a number of constraints formalizing our notion of a veridical imaging process. The design of any lens forms a trade-off between these constraints, leaving us with residual errors that are called *optical aberrations*. Some errors are due to the fact that light coming through different parts of the lens can not be focused onto a single point (spherical aberrations, astigmatism and coma), some errors appear because refraction depends on the wavelength of the light (chromatic aberrations). A third type of error, not treated in the present work, leads to a deviation from a rectilinear projection (image distortion). Camera lenses are carefully designed to minimize optical aberrations by combining elements of multiple shapes and glass types.

However, it is impossible to make a perfect lens, and it is very expensive to make a close-to-perfect lens. A much cheaper solution is in line with the new field of computational photography: correct the optical aberration in software. To this end, we use *non-uniform (non-stationary) blind deconvolution*. Deconvolution is a hard inverse problem, which implies that in practice, even non-blind uniform deconvolution requires assumptions to work robustly. Blind deconvolution is harder still, since we additionally have to estimate the blur kernel, and

non-uniform deconvolution means that we have to estimate the blur kernels as a function of image position. The art of making this work consists of finding the right assumptions, sufficiently constraining the solution space while being at least approximately true in practice, and designing an efficient method to solve the inverse problem under these assumptions. Our approach is based on a *forward model* for the image formation process that incorporates two assumptions:

- (a) The image contains certain elements typical of natural images, in particular, there are sharp edges.
- (b) Even though the blur due to optical aberrations is non-uniform (spatially varying across the image), there are circular symmetries that we can exploit.

Inverting a forward model has the benefit that if the assumptions are correct, it will lead to a plausible *explanation* of the image, making it more credible than an image obtained by sharpening the blurry image using, say, an algorithm that filters the image to increase high frequencies.

Furthermore, we emphasize that our approach is *blind*, i.e., it requires as an input only the blurry image, and not a point spread function that we may have obtained using other means such as a calibration step. This is a substantial advantage, since the actual blur depends not only on the particular photographic lens but also on settings such as focus, aperture and zoom. Moreover, there are cases where the camera settings are lost and the camera may even no longer be available, e.g., for historic photographs.

2 Related Work and Technical Contributions

Correction of optical aberrations: The existing deconvolution methods to reduce blur due to optical aberrations are *non-blind* methods, i.e., they require a time-consuming calibration step to measure the point spread function (PSF) of the given camera-lens combination, and in principle they require this for all parameter settings. Early work is due to Joshi et al. [1], who used a calibration sheet to estimate the PSF. By finding sharp edges in the image, they also were able to remove chromatic aberrations blindly. Kee et al. [2] built upon this calibration method and looked at the problem how lens blur can be modeled such that for continuous parameter settings like zoom, only a few discrete measurements are sufficient. Schuler et al. [3] use point light sources rather than a calibration sheet, and measure the PSF as a function of image location. The commercial software "DxO Optics Pro" (DXO) also removes "lens softness"¹ relying on a previous calibration of a long list of lens/camera combinations referred to as "modules." Furthermore, Adobe's Photoshop comes with a "Smart Sharpener," correcting for lens blur after setting parameters for blur size and strength. It does not require knowledge about the lens used, however, it is unclear if a genuine PSF is inferred from the image, or the blur is just determined by the parameters.

¹ http://www.dxo.com/us/photo/dxo_optics_pro/features/optics_geometry_ corrections/lens_softness

Non-stationary blind deconvolution: The background for techniques of optical aberration deconvolution is recent progress in the area of removing camera shake. Beginning with Fergus et al.'s [4] method for camera shake removal, extending the work of Miskin and MacKay [5] with sparse image statistics, blind deconvolution became applicable to real photographs. With Cho and Lee's work [6], the running time of blind deconvolution has become acceptable. These early methods were initially restricted to uniform (space invariant) blur, and later extended to real world spatially varying camera blur [7,8]. Progress has also been made regarding the quality of the blur estimation [9,10], however, these methods are not yet competitive with the runtime of Cho and Lee's approach.

Technical Contributions: Our main technical contributions are as follows:

- (a) we design a class of PSF families containing realistic optical aberrations, via a set of suitable symmetry properties,
- (b) we represent the PSF basis using an orthonormal basis to improve conditioning, and allow for direct PSF estimation,
- (c) we avoid calibration to specific camera lens combinations by proposing a blind approach for inferring the PSFs, widening the applicability to any photographs (e.g., with missing lens information such as historical images) and avoiding cumbersome calibration steps,
- (d) we extend blur estimation to multiple color channels to remove chromatic aberrations as well, and finally
- (e) we present experimental results showing that our approach is competitive with non-blind approaches.

Spatially varying point spread functions

3



Fig. 1. Optical aberration as a forward model.

Optical aberrations cause image blur that is spatially varying across the image. As such they can be modeled as a non-uniform point spread function (PSF), for which Hirsch et al. [11] introduced the *Efficient Filter Flow* (EFF) framework,

$$y = \sum_{r=1}^{R} a^{(r)} * \left(w^{(r)} \odot x \right),$$
 (1)

where x denotes the ideal image and y is the image degraded by optical aberration. In this paper, we assume that x and y are discretely sampled images, i.e., x and y are finite-sized matrices whose entries correspond to pixel intensities. $w^{(r)}$ is a weighting matrix that masks out all of the image x except for a local patch by Hadamard multiplication (symbol \odot , pixel-wise product). The r-th patch is convolved (symbol *) with a local blur kernel $a^{(r)}$, also represented as a matrix. All blurred patches are summed up to form the degraded image. The more patches are considered (R is the total number of patches), the better the approximation to the true non-uniform PSF. Note that the patches defined by the weighting matrices $w^{(r)}$ usually overlap to yield smoothly varying blurs. The weights are chosen such that they sum up to one for each pixel. In [11] it is shown that this forward model can be computed efficiently by making use of the short-time Fourier transform.

4 An EFF basis for optical aberrations

Since optical aberrations lead to image degradations that can be locally modeled as convolutions, the EFF framework is a valid model. However, not all blurs expressible in the EFF framework do correspond to blurs caused by optical aberrations. We thus define a PSF basis that constrains EFF to physically plausible PSFs only.

To define the basis we introduce a few notions. The image y is split into overlapping patches, each characterized by the weights $w^{(r)}$. For each patch, the symbol l_r denotes the line from the patch center to the image center, and d_r the length of line l_r , i.e., the distance between patch center and image center. We assume that local blur kernels $a^{(r)}$ originating from optical aberrations have the following properties:

- (a) Local reflection symmetry: a local blur kernel $a^{(r)}$ is reflection symmetric with respect to the line l_r .
- (b) **Global rotation symmetry:** two local blur kernels $a^{(r)}$ and $a^{(s)}$ at the same distance to the image center (i.e., $d_r = d_s$) are related to each other by a rotation around the image center.
- (c) **Radial behavior:** along a line through the image center, the local blur kernels change smoothly. Furthermore, the maximum size of a blur kernel is assumed to scale linearly with its distance to the image center.

Note that these properties are compromises that lead to good approximations of real-world lens aberrations.²

For two dimensional blur kernels, we represent the basis by K basis elements b_k each consisting of R local blur kernels $b_k^{(1)}, \ldots, b_k^{(R)}$. Then the actual blur

 $^{^2}$ Due to issues such as decentering, real world lenses may not be absolutely rotationally symmetric. Schuler et al.'s exemplar of the Canon 24mm f/1.4 (see below) exhibits PSFs that deviate slightly from the local reflection symmetry. The assumption, however, still turns out to be useful in that case.

kernel $a^{(r)}$ can be represented as linear combinations of basis elements,

$$a^{(r)} = \sum_{k=1}^{K} \mu_k b_k^{(r)}.$$
 (2)

To define the basis elements we group the patches into overlapping groups, such that each group contains all patches inside a certain ring around the image center, i.e., the center distance d_r determines whether a patch belongs to a particular group. Basis elements for three example groups are shown Figure 2. All patches inside a group will be assigned similar kernels. The width and the overlap of the rings determine the amount of smoothness between groups (see property (c) above).

For a single group we define a series of basis elements as follows. For each patch in the group we generate matching blur kernels by placing a single delta peak inside the blur kernel and then mirror the kernel with respect to the line l_r (see, Figure 3). For patches not in the current group (i.e., in the current ring), the corresponding local blur kernels are zero. This generation process creates basis elements that fulfill the symmetry properties listed above. To increase smoothness of the basis and avoid effects due to pixelization, we place little Gaussian blurs (standard deviation 0.5 pixels) instead of delta peaks.

											+	٠	٠	+					٠	٠	+			+	٠	+	
		+	٠	٠	+					*	٠	٠	٠	٠	٠			٠	٠	٠					+	٠	•
	٠	٠	٠	٠	٠	+			+	٠	+			+	٠	+		٠	٠							٠	٠
	٠	•			٠	+			*	٠					٠	٠		٠	٠							+	*
	٠	•			٠	٠			٠	٠					٠	٠		٠	٠							+	٠
	٠	٠	٠	٠	٠	+			+	٠	+			+	٠	٠		٠	٠							٠	+
		•	٠	٠	+					٠	٠	٠	٠	٠	٠			٠	٠	+					+	٠	
											+	•	•	*					٠	٠	+			+	٠		

Fig. 2. Three example groups of patches, each forming a ring.



Fig. 3. Shifts to generate basis elements for the middle group of Figure 2.

5 An orthonormal EFF basis



Fig. 4. SVD spectrum of a typical basis matrix B with cut-off.

The basis elements constrain possible blur kernels to fulfill the above symmetry and smoothness properties. However, the basis is overcomplete and direct projection on the basis is not possible. Therefore we approximate it with an orthonormal one. To explain this step with matrices, we reshape each basis element as a column vector by vectorizing (operator vec) each local blur kernel $b_k^{(r)}$ and stacking them for all patches r:

$$b_k = \left[\left[\operatorname{vec} b_k^{(1)} \right]^\mathsf{T} \dots \left[\operatorname{vec} b_k^{(R)} \right]^\mathsf{T} \right]^\mathsf{T}.$$
 (3)

Let B be the matrix containing the basis vectors b_1, \ldots, b_K as columns. Then we can calculate the singular value decomposition (SVD) of B,

$$B = USV^{\mathsf{T}}.$$
 (4)

with S being a diagonal matrix containing the singular values of B. Figure 4 shows the SVD spectrum and the chosen cut-off of some typical basis matrix B, with approximately half of the eigenvalues being below numerical precision.

We define an orthonormal EFF basis Ξ that is the matrix that consists of the column vectors of U that correspond to large singular values, i.e., that contains the relevant left singular vectors of B. Properly chopping the column vectors of Ξ into shorter vectors one per patch and reshaping those back to the blur kernel, we obtain an orthonormal basis $\xi_k^{(r)}$ for the EFF framework that is tailored to optical aberrations. This representation can be plugged into the EFF forward model in Eq. (1),

$$y = \mu \diamond x := \sum_{r=1}^{R} \left(\sum_{\mu=1}^{K} \mu_k \xi_k^{(r)} \right) * \left(w^{(r)} \odot x \right).$$
 (5)

Note that the resulting forward model is linear in the parameters μ .

6 Blind deconvolution with chromatic shock filtering

Having defined a PSF basis, we perform blind deconvolution by extending [6] to our non-uniform blur model (5) (similar to [13,8]). However, instead of considering only a gray-scale image during PSF estimation, we are processing the



Fig. 5. Chromatic shock filter removes color fringing (adapted from [12]).

full color image. This allows us to better address chromatic aberrations by an improved *shock filtering* procedure that is tailored to color images: the color channels x_R , x_G and x_B are shock filtered separately but share the same sign expression depending only on the gray scale image z:

$$\begin{aligned} x_R^{t+1} &= x_R^t - \Delta t \cdot \operatorname{sign}(z_{\eta\eta}^t) |\nabla x_R^t| \\ x_G^{t+1} &= x_G^t - \Delta t \cdot \operatorname{sign}(z_{\eta\eta}^t) |\nabla x_G^t| \qquad \text{with } z^t = (x_R^t + x_G^t + x_B^t)/3 \\ x_B^{t+1} &= x_B^t - \Delta t \cdot \operatorname{sign}(z_{\eta\eta}^t) |\nabla x_B^t| \end{aligned}$$
(6)

where $z_{\eta\eta}$ denotes the second derivative in the direction of the gradient. We call this extension *chromatic shock filtering* since it takes all three color channels simultaneously into account. Figure 5 shows the reduction of color fringing on the example of Osher and Rudin [12] adapted to three color channels.

Combining the forward model $y = \mu \diamond x$ defined above and the chromatic shock filtering, the PSF parameters μ and the image x (initialized by y) are estimated by iterating over three steps:

- (a) **Prediction step:** the current estimate x is first denoised with a bilateral filter, then edges are emphasized with chromatic shock filtering and by zeroing flat gradient regions in the image (see [6] for further details). The gradient selection is modified such that for every radius ring the strongest gradients are selected.
- (b) **PSF estimation:** if we work with the overcomplete basis B, we would like to find coefficients τ that minimize the regularized fit of the gradient images ∂y and ∂x ,

$$\left\|\partial y - \sum_{r=1}^{R} (B^{(r)}\tau) * (w^{(r)} \odot \partial x)\right\|^{2} + \alpha \sum_{r=1}^{R} \left\|\partial B^{(r)}\tau\right\|^{2} + \beta \sum_{r=1}^{R} \left\|B^{(r)}\tau\right\|^{2}$$
(7)

where $B^{(r)}$ is the matrix containing the basis elements for the *r*-th patch. Note that τ is the same for all patches. This optimization can be performed iteratively. The regularization parameters α and β are set to 0.1 and 0.01, respectively.

However, the iterations are costly, and we can speed up things by using the orthonormal basis Ξ . The blur is initially estimated *unconstrained* and then projected onto the orthonormal basis. In particular, we first minimize the

fit of the general EFF forward model (without the basis) with an additional regularization term on the local blur kernels, i.e., we minimize

$$\left\|\partial y - \sum_{r=1}^{R} a^{(r)} * (w^{(r)} \odot \partial x)\right\|^{2} + \alpha \sum_{r=1}^{R} \left\|\partial a^{(r)}\right\|^{2} + \beta \sum_{r=1}^{R} \left\|a^{(r)}\right\|^{2}$$
(8)

This optimization problem is approximately minimized using a single step of direct deconvolution in Fourier space, i.e.,

$$a^{(r)} \approx C_r^{\mathsf{T}} \mathcal{F}^{\mathsf{H}} \frac{\overline{\mathcal{F}Z_{\partial x} \partial x} \odot (\mathcal{F}E_r \operatorname{Diag}(w^{(r)}) Z_y y)}{|\mathcal{F}Z_{\partial x} \partial x|^2 + \alpha |\mathcal{F}Z_l l|^2 + \beta} \quad \text{for all } r.$$
(9)

where $l = [-1, 2, -1]^{\mathsf{T}}$ denotes the discrete Laplace operator, \mathcal{F} the discrete Fourier transform, and $Z_{\partial x}$, Z_y , Z_l , C_r and E_r appropriate zero-padding and cropping matrices. |u| denotes the entry-wise absolute value of a complex vector u, \overline{u} its entry-wise complex conjugate. The fraction has to be implemented pixel-wise.

Finally, the resulting unconstrained blur kernels $a^{(r)}$ are projected onto the orthonormal basis Ξ leading to the estimate of the blur parameters μ .

(c) **Image estimation:** For image estimation given the blurry image y and blur parameters μ , we apply Tikhonov regularization with $\gamma = 0.01$ on the gradients of the latent image x, i.e.

$$\left\|y - \mu \diamond x\right\|^2 + \gamma \left\|\partial x\right\|^2. \tag{10}$$

As shown in [8], this expression can be approximately minimized with respect to x using a single step of the following direct deconvolution:

$$x \approx \mathcal{N} \odot \sum_{r} C_{r}^{\mathsf{T}} \mathcal{F}^{\mathsf{H}} \frac{\overline{\mathcal{F}Z_{b}\Xi\mu} \odot (\mathcal{F}E_{r}\operatorname{Diag}(w^{(r)}) Z_{y}y)}{|\mathcal{F}Z_{b}\Xi\mu|^{2} + \gamma|\mathcal{F}Z_{l}l|^{2}}.$$
 (11)

where $l = [-1, 2, -1]^{\mathsf{T}}$ denotes the discrete Laplace operator, \mathcal{F} the discrete Fourier transform, and Z_b , Z_y , Z_l , C_r and E_r appropriate zero-padding and cropping matrices. |u| denotes the entry-wise absolute value of a complex vector u, \overline{u} its entry-wise complex conjugate. The fraction has to be implemented pixel-wise. The normalization factor \mathcal{N} accounts for artifacts at patch boundaries which originate from windowing (see [8]).

Similar to [6] and [8] the algorithm follows a coarse-to-fine approach. Having estimated the blur parameters μ we use a non-uniform version of Krishnan and Fergus' approach [14,8] for the non-blind deconvolution to recover a high-quality estimate of the true image. For the x-sub problem we use the direct deconvolution formula (11).

7 Implementation and running times

The algorithm is implemented on a Graphics Processing Unit (GPU) in Python using PyCUDA³. All experiments were run on 3.0GHz Intel Xeon with an

³ http://mathema.tician.de/software/pycuda

NVIDIA Tesla C2070 GPU with 6GB of memory. The basis elements generated as detailed in Section 4 are orthogonalized using the SVDLIBC library⁴. Calculating the SVD for the occurring large sparse matrices can require a few minutes of running times. However, the basis is independent of the image content, so we can compute the orthonormal basis once and reuse it. Table 1 reports the running times of our experiments for both PSF and final non-blind deconvolution along with the EFF parameters and image dimensions. In particular, it shows that using the orthonormal basis instead of the overcomplete one improves the running times by a factor of about six to eight.

image	image dims	local blur	patches	using B	using Ξ	NBD
bridge	$2601{\times}1733$	19×19	10×8	127 sec	16 sec	1.4 sec
bench	1097×730	81×81	10×6	$85 \ sec$	14 sec	$0.7 \sec$
historical	2191×1464	29×29	10×6	103 sec	13 sec	$1.0 \sec$
facade	2817×1877	29×29	12×8	166 sec	21 sec	$1.7 \sec$
	(a)	(b)	(\mathbf{c})	(d)	(e)	(f)

Table 1. (a) Image sizes, (b) size of the local blur kernels, (c) number of patches horizontally and vertically, (d) runtime of PSF estimation using the overcomplete basis B (see Eq. (7)), (e) runtime of PSF estimation using the orthonormal basis Ξ (see Eq. (8)) as used in our approach, (f) runtime of the final non-blind deconvolution.

8 Results

In the following, we show results on real photos and do a comprehensive comparison with other approaches for removing optical aberrations. Image sizes and blur parameters are shown in Table 1.

8.1 Schuler et al.'s lens 120mm.

Schuler et al. show deblurring results on images taken with a lens that consists only of a single element, thus exhibiting strong optical aberrations, in particular coma. Since their approach is non-blind, they measure the non-uniform PSF with a point source and apply non-blind deconvolution. In contrast, our approach is blind and is directly applied to the blurry image.

To better approximate the large blur of that lens, we additionally assume that the local blurs scale linearly with radial position, which can be easily incorporated into our basis generation scheme. For comparison, we apply Photoshop's "Smart Sharpening" function for removing lens blur. It depends on the blur size and the amount of blur, which are manually controlled by the user. Thus we call this method *semi-blind* since it assumes a parametric form. Even though we choose its parameters carefully, we are not able to obtain comparable results.

⁴ http://tedlab.mit.edu/~dr/SVDLIBC/

Comparing our blind method against the non-blind approach of [3], we observe that our estimated PSF matches their measured PSFs rather well (see Figure 7). However, surprisingly we are getting an image that may be considered sharper. The reason could be over-sharpening or a less conservative regularization in the final deconvolution; it is also conceivable that the calibration procedure used by [3] is not sufficiently accurate. Note that neither DXO nor Kee et al.'s approach can be applied, lacking calibration data for this lens.

8.2 Canon 24mm f/1.4.

The PSF constraints we are considering assume local axial symmetry of the PSF with respect to the radial axis. For a Canon 24mm f/1.4 lens also used in [3], this is not exactly fulfilled, which can be seen in the inset in Figure 8. The wings of the green blur do not have the same length. Nonetheless, our blind estimation with enforced symmetry still approximates the PSF shape well and yields a comparable quality of image correction. We stress the fact that this was obtained blindly in contrast to [3].

8.3 Kee et al.'s image

Figure 9 shows results on an image taken from Kee et al. [2]. The close-ups reveal that Kee's *non-blind* approach is slightly superior in terms of sharpness and noise-robustness. However, our *blind* approach better removes chromatic aberration. A general problem of methods relying on a prior calibration is that optical aberrations depend on the wavelength of the transmitting light continuously: an approximation with only a few (generally three) color channels therefore depends on the lighting of the scene and could change if there is a discrepancy between the calibration setup and a photo's lighting conditions. This is avoided with a blind approach.

We also apply "DxO Optics Pro 7.2" to the blurry image. DXO uses a database for combinations of cameras/lenses. While it uses calibration data, it is not clear whether it additionally infers elements of the optical aberration from the image. For comparison, we process the photo with the options "chromatic aberrations" and "DxO lens softness" set to their default values. The result is good and exhibits less noise than the other two approaches (see Figure 9, however, it is not clear if an additional denoising step is employed by the software.

8.4 Historical Images

A blind approach to removing optical aberrations can also be applied to historical photos, where information about the lens is not available. The left column of Figure 10 shows a photo (and some detail) from the Library of Congress archive that was taken around 1940^5 . Assuming that the analog film used has a sufficiently linear light response, we applied our blind lens correction method

⁵ http://www.loc.gov/pictures/item/fsa1992000018/PP/



Blurred image

Our approach (blind)



Adobe's "Smart Sharpen" (semi-blind)

Schuler et al. [3] (non-blind)

 ${\bf Fig.}\,{\bf 6.}$ Schuler et al.'s lens. Full image and lower left corner.

X	8	*	*	•	•	ø	ø	e di		
			46	R	0	ø	ø	ø	v	s
				×	0	0	0	P	φ	

(a) blindly estimated by our approach

(b) measured by Schuler et al. [3]

Fig. 7. Schuler et al.'s lens. Lower left corner of the PSF.



Fig. 8. Canon 24 mm f1/4 lens. Shown is the upper left corner of the image. PSF inset is three times the original size.



Fig. 9. Comparison between our blind approach and two non-blind approaches of Kee et al.[2] and DXO.

and obtained a sharper image. However, the blur appeared to be small, so algorithms like Adobe's "Smart Sharpen" also give reasonable results. Note that neither DXO nor Kee et al.'s approach can be applied here since lens data is not available.



Blurry image

(blind) (semi-blind)

Fig. 10. Historical image from 1940.

9 Conclusion

We have proposed a method to blindly remove spatially varying blur caused by imperfections in lens designs, including chromatic aberrations. Without relying on elaborate calibration procedures, results comparable to non-blind methods can be achieved. By creating a suitable orthonormal basis, the PSF is constrained to a class that exhibits the generic symmetry properties of lens blurs, while fast PSF estimation is possible.

9.1 Limitations

Our assumptions about the lens blur are only an approximation for lenses with poor rotation symmetry.

The image prior used in this work is only suitable for natural images, and is hence content specific. For images containing only text or patterns, this would not be ideal.

9.2 **Future Work**

While it is useful to be able to infer the image blur from a single image, it does not change for photos taken with the same lens settings. On the one hand, this

implies that we can transfer the PSFs estimated for these settings for instance to images where our image prior assumptions are violated. On the other hand, it suggests the possibility to improve the quality of the PSF estimates by utilizing a substantial database of images.

Finally, while optical aberrations are a major source of image degradation, a picture may also suffer from motion blur. By choosing a suitable basis, these two effects could be combined. It would also be interesting to see if non-uniform motion deblurring could profit from a direct PSF estimation step as introduced in the present work.

References

- Joshi, N., Szeliski, R., Kriegman, D.: PSF estimation using sharp edge prediction. In: Proc. IEEE Conf. Comput. Vision and Pattern Recognition. (June 2008) 1, 2
- Kee, E., Paris, S., Chen, S., Wang, J.: Modeling and removing spatially-varying optical blur. In: Proc. IEEE Int. Conf. Computational Photography. (2011) 1, 2, 10, 12
- Schuler, C., Hirsch, M., Harmeling, S., Schölkopf, B.: Non-stationary correction of optical aberrations. In: Proc. IEEE Intern. Conf. on Comput. Vision. (2011) 1, 2, 10, 11, 12
- 4. Fergus, R., Singh, B., Hertzmann, A., Roweis, S., Freeman, W.: Removing camera shake from a single photograph. ACM Trans. Graph. 25 (2006) 3
- 5. Miskin, J., MacKay, D.: Ensemble learning for blind image separation and deconvolution. Advances in Independent Component Analysis (2000) 3
- Cho, S., Lee, S.: Fast Motion Deblurring. ACM Trans. Graph. 28(5) (2009) 3, 6, 7, 8
- Whyte, O., Sivic, J., Zisserman, A., Ponce, J.: Non-uniform deblurring for shaken images. In: Proc. IEEE Conf. Comput. Vision and Pattern Recognition. (2010) 3
- 8. Hirsch, M., Schuler, C., Harmeling, S., Schölkopf, B.: Fast removal of non-uniform camera shake. In: Proc. IEEE Intern. Conf. on Comput. Vision. (2011) 3, 6, 8
- Krishnan, D., Tay, T., Fergus, R.: Blind deconvolution using a normalized sparsity measure. In: Proc. IEEE Conf. Comput. Vision and Pattern Recognition. (2011) 3
- Levin, A., Weiss, Y., Durand, F., Freeman, W.: Efficient marginal likelihood optimization in blind deconvolution. In: Proc. IEEE Conf. Comput. Vision and Pattern Recognition, IEEE (2011) 3
- Hirsch, M., Sra, S., Schölkopf, B., Harmeling, S.: Efficient filter flow for spacevariant multiframe blind deconvolution. In: Proc. IEEE Conf. Comput. Vision and Pattern Recognition. (2010) 3, 4
- Osher, S., Rudin, L.: Feature-oriented image enhancement using shock filters. SIAM J. Numerical Analysis 27(4) (1990) 7
- Harmeling, S., Hirsch, M., Schölkopf, B.: Space-variant single-image blind deconvolution for removing camera shake. In: Advances in Neural Inform. Processing Syst. (2010) 6
- Krishnan, D., Fergus, R.: Fast image deconvolution using hyper-Laplacian priors. In: Advances in Neural Inform. Process. Syst. (2009) 8