

Computational photography



$$y = Z_y^T \sum_{r=0}^{p-1} C_r^T F^H \text{Diag}(F Z_f f^{(r)}) F \text{Diag}(w^{(r)}) C_r x$$

Figure 1.6: In many real-world imaging applications, the common assumption of stationary blur does not hold. Examples that exhibit non-stationary blur include e.g. camera shake, optical aberrations, and atmospheric turbulence. We derived a mathematically sound and physically well-motivated model, which allows to express and efficiently compute spatially-varying blur [570]. Our “Efficient Filter Flow” framework substantially broadens the application range of image deconvolution methods. In a number of challenging real-world applications [26, 320, 472, 535, 537, 570] we demonstrated both the validity and versatility of our approach.

Digital image restoration as a key area of signal and image processing aims at computationally enhancing the quality of images by undoing the adverse effects of image degradation such as noise and blur. It plays an important role in both scientific imaging and everyday photography. Using probabilistic generative models of the imaging process, our research aims to recover the most likely original image given a low-quality image, or image sequence. In the following we give a number of illustrative examples to highlight some of our work.

Spatially varying blurs: our work on blind deconvolution of astronomical image sequences [620] can recover sharp images through atmospheric turbulence, but is limited to relatively small patches of the sky, since the image defect is modeled as a space-invariant blur. Images that cover larger areas require a convolutional model allowing for space-variant blur. In [570] we proposed such a model based on a generalization of the short-time Fourier transform, called Efficient Filter Flow (EFF), which is illustrated in Fig. 1.6. Deconvolution based on EFF successfully recovers a sharp image from an image sequences distorted by air turbulence (see Fig. 1.7).

Super-resolution: we generalized our online method for the *multi-frame blind deconvolution* problem [620] to account for the adverse effect of saturation and to enable super-resolution given a sequence of degraded images [184].

By drawing a connection between Fraunhofer diffraction and kernel mean maps we are able to show that under certain imaging conditions imaging beyond the physical diffraction limit is in principle possible [423].

Removing camera shake: photographs taken with long exposure times are affected by camera shake creating smoothly varying blur. Real hand-held camera movement involves both translation and rotation, which can be modeled with our EFF framework for space-variant blur [597]. We were also able to recover a sharp image from a single distorted image, using sparsity-inducing image priors and an alternating update algorithm. The algorithm can be made more robust by restricting the EFF to blurs consistent with physical camera shake [537]. This leads to higher image quality (Fig. 1.8) and computational advantages. To foster and simplify comparisons between different algorithms removing camera shake, we created a public benchmark dataset and a comparison of current methods [483].

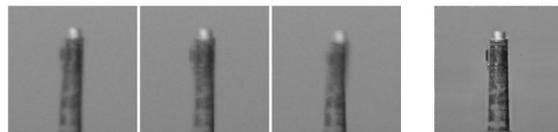


Figure 1.7: Our Efficient Filter Flow model [570] combined with our online multi-frame blind deconvolution [620] allows to restore a single sharp image given a sequence of blurry images degraded by atmospheric turbulence.



Figure 1.8: Example result of our proposed method for fast removal of non-uniform camera shake [537]. The left image shows the original image blurred due to camera shake during exposure, the right image shows the result of our proposed method.

Correcting lens aberration: even good lenses exhibit optical aberration when used with wide apertures. Similar to camera shake, this creates a certain type of blur that can be modeled with our EFF framework; but optical aberrations affects each color channel differently (chromatic aberration). We measured these effects with a robotic setup and corrected them using a non-blind deconvolution based on the EFF framework [535]. We were also able to implement a blind method rectifying optical aberrations in single images [472]. An example is shown in Fig. 1.9. The key was to constrain the EFF framework to rotationally symmetric blurs varying smoothly from the image center to the edges. We have recently been able to formulate this as non-parametric kernel regression, enabling faithful optical aberration estimation and correction [320], which might lead to new approaches in lens design.



Figure 1.9: To test the capability of our approach to optical aberration correction, we built a camera lens containing a single glass element only. An example image is shown on the left, which exhibits strong blur artifacts especially towards the corners of the image. In contrast, the restored image (right) is of high quality and demonstrates that our method [472] is able to correct for most degradations caused by optical aberrations from a single input image only.

Denoising: another classical image distortion is *noise*. Noise can exhibit different structure, e.g., additive white Gaussian, salt-and-pepper, JPEG-artifacts, and stripe noise, depending on the application.

More information: <https://ei.is.tuebingen.mpg.de/project/computational-imaging>

In astronomical imaging, dim celestial objects require very long exposure times, which causes high sensor noise. It is common to subtract a dark frame from the image — an image taken with covered lens, containing only sensor noise. The difficulty with this is that the sensor noise is stochastic, and image information is not taken into account. We studied the distribution of sensor noise generated by a specific camera sensor and proposed a parameterized model [538]. Combined with a simple image model for astronomical images, this gives superior denoising.

Multi-scale denoising: noise usually has a more damaging effect on the higher spatial frequencies, so most algorithms focus on those. However, if the noise variance is large, also lower frequencies are distorted. We were able to significantly improve the performance of many methods for that setting by defining a multi-scale meta-procedure, leading to a DAGM 2011 prize [539].

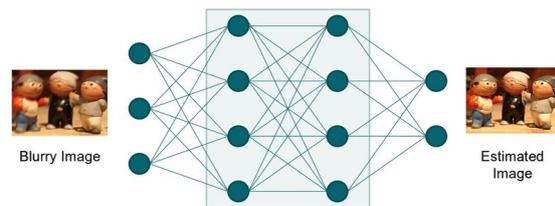


Figure 1.10: Schematic illustration of casting image restoration as a learning problem. A neural network (middle) is trained with degraded/clean image pairs to learn a mapping from a degraded image (left) to a restored image with enhanced quality (right).

Image restoration as a learning problem: since denoising can be also seen as a non-trivial mapping from noisy to clean images as sketched in Fig. 1.10, we were particularly interested whether a learning-based approach can be applied. Using very large data sets, we trained a multi-layer perceptron (MLP) that is able to denoise images better than all existing methods, leading to the new state-of-the-art denoising method [450].

Such a discriminative approach turned out to be effective also for other image restoration tasks such as inpainting [402], non-blind deconvolution [424], as well as blind image deconvolution [26, 320].