Learning to Push the Limits of Efficient FFT-based Image Deconvolution

Jakob Kruse[1] Carsten Rother[1] Uwe Schmidt[2]

[1]Heidelberg University [2]MPI-CBG Dresden

<CODE>
goo.gl/7MvKZy

Motivation & Contributions

- *Non-blind deconvolution* is an important component for removing image blur (e. g., due to camera shake)
- High-quality methods are often slow and do not scale to large megapixel-sized images
- Fast Fourier-based methods are lacking in restoration quality

Our Contributions

1. **Generalize** discriminative FFT-based deconvolution by using regularization based on **convolutional neural networks**
2. Propose a **simple and effective boundary adjustment** to adhere to the circular convolution assumption imposed by FFTs
3. Obtain **state-of-the-art results** on two deconvolution benchmarks, even compared to much slower high-quality methods

FFT-based Image Deconvolution

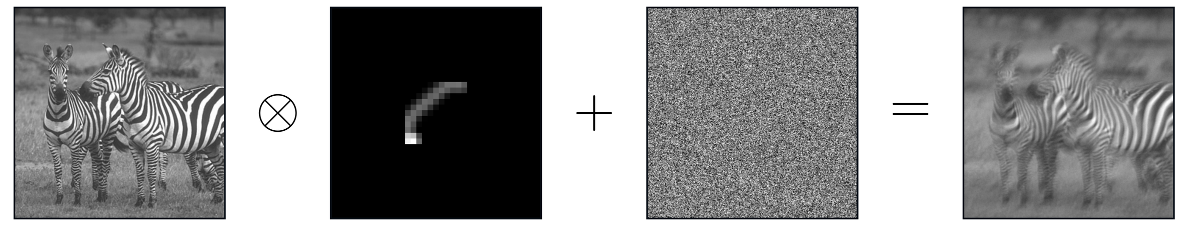

sharp image $\mathbf{x}$ blur kernel $\mathbf{k}$ $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}/\lambda)$ observation $\mathbf{y}$

Restoring $\mathbf{x}$ from $\mathbf{y}, \mathbf{k}$ with regularization requires optimization
→ often by solving a sequence of linear equation systems (EqS)

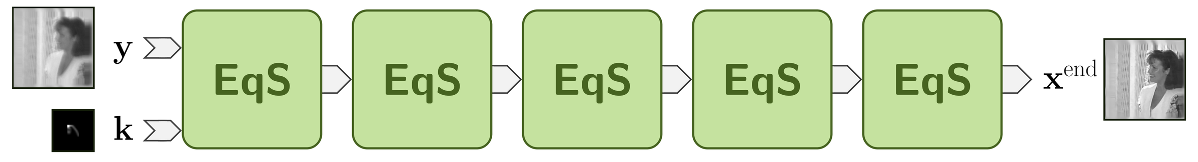

Circular convolution assumption allows direct solution of EqS:

$$\mathbf{x}^{t+1} = \mathcal{F}^{-1}\left(\frac{\mathcal{F}(\mathbf{k} \circledast \mathbf{y} + \boxed{A})}{|\mathcal{F}(\mathbf{k})|^2 + \boxed{B}}\right) \qquad \mathcal{F}: \text{Fourier transform}$$

- Wiener filter:

  $\boxed{A} = \mathbf{0}$ and $\boxed{B}$ = noise/image spectrum ratio

- Shrinkage Fields (CSF) [1]:

  $$\boxed{A} = \frac{\beta_t}{\lambda}\sum_i \mathbf{f}_{it} \circledast \psi_{it}(\mathbf{f}_{it} \otimes \mathbf{x}^t) \quad \text{and} \quad \boxed{B} = \frac{\beta_t}{\lambda}\sum_i |\mathcal{F}(\mathbf{f}_{it})|^2,$$

  where $\mathbf{f}_{it}$ are linear filters and ψ_{it} are 1D shrinkage functions

FDN Fourier Deconvolution Network

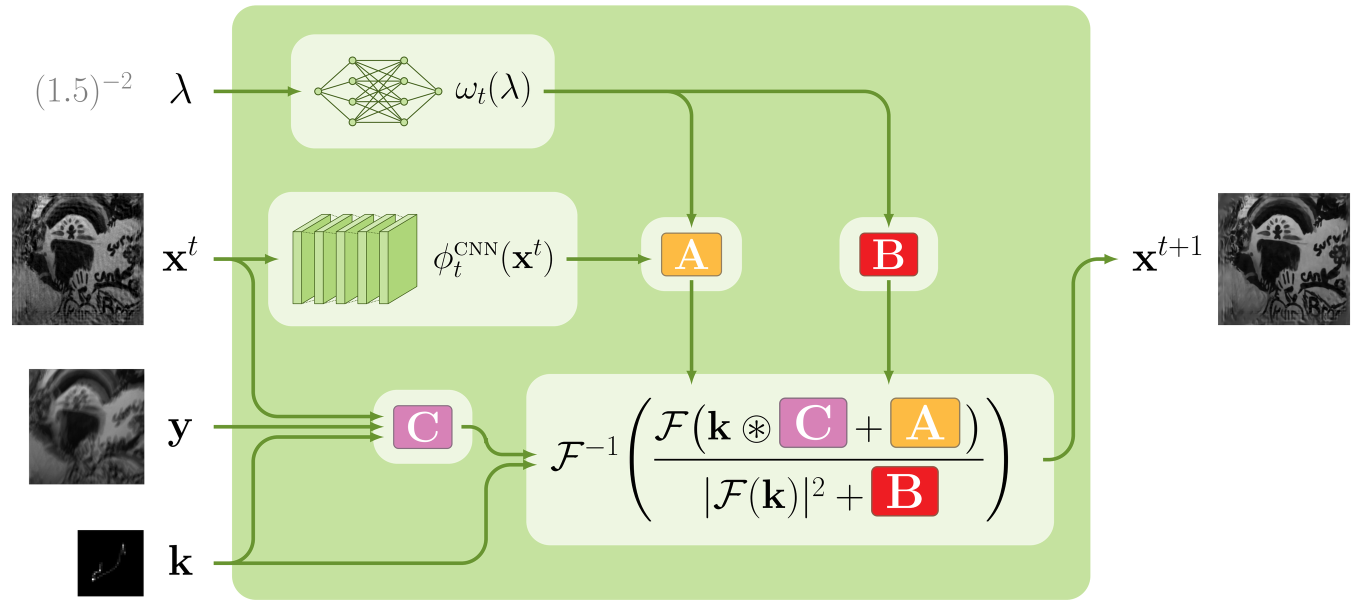

We generalize Shrinkage Fields [1] by choosing

$$\boxed{A} = \frac{1}{\omega_t(\lambda)} \cdot \phi_t^{\text{CNN}}(\mathbf{x}^t) \quad \text{and} \quad \boxed{B} = \frac{1}{\omega_t(\lambda)} \cdot \sum_i |\mathcal{F}(\mathbf{f}_{it})|^2$$

- More powerful: CNNs instead of pixel-wise shrinkage functions
- More flexible: Filters $\mathbf{f}_{it}$ in $\boxed{B}$ are decoupled from $\boxed{A}$
- Noise generalization: Noise-adaptive regularization weight $\omega_t(\lambda)$ allows one model to be used for images with varying noise levels

- CNNs modulate smoothness → strong response at sharp edges

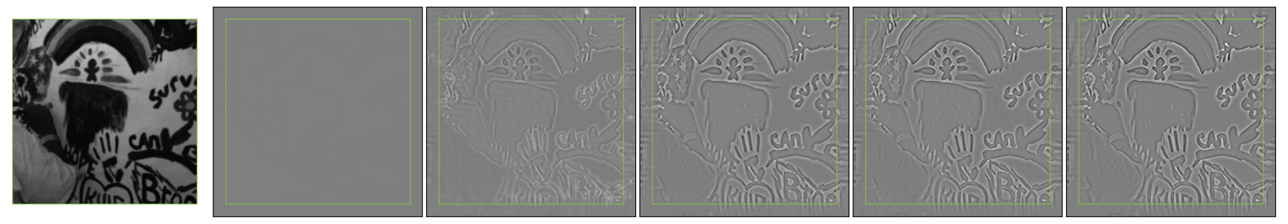

CNN output: Sharp image (left) and output of CNNs for the first 5 model stages.

- More versatile noise-adaptive models trained for a range of noise levels perform similar to noise-specialized ones

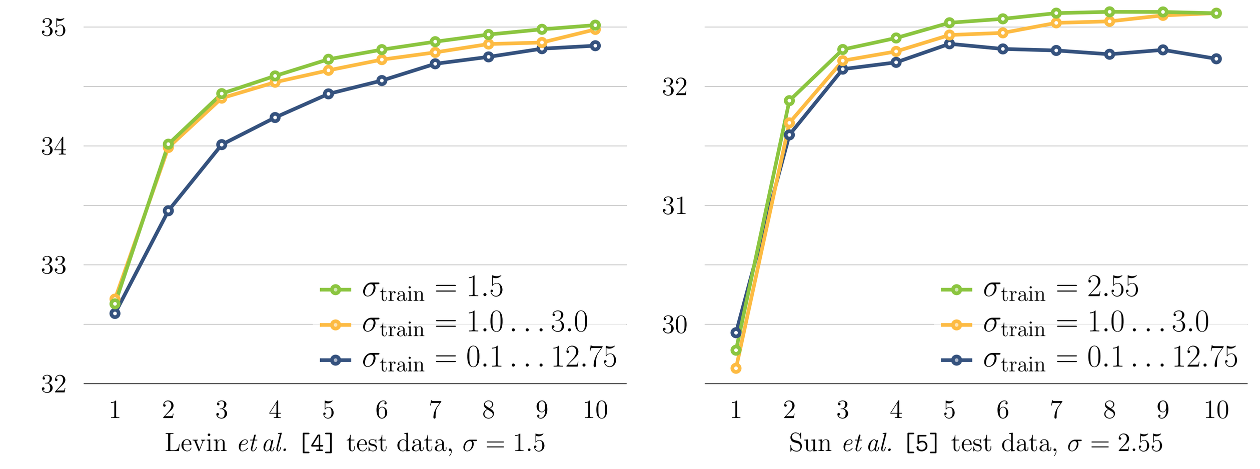

Noise-specialized vs. noise-adaptive models: PSNRs after each model stage.

New Boundary Adjustment Method

- Fourier-based deconvolution assumes **circular boundary conditions**, which is inaccurate for blurred natural images
- Common `edgetaper` operation is used for pre-processing to approximate missing boundary regions of blurred image $\mathbf{y}$

From analysis of [6,7], we propose to iteratively update the boundary region based on current deblurred image estimate $\mathbf{x}^t$:

$$\boxed{C} = \mathbf{y} + \mathtt{boundary}(\mathbf{k} \otimes \mathbf{x}^t)$$

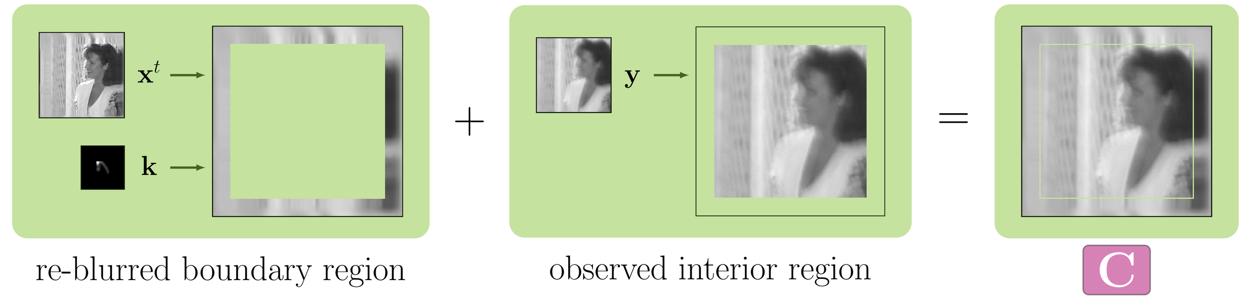

- Simple, yet effective: Consistently improves upon `edgetaper` to alleviate restoration artifacts and obtain better results
- Trivial to implement, works with all FFT-based methods
- Parameter-free, negligible computational cost per iteration

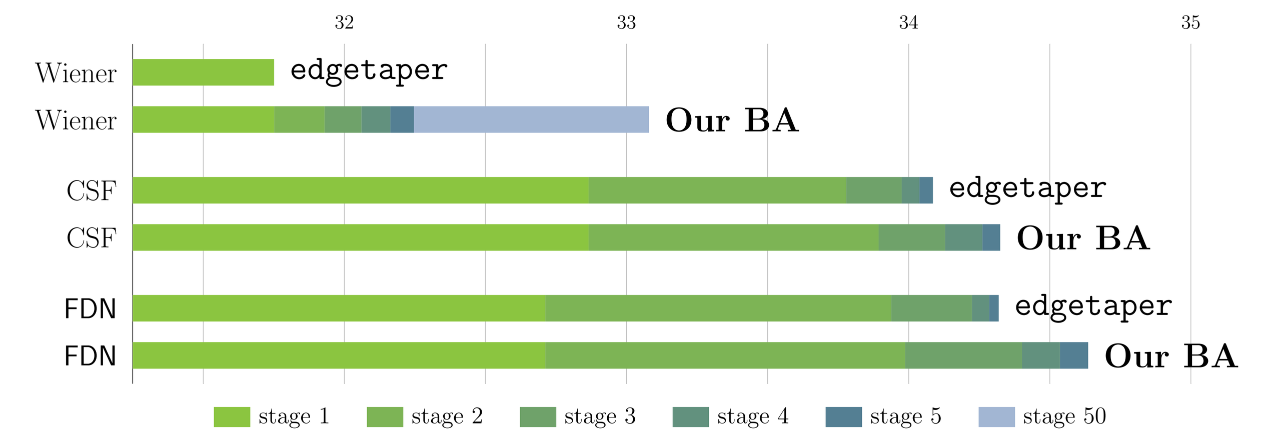

Comparison of our boundary adjustment and standard `edgetaper` (Levin *et al.* data [4]).

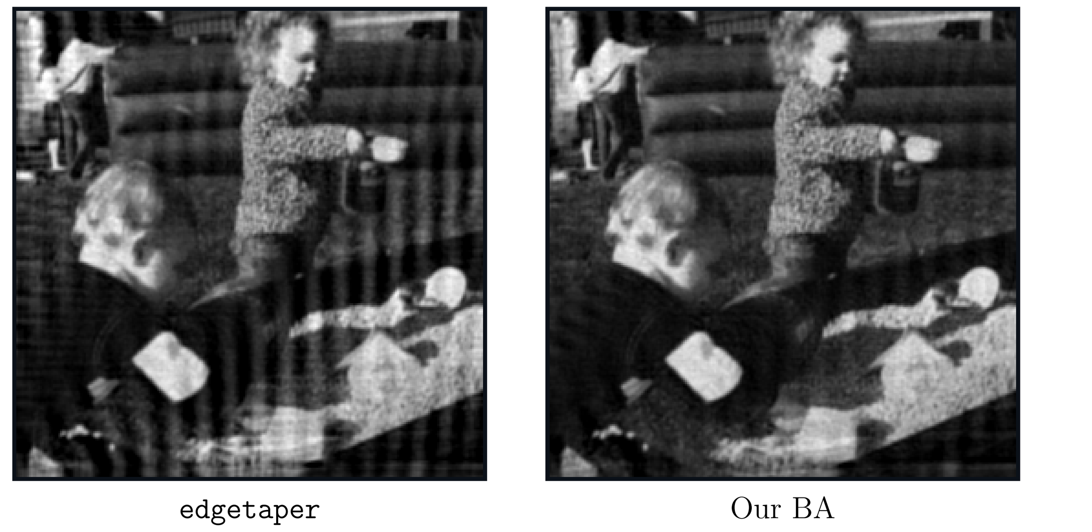

Example (Wiener filter): Our boundary adjustment can reduce strong restoration artifacts.

High Resolution Example

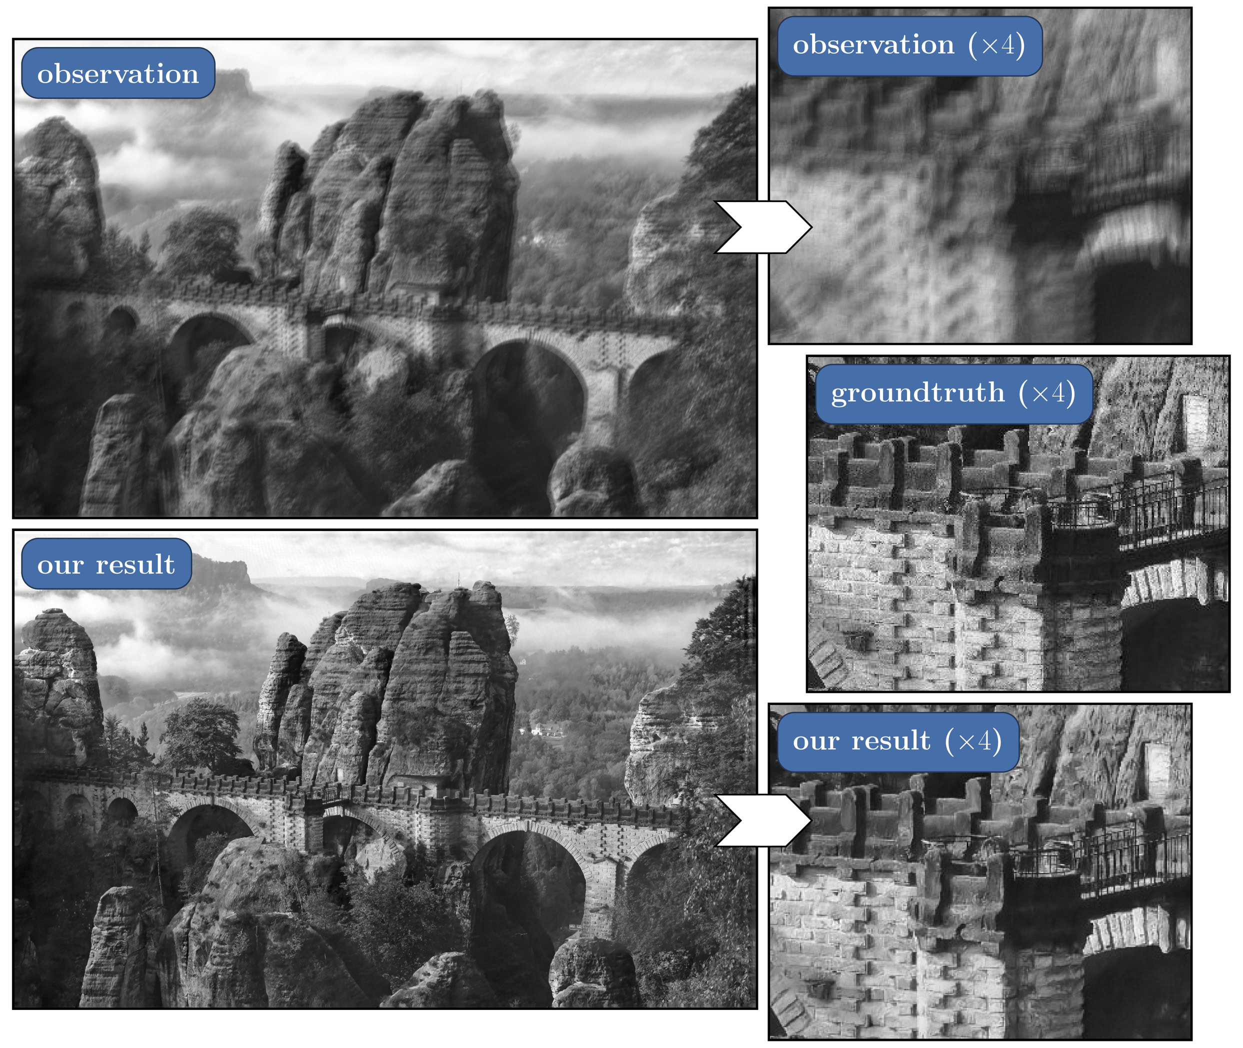

- Less than 10s to restore 4 megapixel with unoptimized code on a GPU (much faster if image size is known in advance)
- Other high-quality methods too slow for images of this size

Quantitative Results

- Two common benchmarks for non-blind deconvolution
- 10-stage FDN trained with noise range $\sigma_{\text{train}} = 1.0 \ldots 3.0$
- We outperform CSF and the much slower high-quality methods RTF and EPLL

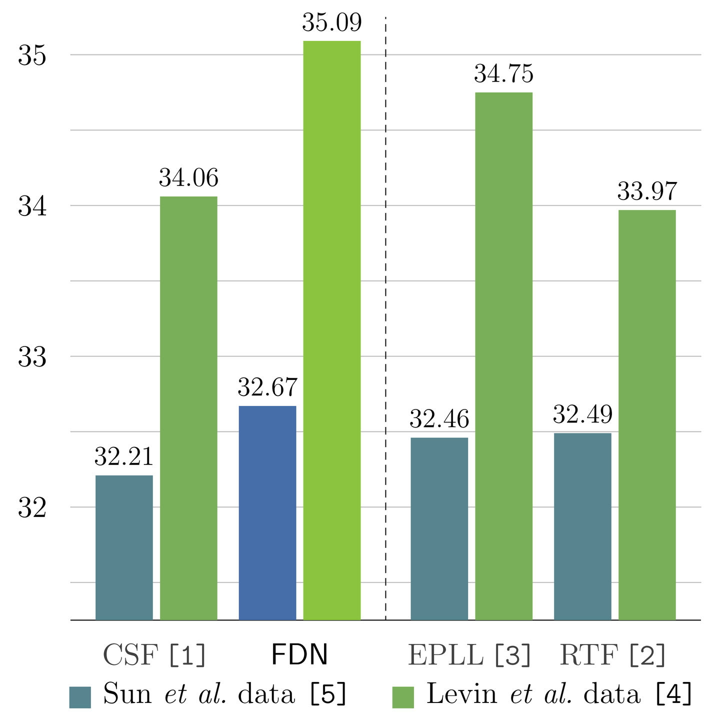

Average PSNR in dB on two benchmarks.

CBG Max Planck Institute of Molecular Cell Biology and Genetics
UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386
erc

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 647769).

Contact

uschmidt@mpi-cbg.de

Code

Keras/TensorFlow
https://goo.gl/7MvKZy

References

[1] U. Schmidt and S. Roth. **Shrinkage fields for effective image restoration**. CVPR 2014.
[2] U. Schmidt, C. Rother, S. Nowozin, J. Jancsary, and S. Roth. **Discriminative non-blind deblurring**. CVPR 2013.
[3] D. Zoran and Y. Weiss. **From learning models of natural image patches to whole image restoration**. ICCV 2011.
[4] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. **Understanding and evaluating blind deconvolution algorithms**. CVPR 2009.
[5] L. Sun, S. Cho, J. Wang, and J. Hays. **Edge-based blur kernel estimation using patch priors**. ICCP 2013.
[6] M. S. C. Almeida and M. A. T. Figueiredo. **Deconvolving images with unknown boundaries using the alternating direction method of multipliers**. IEEE Trans. Image Process. 2013.
[7] A. Matakos, S. Ramani, and J. A. Fessler. **Accelerated edge-preserving image restoration without boundary artifacts**. IEEE Trans. Image Process. 2013.