Spot and Edge Feature Based Estimation of Point-Spread Functions for Image Deconvolution

Martin Welk
Institute of Biomedical Image Analysis
UMIT TIROL – Private University for Health Sciences and Health Technology
6060 Hall in Tirol, Austria
martin.welk@umit-tirol.at

Abstract

We consider the extraction of point-spread function (PSF) information for blind image deconvolution from blurred images in a way that preserves phase information, in contrast to using cues like autocorrelation that reveal only spectral information. Our approach is based on extracting suitable feature patches, depending on the type of images either spot highlights or edge segments. We discuss how edge patches in fact constitute a tomographic representation of the PSF. In integrating information from spot or edge patches into a PSF estimate, it is essential to compensate spatial misalignments. We achieve this by an iterative update rule that combines Fourier transformation with a nonlinear intensity transformation to achieve shift invariance. Although designed with the goal of integration into alternating minimisation schemes, the two-step procedure of PSF estimation followed by non-blind deconvolution developed here performs surprisingly well as a fast blind deconvolution method in its own right on suitable image classes.

1 Introduction

Spatially invariant blur in an image can be described as $f = u * h + n$ where f, u are functions on a planar domain representing the observed blurred and unobservable sharp image, respectively, h is a point-spread function (PSF) given as a convolution kernel, and n denotes some additive noise. The task to recover a sharp image can either be a non-blind deconvolution problem in which f and h are given, and u is to be estimated, or a blind deconvolution problem in which only f is known, and both u and h need to be estimated. Whereas both problems are ill-posed, for non-blind deconvolution there exist a wide-range of well-established techniques to regularise the problem sufficiently to obtain high-quality solutions, e.g. [4, 7, 13, 15, 21, 25, 28, 30, 31, 32, 33, 34, 35, 38, 39].

Blind deconvolution is substantially more difficult and a long-standing field of intensive research [1, 2, 9, 14, 18, 29]. Many existing methods are either based on statistical or classic machine-learning frameworks [8, 19] or variational models [3, 5, 6, 20, 36, 42]. End-to-end-trained neural network approaches have been considered [22, 27, 41, 43] but do not dominate the field, probably since they are not well-suited to cope with the strict algebraic condition imposed by the blur model and the difficulties to generalise beyond image classes represented by the training data. However, hybrid approaches where the regulariser of a variational model is formulated as a deep-learning component [23, 41] hold some potential.

Model-based approaches to blind deconvolution mostly fall into two categories: on one hand, two-step procedures in which one first estimates the PSF and applies non-blind deconvolution afterwards, see e.g. [8, 11, 12, 40], and on the other hand approaches that perform a joint estimation of u and h, often in the form of alternating minimisation of a variational functional, e.g. [10, 20, 36, 42]. In some cases a joint estimation is performed but the final estimate of the sharp image is obtained by a separate

The Third Austrian Symposium on AI, Robotics, and Vision (AIROV26).

non-blind deconvolution step using the PSF from the joint estimation, whereas the image estimate from the joint estimation is discarded [9].

Several approaches to extract PSF information from the blurred image either directly [12] or in the form of a regulariser for PSFs [20] heavily rely on the autocorrelation of the blurred image. Generally the autocorrelation of the blurred image f, in the form of its convolution with its point-reflected counterpart $f^*(\boldsymbol{x}) := f(-\boldsymbol{x})$, enters many methods naturally via the gradient descent of the relevant data fidelity terms. Autocorrelation is a good source for the Fourier spectrum of the point spread function but entirely loses the phase information, and thus any spatial asymmetry of h, because one has (neglecting noise) $f * f^* = (u * u^*) * (h * h^*)$, and the Fourier transform $\mathcal{F}(h * h^*)$ equals $|\mathcal{F}(h)|^2$. Full information on h then requires some kind of phase retrieval, either explicitly [12] or via the outer iteration of alternating minimisation [20] which makes it necessary to perform many computationally expensive iterations until the asymmetric structure of h is established. Indeed, papers reporting fast and efficient blind deconvolution often focus on symmetric PSFs.

In contrast, f itself often contains the asymmetry of the PSF in ways fairly obvious to the human observer. For example, night-view images (see Fig. 1 later in this paper) contain impulse responses from street lights that clearly show the PSF shape. The intensity profiles of blurred edges are also clearly asymmetric when the PSF is asymmetric w.r.t. the edge-normal direction. If PSF information could be extracted from the blurred image directly in a way that preserves the phase, feeding this into the PSF estimation or alternating minimisation could speed up these methods considerably.

In this paper, we explore therefore how this can be done. Our overarching goal is to allow the integration of such information into alternating minimisation, but for the scope of this paper we focus on designing a two-step procedure, i.e. direct PSF estimation for subsequent non-blind deconvolution. For some image classes, particularly those rich in spot highlights, the two-step method itself already yields competitive results.

Our contribution We design methods to extract spot and edge patches suitable for PSF reconstruction from blurred images. To this end, we fit foreground-background templates to image patches using the ratio of inter-segment and intra-segment variances as a criterion analogous to Otsu segmentation [26]. We devise an iterative update rule to integrate PSF phase information from patches into an overall PSF estimate in a way insensitive w.r.t. misalignments of patches. This is achieved by a shift-invariant combination of Fourier transforms and nonlinear intensity transforms. We discuss the relation between blurred edge patches and PSFs which can be cast in the form of a tomographic (Radon) transform. We demonstrate the validity and viability of the approach by proof-of-concept experiments on synthetic and real-world image examples.

Structure of the paper Section 2 describes the extraction of three types of features, spot, edge and crossing edge features, from image patches and the process of estimating PSFs from these, along with some analysis. Experiments on all three feature types are presented in Section 3. A short discussion and outlook to future work in Section 4 concludes the paper.

2 Patch-Based PSF Estimation

2.1 Spot Features

In some classes of images such as astronomical images containing isolated stars, or urban night views containing isolated highlights such as street lanterns, these highlights act as approximate Dirac impulses probing the point-spread function, i.e. the blurred image of such a highlight yields an approximation of the PSF. This approach can be used to manually extract suitable PSF candidates, enabling subsequent non-blind deconvolution; it was also used in the 1990s to extract PSFs from blurred images produced by the Hubble space telescope in its early phase before the insertion of correcting optics [16].

For a more robust approximation of the PSF it is, however, desirable to rely not on a single highlight but instead to combine information of several of these to reduce effects of noise, background structures, and saturation. Moreover, automated extraction of suitable patches will generally yield some patches approximating the correct PSF along with less accurate ones resulting from proximity of several highlights, strong background, or occasional moving objects (like car headlights). Thus, we will aim at automatically extracting a number of spot patches and estimate the PSF from all of them.

Extraction of spot patches We use a sliding window with radius r corresponding to the desired PSF size with concentric foreground and background regions; the foreground part is a discretised disc of radius $0.6\,r$ whereas the background part is a discretised circular annulus of inner radius $0.8\,r$ and outer radius r, yielding approximately equal foreground and background pixel counts n_f, n_b. We now borrow an optimisation criterion from Otsu segmentation [26] (where, however, foreground and background are distinguished by their intensities instead of their locations): For each centre pixel location we compute the mean value μ_f and variance σ_f^2 of the intensities in the foreground region, and analogously μ_b, σ_b^2 of the background region. Denoting by $\mu = (n_\mathrm{f}\mu_\mathrm{f} + n_\mathrm{b}\mu_\mathrm{b})/(n_\mathrm{f} + n_\mathrm{b})$ the overall mean, we compute the ratio of inter-region variance and intra-region variance as $q = (n_\mathrm{f}(\mu_\mathrm{f} - \mu)^2 + n_\mathrm{b}(\mu_\mathrm{b} - \mu)^2)/(n_\mathrm{f}\sigma_\mathrm{f}^2 + n_\mathrm{b}\sigma_\mathrm{b}^2)$. Maxima of q are considered candidates for spot locations. These candidates are pruned using thresholds on contrast between the foreground and background mean intensities (discarding patches with less than $1/4$ the greatest contrast among patch candidates) and moments of inertia (discarding patches with more than 2 times the least momen of inertia among candidates), and remaining patches are finally normalised. The PSF update iteration (1) is initialised with a Gaussian of standard deviation $0.5\,r$, and typically converges numerically after a few hundred iterations. The final PSF is pruned by replacing entries below a small positive threshold (in our experiments, about 0.2 times the maximal value) with zero.

Shift-invariant superposition Whereas spot-feature patches provide more or less reasonable approximations to the PSF, the exact position of the PSF in each of them is unknown. If a highly accurate PSF estimate were to be derived from those patches alone, a sub-pixel accurate alignment would be necessary. This poses a challenge analogous to the application of point-based measures like PSNR for the quality assessment of blind deconvolution results, a problem that has been discussed e.g. in [17, 37] but for which, unfortunately, no fully satisfactory solution is available so far. We made an attempt to estimate sub-pixel displacements by a Fourier-based procedure similar to [17] but averaging the so aligned patches yielded a very blurred PSF estimate.

We follow therefore a different approach. Starting from some well-centred initial estimate (like a Gaussian), we refine the PSF not directly with the spot patches but using quantities that are invariant under shifts and therefore insensitive to any misalignment.

To construct a spectral shift-invariant quantity, we recall the Fourier shift theorem: shifting an image patch p by some displacement vector $\boldsymbol{v}$ multiplies its Fourier transform $\hat{p} = \mathcal{F}(p)$ by multiplication with a frequency-dependent phase function, $\hat{p}(\boldsymbol{\omega}) \mapsto \hat{p}(\boldsymbol{\omega})\exp(2\pi\mathrm{i}\langle\boldsymbol{v},\boldsymbol{\omega}\rangle)$. As a consequence, for any frequencies $\boldsymbol{\xi}$, $\boldsymbol{\eta}$ the quantity $\hat{p}(\boldsymbol{\xi})\hat{p}(\boldsymbol{\eta})\bar{\hat{p}}(\boldsymbol{\xi}+\boldsymbol{\eta})$ is shift-invariant.

Given a discrete PSF estimate $\boldsymbol{h} = (h_{\boldsymbol{x}})$ where $\boldsymbol{x}$ runs over $\mathbb{Z}_{m_x} \times \mathbb{Z}_{m_y}$ (thus, with periodic boundary conditions) and a set $\{\boldsymbol{p}^k \mid k = 1,\dots,K\}$ of discrete patches $\boldsymbol{p}^k = (p^k_{\boldsymbol{x}})$, we use therefore the conditions $\hat{h}^{\mathrm{new}}_{\boldsymbol{\xi}} = \hat{h}_{\boldsymbol{\eta}}\hat{h}_{\boldsymbol{\xi}-\boldsymbol{\eta}}\bar{\hat{p}}^k_{\boldsymbol{\eta}}\bar{\hat{p}}^k_{\boldsymbol{\xi}-\boldsymbol{\eta}}\hat{p}^k_{\boldsymbol{\xi}}$ where $\boldsymbol{h}$ and $\boldsymbol{p}^k$ are normalised to $\hat{h}_{\boldsymbol{0}} = \hat{p}^k_{\boldsymbol{0}} = 1$ as an overcomplete system of equations from which $\boldsymbol{h}^{\mathrm{new}}$ is obtained as a least-squares estimate, which yields with an additional regularisation the update rule

$$\hat{h}^{\mathrm{new}}_{\boldsymbol{\xi}} = \frac{\frac{1}{K}\sum_k \mathcal{F}\big((h * (p^k)^*)^2\big)_{\boldsymbol{\xi}}\hat{p}^k_{\boldsymbol{\xi}} + \theta\hat{h}_{\boldsymbol{\xi}}}{\mathcal{F}\big((h * h^*)^2\big)_{\boldsymbol{\xi}} + \theta} \,. \tag{1}$$

This iteration can also be understood as a discrete minimisation procedure for the cost function

$$R_{hp}[h] = \int_\Omega (h * h^*)^3 - \frac{1}{K}\sum_k \big(h * (p^k)^*\big)^3 \,\mathrm{d}\boldsymbol{x} \,. \tag{2}$$

Note the combination of non-linear operations with Fourier transforms and convolutions.

Originally designed as a tool for phase retrieval for PSF spectra obtained from autocorrelation-related approaches this update rule performs surprisingly well as a stand-alone PSF estimator as we will demonstrate in the experimental section.

2.2 Edge Features

Unfortunately, the retrieval of PSFs as impulse responses as discussed before is restricted to images containing well-defined isolated impulses. In absence of these, such as in blurred daylight images, other image structures need to be exploited for PSF estimation. The most obvious candidate are blurred images of sharp edges.

Point-spread function tomography Let us turn to analyse what information about a PSF can be expected from an edge patch, assuming it shows the blurred image of a straight edge segment and no other structure within the patch. Modelling sharp edges as step functions, we notice that the underlying sharp patch would be given, up to linear intensity rescaling, by $u(\boldsymbol{x}) = \mathrm{H}(\langle \boldsymbol{x} - \boldsymbol{x}_0, \boldsymbol{v} \rangle)$ with H denoting the Heaviside function, $\boldsymbol{x}$ the location within the patch, $\boldsymbol{x}_0$ a reference point located on the edge, and $\boldsymbol{v}$ an edge-normal unit vector. Taking the directional derivative in edge-normal direction yields the single-layer distribution $\partial_{\boldsymbol{v}} u(\boldsymbol{x}) = \delta(\langle \boldsymbol{x} - \boldsymbol{x}_0, \boldsymbol{v} \rangle)$ where δ denotes the Dirac distribution. Since the derivative and the blur, both being linear operations, commute, applying the directional derivative to the observed blurred image $f = u * h$ yields $\partial_{\boldsymbol{v}} u(\boldsymbol{x}) = ((\delta(\langle \cdot - \boldsymbol{x_0}, \boldsymbol{v} \rangle) * h)(\boldsymbol{x})$ which is an integral of h along the straight line through $\boldsymbol{x}$ perpendicular to $\boldsymbol{v}$. For some fixed location $\boldsymbol{x}_0$, the values of $p_{\boldsymbol{v}} h(t) := \partial_{\boldsymbol{v}} u(\boldsymbol{x}_0 + t\boldsymbol{v})$ along the edge-normal line $\boldsymbol{x} = \boldsymbol{x}_0 + t\boldsymbol{v}$ therefore form a projection of h onto this line. Projections $p_{\boldsymbol{v}}(t)$ along all directions $\boldsymbol{v}$ within an angular range of π (provided there are sufficiently many edge patches to cover this range of directions) therefore constitutes a Radon transform of h. Invoking the Fourier slice theorem, the Fourier transform of h can then be obtained by combining the 1D Fourier transforms of directional derivatives $\partial_{\boldsymbol{v}} u(\boldsymbol{x}_0 + t\boldsymbol{v})$, where each of these 1D Fourier transforms yields one radial line in the 2D Fourier transform of h. Thus we obtain in fact h by a tomographic reconstruction.

Computation of directional derivatives In the computation of directional derivatives of edge patches it is relevant to minimise additional blur introduced by finite-difference numerics. To this end, we strive to compute the directional derivative via the Fourier domain using the identity $\mathcal{F}(\partial_{\boldsymbol{v}} h)(\boldsymbol{\xi}) = 2\pi \mathrm{i} \langle \boldsymbol{\xi}, \boldsymbol{v} \rangle \mathcal{F}(\boldsymbol{\xi})$. However, we have to work with the discrete Fourier transform on patches, which interprets the patch as a periodic function. The mismatch of intensities between opposite boundaries of the patch then creates additional jumps in the function being differentiated. To avoid this, we use the periodic-plus-smooth image decomposition [24] to decompose the patch f into a periodic function f_{p} which is consistent with the Fourier derivative, and a smooth function f_{s} absorbing the intensity difference between opposite boundaries. The derivative of f is then obtained as the sum of the derivative of f_{p} computed via the Fourier domain and the derivative of f_{s} computed in the spatial domain.

Shift uncertainty Unfortunately, the tomographic reconstruction of h rests on the assumption that the edge patches can be aligned such that all edges go through one and the same centre point $\boldsymbol{x}_0$. As this alignment meets similar difficulties as in the case of spot patches, the direct superposition again needs to be replaced with a shift-invariant approach. The iteration (1) can be used to this end. However, as we will demonstrate in the experimental section, the alignment does not work very well due to the lack of interaction between the different radial lines in the Fourier domain.

Extraction of edge features To compute edge features, we rely on the same Otsu-like criterion as for spot patches but the foreground and background regions within a patch of radius r are now chosen as half-discs separated by a line orthogonal to the local gradient direction, with a narrow neutral zone of width $0.1\,r$ between them. Again we discard patches with low contrast (in this case, below 0.05 times the greatest contrast among candidates). Retained patches are then post-processed by computing the directional derivative in local gradient direction, spatial windowing, suppressing frequencies away from the gradient direction in the Fourier domain and normalisation.

2.3 Crossing Edge Features

In experiments it is observed that the alignment of spot patches via (1) works considerably better than the alignment of the edge features obtained from directional derivatives of edge patches. We attribute this to the fact that spot features offer rich information across the 2D Fourier domain whereas edge features contain only information on largely decoupled 1D lines in the Fourier domain. We are therefore interested in making the information of edge features available in a way that creates stronger interaction across the 2D Fourier domain.

The commutativity of linear operations such as the convolution in the blur model, directional derivatives involved in the edge feature extraction, and linear combination of image patches offers a way to achieve this: Assuming we choose two edge-features with substantially different edge orientations, their superposition will contain information on two radial lines in the Fourier domain. Moreover, such a superposed feature from the blurred image is equivalent to the same superposition

of directional derivatives of the corresponding patches of the latent sharp image, blurred with h. The directional derivatives of sharp edges are sharp straight lines, and their superposition creates a single sharp dot at their intersection, which is similar to a spot highlight. It can therefore be expected that in a large set of such crossing edge features sufficient information will be contained to reconstruct the entire Fourier transform of h in a phase-consistent way.

Note that the superposition should not be extended to more than two edge features because the superposition of the corresponding sharp image patches would contain more than two straight line segments that will in general not have a common intersection point.

Generation of crossing edge features We use the edge features computed as in the previous subsection and linearly superpose each pair of edge features whose orientations enclose an angle between 60 and 120 degrees. This yields a large, highly redundant set of crossing edge features, introducing some inefficiency into the process. For the moment we leave it at that but will consider suitable pruning procedures for the patch set in future work. Again, the PSF estimation via (1) takes a few hundred iterations after which the PSF is pruned using a small positive threshold.

3 Experiments

Computations were performed on a workstation with an Intel i9-14900KF processor (24 cores (8p/16e), clock frequencies 2.4–5.8 GHz), 64 GB main memory, and an Nvidia RTX Pro 4000 graphics card (8,960 cores, 24 GB) running Ubuntu Linux 24.04. Implementations were made in C++ (using gcc 13.3.0) and partially CUDA v13.1.

For the patch detection and patch-based PSF estimation currently a single-threaded CPU implementation is used. The PSF/patch size is set manually as a parameter in all methods. In the RRRL method [35] used for non-blind deconvolution as well as the blind deconvolution method (modified from [36]) used for comparisons all essential parts of the algorithms were parallelised on the GPU. For RRRL, we use a fixed parameter set similar to the values used in [35]; results are not highly sensitive to these parameters.

3.1 Spot Features for Night-View Images: Real-World Example

We demonstrate the reconstruction of a point-spread function on a night-view photograph blurred during acquisition (Fig. 1). In this case, 30 spot patches are found in the image from which a visually convincing PSF is reconstructed. Non-blind deconvolution confirms a high accuracy of the PSF. The patch-based PSF estimation (1000 iterations) takes 87 seconds (CPU, single-threaded); the non-blind deconvolution is performed within a few seconds (GPU). With a future full GPU implementation, the complete blind deconvolution of the three-megapixel colour photograph can be expected to be done within a few seconds.

For comparison we also show results (PSF and deblurred image) obtained by a modification of the variational blind deconvolution method from [36] (modified by a coarse-to-fine approach not present in the original paper) which took about 42 minutes computing time with a full GPU implementation. Despite the significant lower computational effort, the restoration quality of our approach is very similar to the variational method.

3.2 Edge Features: Simplistic Synthetic Examples

To test the validity of the tomographic reconstruction of h as such, we use a specific synthetic setting (Fig. 2): A synthetic test image featuring straight edge segments covering all orientations is blurred with a known PSF, and we prescribe the patch locations exactly on the edges. Instead of the shift-invariant iteration (1) we reconstruct the PSF directly in the Fourier domain by placing the Fourier coefficients from each patch along the respective radial line. After a final pruning of small entries the resulting PSF is visually indistinguishable from the original, and non-blind deconvolution yields a near-perfect sharpened image.

However, despite this encouraging finding, our next experiment (Fig. 3) reveals the severe limitations of edge-feature-based PSF estimation in a more realistic context. The test image contains sharp circles, thus again providing good coverage of the orientation space, and we use the same PSF for

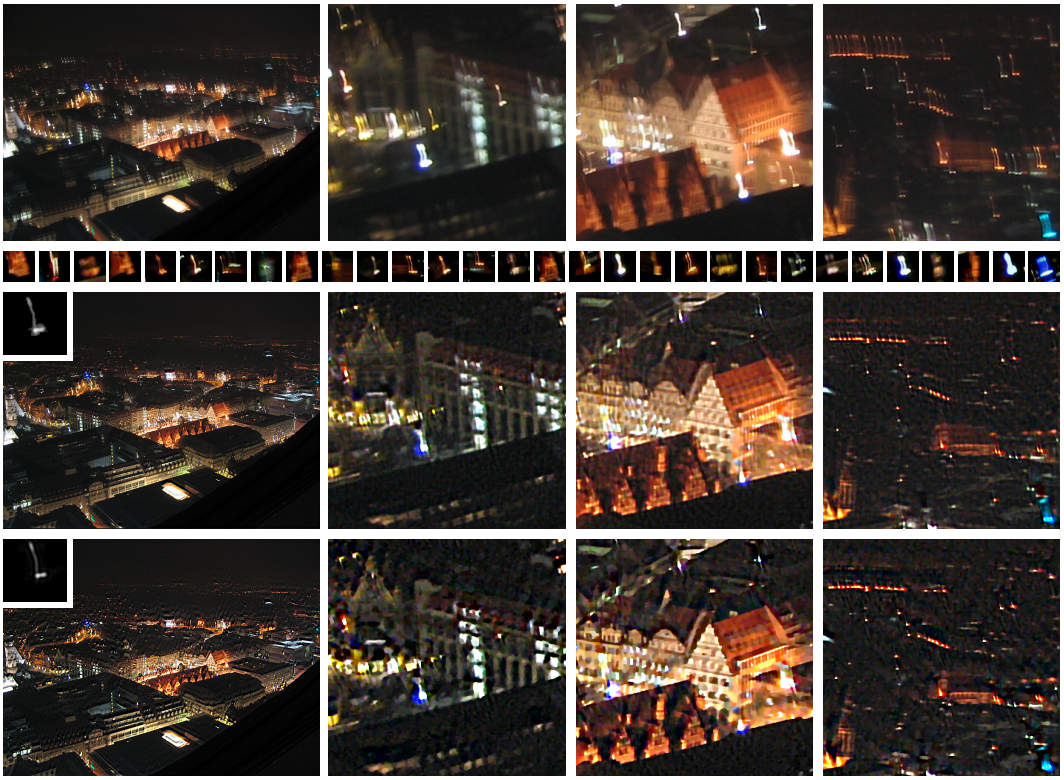

Figure 1: *Top row:* Urban night-view photograph, 2048×1536 pixels, blurred during acquisition, with three zoom-ins, 320×320 pixels each. *Second row:* 30 spot patches, 51×51 pixels each, automatically detected. *Third row:* PSF estimated from spot patches, 51×51 pixels, and deblurred photograph (RRRL, 100 iterations, $\alpha = 0.005$) with three zoom-ins; *Bottom row:* PSF and deblurred photograph obtained by a modified version of the blind deconvolution method from [36] with three zoom-ins. – *Author of original image:* Gregor Peltri.

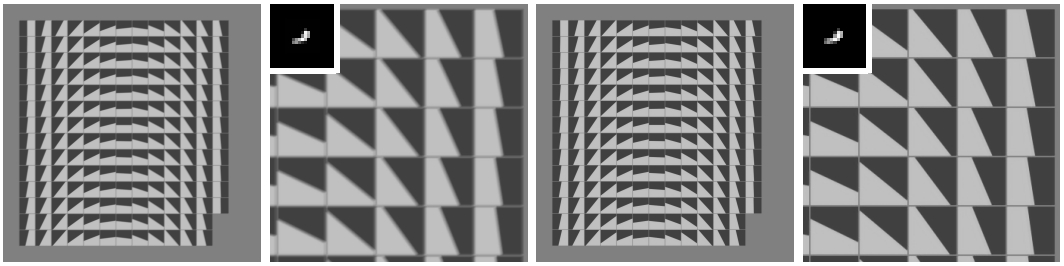

Figure 2: *Left to right:* Synthetic image containing 180 tiles with straight edges in 1 degree angular steps blurred with a known PSF; zoom-in with insert showing the PSF (21×21 pixels); deblurred (RRRL, 100 iterations, $\alpha = 0.005$); zoom-in with insert showing PSF used for deblurring, which was reconstructed from edge patches with exactly prescribed positions centred on the straight edges

blurring as in the previous example. The PSF reconstructed from edge features using (1) looks visually plausible but the result of the subsequent non-blind deconvolution (third frame of Fig. 3) is severely distorted; note that the PSF obviously displaces edge segments dependent on their direction. This is most likely caused by the decoupling of radial components in the Fourier domain: each edge patch provides only information on one radial line of the Fourier transform of h, and the iteration rule (1) is incapable of aligning the phases on different radial lines.

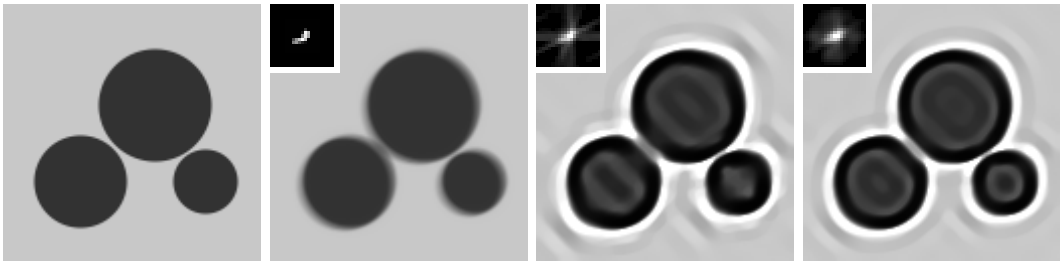

Figure 3: *Left to right:* Synthetic image showing three circles; same image blurred with the PSF shown as insert; deblurred by RRRL (100 iterations, $\alpha = 0.005$) using the PSF estimated using edge patches (shown as insert); deblurred by RRRL (same parameters) using the PSF estimated from crossed edge patches (shown as insert).

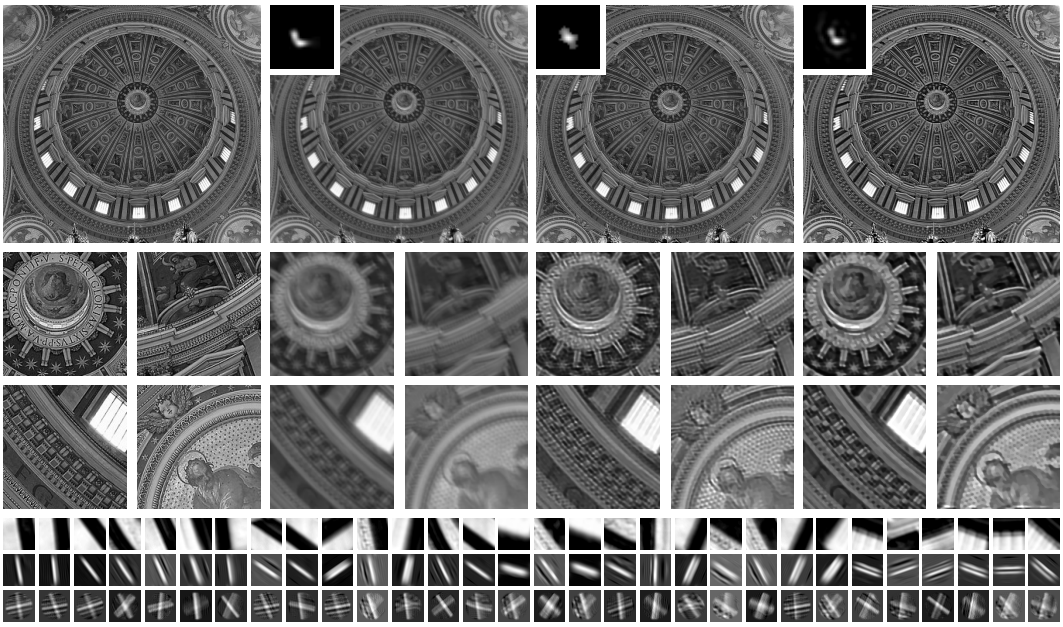

Figure 4: *Top rows, left to right:* Photograph (1280×1181 pixels) with zoom-ins; synthetically blurred using, PSF (31×31 pixels) shown as insert; deblurred (RRRL, 100 iterations, $\alpha = 0.005$), PSF reconstructed from crossed edge patches; deblurred using a modification of the blind deconvolution method from [36]. *Bottom rows:* 30 (from a total of 149) extracted edge patches; corresponding edge features; 30 (from a total of 3386) crossing edge features. – *Source of original image:* https://commons.wikimedia.org/wiki/File:Dome_of_Saint_Peter%27s_Basilica_(Interior).jpg, author: LivioAndronico, licence: CC BY-SA 4.0

3.3 Crossing Edge Features for Daylight Images: Synthetic Example

Our first test for crossing edge features is the continuation of the experiment shown in Fig. 3: The last frame in the figure shows the PSF (inserted) and non-blind deconvolution result obtained with crossing edge features. Indeed, the distortion is reduced (albeit not completely gone) which indicates that crossing-edge features allow for a better alignment of phases in the PSF reconstruction iteration.

Finally, we demonstrate PSF reconstruction from crossing edge features on a synthetically blurred test image based on a real-world daylight photograph (Fig. 4). The PSF estimated from crossing edge features is visually similar to the ground-truth PSF, but somewhat blurred in comparison. Still, the non-blind deconvolution result is favourable and of overall similar visual quality to that obtained by the variational blind deconvolution method, despite the visually more accurate PSF estimate achieved by the latter. Some fine details appear even clearer in our result than in that of the variational blind deconvolution (see e.g. the lower right zoom-in).

Table 1: Shift-compensated PSNR measurements for the blurred and deblurred images from Fig. 4.

Measurement method	Blurred image	Feature-based deblurring	Variational blind deconvolution
Shift-compensated PSNR [17]	16.71 dB	16.78 dB	17.49 dB
Innocence-alignment PSNR [37]	18.29 dB	18.64 dB	18.92 dB

To complement the visual assessment, we also performed shift-compensated PSNR measurements using the methods from [17] (with margin width parameters $T_1 = 16$, $T_2 = 48$) and [37] (with margin width 48), see Table 1. Note that PSNR measurements and visual assessment often diverge since PSNR, being derived from mean square error, is more sensitive to average intensity deviations in larger areas than to the visually important sharpness of details. Moreover, it is important to remark that the method from [37] yields systematically higher PSNR values (as it computes a lower bound for the underlying MSE).

Regarding computation times, the variational coarse-to-fine method is faster here with about 4 min, whereas the patch-based estimation takes about 20 min but this is mainly due to the overly large patch set and the single-threaded CPU computation currently used for it.

4 Discussion and Outlook

Our experiments confirm on a proof-of-concept level that phase-preserving extraction of PSF information from patches of blurred images is possible. For urban night-view images with sufficiently many spot highlights the two-phase approach using PSF estimation from spot patches with subsequent non-blind deconvolution performed surprisingly good and can stand as a simple and fast blind deconvolution method in its own right. A similar method based on crossing edge features for blurred daylight images yields promising results coming close to those of a variational blind deconvolution method used for comparison, but requires further work e.g. to improve its algorithmic efficiency. For pure edge features, the lack of interaction between radial directions in the Fourier domain prevents the approach from working in a stand-alone manner.

Integration into blind deconvolution frameworks However, our approach is well-suited to serve as an additional component in alternating minimisation frameworks for blind deconvolution. First, the patch-based estimate can be used to initialise the PSF in such methods in a way that already contains valuable phase information, which can be expected to speed up the phase retrieval which is often slow as it is mostly achieved in the outer loop of alternating minimisation. Second, a patch-based component can be integrated into the PSF estimation step of an alternating minimisation, either by pre- or post-processing the PSF estimates with the update rule (1), or by using the cost term from (2) in the variational framework itself.

An obvious limitation of our approach is that it requires the presence of sufficiently many suitable features, either spot highlights or edge segments covering a wide range of directions, in the input image, thus restricting its effectivity as a stand-alone method to suitable image classes. Built into an alternating minimisation framework, it will not speed up the latter in the case of unsuitable images.

Future work A more rigorous evaluation with larger image sets and more extensive quantitative measurements will be important for a fuller assessment of the performance of the approach. As the purpose of the present paper is mainly to introduce the theoretical basis of the approach and demonstrate its feasibility, this is left for future work. Note that the quantitative evaluation of blind deconvolution results generally meets the problem of underdetermination of the spatial location of PSF and sharpened image, see [14, 17, 37].

Ongoing and future work on the algorithmic level is directed at improving the efficiency of the crossing edge features by pruning the large set of crossed features, and on detecting automatically which feature type should be used in a given image. A full GPU (CUDA) implementation of the patch extraction and PSF estimation procedures is underway. Moreover, we will investigate the integration of the approach into alternating minimisation approaches such as [9, 36] along the lines mentioned.

Last but not least, we work on theoretical foundations for an extension to more general types of features in order to overcome the limitation to specific image classes.

Acknowledgements

Funding disclosure No third-party funding was involved in this work.

References

[1] M. S. C. Almeida and L. B. Almeida. Blind and semi-blind deblurring of natural images. *IEEE Transactions on Image Processing*, 19(1):36–52, 2010.

[2] G. R. Ayers and J. C. Dainty. Iterative blind deconvolution method and its applications. *Optics Letters*, 13(7):547–549, 1988.

[3] L. Bar, N. Sochen, and N. Kiryati. Variational pairing of image segmentation and blind restoration. In T. Pajdla and J. Matas, editors, *Computer Vision – ECCV 2004, Part II*, volume 3022 of *Lecture Notes in Computer Science*, pages 166–177. Springer, Berlin, 2004.

[4] L. Bar, N. Sochen, and N. Kiryati. Image deblurring in the presence of salt-and-pepper noise. In R. Kimmel, N. Sochen, and J. Weickert, editors, *Scale Space and PDE Methods in Computer Vision*, volume 3459 of *Lecture Notes in Computer Science*, pages 107–118. Springer, Berlin, 2005.

[5] T. F. Chan and C. K. Wong. Total variation blind deconvolution. *IEEE Transactions on Image Processing*, 7:370–375, 1998.

[6] T. F. Chan, A. M. Yip, and F. E. Park. Simultaneous total variation image inpainting and blind deconvolution. *International Journal of Imaging Systems and Technology*, 15(1):92–102, 2005.

[7] N. Dey, L. Blanc-Feraud, C. Zimmer, P. Roux, Z. Kam, J.-C. Olivo-Marin, and J. Zerubia. Richardson-Lucy algorithm with total variation regularization for 3D confocal microscope deconvolution. *Microscopy Research and Technique*, 69:260–266, 2006.

[8] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. In *SIGGRAPH '06 – Special Interest Group on Computer Graphics and Interactive Techniques*, pages 787–794. ACM, 2006.

[9] X. Ge, J. Liu, D. Hu, and J. Tan. An extended sparse model for blind image deblurring. *Signal, Image and Video Processing*, 18:1863–1877, 2024.

[10] D. Hu, J. Tan, L. Zhang, X. Ge, and J. Liu. Image deblurring via enhanced local maximum intensity prior. *Signal Processing: Image Communication*, 96:116311, 2021.

[11] H. Hu and G. de Haan. Low cost robust blur estimator. In *IEEE International Conference on Image Processing (ICIP 2006)*, pages 617–620. IEEE, 2006.

[12] W. Hu, J. Xue, and N. Zheng. PSF estimation via gradient domain correlation. *IEEE Transactions on Image Processing*, 21(1):386–392, 2012.

[13] M. Jung and L. A. Vese. Nonlocal variational image deblurring models in the presence of Gaussian or impulse noise. In X.-C. Tai, K. Mørken, M. Lysaker, and K.-A. Lie, editors, *Scale-Space and Variational Methods in Computer Vision*, volume 5567 of *Lecture Notes in Computer Science*, pages 402–413. Springer, Berlin, 2009.

[14] 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 A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, editors, *Computer Vision – ECCV 2012, Part VII*, volume 7578 of *Lecture Notes in Computer Science*, pages 27–40. Springer, Berlin, 2012.

[15] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-Laplacian priors. In *Advances in Neural Information Processing Systems*, pages 1033–1041, 2009.

[16] J. E. Krist. Deconvolution of Hubble Space Telescope images using simulated point spread functions. *Astronomical Data Analysis Software and Systems I, A.S.P. Conference Series*, 25:226–228, 1992.

[17] G. Laribi and M. Welk. Towards quality assessment of blind deconvolution with shift compensation. In *26th International Conference on Pattern Recognition (ICPR 2022)*, pages 421–427. IEEE, 2022.

[18] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Efficient marginal likelihood optimization in blind deconvolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011)*, pages 2657–2664. IEEE, 2011.

[19] D. Li, R. M. Mersereau, and S. Simske. Blind image deconvolution through support vector regression. *IEEE Transactions on Neural Networks*, 18(3):931–935, 2007.

[20] G. Liu, S. Chang, and Y. Ma. Blind image deblurring using spectral properties of convolution operators. *IEEE Transactions on Image Processing*, 23(12):5047–5056, 2014.

[21] L. B. Lucy. An iterative technique for the rectification of observed distributions. *The Astronomical Journal*, 79(6):745–754, June 1974.

[22] M. Makarkin and D. Bratashov. State-of-the-art approaches for image deconvolution problems, including modern deep learning architectures. *MDPI Micromachines*, 12:1558:1–25, 2021.

[23] T. Meinhardt, M. Moeller, C. Hazirbas, and D. Cremers. Learning proximal operators: Using denoising networks for regularizing inverse imaging problems. In *IEEE International Conference on Computer Vision (ICCV 2017)*, pages 1799–1808. IEEE, 2017.

[24] L. Moisan. Periodic plus smooth image decomposition. *Journal of Mathematical Imaging and Vision*, 39:161–179, 2011.

[25] J. G. Nagy and D. P. O'Leary. Fast iterative image restoration with a spatially varying PSF. In F. T. Luk, editor, *Advanced Signal Processing Algorithms, Architectures, and Implementations*, volume 3162 of *Proceedings of SPIE*, pages 388–399. SPIE Press, Bellingham, 1997.

[26] N. Otsu. A threshold selection method from grey-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-9(1):62–66, 1979.

[27] D. Ren, K. Zhang, Q. Wang, Q. Hu, and W. Zuo. Neural blind deconvolution using deep priors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2020)*, pages 3341–3350. IEEE, 2020.

[28] W. H. Richardson. Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America*, 62(1):55–59, 1972.

[29] K. Schelten, S. Nowozin, J. Jancsary, C. Rother, and S. Roth. Interleaved regression tree field cascades for blind image deconvolution. In *IEEE Winter Conference on Applications of Computer Vision*, pages 494–501. IEEE, 2015.

[30] C. J. Schuler, H. C. Burger, S. Harmeling, and B. Schölkopf. A machine learning approach for non-blind image deconvolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2013)*, pages 1067–1074. IEEE, 2013.

[31] D. Snyder, T. J. Schulz, and J. A. O'Sullivan. Deblurring subject to nonnegativity constraints. *IEEE Transactions on Image Processing*, 40(5):1143–1150, 1992.

[32] P. H. van Cittert. Zum Einfluß der Spaltbreite auf die Intensitätsverteilung in Spektrallinien. II. *Zeitschrift für Physik*, 65:298–308, 1933.

[33] C. R. Vogel and M. E. Oman. Fast, robust total variation-based reconstruction of noisy, blurred images. *IEEE Transactions on Image Processing*, 7:813–824, 1998.

[34] Y. Wang, J. Yang, W. Yin, and Y. Zhang. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences*, 1(3):248–272, 2008.

[35] M. Welk. A robust variational model for positive image deconvolution. *Signal, Image and Video Processing*, 10(2):369–378, 2016.

[36] M. Welk. Robust blind deconvolution with convolution-spectrum-based kernel regulariser and Poisson-noise data terms. In F. Lauze, Y. Dong, and A. B. Dahl, editors, *Scale Space and Variational Methods in Computer Vision*, volume 10302 of *Lecture Notes in Computer Science*, pages 159–171. Springer, Cham, 2017.

[37] M. Welk. Superresolution alignment with innocence assumption: Towards a fair quality measurement for blind deconvolution. In P. M. Roth, M. Vincze, W. Kubinger, A. Müller, B. Blaschitz, and S. Stolc, editors, *Proceedings of the OAGM-ARW Joint Workshop: Vision, Automation and Robotics*, pages 145–150. Verlag der Technischen Universität Graz, Graz, 2017.

[38] M. Welk, D. Theis, T. Brox, and J. Weickert. PDE-based deconvolution with forward-backward diffusivities and diffusion tensors. In R. Kimmel, N. Sochen, and J. Weickert, editors, *Scale Space and PDE Methods in Computer Vision*, volume 3459 of *Lecture Notes in Computer Science*, pages 585–597. Springer, Berlin, 2005.

[39] N. Wiener. *Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engineering Applications*. MIT Press, Cambridge, MA, 1949.

[40] L. Xu and J. Jia. Two-phase kernel estimation for robust motion deblurring. In K. Daniilidis, P. Maragos, and N. Paragios, editors, *Computer Vision – ECCV 2010, Part I*, volume 6311 of *Lecture Notes in Computer Science*, pages 157–170. Springer, Berlin, 2010.

[41] L. Xu, J. Ren, C. Liu, and J. Jia. Deep convolutional neural network for image deconvolution. In *Advances in Neural Information Processing Systems*, pages 1790–1798, 2014.

[42] Y.-L. You and M. Kaveh. Anisotropic blind image restoration. In *IEEE International Conference on Image Processing (ICIP 1996)*, volume 2, pages 461–464. IEEE, 1996.

[43] J. Zhang, Z. Yue, H. Wang, Q. Zhao, and D. Meng. Blind image deconvolution by generative-based kernel prior and initializer via latent encoding. In A. Leonardis, E. Ricci, S. Roth, O. Russakovsky, T. Sattler, and V. G., editors, *Computer Vision – ECCV 2024, Part XLVI*, volume 15104 of *Lecture Notes in Computer Science*, pages 73–92. Springer, Cham, 2024.