

Learning to Deblur

Christian J. Schuler, Michael Hirsch, Stefan Harmeling, Bernhard Schölkopf
{cschuler,mhirsch,harmeling,bs}@tuebingen.mpg.de

Max Planck Institute for Intelligent Systems

We describe a learning-based approach to blind image deconvolution. It uses a deep layered architecture, parts of which are borrowed from recent work on neural network learning, and parts of which incorporate computations that are specific to image deconvolution. The system is trained end-to-end on a set of artificially generated training examples, enabling competitive performance in blind deconvolution, both with respect to quality and runtime.

1. Introduction

Blind image deconvolution is the task of recovering the underlying sharp image from a recorded image corrupted by blur and noise. Examples of such distortions are manifold: In photography, long exposure times might result in camera shake, often in combination with image noise caused by low light conditions. Lens aberrations can create blur across images in particular for wide apertures. In astronomy and remote sensing, atmospheric turbulence blurs images. All these image reconstruction tasks are examples of *inverse problems*, which are particularly hard to solve since not only the underlying image is unknown, but also the blur.

We assume that the blurred image \mathbf{y} is generated by linearly transforming the underlying image \mathbf{x} (sometimes called the “true image” or “latent image”) by a convolution (denoted by $*$), and additive noise \mathbf{n} ,

$$\mathbf{y} = \mathbf{k} * \mathbf{x} + \mathbf{n}. \quad (1)$$

The task of blind image deconvolution is to recover \mathbf{x} given only the blurry image \mathbf{y} , without knowing \mathbf{k} . A number of approaches have recently been proposed to recover the true image from blurred photographs, e.g., [Fer+06; CL09; XJ10]. Usually these methods assume some sparsity inducing image prior for \mathbf{x} , and follow an iterative, multi-scale estimation scheme, alternating between blur and latent image estimation.

The idea of the proposed method is to “unroll” this reconstruction procedure and pose it as a nonlinear regression problem, where the optimal parameters are learned from artificially generated data. As a function approximator, we use a layered architecture akin to a deep neural network

or multilayer perceptron. Some of the layers are convolutional, as popular in many recent approaches, while others are non-standard and specific to blind deconvolution. Overall, the system is inspired by [BG91] who formulated the idea of neural networks (NNs) as general processing schemes with large numbers of parameters.

Using extensive training on a large image dataset in combination with simulated camera shakes, we train the blind deconvolution NN to automatically obtain an efficient procedure to approximate the true underlying image, given only the blurry measurement.

A common problem with NNs, and in particular with large custom built architectures, is that the devil is in the details and it can be nontrivial to get systems to function as desired. We thus put particular emphasis on describing the implementation details, and we make the code publicly available.

Main contributions: We show how a trainable model can be designed in order to be able to *learn* blind deconvolution. Results are comparable to the state of the art of hand-crafted approaches and go beyond when the model is allowed to specialize to a particular image category.

2. Related work

Some recent approaches to blind deconvolution in photography have already been mentioned above. We refer the interested reader to [WZ13] for an overview of the subject, and focus only on how NNs have been previously used for deconvolution.

Neural networks have been used extensively in image processing. Comprehensive reviews are [ERH02; De +03], both of which present broad overviews of applying NNs to all sorts of image processing operations, including segmentation, edge detection, pattern recognition, and nonlinear filtering.

Image deconvolution, often also called image reconstruction, has been approached with NNs for a long time. However, these approaches are quite different from our proposed work:

- *NN to identify the type of blur:* [Aiz+06] and similarly [KN11] apply NNs to identify the type of blur from a restricted set of parametrized blurs, e.g. linear-motion, out-of-focus and Gaussian blur, and possibly their parameters.
- *NN to model blurry images:* [CD91] models the blurry image as the result of a neural network where at the different layers the blur kernel and the true image appear.
- *NN to inversely filter blurry images:* [TOM96] learns an inverse filter represented by a NN to deblur text.
- *NN to optimize a regularized least squares objective:* [SF94] proposes also the common two-stage procedure to first estimate the blur kernel and then to recover the image. For both tasks Hopfield networks are employed to solve the optimization problems.

- *NN to remove colored noise*: The previous chapter presented a method for non-blind deblurring that starts with a straight-forward division in Fourier space and then removes the resulting artifacts (mainly colored noise) with large neural networks.

Other learning-based approaches for blind deconvolution try to learn the deconvolution solution for a single image, in particular they attempt to learn an appropriate sparse representation for a given blurry image, e.g. [HHY10]. In our work, we follow a different strategy: instead of learning the solution or part of a solution for a single fixed image, we use a neural network to learn a general procedure that is directly applicable to other images and to different blurs. Closest to our approach is the work of [Sch+13], who train a deblurring procedure based on regression tree fields. However, even though their approach is not limited to a specific blur or a specific image, they consider only the problem of *non-blind* deblurring, i.e. their method assumes that the blur kernel is known. This is in contrast to our contribution, which demonstrates how to train a NN to solve the *blind* deconvolution problem. This is a much harder problem, since not only do we have to solve an underdetermined *linear* problem (originating from the least-squares formulation of non-blind deconvolution), but an underdetermined *bilinear* problem, which appears since the entries of the unknown blur kernel k and the pixel values of the unknown image x are multiplied by the convolution, see Eq. (1).

Finally we note that *deconvolutional networks*, introduced in [Zei+10] and further extended in [ZTF11], are not architectures for image deconvolution. Instead, they use convolutions to link sparse features to given images. Their goals are good image representations for tasks such as object recognition and image denoising.

3. Blind deconvolution as a layered network

Existing fast blind deconvolution methods work by alternating between the following three steps [XZJ13; XJ10; CL09]:

1. *Feature extraction* computes image representations that are useful for kernel estimation. These representations may simply be the gradient of the blurry image and the gradient of the current estimate, possibly thresholded in flat regions [XJ10]; they may also be preliminary estimates of the sharp image, for instance computed by heuristic nonlinear filtering [CL09].
2. *Kernel estimation* uses the extracted features to estimate the convolution kernel.
3. *Image estimation* attempts to compute an approximation of the sharp image using the current kernel estimate.

We will represent these steps by a trainable deep neural network, thus adding more flexibility to them and allowing them to optimally adapt to the problem. The layers of the network alternate between (1) a local convolutional estimation to extract good features, (2) the estimation of the blur kernel, globally combining the extracted features, and finally (3) the estimation of the sharp image. Parts (2) and (3) are fixed (having only one hyper-parameter for regularization). The free

parameters of the network appear in part (1), the feature extraction module. Thus, instead of having to learn a model on the full dimensionality of the input image, which would not be doable using realistic training set sizes, the learning problem is naturally reduced to learning filters with a limited receptive field.

3.1. Architecture layout

Below, we describe the different parts of the network.

3.1.1. Feature extraction module

What makes a good feature for kernel estimation can reasonably be assumed to be a translation invariant problem, i.e., independent from the position within the image. We therefore model the feature extractors using shared weights applying across all image locations, i.e., as a convolutional NN layer,¹ creating several feature representations of the image. This is followed by two layers that introduce nonlinearity into the model. First, every value is transformed by a tanh-unit, then the feature representations are pixel-wise linearly combined to new hidden images, formally speaking

$$\tilde{\mathbf{y}}_i = \sum_j \alpha_{ij} \tanh(\mathbf{f}_j * \mathbf{y}) \quad \text{and} \quad \tilde{\mathbf{x}}_i = \sum_j \beta_{ij} \tanh(\mathbf{f}_j * \mathbf{y}) \quad (2)$$

where \mathbf{f}_j are the filters of the convolution layer (shared for $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{y}}_i$), the function tanh operates coordinate-wise, and α_{ij} and β_{ij} are the coefficients to linearly combine the hidden representations. Note that we usually extract several gradient-like images $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{y}}_i$. Depending on the desired nonlinearity, these two layers can be stacked multiple times, leading to the final image representations based on features useful for kernel estimation.

3.1.2. Kernel estimation module

Given $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{y}}_i$ which contain features tuned for optimal kernel estimation, the kernel $\tilde{\mathbf{k}}$ can be estimated by minimizing

$$\sum_i \|\tilde{\mathbf{k}} * \tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i\|^2 + \beta_k \|\tilde{\mathbf{k}}\|^2 \quad (3)$$

for $\tilde{\mathbf{k}}$ given the results from the previous step $\tilde{\mathbf{x}}_i$ and their blurry counterparts $\tilde{\mathbf{y}}_i$. Assuming no noise and a kernel without zeros in its power spectrum, the true gradients of the sharp image \mathbf{x} and its blurred version \mathbf{y} would return the true kernel for $\beta_k = 0$. Typically, in existing methods $\tilde{\mathbf{y}}_i$ are just the gradients of the blurry image, while here these can also be learned representations

¹Note that this convolution has nothing to do with the convolution appearing in our image formation model (Eq. (1)) — causally, it goes in the opposite direction, representing one step in the inverse process turning the blurry image into an estimate of the underlying image.

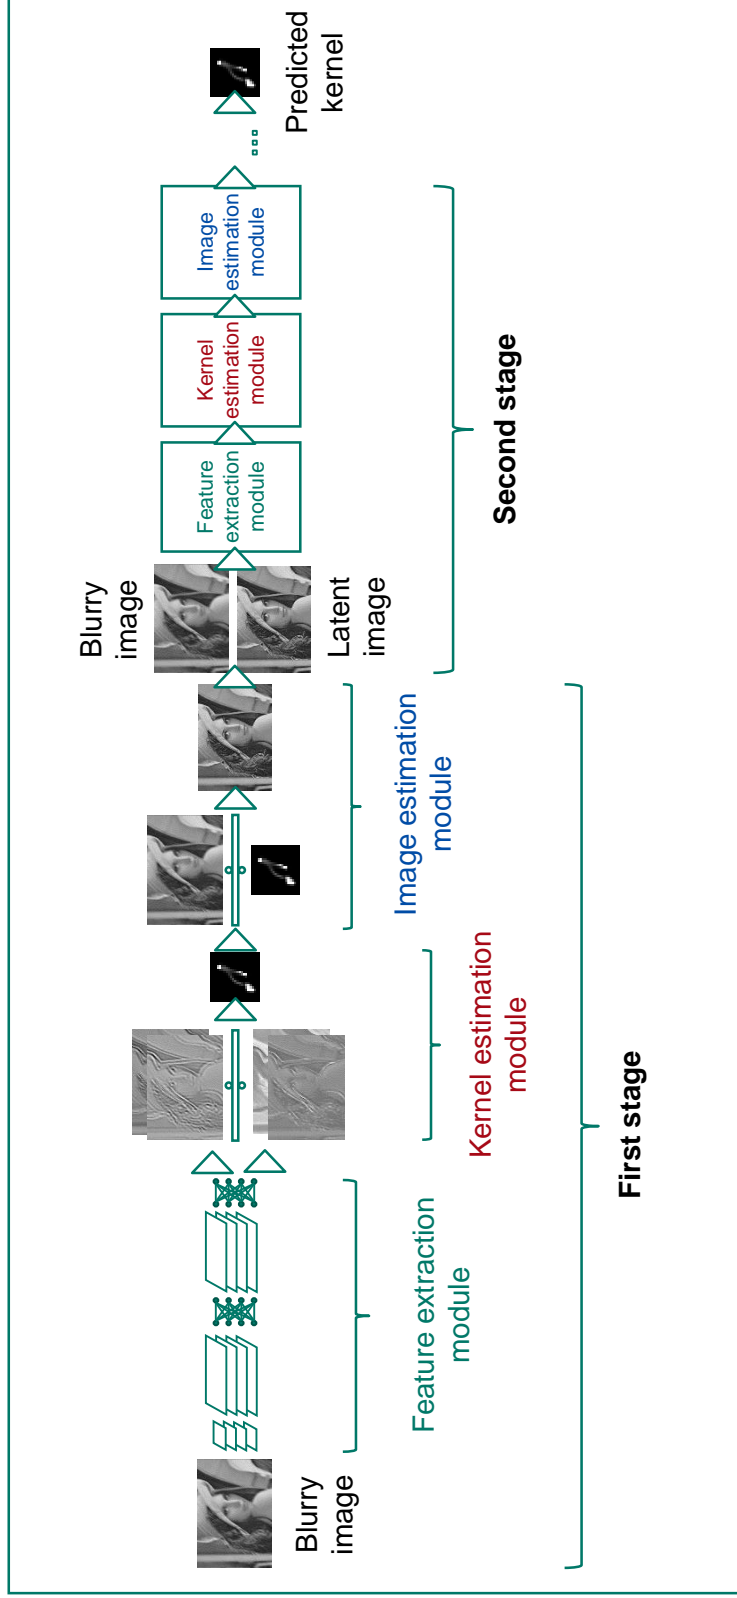


Figure 1: Architecture of our proposed blind deblurring network. First the *feature extraction module* transforms the image to a learned gradient-like representation suitable for kernel estimation. Next, the kernel is estimated by division in Fourier space, then similarly the latent image. The next stages, each consisting of these three operations, operate on both the blurry image and the latent image.

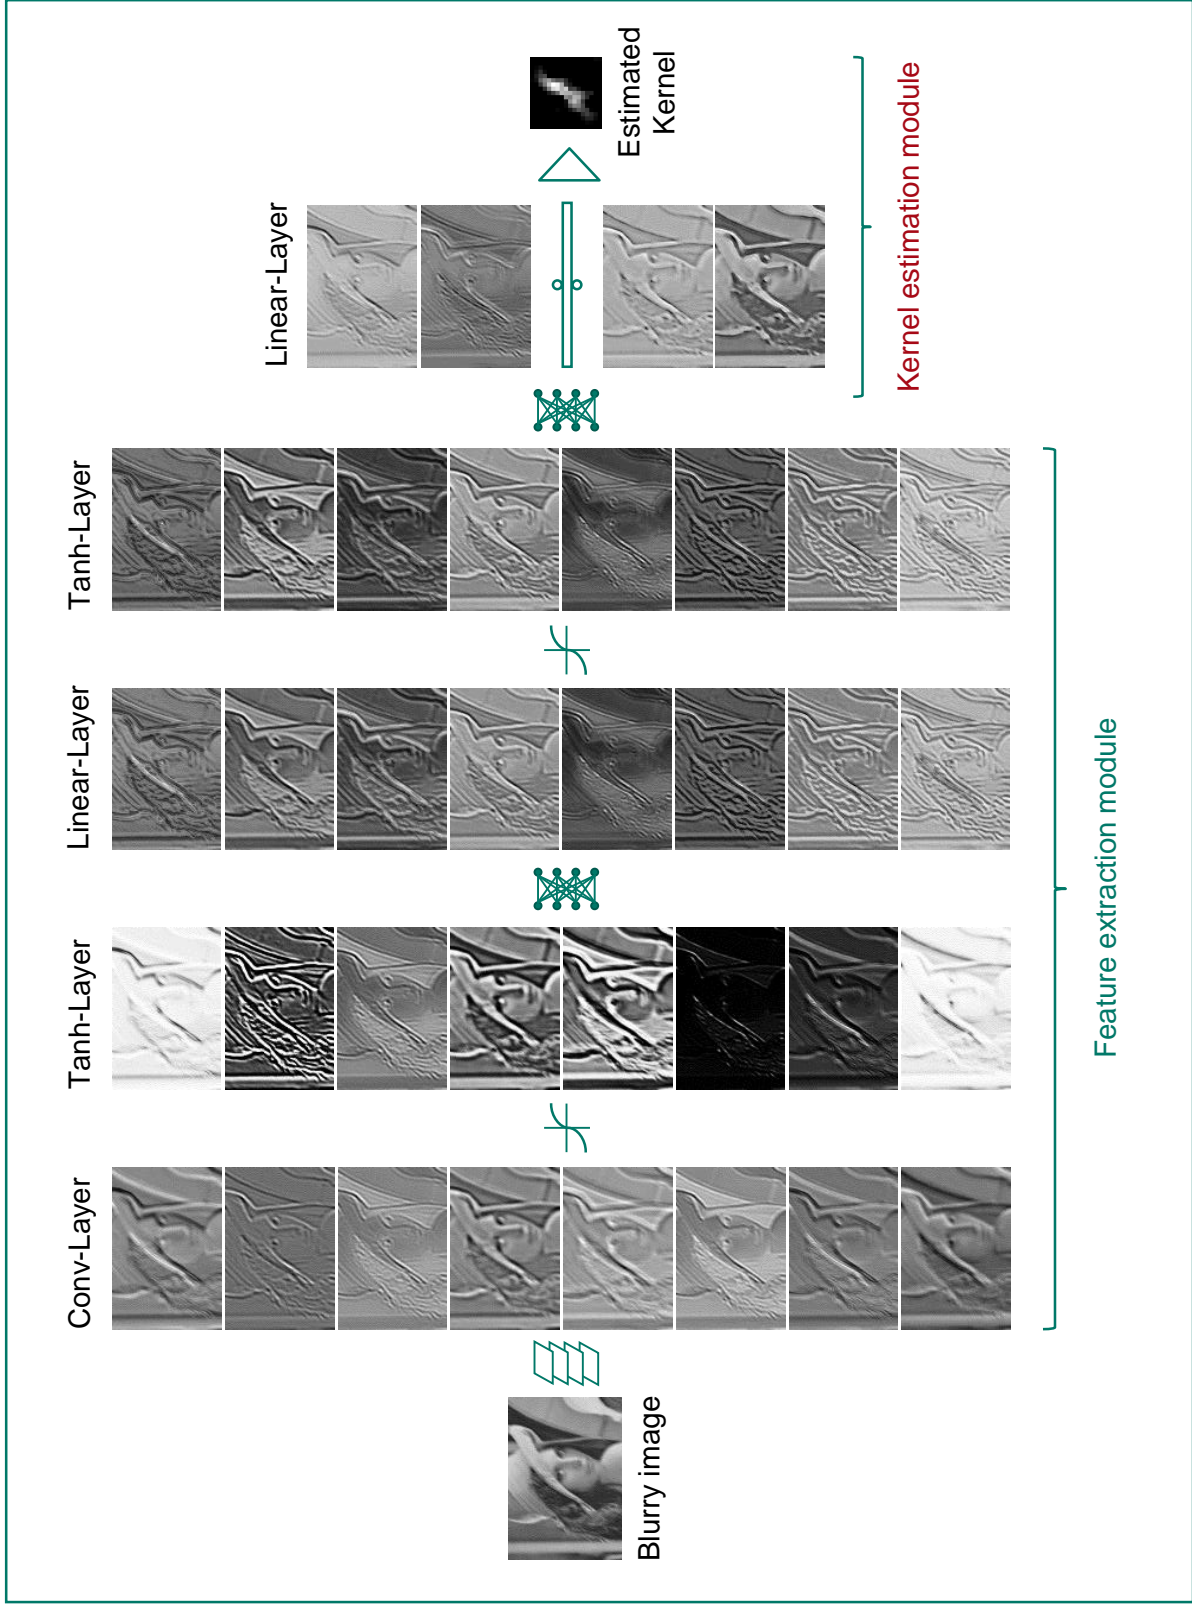


Figure 2: Intermediary outputs of a single-stage NN with architecture $8 \times \text{Conv}$, 8×8 , Tanh , 8×8 , Tanh , 8×4 .

predicted from the previous layer. The minimization problem can be solved in one step in Fourier space if we assume circular boundary conditions [CL09]:

$$\tilde{\mathbf{k}} = F^H \frac{\sum_i \overline{F\tilde{\mathbf{x}}_i} \odot F\tilde{\mathbf{y}}_i}{\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k}. \quad (4)$$

Here F is the discrete Fourier transform matrix, \odot the Hadamard product, $\bar{\cdot}$ the complex conjugate of a vector \mathbf{v} , and the division and absolute value are performed element-wise. The one step solution is only possible because we use a simple Gaussian prior on the kernel. We call this step the *quotient layer*, which is an uncommon computation in NNs that usually only combine linear layers and nonlinear thresholding units. The final kernel is returned by cropping to the particular kernel size and thresholding negative values. To reduce artifacts from the incorrect assumption of circular boundary conditions, the borders of the image representations are weighted with a Barthann window such that they smoothly fade to zero.

Due to varying size of the input image, we set $\beta_k = 10^{-4}$ for numerical stability only. This forces the network to not rely on the prior, which would lose importance for a larger input image relative to the likelihood term (since the kernel size is fixed).

3.1.3. Image estimation module

Before adding another *feature extraction module*, the estimated kernel is used to obtain an update of the sharp latent image: analogously to Eq. (3), we solve

$$\|\tilde{\mathbf{k}} * \tilde{\mathbf{x}} - \mathbf{y}\|^2 + \beta_x \|\tilde{\mathbf{x}}\|^2 \quad (5)$$

for $\tilde{\mathbf{x}}$, which can also be performed with a quotient layer. This can be done in one step (which would not be possible when using a sparse prior on $\tilde{\mathbf{x}}$). The following convolution layer then has access to both the latent image and the blurry image, which are stacked along the third dimension (meaning that the learned filters are also three-dimensional). The hyper-parameter β_x is also learned during training.

3.2. Iterations as stacked networks

The *feature extraction module*, *kernel estimation module* and the *image estimation module* can be stacked several times, corresponding to multiple iterations in non-learned blind deconvolution methods. This leads to a single NN that can be trained end-to-end with back-propagation by taking the derivatives of all steps (see Appendix A for the derivatives of the solutions to Eq. (3) and the analogous Eq. (5)), increasing the performance as shown in Fig. 3 (but at the same time increasing runtime).

Similar to other blind deconvolution approaches, we also use a multi-scale procedure. The first (and coarsest) scale is just a network as described above, trained for a particular blur kernel size. For the second scale, the previous scale network is applied to a downsampled version of the blurry image, and its estimated latent image is upsampled again to the second scale. The second scale network then takes both the blurry image and the result of the previous scale as input. We repeat these steps until the final scale can treat the desired blur kernel size on the full resolution image.




Figure 3: Deeper networks are better at kernel prediction. One stage for kernel prediction consists of a convolutional layer, two hidden layers and a kernel estimation module.

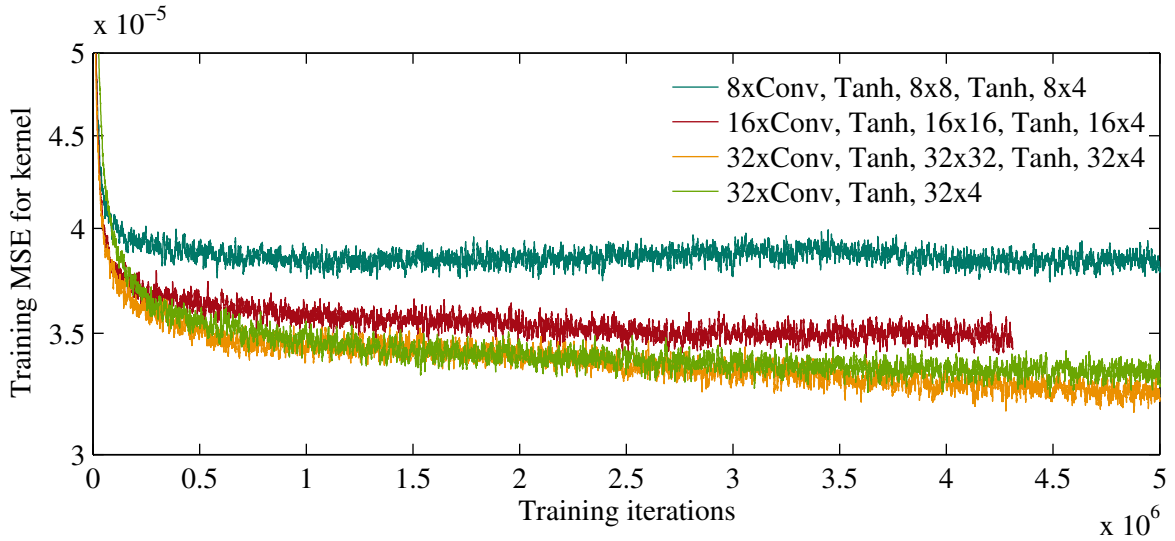


Figure 4: The performance of the network for kernel estimation depends on the architecture. More filters in the convolutional layer and more hidden layers are better.

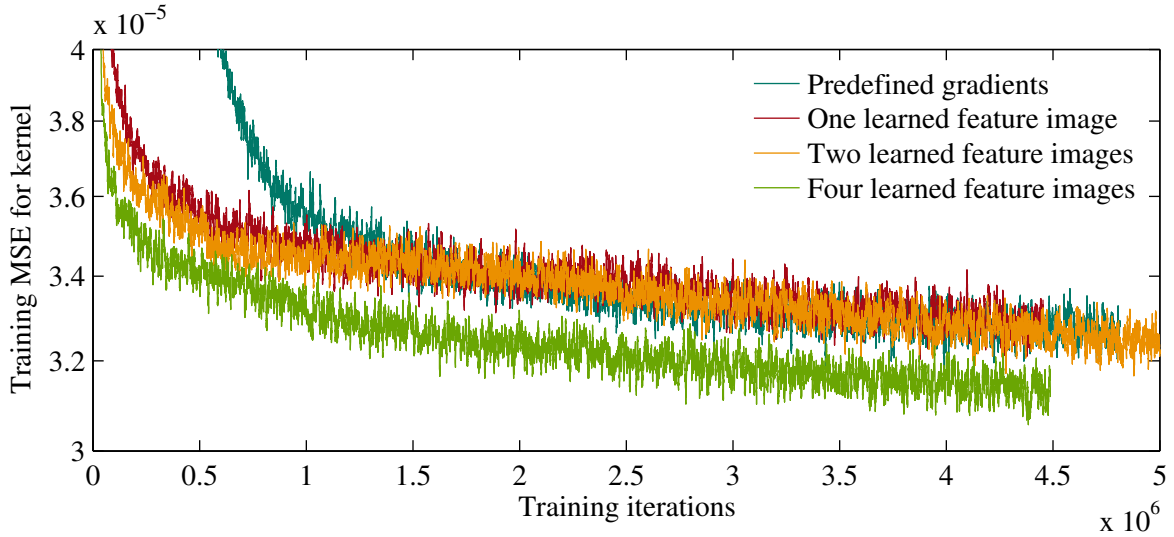


Figure 5: The performance of the network depends on the number of the predicted gradient-like images used in the kernel estimation module. Predefining \tilde{y}_i to x- and y-gradients and not learning these representations slows down training.




Figure 6: Influence of the parameters of ADDELTA [Zei12] or stochastic gradient descent (SGD) on the convergence of the network.




Figure 7: Examples of blurs sampled from a Gaussian process (left: 33 px, right: 17 px).

3.3. Training

To train the network, we generate pairs of sharp and blurry images. The sharp image is sampled from a subset of about 1.6 million images of the ImageNet [Den+09] dataset, randomly cropped to a size of 256×256 . Next, a blur trajectory is generated by sampling both its x- and y-coordinates separately from a Gaussian Process

$$f_x(t), f_y(t) \sim \mathcal{GP}(0, k(t, t')), \quad k(t, t') = \sigma_f^2 \left(1 + \frac{\sqrt{5}|t - t'|}{l} + \frac{5(t - t')^2}{3l^2} \right) \exp \left(-\frac{\sqrt{5}|t - t'|}{l} \right), \quad (6)$$

which is a Matérn covariance function with $\nu = 3/2$ [RW06]. The length scale l is set to 0.3, the signal standard deviation σ_f to $1/4$. The coordinates are then scaled to a fixed blur kernel size and the final kernel is shifted to have its center of mass in the middle. This simple procedure generates realistic looking blur kernels, examples for both small and large kernel sizes are illustrated in Fig. 7. For every setting, we generate 1 million noise-free training examples, and add Gaussian noise during training (by default, $\sigma = 0.01$).

To avoid that the training process is disturbed by examples where the blur kernel cannot be estimated, e.g., a homogeneous area of sky without any gradient information, we reject examples where less than 6% pixels have gradients in x- and y-direction with an absolute value 0.05 or above.

As described in the previous subsection, a network for a certain blur kernel, i.e., a particular scale, consists of several stages, each iterating between the *feature extraction*, *kernel estimation* and the *image estimation module*.

We use pre-training for our network: we start by training only one stage minimizing the L2 error between the estimated and the ground truth kernel, then add a second stage after about 1 million training steps. For the next 1000 steps, the parameters of the first stage stay fixed and only the parameters of the second stage are updated. After that, the network is trained end-to-end, until a potential next stage is added.

For the update of the parameters, convergence proved to be best with ADADELTA [Zei12], a heuristic weighting scheme of parameter updates using gradients and updates from previous training steps. For the influence of its parameters on the convergence speed see Fig. 6. For our experiments, we choose a learning rate of 0.01 and a decay rate of 0.95. Moreover, it makes training more robust to outliers with strong gradients since it divides by the weighted root-mean-square of the seen gradients, including the current one. Responsible for the mentioned strong gradients are typically images dominated by abrupt step-edges, which create ringing artifacts in

Blur size	255×255	800×800	1024×1024	2048×2048
17×17	1.1	5.2	9.5	53.1
25×25	1.2	14.3	18.5	89.1
33×33	1.6	10.7	22.6	91.9

Table 1: The method is very fast: runtime in seconds for kernel estimation with varying image size on an Intel i5 in Matlab.

the deconvolution Eq. (5). In ImageNet these often are photos of objects with trimmed background. To make the training even more robust, we don’t back-propagate examples with an error above 10 times the current average error.

4. Implementation

We make the code for both training and testing our method available for download. For training, we use our own C++/CUDA-based neural network toolbox. After training once for a certain blur class (e.g., camera shake), which takes about two days per stage, applying the network is very fast and can be done in Matlab without additional dependencies.

The runtimes on an Intel i5 using only Matlab are shown in Table 1. The most expensive calculation is the creation of the multiple hidden representations in the convolutional layer of the NN (in this example: 32).

5. Experiments

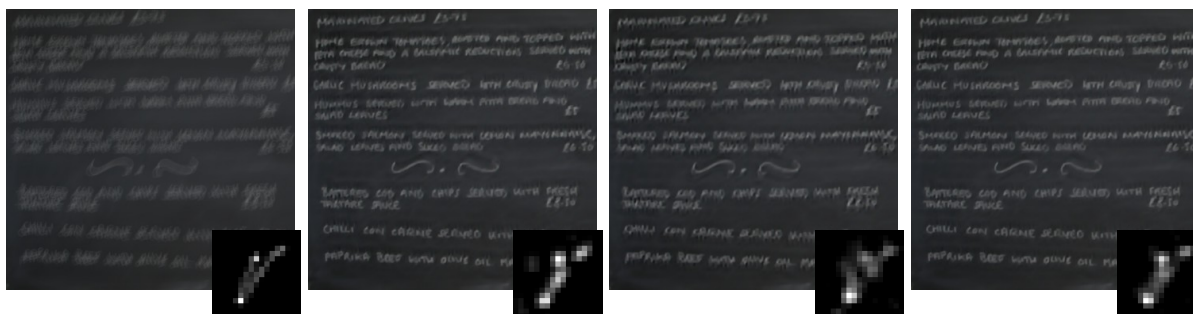

If not otherwise stated, all experiments were performed with a multi-scale, triple stage architecture. We use up to three scales for kernels of size 17×17, 25×25, 33×33. On each scale each *feature extraction module* consists of a convolution layer with 32 filters, a tanh-layer, a linear recombination to 32 new hidden images, a further tanh-layer, and a recombination to four gradient-like images, two for \tilde{x}_i and \tilde{y}_i each (in the third stage: eight gradient-like images). In the case of the network with blur kernels of size 33×33, we deconvolve the estimated kernels with a small Gaussian with $\sigma = 0.5$ to counter the over-smoothing effect of the L2 norm used during training. For the specific choice of the architecture, we refer to the influence of model parameters on the kernel estimation performance in Figs. 3 to 4 .

5.1. Image content specific training

A number of recent works [HY12; Sun+13; Wan+13] have pointed out the shortcoming of state-of-the-art algorithms [CL09; XJ10] to depend on the presence of strong salient edges and their diminished performance in the case of images that contain textured scenes such as natural landscape images. The reason for this is the deficiency of the so-called image prediction step, which



Figure 8: Typical example images of *valley* (top row) and *blackboard* (bottom row) categories from ImageNet used for content specific training.



Blurry image with ground truth kernel

Deblurring result of Xu et al. [XZJ13]

Deblurring result with content *agnostic* training

Deblurring result with content *specific* training

Figure 9: Comparison of deblurring results for NNs that have been trained with image examples from the entire ImageNet dataset (content *agnostic*) and from particular subsets (content *specific*), i.e., image categories *valley* (top row) and *blackboard* (bottom row). We also show the results of the state-of-the-art method [XZJ13].

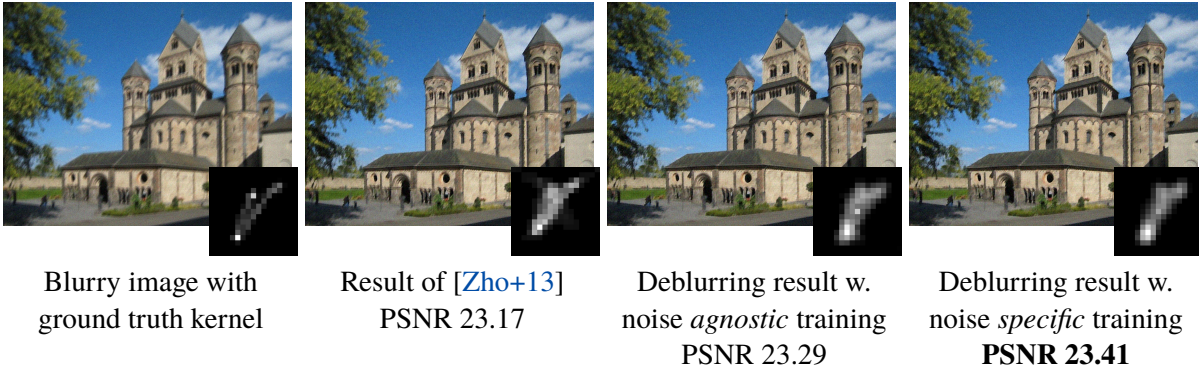


Figure 10: Comparison of deblurring results for NNs that have been trained with different amounts of noise added to the sample images during training. The network that has been trained with the same amount of noise as the input blurry image (5% noise) performs best. We also show the results of a recently proposed deblurring method tailored for increased levels of noise [Zho+13].

applies a combination of bilateral and shock filtering to restore latent edges that are used for subsequent kernel estimation.

In this context, learning the latent image prediction step offers a great advantage: by training our network with a particular class of images, it is able to focus on those features that are informative for the particular type of image. In other words, the network learns *content-specific* nonlinear filters, which yield improved performance.

To demonstrate this, we used the same training procedure as described above, however, we reduced the training set to images from a specific image category within the ImageNet dataset. In particular, we used the image category *valley*² containing a total of 1395 pictures. In a second experiment, we trained a network on the image category *blackboard*³ with a total of 1376 pictures. Figure 8 shows typical example images from these two classes. Fig. 9 compares deblurring results of the state-of-the-art algorithm described in [XZJ13] with our approach trained on images sampled from the entire ImageNet dataset, and trained on the aforementioned image categories only. Note that the particular images shown in Fig. 9 were not used for training. We see that content specific training outperforms both content agnostic training and the generic state-of-the-art method.

5.2. Noise specific training

Typically, image noise impedes kernel estimation. To counter noise in blurry images, current state-of-the-art deblurring algorithms apply a denoising step during latent image prediction such

²ImageNet 2011 Fall Release > Geological formation, formation > Natural depression, depression > Valley, vale (<http://www.image-net.org/synset?wnid=n09468604>)

³ImageNet 2011 Fall Release > Artifact, artifact > Sheet, flat solid > Blackboard, chalkboard (<http://www.image-net.org/synset?wnid=n02846511>)

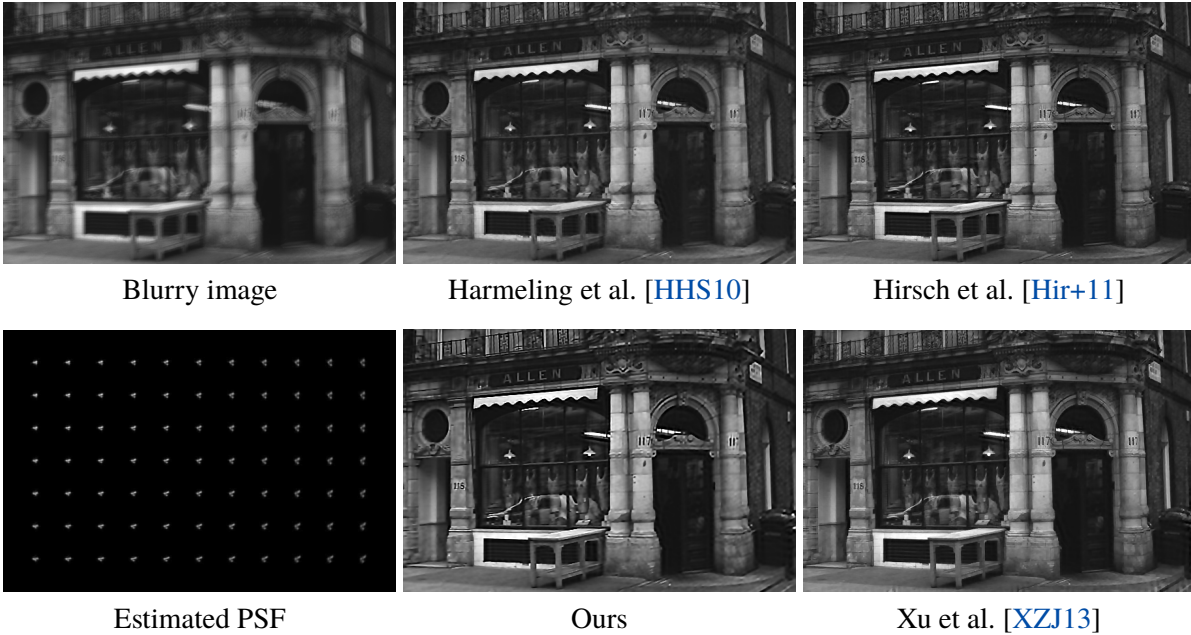


Figure 11: Comparison on *Butcher Shop* example [HHS10] of state-of-the-art deblurring methods for removing non-uniform blur together with our estimated PSF.

as bilateral filtering [CL09] or Gaussian filtering [XJ10]. However, in a recent work [Zho+13], the authors show that for increased levels of noise current methods fail to yield satisfactory results and propose a novel robust deblurring algorithm. Again, if we include image noise in our training phase, our network is able to adapt and learn filters that perform better in the presence of noise. In particular, we trained a network on images with Gaussian noise of 5% added during the training phase. Figure 9 compares the results for an image taken from [Zho+13] with 5% Gaussian noise for a network trained with 1% and 5% of added Gaussian noise during training, respectively. We also show the result of [Zho+13] and compare peak signal-to-noise ratio (PSNR) for objective evaluation. All results use the same non-blind deconvolution of [Zho+13]. The noise specific training is most successful, but even the noise agnostic NN outperforms the non-learned method on this example.

5.3. Spatially-varying blur

Since the prediction step of our trained NN is independent of the convolution model, we can also use it in conjunction with the recently proposed fast forward model of [Hir+11] to restore images with spatially-varying blur. To this end, we replace the objective function Eq. (5) with Eq. (8) of [Hir+11] in our kernel estimation module in a network trained for spatially invariant deconvolution. We solve for k in a two-step procedure: first we compute local blur kernels using the efficient filter flow (EFF) model of [Hir+10]; in a second step we project the blur kernels onto a motion basis aka [Hir+11], as explained below. Figure 11 shows a comparison between recent

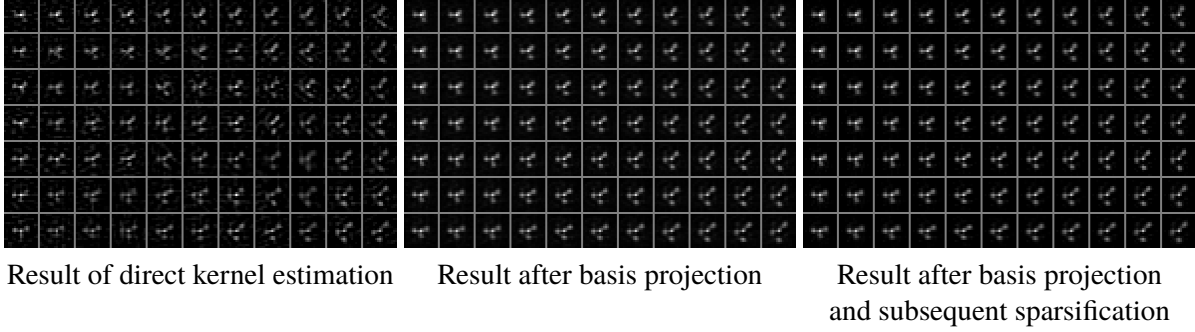


Figure 12: Visualisation of our kernel estimation in the case of spatially-varying blur for the *Butcher Shop* example shown in Fig. 11. The left panel shows the kernel estimated with Eq. (8), the middle and right panels show the kernel after applying a subsequent projection step to our motion basis, i.e. the result of Eq. (9) with η set to 1.0 and 0.0314 respectively.

state-of-the-art algorithms for spatially-varying blur along with our deblurring result that features comparable quality.

For the estimation of spatially-varying blur we solve the following objective

$$\sum_i \|\tilde{X}_i \tilde{\mathbf{k}} - \tilde{\mathbf{y}}_i\|^2 + \beta_k \|\tilde{\mathbf{k}}\|^2 \quad (7)$$

in our kernel estimation module. Here \tilde{X}_i denotes the EFF matrix of $\tilde{\mathbf{x}}$ (cf. Eq. (9) in [Hir+10]) and $\tilde{\mathbf{k}}$ the stacked sequence of local kernels $\tilde{\mathbf{k}}^{(r)}$, one for each patch that are enumerated by index r . Since Eq. (7) is quadratic in $\tilde{\mathbf{k}}$, we can solve for a local blur $\tilde{\mathbf{k}}^{(r)}$ in a single step

$$\tilde{\mathbf{k}}_{direct}^{(r)} \approx F^H \frac{\sum_i \overline{FC_r \text{Diag}(\mathbf{w}^{(r)}) \tilde{\mathbf{x}}_i} \odot (FC_r \text{Diag}(\mathbf{w}^{(r)}) \tilde{\mathbf{y}}_i)}{\sum_i |FC_r \text{Diag}(\mathbf{w}^{(r)}) \tilde{\mathbf{x}}_i|^2 + \beta_k}, \quad (8)$$

where C_r are appropriately chosen cropping matrices, and $\mathbf{w}^{(r)}$ are window functions matching the size of $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{y}}_i$. Note that Eq. (8) is only approximately true and is motivated by Eq. (8) in [Hir+11]. Subsequently, we project the estimated kernel computed by Eq. (8) to a motion blur kernel basis. In our experiments we use the same basis as [Hir+11] comprising translations within the image plane and in-plane rotations only. This additional projection step constrains the estimated blur to physically plausible ones. Formally, we compute

$$\tilde{\mathbf{k}}_{est}^{(r)} \approx B^{(r)} T_\eta \underbrace{\sum_r (B^{(r)})^T \tilde{\mathbf{k}}_{direct}^{(r)}}_{\boldsymbol{\mu}}, \quad (9)$$

where again we make use of the notation chosen in [Hir+11], i.e. $B^{(r)}$ denotes the motion blur kernel basis for patch r . Then $\boldsymbol{\mu}$ are the coefficients in the basis of valid motion blurs. T_η denotes a thresholding operator that sets all elements to zero below a certain threshold whereby the

threshold is chosen such that only η percent of entries remain non-zero. This thresholding step is motivated by [CL09], who also apply a hard thresholding step to the estimated kernels in order to get rid of spurious artefacts. Figure 12 shows the intermediate results of our kernel estimation procedure in the case of spatially-varying blur.

5.4. Comparisons

5.4.1. Benchmark Datasets

We evaluate our method on the standard test sets from [Lev+09; Sun+13]. The four images of [Lev+09] are 255×255 pixels in size and are artificially blurred each with eight different blur kernels and contain 1% additive Gaussian noise. The performance is illustrated in Fig. 13 on the left, with blur kernels sorted according to increasing size. We compare with Levin et al. [Lev+11], Cho and Lee [CL09], and Xu et al. [XZJ13]. While our method is competitive with the state of the art for small blur kernels, our method falls short in performance for blur kernel sizes above 25×25 pixels. We discuss reasons for this in Section 6.3. The second benchmark from [Sun+13] extends this dataset to 80 new images with about one megapixel in size each, using the same blur kernels as in [Lev+09]. Results are shown in Fig. 13 on the right. Here we compare with Levin et al. [Lev+11], Cho and Lee [CL09], Krishnan et al. [KTF11], Sun et al. [Sun+13], and Xu and Jia [XJ10], where however [Sun+13] has runtime in the order of hours. Again, we see competitive performance for small blur kernels.

5.4.2. Real-World Images

In Figs. 14 and 15 we show results of our method on real-world images. Figure 14 shows examples for invariant blur, while Fig. 15 depicts images with spatially-varying camera shake. In both examples, our approach is able to recover images comparable in quality with the state of the art.

6. Discussion

6.1. Learned filters

The task of the *Feature Extraction Module* is to emphasize and enhance those image features that contain information about the unknown blur kernel. Figure 16 shows the learned filters of the convolution layer for each of the three stages within a single scale of a trained NN for kernel size 17×17 pixels. While the first stage takes a single (possibly down-sampled) version of a blurry image as input, the subsequent stages take both the restored latent image (obtained by non-blind deconvolution with the current estimate of the kernel) and the blurry image as input. The outputs of each stage are nonlinearly filtered versions of the input images. In Fig. 18 we visualize the effect of the first stage of a NN with two predicted output images on both the Lena image and a toy example image consisting of four disks blurred with Gaussians of varying size. Note that our feature extraction module outputs nonlinearly filtered images for both the blurry and the

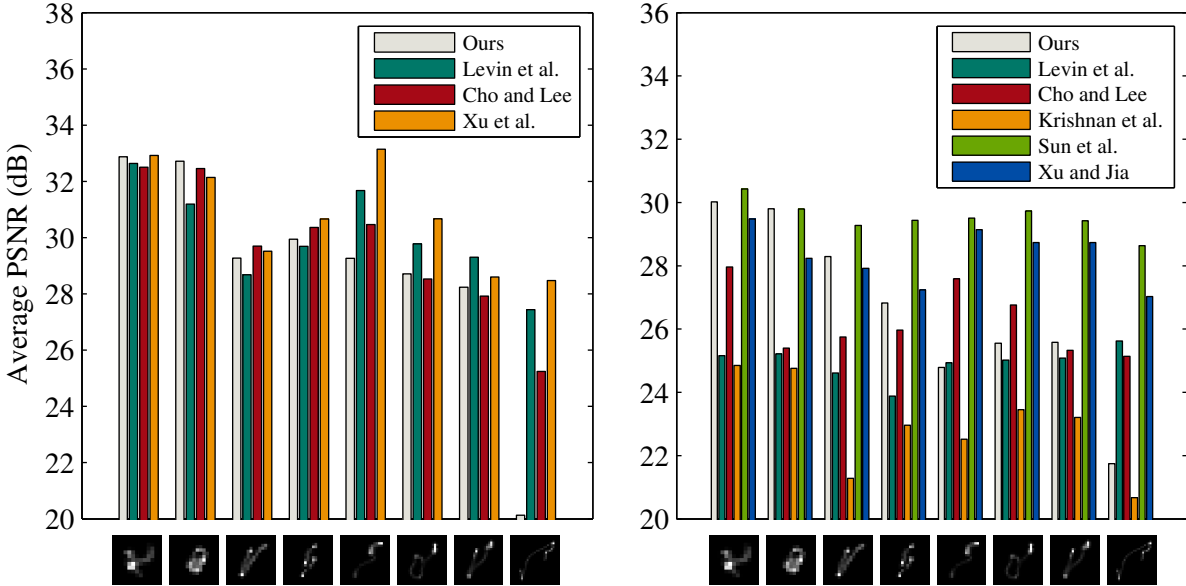


Figure 13: Results of the benchmark dataset of Levin et al. [Lev+09] and the extended benchmark of Sun et al. [Sun+13]. The results are sorted according to blur kernel size. While for kernels up to a size of 25×25 pixels, our approach yields comparable results, it falls short for larger blur kernels. Reasons for the performance drop are discussed in Section 6.3.

latent sharp image, both of which serve as input to the subsequent quotient layer, which in turn computes an estimate of the blur kernel. This is in contrast to other existing approaches [CL09; XJ10], which apply a *nonlinear* filter to the current estimate of the latent image, but use only a *linearly* filtered version of the blurry input image for kernel estimation.

Once these feature images have passed the subsequent tanh and recombination layer, they serve as input to the *quotient layer*, which computes an estimate of the unknown blur kernel. Figure 2 shows the intermediary results directly after the convolution layer and how they progressively change after passing through tanh and linear recombination layer two times, which seem to emphasize strong edges.

While the learned filters of the first stage are reminiscent of gradient filters of various extent and orientations including Gabor and Laplace-like filters, the filters of the subsequent stages are much more intricate and more difficult to interpret.

As discussed in Section 5, the learned filters depend on the image set that the NN was trained with, i.e., the *feature extraction module* learns image content specific features that are informative about the unknown blur kernel. In Fig. 17, we show the learned filters of the experiments in Section 5.1 for the *valley* and *blackboard* image category of the ImageNet dataset; they indeed differ from the ones trained on all images. For example, most of the filters learned for valley images are mirror or rotational symmetric, unlike many filters of the generic NN.



Blurry image

Cho et Lee [CL09]

Our results



Blurry image

Fergus et al. [Fer+06]

Our result



Blurry image

Shan et al. [SJA08]

Our result


Figure 14: Comparison on real-world example images taken from the literature with spatially invariant blur.



Blurry image

Joshi et al. [Jos+10]

Our result



Blurry image [Why+10]

Whyte et al. [Why+10]

Our result



Blurry image [Gup+10]

Gupta et al. [Gup+10]


Hirsch et al. [Hir+11]




Xu et al. [XZJ13]

Our result


Figure 15: Comparison on real-world example images taken from the literature with spatially-varying blur.



First iteration




Second iteration




Third iteration


Figure 16: Learned filters of the convolution layer for each of the three iterations within a single scale of a trained NN for kernel size 17×17 pixels. See text for details.



NN trained on all images



NN trained on valley images



NN trained on blackboard images

Figure 17: Learned filters of the convolution layer for three different types of images.




Figure 18: Visualization of the effect of the first stage of a network with two predicted output images on toy example with disks blurred with Gaussians of varying size (top row) and motion blurred Lena image (bottom row). While for the circles in \tilde{y}_i the different sizes of the Gaussian blurs are clearly visible, the NN replaces them in \tilde{x}_i with shapes of comparable sharpness.

6.2. Dependence on the size of the observed image

As noted by Hu and Yang [HY12], blind deblurring methods are most successful in predicting the kernel in regions of an image that exhibit strong salient edges. Other regions are less informative about the kernel, and have been shown to even hurt kernel estimation when included in the input to the blind deconvolution algorithm. Ideally, an estimation procedure should weight its input according to its information content. In the worst case, a larger input would not improve the results, but would not cause a deterioration either.

We study the behavior of our method with respect to the size of the observed image. It is possible that the NN learned to ignore image content detrimental to the kernel estimation. In Fig. 19 the predicted kernels for different sized crops of a blurry image are shown. Indeed, this example suggests that for our learned algorithm the kernel converges with input images increasing in size, while a non-learned state-of-the-art algorithm [XZJ13] exhibits no such trend. The more thorough analysis in Fig. 20 confirms this behavior: when evaluating the MSE for three different kernels, averaged over the 52 largest images from [Sun+13], it monotonously decreases for larger crops of the blurry image. We also see that the NN outperforms the competing method for the two small blur kernels.




Figure 19: Results for kernel estimation for different sizes of the observed image. *Top row:* Differently sized inputs to the blind deblurring algorithm. *Middle row:* Estimated kernels of our method. Larger inputs lead to better results. *Bottom Row:* Estimated kernels of [XZJ13]. No clear trend is visible.




Figure 20: Dependence of the estimated kernel on the size of the observed image. We show the MSE of predictions of three kernels for different sized inputs (cf. Fig. 19), averaged over the 52 largest images from [Sun+13]. The results of the NN decrease monotonously.



Figure 21: For larger blur kernels our approach falls short in yielding acceptable deblurring results. Shown is an example from [Lev+09] with a kernel of size 27×27 .

6.3. Limitations

A limitation of our current approach is the performance drop in the case of larger blur kernels. Figure 21 shows an example failure case from the benchmark dataset of [Lev+09]. We believe the reason for this is the suboptimal architecture of our multi-scale approach at higher resolution scales. While a multi-scale approach exhibits better performance compared to a single scale network, the observed performance drop for larger blurs suggests that using the same architecture at all scales is not optimal.

7. Conclusion

We have shown that it is possible to automatically learn blind deconvolution by reformulating the task as a single large nonlinear regression problem, mapping between blurry input and predicted kernel. The key idea is to incorporate the properties of the generative forward model into our algorithm, namely that the image is convolved with the same blur kernel everywhere. While features are extracted locally in the image, the *kernel estimation module* combines them globally. Next, the *image estimation module* propagates the information to the whole image, reducing the difficulty of the problem for the following iteration.

Our approach can adapt to different settings (e.g., blurry images with strong noise, or specific image classes), and it could be further extended to combine deblurring with other steps of the imaging pipeline, including over-saturation, Bayer filtering, HDR, or super-resolution.

The blur class also invites future research: instead of artificially sampling from a stochastic process, one could use recorded spatially-varying camera shakes, or a different source of unsharpness, like lens aberrations or atmospheric turbulences in astrophotography. Additionally, the insights gained from the trained system could be beneficial to existing hand-crafted methods. This includes using higher-order gradient representations and extended gradient filters.

Scalability of our method to large kernel sizes is still an issue, and this may benefit from future improvements in neural net architecture and training.

Appendix A Quotient Layer

As introduced in Eq. (4), the quotient layer performs the operation

$$\tilde{\mathbf{k}} = F^H \frac{\sum_i \overline{F\tilde{\mathbf{x}}_i} \odot F\tilde{\mathbf{y}}_i}{\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k} \quad (10)$$

to estimate the kernel $\tilde{\mathbf{k}}$ from images $\tilde{\mathbf{x}}_i$ and their blurry counterparts $\tilde{\mathbf{y}}_i$, both predicted by the previous layers of the NN.

The quotient layer also includes a learned regularization parameter β_k . To train the NN, we need the gradients of the output with respect to its parameters (in this case $\tilde{\mathbf{x}}_i$, $\tilde{\mathbf{y}}_i$ and β_k) in every layer. From these we obtain the gradient steps $\Delta\tilde{\mathbf{x}}_j$, $\Delta\tilde{\mathbf{y}}_k$ and $\Delta\beta_k$ in terms of $\Delta\tilde{\mathbf{k}}$, where $\Delta\tilde{\mathbf{k}}$ is determined by the loss for the current training example, back-propagated through the layers subsequent to the quotient layer.

A.1 Derivative with respect to sharp images

To obtain $\Delta\tilde{\mathbf{x}}_j$, we first determine the differential form in numerator layout, which means for vectors \mathbf{u} , \mathbf{v} and a matrix M that

$$d\mathbf{u} = \begin{pmatrix} du_1 \\ du_2 \\ \vdots \end{pmatrix}, \quad \frac{d\mathbf{u}}{d\mathbf{v}} = \begin{pmatrix} \frac{du_1}{dv_1} & \frac{du_1}{dv_2} & \cdots \\ \frac{du_2}{dv_1} & \frac{du_2}{dv_2} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \quad \text{and} \quad dM = \begin{pmatrix} dm_{11} & dm_{12} & \cdots \\ dm_{21} & dm_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}. \quad (11)$$

Therefore, assuming that only $\tilde{\mathbf{x}}_j$ is variable, with the rules of matrix calculus [PP12] we arrive at

$$\begin{aligned} d\tilde{\mathbf{k}} &= F^H \frac{\overline{F d\tilde{\mathbf{x}}_j} \odot F\tilde{\mathbf{y}}_j}{\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k} - F^H \frac{(\sum_i \overline{F\tilde{\mathbf{x}}_i} \odot F\tilde{\mathbf{y}}_i) \odot d(F\tilde{\mathbf{x}}_j \odot \overline{F\tilde{\mathbf{x}}_j})}{(\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k)^2} = \\ &= F^H \underbrace{\text{Diag} \left(\frac{F\tilde{\mathbf{y}}_j}{\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k} \right)}_A \overline{F} d\tilde{\mathbf{x}}_j \\ &\quad - F^H \underbrace{\text{Diag} \left(\frac{(\sum_i \overline{F\tilde{\mathbf{x}}_i} \odot F\tilde{\mathbf{y}}_i) \odot \overline{F\tilde{\mathbf{x}}_j}}{(\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k)^2} \right)}_B F d\tilde{\mathbf{x}}_j - F^H \underbrace{\text{Diag} \left(\frac{(\sum_i \overline{F\tilde{\mathbf{x}}_i} \odot F\tilde{\mathbf{y}}_i) \odot F\tilde{\mathbf{x}}_j}{(\sum_i |F\tilde{\mathbf{x}}_i|^2 + \beta_k)^2} \right)}_C \overline{F} d\tilde{\mathbf{x}}_j, \end{aligned} \quad (12)$$

with the Hadamard product $\mathbf{a} \odot \mathbf{b} = \text{Diag}(\mathbf{a})\mathbf{b}$ and the operation ‘‘Diag’’ that creates a diagonal matrix from a vector. Additionally, $\tilde{\mathbf{x}}_i$ is real, and therefore $\overline{d\tilde{\mathbf{x}}_j} = d\tilde{\mathbf{x}}_j$.

With the differential $d\tilde{\mathbf{k}} = M d\tilde{\mathbf{x}}_j$ (where $M = A + B + C$) in numerator layout, we note that the derivative $\frac{d\tilde{\mathbf{k}}}{d\tilde{\mathbf{x}}_j} = M$ and the gradient step $\Delta\tilde{\mathbf{x}}_j = M^T \Delta\tilde{\mathbf{k}}$. Additionally, we use that $B^T \Delta\tilde{\mathbf{k}}$ is

real since it is the inverse Fourier transform of a point-wise product of vectors all with Hermitian symmetry (they are themselves purely real vectors that have been Fourier transformed). Hermitian symmetry for a vector \mathbf{v} of length n means $v_{n-i+1} = \overline{v_i}$ for all its components i . Therefore,

$$\begin{aligned}
\Delta \tilde{\mathbf{x}}_j &= A^T \Delta \tilde{\mathbf{k}} - B^T \Delta \tilde{\mathbf{k}} - C^T \Delta \tilde{\mathbf{k}} = A^T \Delta \tilde{\mathbf{k}} - \overline{B^T \Delta \tilde{\mathbf{k}}} - C^T \Delta \tilde{\mathbf{k}} = \\
&= F^H \frac{\overline{F \Delta \tilde{\mathbf{k}}} \odot F \tilde{\mathbf{y}}_j}{\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k} \\
&\quad - F^H \frac{\sum_i \overline{F \tilde{\mathbf{x}}_i} \odot F \tilde{\mathbf{y}}_i \odot F \tilde{\mathbf{x}}_j \odot F \Delta \tilde{\mathbf{k}}}{(\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k)^2} - F^H \frac{(\sum_i \overline{F \tilde{\mathbf{x}}_i} \odot F \tilde{\mathbf{y}}_i) \odot F \tilde{\mathbf{x}}_j \odot \overline{F \Delta \tilde{\mathbf{k}}}}{(\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k)^2} \\
&= F^H \left(\frac{\overline{F \Delta \tilde{\mathbf{k}}} \odot F \tilde{\mathbf{y}}_j}{\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k} - \frac{2\Re \left(\sum_i \overline{F \tilde{\mathbf{x}}_i} \odot F \tilde{\mathbf{y}}_i \odot F \Delta \tilde{\mathbf{k}} \right) \odot F \tilde{\mathbf{x}}_j}{(\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k)^2} \right), \tag{13}
\end{aligned}$$

since the transpose has no effect on a diagonal matrix. The real part of a vector \mathbf{v} is denoted as $\Re(\mathbf{v}) = \frac{1}{2}(\mathbf{v} + \overline{\mathbf{v}})$.

A.2 Derivative with respect to blurry images

The derivation for the gradient $\Delta \tilde{\mathbf{y}}_j$ is similar. We again start with the differential form

$$d\tilde{\mathbf{k}} = F^H \underbrace{\frac{\overline{F \tilde{\mathbf{x}}_j} \odot F d\tilde{\mathbf{y}}_j}{\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k}}_A = F^H \text{Diag} \left(\frac{\overline{F \tilde{\mathbf{x}}_j}}{\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k} \right) F d\tilde{\mathbf{y}}_j, \tag{14}$$

this time assuming $d\tilde{\mathbf{x}}_j$ and $d\beta_k$ to be zero. Next, we use again that a real vector is unchanged by conjugation, and we obtain

$$\Delta \tilde{\mathbf{y}}_j = A^T \Delta \tilde{\mathbf{k}} = \overline{A^T \Delta \tilde{\mathbf{k}}} = F^H \text{Diag} \left(\frac{\overline{F \tilde{\mathbf{x}}_j}}{\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k} \right) F \Delta \tilde{\mathbf{k}} = F^H \left(\frac{F \tilde{\mathbf{x}}_j \odot F \Delta \tilde{\mathbf{k}}}{\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k} \right). \tag{15}$$

A.3 Derivative with respect to the regularization parameter

Following the previous procedure for the regularization parameter β_k we arrive at

$$d\tilde{\mathbf{k}} = F^H \underbrace{\frac{\sum_i \overline{F \tilde{\mathbf{x}}_i} \odot F \tilde{\mathbf{y}}_i}{(\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k)^2}}_A d\beta_k, \tag{16}$$

assuming only β_k to be variable. Then, multiplying $\Delta \tilde{\mathbf{k}}$ by the transpose of $\frac{d\tilde{\mathbf{k}}}{d\beta_k} = A$ we get

$$\Delta \beta_k = A^T \Delta \tilde{\mathbf{k}} = \left(F^H \frac{\sum_i \overline{F \tilde{\mathbf{x}}_i} \odot F \tilde{\mathbf{y}}_i}{(\sum_i |F \tilde{\mathbf{x}}_i|^2 + \beta_k)^2} \right)^T \Delta \tilde{\mathbf{k}}. \tag{17}$$

References

- [Aiz+06] I. Aizenberg, D. Paliy, C. Moraga, and J. Astola. “Blur identification using neural network for image restoration”. In: *Computational Intelligence, Theory and Applications*. Vol. 38. Springer, 2006, pp. 441–455. DOI: [10.1007/3-540-34783-6_45](https://doi.org/10.1007/3-540-34783-6_45).
- [BG91] L. Bottou and P. Gallinari. “A Framework for the Cooperation of Learning Algorithms”. In: *Advances Neural Information Processing Systems*. 1991, pp. 781–788.
- [CD91] C. M. Cho and H. S. Don. “Blur identification and image restoration using a multi-layer neural network”. In: *IEEE Int. Joint Conf. Neural Networks*. 1991, pp. 2558–2563. DOI: [10.1109/ijcnn.1991.170774](https://doi.org/10.1109/ijcnn.1991.170774).
- [CL09] S. Cho and S. Lee. “Fast motion deblurring”. In: *ACM Trans. Graphics* 28.5 (2009), p. 145. DOI: [10.1145/1618452.1618491](https://doi.org/10.1145/1618452.1618491).
- [De +03] D. De Ridder et al. “Nonlinear image processing using artificial neural networks”. In: *Advances Imaging and Electron Physics* 126 (2003), pp. 352–450.
- [Den+09] J. Deng et al. “ImageNet: A large-scale hierarchical image database”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2009. DOI: [10.1109/cvpr.2009.5206848](https://doi.org/10.1109/cvpr.2009.5206848).
- [ERH02] M. Egmont-Petersen, D. de Ridder, and H. Handels. “Image processing with neural networks—a review”. In: *Pattern recognition* 35.10 (2002), pp. 2279–2301. DOI: [10.1016/S0031-3203\(01\)00178-9](https://doi.org/10.1016/S0031-3203(01)00178-9).
- [Fer+06] R. Fergus et al. “Removing camera shake from a single photograph”. In: *ACM Trans. Graphics* 25.3 (2006), pp. 787–794. DOI: [10.1145/1141911.1141956](https://doi.org/10.1145/1141911.1141956).
- [Gup+10] A. Gupta et al. “Single Image Deblurring Using Motion Density Functions”. In: *IEEE Europ. Conf. Computer Vision*. 2010.
- [HHS10] S. Harmeling, M. Hirsch, and B. Schölkopf. “Space-Variant Single-Image Blind Deconvolution for Removing Camera Shake.” In: *Advances Neural Information Processing Systems*. 2010, pp. 829–837.
- [HHY10] Z. Hu, J.-B. Huang, and M.-H. Yang. “Single image deblurring with adaptive dictionary learning”. In: *IEEE Int. Conf. Image Processing*. 2010, pp. 1169–1172. DOI: [10.1109/icip.2010.5651892](https://doi.org/10.1109/icip.2010.5651892).
- [Hir+10] M. Hirsch, S. Sra, B. Schölkopf, and S. Harmeling. “Efficient filter flow for space-variant multiframe blind deconvolution.” In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2010, pp. 607–614. DOI: [10.1109/cvpr.2010.5540158](https://doi.org/10.1109/cvpr.2010.5540158).
- [Hir+11] M. Hirsch, C. J. Schuler, S. Harmeling, and B. Schölkopf. “Fast removal of non-uniform camera shake”. In: *IEEE Int. Conf. Computer Vision*. 2011, pp. 463–470. DOI: [10.1109/iccv.2011.6126276](https://doi.org/10.1109/iccv.2011.6126276).

- [HY12] Z. Hu and M.-H. Yang. “Good Regions to Deblur”. In: *Computer Vision – ECCV 2012*. Lecture Notes in Computer Science. Springer, 2012, pp. 59–72. DOI: [10 . 1007/978-3-642-33715-4_5](https://doi.org/10.1007/978-3-642-33715-4_5).
- [Jos+10] N. Joshi, S. B. Kang, C. L. Zitnick, and R. Szeliski. “Image deblurring using inertial measurement sensors”. In: *ACM Trans. Graphics* 29.4 (2010), p. 30. DOI: [10 . 1145/1833351.1778767](https://doi.org/10.1145/1833351.1778767).
- [KN11] C. Khare and K. K. Nagwanshi. “Image Restoration in Neural Network Domain using Back Propagation Network Approach”. In: *Int. J. Computer Information Systems* 2.5 (2011), pp. 25–31.
- [KTF11] D. Krishnan, T. Tay, and R. Fergus. “Blind deconvolution using a normalized sparsity measure”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2011. DOI: [10 . 1109/cvpr.2011.5995521](https://doi.org/10.1109/cvpr.2011.5995521).
- [Lev+09] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. “Understanding and evaluating blind deconvolution algorithms”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2009, pp. 1964–1971. DOI: [10 . 1109/cvpr.2009.5206815](https://doi.org/10.1109/cvpr.2009.5206815).
- [Lev+11] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. “Efficient marginal likelihood optimization in blind deconvolution”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2011. DOI: [10 . 1109/cvpr.2011.5995308](https://doi.org/10.1109/cvpr.2011.5995308).
- [PP12] K. B. Petersen and M. S. Pedersen. *The matrix cookbook*. 2012.
- [RW06] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. 2006. ISBN: 978-0262182539.
- [Sch+13] U. Schmidt et al. “Discriminative Non-blind Deblurring”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2013, pp. 604–611. DOI: [10 . 1109/cvpr.2013.84](https://doi.org/10.1109/cvpr.2013.84).
- [SF94] R. J. Steriti and M. A. Fiddy. “Blind deconvolution of images by use of neural networks”. In: *Optics letters* 19.8 (1994), pp. 575–577. DOI: [10 . 1364/ol.19.000575](https://doi.org/10.1364/ol.19.000575).
- [SJA08] Q. Shan, J. Jia, and A. Agarwala. “High-quality motion deblurring from a single image”. In: *ACM Trans. Graphics* 27.3 (2008), p. 73. DOI: [10 . 1145 / 1360612 . 1360672](https://doi.org/10.1145/1360612.1360672).
- [Sun+13] L. Sun, S. Cho, J. Wang, and J. Hays. “Edge-based blur kernel estimation using patch priors”. In: *IEEE Int. Conf. Computational Photography*. 2013. DOI: [10 . 1109/ iccpht.2013.6528301](https://doi.org/10.1109/iccpht.2013.6528301).
- [TOM96] J. E. Tansley, M. J. Oldfield, and D. J. MacKay. “Neural network image deconvolution”. In: *Maximum Entropy and Bayesian Methods*. Fundamental Theories of Physics. Springer, 1996, pp. 319–325. DOI: [10 . 1007/978-94-015-8729-7_25](https://doi.org/10.1007/978-94-015-8729-7_25).
- [Wan+13] C. Wang et al. “Nonedge-specific adaptive scheme for highly robust blind motion deblurring of natural images”. In: *IEEE Trans. Image Processing* 22.3 (2013), pp. 884–897. DOI: [10 . 1109/tip.2012.2219548](https://doi.org/10.1109/tip.2012.2219548).

- [Why+10] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce. “Non-uniform Deblurring for Shaken Images”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2010, pp. 491–498. DOI: [10.1109/cvpr.2010.5540175](https://doi.org/10.1109/cvpr.2010.5540175).
- [WZ13] D. Wipf and H. Zhang. “Revisiting Bayesian Blind Deconvolution”. In: *ArXiv e-prints* (2013). arXiv: [1305.2362](https://arxiv.org/abs/1305.2362) [cs.CV].
- [XJ10] L. Xu and J. Jia. “Two-Phase Kernel Estimation for Robust Motion Deblurring”. In: *Computer Vision – ECCV 2010*. Lecture Notes in Computer Science. Springer, 2010, pp. 157–170. DOI: [10.1007/978-3-642-15549-9_12](https://doi.org/10.1007/978-3-642-15549-9_12).
- [XZJ13] L. Xu, S. Zheng, and J. Jia. “Unnatural l0 sparse representation for natural image deblurring”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2013, pp. 1107–1114. DOI: [10.1109/cvpr.2013.147](https://doi.org/10.1109/cvpr.2013.147).
- [Zei+10] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus. “Deconvolutional networks”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2010, pp. 2528–2535. DOI: [10.1109/cvpr.2010.5539957](https://doi.org/10.1109/cvpr.2010.5539957).
- [Zei12] M. D. Zeiler. “ADADELTA: An Adaptive Learning Rate Method”. In: *ArXiv e-prints* (2012). arXiv: [1212.5701](https://arxiv.org/abs/1212.5701) [cs.LG].
- [Zho+13] L. Zhong et al. “Handling Noise in Single Image Deblurring Using Directional Filters”. In: *IEEE Conf. Computer Vision and Pattern Recognition*. 2013, pp. 612–619. DOI: [10.1109/cvpr.2013.85](https://doi.org/10.1109/cvpr.2013.85).
- [ZTF11] M. D. Zeiler, G. W. Taylor, and R. Fergus. “Adaptive deconvolutional networks for mid and high level feature learning”. In: *IEEE Int. Conf. Computer Vision*. 2011, pp. 2018–2025. DOI: [10.1109/iccv.2011.6126474](https://doi.org/10.1109/iccv.2011.6126474).