

# End-to-End Learning for Image Burst Deblurring

Patrick Wieschollek<sup>1,2</sup>, Bernhard Schölkopf<sup>1</sup>,  
Hendrik P.A. Lensch<sup>2</sup>, and Michael Hirsch<sup>1</sup>

<sup>1</sup> Max Planck Institute for Intelligent Systems, <sup>2</sup> University of Tübingen

**Abstract.** We present a neural network model approach for multi-frame blind deconvolution. The discriminative approach adopts and combines two recent techniques for image deblurring into a single neural network architecture. Our proposed hybrid-architecture combines the explicit prediction of a deconvolution filter and non-trivial averaging of Fourier coefficients in the frequency domain. In order to make full use of the information contained in all images in one burst, the proposed network embeds smaller networks, which explicitly allow the model to transfer information between images in early layers. Our system is trained end-to-end using standard backpropagation on a set of artificially generated training examples, enabling competitive performance in multi-frame blind deconvolution, both with respect to quality and runtime.

## 1 Introduction

Nowadays, consumer cameras are able to capture an entire series of photographs in rapid succession. Hand-held acquisition of a burst of images is likely to cause blur due to unwanted camera shake during image capture. This is particularly true along with longer exposure times needed in low-light environments.

Motion blurring due to camera shake is commonly modeled as a spatially invariant convolution of a latent sharp image  $X$  with an unknown blur kernel  $k$

$$Y = k * X + \varepsilon, \quad (1)$$

where  $*$  denotes the convolution operator,  $Y$  the blurred observation and  $\varepsilon$  additive noise. Single image blind deconvolution (BD), i.e. recovering  $X$  from  $Y$  without knowing  $k$ , is a highly ill-posed problem for a variety of reasons. In contrast, multi-frame blind deconvolution or burst deblurring methods aim at recovering a single sharp high-quality image from a sequence of blurry and noisy observed images  $Y_1, Y_2, \dots, Y_N$ . Accumulating information from several observations can help to solve the reconstruction problem associated with Eq.(1) more effectively.

Traditionally, generative models are used for blind image deconvolution. While they offer much flexibility they are often computationally demanding and time-consuming.

Discriminative approaches on the other hand keep the promise of fast processing times, and are particularly suited for situations where an exact modeling of the image formation process is not possible. A popular choice in this context are neural networks. They gained some momentum due to the great success of deep learning in many supervised computer vision tasks, but also for a number of low-level vision tasks state-of-the-art results have been reported [1, 2]. Our proposed method lines up with the latter approaches and comes along with the following main contributions:

1. A robust state-of-the-art method for multi-image blind-deconvolution for both invariant and spatially-varying blur.
2. A hybrid neural network architecture as a discriminative approach for image deblurring supporting end-to-end learning in the fashion of deep learning.
3. A neural network layer version of Fourier-Burst-Accumulation [3] with learnable weights.
4. The proposed embedding of a small neural network allowing for information sharing across the image burst early in the processing stage.

## 2 Related Work

Blind image deconvolution (BD) has seen considerable progress in the last decade. A comprehensive review is provided in the recent overview article by Wang and Tao [4].

**Single image blind deconvolution.** Approaches for single image BD that report state-of-the-art results include the methods of Sun et al. [5], and Michaeli and Irani [6] that use powerful patch-based priors for image prediction. Following the success of deep learning methods in computer vision, also a number of neural network based methods have been proposed for image restoration tasks including *non-blind* deconvolution [2, 7, 8] which seeks to restore a blurred image when the blur kernel is known, but also for the more challenging task of *blind* deconvolution [9–14] where the blur kernel is not known a priori. Most relevant to our work is the recent work of Chakrabarti [11] which proposes a neural network that is trained to output the complex Fourier coefficients of a deconvolution filter. When applied to an input patch in the frequency domain, the network returns a prediction of the Fourier transform of the corresponding latent sharp image patch. For whole image restoration, the input image is cut into overlapping patches, each of which is independently processed by the network. The outputs are recomposed to yield an initial estimate of the latent sharp image, which is then used together with the blurry input image for the estimation of a single invariant blur kernel. The final result is obtained using the state-of-the-art non-blind deconvolution method of Zoran and Weiss [15].

**Multi-frame blind deconvolution.** Splitting an exposure budget across many photos can lead to a significant quality advantage [16]. Thus, it has been shown that multiple captured images can help alleviating the illposedness of the BD problem [17]. This has been exploited in several approaches [18–22] for *multi-frame* BD, i.e. combining multiple, differently blurred images into a single latent

sharp image. Generative methods, that make explicit use of an image formation model, mainly differ in the prior and/or the optimization procedure they use. State-of-the-art methods use sparse priors with fast Bregman splitting techniques for optimization [20], or within a variational inference framework [18], cross-blur penalty functions between image pairs [21, 22], also in combination with robust cost functions [19].

More recently proposed methods [23–26] also model the inter-frame motion and try to exploit the interrelation between camera motion, blur and image mis-alignment. Camera motion, when integrated during the exposure time of a single frame will produce intra-frame motion blur while leading to inter-frame mis-alignment during readout time between consecutive image capture.

All of the above-mentioned methods employ generative models and try to explicitly estimate one unknown blur kernel for each blurry input frame along with predicting the latent sharp image. A common shortcoming is the large computational burden with typical computation times in the order of tens of minutes, which hinders their wide-spread use in practice.

Recently, Delbracio and Sapiro have presented a fast method that aggregates a burst of images into a single image that is both sharper and less noisy than all the images in the burst [3]. The approach is inspired by a recently proposed Lucky Imaging method [27] targeted for astronomical imaging. Traditional Lucky Imaging approaches would select only a few “lucky” frames from a stack of hundreds to thousands recorded short-exposure images and combine them via non-rigid shift-and-add techniques into a single sharp image. In contrast, the authors of [27] propose to take all images into account (rather than a small subset of carefully chosen frames) taking and combining what is less blurred of each frame to form an improved image. It does so by computing a weighted average of the Fourier coefficients of the registered images in the burst. In [3, 28] it has been demonstrated that this approach can be adapted successfully to remove camera shake originated from hand tremor vibrations. Their Fourier Burst Accumulation (FBA) approach allows for fast processing even for Megapixel images while at the same time yielding high-quality results provided that a “lucky”, i.e. almost sharp frame is amongst the captured image burst.

In our work, we not only present a learning-based variant of FBA but also show how to alleviate the drawback of requiring a sharp frame amongst the input sequence of images. To this end we combine the single image BD method of [11] with FBA in a single network architecture which facilitates end-to-end learning. To the best of our knowledge this is the first time that a fully discriminative approach has been presented for the challenging problem of multi-frame BD.

### 3 Method

Let’s assume we have given a burst of observed color images  $Y_1, Y_2, \dots, Y_N \in \mathcal{I}$  capturing the same scene  $X \in \mathcal{I}$ . Assuming each image in the captured sequence is blurred differently, our image formation model reads

$$Y_t = k_t * X + \varepsilon_t, \quad (2)$$

where  $*$  denotes the convolution operator,  $k_t$  the blur kernel for observation  $Y_t$  and  $\varepsilon_t$  additive zero-mean Gaussian noise.

We aim at predicting a latent single sharp image  $\hat{X}$  through a deep neural network architecture, i.e.

$$\pi^{(\theta)}: \mathcal{I}_p^N \rightarrow \mathcal{I}_p, \quad (y_1, y_2, \dots, y_N) \mapsto \hat{x} = \pi^{(\theta)}(y_1, y_2, \dots, y_N).$$

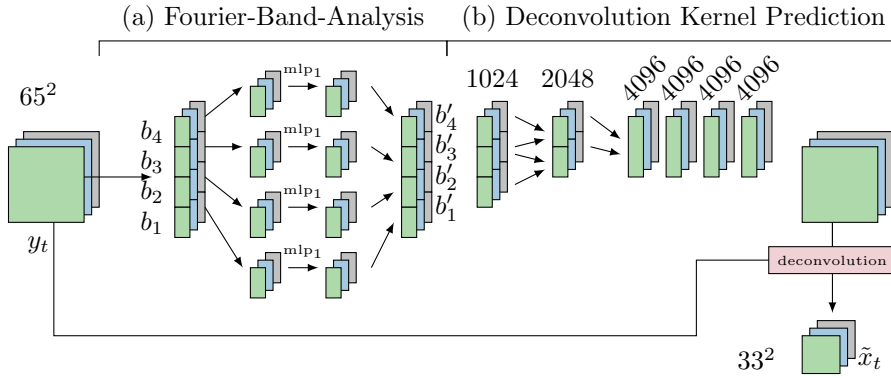
The network operates on a patch-by-patch basis, here  $y_t \in \mathcal{I}_p$  and  $\hat{x} \in \mathcal{I}_p$  denote a patch in  $Y_t$  and  $X$  respectively. The patches are chosen to be overlapping. Our network predicts a single sharp patch  $\hat{x} \in \mathcal{I}_p$  from multiple input patches  $y_t \in \mathcal{I}_p$ . All predicted patches are recomposed to form the final prediction  $\hat{X}$  by averaging the predicted pixel values. During the training phase we optimize the learning parameters  $\theta$  by directly minimizing the objective

$$\|\pi^{(\theta)}(y_1, y_2, \dots, y_N) - x\|_2^2. \quad (3)$$

In the following we will describe the construction of  $\pi^{(\theta)}(\cdot)$ , the optimization of network parameters  $\theta$  during the training of the neural network and the restoration of an entire sharp image.

### 3.1 Network Architecture

The architecture  $\pi^{(\theta)}(\cdot)$  consists of several stages: (a) frequency band analysis with Fourier coefficient prediction, (b) a deconvolution part and (c) image fusion. Figure 1 illustrates the first two stages of our proposed system.




**Fig. 1.** Frequency band analysis and deconvolution for an image burst with 3 patches  $y_1, y_2, y_3$ . Following the work of Chakrabarti [11] we separate the Fourier spectrum in 4 different bands  $b_1, \dots, b_4$ . In addition, we allow each band separately to interact across all images in one burst to support early information sharing. The predicted output of the deconvolution step are smaller patches  $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3$ .

**(a) Frequency band analysis.** The frequency band analysis computes the discrete Fourier transform of the observed patch  $y_t$  according to the neural network approach in [11] at three different sizes ( $17 \times 17, 33 \times 33, 65 \times 65$ ) using different sample sizes, which we will refer to bands  $b_1, b_2, b_3$ . In addition, band  $b_4$  represents a low-pass band containing all coefficients with  $\max |z| \leq 4$  from band  $b_3$ . This is depicted in Fig. 1. To enable early information sharing within one burst of patches, we allow the neural network to spread the per band information extracted from one patch across all images of the burst using  $1 \times 1$  convolution.

This essentially embeds a fully connected neural network for each Fourier coefficient  $(f_{ij})_t$  with weight sharing. Since we use these operations in the image fusion stage again, we elaborate on this idea in more detail.

The values of one Fourier coefficient  $(f_{ij})_t$  at frequency position  $(i, j)$  across the entire burst  $t = 1, 2, \dots, N$  can be considered as a single vector  $(f_{ij})_{t=1,2,\dots,N}$  of dimension  $N$  (compare Fig. 2). Each of these vectors is fed through a small



**Fig. 2.** For arbitrary inputs (bands  $b_1, b_2, b_3, b_4$  or later FBA weights) we interpret each coefficient across one burst as a single vector. A transformed version of this excerpt will be placed at the same location in the output patch again. To reduced the number of learnable parameters, we employ weight sharing independent of the position.

network of fully connected layers, labeled by `mlp1` in Fig. 1. This allows the neural network to adjust the extracted Fourier coefficients right before a dimensionality reduction occurs. These modified values  $(f'_{ij})_{t=1,2,\dots,N}$  give rise to adjusted Fourier bands  $b'_1, b'_2, b'_3, b'_4$ .

**(b) Deconvolution.** Pairwise merging of the resulting bands  $b'_1, b'_2, b'_3, b'_4$  with modified Fourier coefficients using fully connected layers with ReLU activation units entails a dimensionality reduction. The produced 4096 feature vector encoding is then fed through several fully connected layers producing a 4225 dimensional prediction of the filter coefficients of the deconvolution kernel. Applying the deconvolution kernel predicts a sharp patch  $\hat{x}$  of size  $33 \times 33$  from each input sequence of patches. This step is implemented as a multiplication of the predicted Wiener Filter with the Fourier transform of the input patch.

**(c) Image fusion.** In the last part of our pipeline we fuse all available sharp patches  $y_1, y_2, \dots, y_N$  by adopting the FBA approach described in [3] as a neural network component with learnable weights. The vanilla FBA algorithm applies

the following weighted sum to a Fourier transform  $\hat{\alpha}$  of a patch  $\alpha$ :

$$u(\hat{\alpha}) = \mathcal{F}^{-1} \left( \sum_{i=1}^N w_i(\zeta) \hat{\alpha}_i(\zeta) \right) (x) \quad (4)$$

$$w_i(\zeta) = \frac{|\hat{\alpha}_i(\zeta)|^p}{\sum_{j=1}^N |\hat{\alpha}_j(\zeta)|^p}, \quad (5)$$

where  $w_i$  denotes the contribution of frequency  $\zeta$  of a patch  $\alpha_i$ . Note, that  $u(\hat{\alpha})$  is differentiable in  $\hat{\alpha}$  allowing to pass gradient information to previous layers through back-propagation. To incorporate this algorithm as a neural network layer into our pipeline, we replace Equation (4) by a parametrized version

$$u(\hat{\alpha}) = \mathcal{F}^{-1} \left( \sum_{i=1}^N h_\phi(\zeta) \hat{\alpha}_i(\zeta) \right) (x). \quad (6)$$


Hence, instead of a hard-coded weight-averaging (using  $w_i$ ) the network is able to learn a data-dependent weighted-averaging scheme. Again, the function  $h_\phi(\cdot)$  represents two 1x1 convolutional layers with trainable parameters  $\phi$  following the same idea of considering the Fourier coefficient across one burst as a single vector (compare Fig. 2).

### 3.2 Training

The network is trained on an artificially generated dataset obtained by applying synthetic blur kernels to patches extracted from the MS COCO dataset [29]. This dataset consists of real-world photographs collected from the internet. To increase the quality of ground-truth patches guiding the training process we reject patches with too small image gradients. This process gives us 542217 sharp patches. For a fair evaluation we use a splitting<sup>1</sup> in training and validation set. Optimizing the neural network parameters is done on the training set only. The input bursts of 14 blurry images are generated on-the-fly by applying synthetic blur kernels to the ground-truth patches. These synthetic blur kernels of sizes  $17 \times 17$  and  $7 \times 7$  pixels are generated using a Gaussian process with a Matérn covariance function following [9], a random subset of which is shown in Fig. 3. In addition, we apply standard data augmentation methods like rotating and mirroring to the ground-truth data. Hence, this approach gives nearly an infinite amount of training data. We also add zero-mean Gaussian noise with variance 0.1. The validation data is precomputed to ensure fair evaluation during training.

Unfortunately, sophisticated stepsize heuristics like Adam [30] or Adagrad [31] failed to guarantee a stable training. We suspect the large range of values in the Fourier space to mislead those heuristics. Instead, we use stochastic gradient descent with momentum ( $\beta = 0.9$ ), batchsize 32 and an initial learning rate of  $\eta = 2$  which decreases every 5000 steps by a factor of 0.8. Training the neural network took 6 days using TensorFlow [32] on a NVIDIA Titan X.

<sup>1</sup> Provided by [29]



**Fig. 3.** Some of the synthetically generated PSFs using a Gaussian process for generating training examples on-the-fly.

The FBA approach [3] applies Gaussian smoothing to the weights  $w_i$  to account for the fact that small camera shakes are likely to vary the Fourier spectrum in a smooth way. While this removes strong artefacts in the restored recomposed image, it prevents the network to convergence during training. Following this idea we tried a fixed Gaussian blur with parameters set to the reported values of [3] as well as learning a blur kernel (initialized by a Gaussian) during training. In both cases we observed no convergence during training. Therefore, we apply this smoothing only for the final application of the neural network.

### 3.3 Deployment

During deployment we feed input patches of size  $65 \times 65$  into our neural network with stride 5. Using overlapping patches helps to average multiple predictions. For recombination of overlapping patches we apply a 2-dimensional Hanning window to each patch to favour pixel values in the patch center and devalue information at the border of the patch.

While the predicted images  $\hat{X}$  generated by our neural network contain well-defined sharp edges we observed desaturation in color contrast. To correct the color of the predicted image we replace its *ab*-channel in the *Lab* color space by the *ab*-channel of the FBA results (compare Figure 4).

Regarding runtime the most expensive step is the frequency band analysis. Given a burst of 14 images of size  $1000 \times 700$  pixels the entire reconstruction process takes roughly 5 minutes per channel with our unoptimized implementation.




**Fig. 4.** Deblurring a burst of degraded images from a groundtruth image (left) results in a desaturated image (middle). Therefore we correct those colors (right image) using color transfer.

## 4 Experiments

To evaluate and validate our approach we conduct several experiments including a comprehensive comparison with state-of-the-art techniques on a real-world dataset, and a performance evaluation on a synthetic dataset to test the robustness of our approach with varying image quality of the input sequence.

### 4.1 Comparison on real-world dataset

We compare the restored images with other state-of-the-art multi-image blind deconvolution algorithms. In particular, we compare with the multichannel blind deconvolution method from Šroubek et al. [21], the sparse-prior method of [33] and the FBA method proposed in [3]. We used the data provided by [3], which contains typical photographs captured with hand-held cameras (iPad back camera, Canon 400D). As they are captured under various challenging lighting conditions they exhibit both noise and saturated pixels. As shown in [3] the FBA



**Fig. 5.** Comparison to state-of-the-art multi-frame blind deconvolution algorithms on real-world data. See the supplementary material for high-resolution images. Note that our approach produces the sharpest results except for the last scene, which could be caused by the color transfer described in Section 3.3.



algorithms demonstrated superior performance compared to previous state-of-the-art multi-image blind deconvolution algorithms [21, 33] in both reconstruction quality and runtime. Figure 5 shows crops of the deblurred results on these images. The high-resolution images are enclosed in the supplemental material. Our trained neural network featuring the FBA-like averaging yields comparable if not superior results compared to previous approaches [21, 33, 3]. In direct comparison to the FBA results, our method is better removing blur due to our additional prepended deconvolution module.

## 4.2 Deblurring bursts with varying number of frames and quality

Here, we analyse the performance of our approach depending on the burst “quality”. Sorting all images provided by [3] within one burst according to their PSNR beginning with images of strong blur and consequently adding sharper shots to the burst gives a series of bursts starting with images of poor quality up to bursts with at least one close-to-sharp shot. Since our architecture is trained for deblurring bursts with exactly 14 input images, we duplicated images of bursts with fewer frames. Figure 6 clearly indicates good performance of our neural network even for a relative small number of input images with strong blur artifact.


## 4.3 Deblurring image bursts without reasonable sharp frames

To further challenge our neural network approach, we artificially sampled image bursts from unseen images taken from the MS COCO validation set and blurred them by applying synthetic blur kernels of size  $14 \times 14$ . The restored sharp images from the input bursts of 14 artificially blurred images under absence of a close-to-sharp frame (best shot) are depicted in Figure 7. As the experiments indicate the explicit deconvolution step in our approach is absolutely necessary to handle these kind of snapshots and to remove blur artifacts. In contrast, while FBA [3] stands out in small memory footprint and fast processing times it clearly failed to recover sharp images for cases where no reasonably sharp frame is available amongst the input sequence.

## 4.4 Comparing to a baseline version

One might ask, how our trained neural network compares to an approach that applies the methods of Chakrabarti [11] and Delbracio and Sapiro [3] subsequently, each in a separate step. We fine-tuned the provided weights from [11] in combination with our FBA-layer. Figure 8 shows the training progress for an exemplar patch, where the improvement in sharpness is clearly visible.


In addition, we run the *entire* pipeline of Chakrabarti [11] including the costly non-blind deconvolution EPLL step and afterwards FBA. The approach is significantly slower and results in less sharp reconstructions (see Fig. 9).



**Fig. 6.** FBA and our algorithm are compared on bursts with a growing number of images of increasing quality. The individual images are sorted according to their PSNR starting with the most blurry images. The input images were taken from [3].

#### 4.5 Spatially-varying blur

To test whether our network is also able to deal with spatially-varying blur we generated a burst of images degraded by non-stationary blur. To this end, we took one of the recorded camera trajectories of [34] that are provided on



**Fig. 7.** Comparing FBA (third column) and our trained neural network (fourth column) against best shot and a typical shot. These images are taken from the validation set. For image bursts without a single sharp frame lucky imaging approaches fail due to a missing explicit deconvolution step, while our approach gives reasonable results.

the project webpage<sup>2</sup>. The camera trajectory has been recorded with a Vicon system at 500 fps and represents the camera motion during a slightly longer-exposed shot (1/30s). The trajectory comprises a 6-dimensional time series with 167 time samples. We divided this time series into 8 fragments of approximately equal lengths.

With a Matlab script (see Supplemental material) 8 spatially-varying PSFs are generated as shown at the bottom of Fig. 10.

<sup>2</sup> <http://webdav.is.mpg.de/pixel/benchmark4camerashake/>



**Fig. 8.** The combination of the work of Chakrabarti [11] and Delbracio et al. [3] can be considered as a baseline version of our neural network. We fine-tuned the published weights from the work of Chakrabarti [11] in an end-to-end fashion in combination with our FBA-layer. The left-most patch is the ground-truth patch. Note how the sharpness continuously increases with training.



**Fig. 9.** Comparison to a baseline approach of simply stacking [11] and [3]. Without end-to-end training ringing-artifacts are clearly visible on the blue roof. They are significantly dampened after training.

The spatially varying kernels of size  $17 \times 17$  pixels are applied using the Efficient Filter Flow model of [35]. The results of our network along with re-



sults of FBA for three example images are shown in Fig. 10. Our results are consistently sharper and demonstrate that our approach is also able to correct for spatially-varying blur.



**Fig. 10.** Comparison to FBA on image sequences with spatially-varying blur. Our approach is able to reconstruct consistently sharper images.

## 5 Conclusion, Limitations and Future Work

We presented a discriminative approach for multi-frame blind deconvolution (BD) by posing it as a nonlinear regression problem. As a function approximator, we use a deep layered neural network, whose optimal parameters are learned from artificially generated data. Our proposed network architecture draws inspiration from two recent works as (a) a neural network approach to single image blind deconvolution of Chakrabarti [11], and (b) the Fourier Burst Accumulation (FBA) algorithm of Delbracio and Sapiro [3]. The latter takes a burst of images as input and combines them through a weighted average in the frequency domain to a single sharp image. We reformulated FBA as a learning method and casted it into a deep layered neural network. Instead of resorting to heuristics and hand-tuned parameters for weight computation, we learn optimal weights as network parameters through end-to-end training.

By prepending parts of the network of Chakrabarti to our FBA network we are able to extend its applicability by alleviating the necessity of a close-to-sharp frame being amongst the image burst. Our system is trained end-to-end on a set of artificially generated training examples, enabling competitive performance in multi-frame BD, both with respect to quality and runtime. Due to its novel information sharing in the frequency band analysis stage and its explicit deconvolution step, our network outperforms state-of-the-art techniques like FBA [3] especially for bursts with few severely degraded images.

Our contribution resides at the experimental level and despite competitive results with state-of-the-art, our proposed approach is subject to a number of limitations. However, at the same time it opens up several exciting directions for future research:

- Our proposed approach doesn’t exploit the temporal structure of the input image sequence, which encodes valuable information about intra-frame blur and inter-frame image mis-alignment [24–26]. Embedding our described network into a network architecture akin to the spatio-temporal auto-encoder of Pătrăucean et al. [36] might enable such non-trivial inference.
- Our current model assumes a static scene and is not able to handle object motion. Inserting a Spatial Transformer Network Layer [37] which also facilitates optical flow estimation [36] could be an interesting avenue to capture and correct for object motion occurring between consecutive frames.

**Acknowledgement.** This work has been partially supported by the DFG Emmy Noether fellowship Le 1341/1-1 and an NVIDIA hardware grant.

## References

1. Burger, H.C., Schuler, C.J., Harmeling, S.: Image denoising: Can plain neural networks compete with bm3d? In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society (2012)
2. Schuler, C., Burger, H., Harmeling, S., Scholköpf, B.: A machine learning approach for non-blind image deconvolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society (2013)
3. Delbracio, M., Sapiro, G.: Burst deblurring: Removing camera shake through fourier burst accumulation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society (2015)
4. Wang, R., Tao, D.: Recent progress in image deblurring. arXiv preprint arXiv:1409.6838 (2014)
5. Sun, L., Cho, S., Wang, J., Hays, J.: Edge-based blur kernel estimation using patch priors. In: IEEE International Conference in Computational Photography (ICCP), IEEE (2013) 1–8
6. Michaeli, T., Irani, M.: Blind deblurring using internal patch recurrence. In: Proceedings of the European Conference on Computer Vision (ECCV). Springer (2014) 783–798
7. Xu, L., Ren, J.S., Liu, C., Jia, J.: Deep convolutional neural network for image deconvolution. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q., eds.: Advances in Neural Information Processing Systems (NIPS). Curran Associates, Inc. (2014) 1790–1798
8. Rosenbaum, D., Weiss, Y.: The return of the gating network: Combining generative models and discriminative training in natural image priors. In: Advances in Neural Information Processing Systems (NIPS). (2015) 2665–2673
9. Schuler, C.J., Hirsch, M., Harmeling, S., Schölkopf, B.: Learning to deblur. IEEE Transactions on Pattern Analysis and Machine Intelligence (2015)
10. Sun, J., Cao, W., Xu, Z., Ponce, J.: Learning a convolutional neural network for non-uniform motion blur removal. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society (2015)
11. Chakrabarti, A.: A neural approach to blind motion deblurring. In: Proceedings of the European Conference on Computer Vision (ECCV). (2016)
12. Hradiš, M., Kotera, J., Zemčík, P., Šroubek, F.: Convolutional neural networks for direct text deblurring. In: Proceedings of BMVC. (2015) 2015–10
13. Svoboda, P., Hradi, M., Mark, L., Zemck, P.: Cnn for license plate motion deblurring. In: IEEE International Conference on Image Processing (ICIP). (2016) 3832–3836
14. Loktyushin, A., Schuler, C., Scheffler, K., Schölkopf, B.: Retrospective motion correction of magnitude-input mr images. In: Machine Learning Meets Medical Imaging. Springer (2015) 3–12
15. Zoran, D., Weiss, Y.: From learning models of natural image patches to whole image restoration. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV), IEEE Computer Society (2011) 479–486
16. Hasinoff, S.W., Kutulakos, K.N., Durand, F., Freeman, W.T.: Time-constrained photography. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV), IEEE (2009) 333–340

17. Rav-Acha, A., Peleg, S.: Two motion-blurred images are better than one. *Pattern recognition letters* **26** (2005) 311–317
18. Zhang, H., Wipf, D.P., Zhang, Y.: Multi-observation blind deconvolution with an adaptive sparse prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **36** (2014) 1628–1643
19. Chen, J., Yuan, L., Tang, C.K., Quan, L.: Robust dual motion deblurring. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, IEEE Computer Society (2008) 1–8
20. Cai, J.F., Ji, H., Liu, C., Shen, Z.: Blind motion deblurring using multiple images. *Journal of computational physics* **228** (2009) 5057–5071
21. Šroubek, F., Milanfar, P.: Robust multichannel blind deconvolution via fast alternating minimization. *Image Processing, IEEE Transactions on* **21** (2012) 1687–1700
22. Zhu, X., Šroubek, F., Milanfar, P.: Deconvolving psfs for a better motion deblurring using multiple images. In: *Computer Vision—ECCV 2012*. Springer (2012) 636–647
23. Zhang, H., Carin, L.: Multi-shot imaging: joint alignment, deblurring and resolution-enhancement. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE Computer Society (2014)
24. Zhang, H., Yang, J.: Intra-frame deblurring by leveraging inter-frame camera motion. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE Computer Society (2015)
25. Kim, T.H., Nah, S., Lee, K.M.: Dynamic scene deblurring using a locally adaptive linear blur model. *arXiv preprint arXiv:1603.04265* (2016)
26. Ito, A., Sankaranarayanan, A.C., Veeraraghavan, A., Baraniuk, R.G.: Blurburst: Removing blur due to camera shake using multiple images. *ACM Trans. Graph.*, Submitted **3** (2014)
27. Garrel, V., Guyon, O., Baudoz, P.: A highly efficient lucky imaging algorithm: Image synthesis based on fourier amplitude selection. *Publications of the Astronomical Society of the Pacific* **124** (2012) 861
28. Delbracio, M., Sapiro, G.: Hand-held video deblurring via efficient fourier aggregation. *Computational Imaging, IEEE Transactions on* **1** (2015) 270–283
29. Lin, T., Maire, M., Belongie, S.J., Bourdev, L.D., Girshick, R.B., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. *arXiv preprint arXiv:1405.0312* (2014)
30. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
31. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.* **12** (2011) 2121–2159
32. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015) Software available from tensorflow.org.
33. Zhang, H., Wipf, D., Zhang, Y.: Multi-image blind deblurring using a coupled adaptive sparse prior. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE Computer Society (2013) 1051–1058



34. Köhler, R., Hirsch, M., Mohler, B., Schölkopf, B., Harmeling, S.: Recording and playback of camera shake: Benchmarking blind deconvolution with a real-world database. In: Proceedings of the European Conference on Computer Vision (ECCV). Springer (2012) 27–40
35. Hirsch, M., Sra, S., Schölkopf, B., Harmeling, S.: Efficient filter flow for space-variant multiframe blind deconvolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE (2010)
36. Pătrăucean, V., Handa, A., Cipolla, R.: Spatio-temporal video autoencoder with differentiable memory. arXiv preprint arXiv:1511.06309 (2015)
37. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in Neural Information Processing Systems (NIPS). (2015) 2008–2016