This ICCV paper is the Open Access version, provided by the Computer Vision Foundation.
Except for this watermark, it is identical to the accepted version;
the final published version of the proceedings is available on IEEE Xplore.

Universal and Flexible Optical Aberration Correction Using Deep-Prior Based Deconvolution

Xiu Li[1*], Jinli Suo[1], Weihang Zhang[1], Xin Yuan[2], Qionghai Dai[1]

[1]Tsinghua University, [2]Westlake University

Abstract

High quality imaging usually requires bulky and expensive lenses to compensate geometric and chromatic aberrations. This poses high constraints on the optical hash or low cost applications. Although one can utilize algorithmic reconstruction to remove the artifacts of low-end lenses, the degeneration from optical aberrations is spatially varying and the computation has to trade off efficiency for performance. For example, we need to conduct patch-wise optimization or train a large set of local deep neural networks to achieve high reconstruction performance across the whole image. In this paper, we propose a PSF aware deep network, which takes the aberrant image and PSF map as input and produces the latent high quality version via incorporating deep priors, thus leading to a universal and flexible optical aberration correction method. Specifically, we pre-train a base model from a set of diverse lenses and then adapt it to a given lens by quickly refining the parameters, which largely alleviates the time and memory consumption of model learning. The approach is of high efficiency in both training and testing stages. Extensive results verify the promising applications of our proposed approach for compact low-end cameras. The code is available at https://github.com/leehsiu/UABC

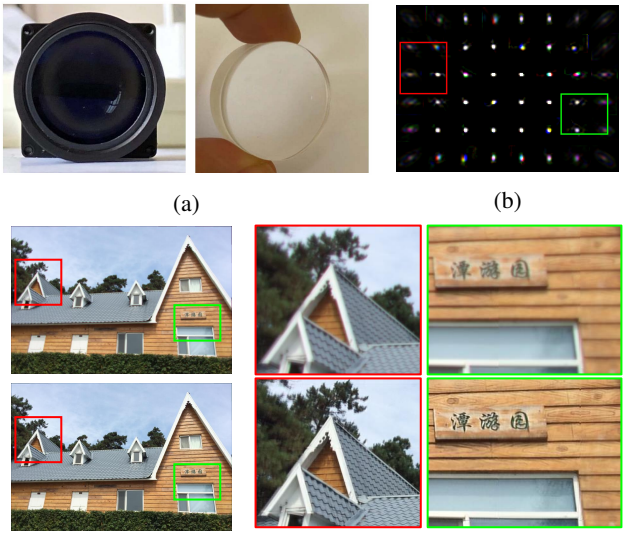

Figure 1: One example of computationally reconstructing high quality image with a simple lens. (a) A camera with a simple double glued lens (Thorlabs, AC254-075-A-ML). (b) The calibrated PSF of the camera in (a). (c) The input degenerated image (upper row) and our reconstruction result (bottom row).

1. Introduction

Optical aberration is one of the most common degeneration in real lens-based imaging systems. Due to the deviation from ideal thin-lens model, the simple/single lens elements suffer from chromatic, spherical aberration and coma aberrations, and degenerate the imaging quality significantly. To cancel out these artifacts, modern camera lenses are usually made of a complex combination of several (even dozens of) lens elements with carefully designed parameters (aka. lens prescription). In a nutshell, existing techniques achieve high imaging quality via such complex design, at the expense of high cost, bulky weight and inevitable lens flare. With the rapid development of large pixel count digital sensors (*e.g.* 100 mega-pixel scale), effective compensation of lens aberration is highly desired.

To achieve *light-weighted* and *low-cost* high-quality imaging, computational optical aberration correction has been exploited during the past decades. Different from advanced optical design, computational methods employ a simple lens for imaging and remove the aberration afterwards with algorithms. Mathematically, the lens aberration can be formulated as convolution with spatially varying kernels, and the compensation is conducted by deconvolution with the assistance of image priors. These computa-

*lixiu15@mails.tsinghua.edu.cn

tional methods can be either *non-blind* or *blind* depending on whether the aberration model, *i.e.*, Point Spread Function (PSF) is calibrated beforehand. The former measure the PSF using specially designed systems [26, 7, 11] or estimating it from degenerated images using some developed algorithms explicitly [23, 30, 28], deconvolution is conducted afterwards. The latter usually jointly estimates the PSF and the latent sharp image in an iterative manner [35]. Considering the fact that the PSF is fixed for a given lens, and the estimation of PSF has already been extensively studied, we focus on the non-blind case in this paper.

As illustrated in Fig. 1(b), the PSF by lens aberration is spatially non-uniform and varying across color channels as well. Previously, researchers have proposed a series of algorithms for non-uniform deconvolution, includes pixel-wise deconvolution [6, 33, 5] and patch-wise deconvolution [7, 23, 36, 8]. The later one assumes that the PSF varies smoothly in the spatial dimension thus can be approximated locally uniform. The deconvolution in this situation is usually faster than the pixel-wise methods. This assumption also holds for the PSF caused by the aberration of a low-end lens.

Recently, deep neural network has also been exploited for image deconvolution and has shown encouraging results [27, 3]. However, handling non-uniformity is still challenging for end-to-end deep learning. Using a single network weight setting for different sub-regions with different PSF will either requires a large model capacity or leads to an average performance. While training a series of models specifically for each sub-region is quite time consuming. In this paper, we propose to combine deep network and model-based deconvolution into an iterative framework. Prior-guided deconvolution is conducted locally and explicitly models the physical degeneration process, which ensures the high fidelity and easy adaption to different lenses. A single global deep projection network is applied to the whole image to suppress ringing artifacts and eliminate blocking artifacts. Meanwhile, we propose to pretrain a base model from a large set of diverse lenses first and then adapt it to a specific lens quickly, greatly simplify the training process for different lenses.

The main contributions of this paper are as follows.

- We propose a PSF-aware deep network to address the spatially varying degeneration caused by lens aberration, by *decoupling* the physical imaging model and deep network prior.
- The model training is built on a pre-trained base model plus a fast *adaption* procedure to different lenses. Thus, it is of high efficiency and high feasibility.
- The approach achieves state-of-the-art performance but with much higher running efficiency.

We conduct synthetic and qualitative experiments to demonstrate the training efficiency and performance of the proposed method. We also get promising results on real data captured by low end lenses, please refer to Fig. 1, and more results in Fig. 9.

2. Related work

2.1. Computational aberration correction

Computational aberration correction can be roughly divided into non-blind and blind methods. For non-blind method, an explicit calibration stage is required to estimate the lens aberration model. Shih et al.[26] literately measures the point-spread function by imaging a pinhole grid pattern in a dark room. Heide et al.[7] calibrate the PSF by using a framed random pattern. Specially designed calibration chart [10, 11, 23] is also widely used due to its lower cost. Once the blur kernel is known, the latent clear image can be solved by image deconvolution [7, 23]. Blind methods skips this calibration stage and directly estimate the blur kernel from the images based on the natural image statistics and knowledge of the aberration model. Rahbar and Faez [19] used a Zernike model to describe the lens aberration and estimated the Zernike polynomial coefficients statistically. Schuler et al. [24] proposed an Efficient Filter Flow (EFF) basis method to describe the non-uniform blurry kernels by exploiting the reflection symmetry, rotational symmetry and radial distributions of the PSF. This prior knowledge is also used by Yue et al. [36] while they also propose a radial-splitting technique (representing the image in polar coordinates) and a sharp-to-blur estimation strategy, which is further extensively studied by[29].

2.2. Image deconvolution

Mathematically, optical aberration correction is a special case of non-uniform image deconvolution. Usually, the modelling of the non-uniformity can be done fully paramatrically, which results in a pixel-wise deconvolution [6, 33, 5, 9]. To accelerate the calculation, the PSF can also be assumed to vary smoothly locally, thus the non-uniform deconvolution can approximated with patch-wise uniform deconvolution. [7, 23, 36, 8]. As image deconvolution is usually ill-posed, regularization terms (or priors) are essential to produce reasonable result. General prior includes TV-norm minimization [31], hyper-Laplacian [13, 16] and self-similarity [17]. In the recent years, deep learning has been widely exploited for deconvolution problems for its ability in modelling complex data distribution. Although it is possible to directly learn to map the blurred image to its corresponding clear version [37, 15], is has been shown that non-blind deep deconvolution performs better than direct end-to-end mapping [38, 14, 25, 27, 3]. These works utilize the PSF by convert the deep deconvolution problem into de-

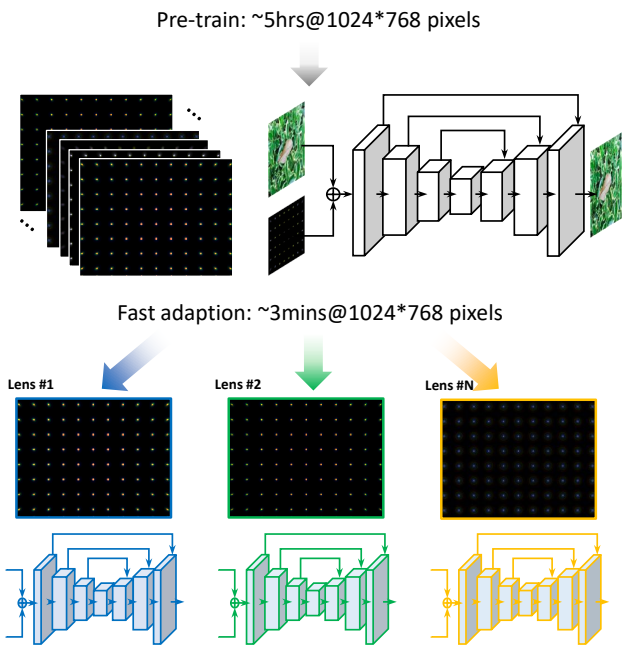

Figure 2: The basic idea of building fast and high quality lens aberration correction model, via pre-training a base model and fast adapting to different lenses.

convolution and network based refinement. This way the network is capable to handle various degeneration kernels.

2.3. Plug-and-Play algorithms

Another trending way to utilize PSF-awareness is plug-and-play(PnP) [22, 20]. A straightforward way is replacing the general or handcraft prior with a learned deep generative model [18] solve the optimization problem using gradient decent. To avoid the heavy computational load of tracking the network gradient, the deep network can also serves as a proximal operator, and the image restoration problem is solved by iteratively running model-based optimization and deep network projection. This modeling has shown promising results in linear inverse problems[2], for image super-resolution [39], snapshot compressive imaging [34] and high-spectral imaging [9]. Our work also follows this formulation.

3. Method

In this section, we first mathematically formulate the lens aberration model and then propose our deep-prior based approach for optical aberration correction.

3.1. Problem formulation

As mentioned before, lens aberration can be interpreted as non-uniform channel-wise blur as the PSF varies both spatially and spectrally (*i.e.*, RGB channel). Thanks to the fact the neighbouring PSFs are highly correlated, we can approximate the non-uniform PSF as patch-wise uniform. For each small image patch p at channel c, the recorded degraded image patch $\mathbf{y}_{c,p}$ and the latent clear patch $\mathbf{x}_{c,p}$ is related by

$$\mathbf{y}_{c,p} = \mathbf{k}_{c,p} \otimes \mathbf{x}_{c,p} + \mathbf{n}_{c,p}, \tag{1}$$

where $\mathbf{k}_{c,p}$ is the corresponding PSF (blur kernel), $\mathbf{n}_{c,p}$ is the noise and $\otimes$ denotes the 2D convolution. For simplicity, we will drop this c index if we consider $\otimes$ is carried channel-wisely.

To solving Eq. (1), a prior term is necessary to constrain the space of $\mathbf{x}$. This leads to

$$\mathbf{x} = \arg\min_{\mathbf{x}} \sum_p \|\mathbf{k}_p \otimes \mathbf{x}_p - \mathbf{y}_p\|_2^2 + \lambda\Phi(\mathbf{x}), \tag{2}$$

where the first term enforces the data term residual and the second one $\Phi(\mathbf{x})$ denotes the prior(statistical distribution) of $\mathbf{x}$.

3.2. Data/Prior splitting

We follow [9, 4, 39] to use Half-Quadratic Splitting (HQS) to solve Eq. (2). By introducing an auxiliary variable $\mathbf{z}$, Eq. (2) can be re-modeled as

$$\arg\min_{\mathbf{x},\mathbf{z}} \sum_p \|\mathbf{k}_p \otimes \mathbf{z}_p - \mathbf{y}_p\|_2^2 + \lambda\Phi(\mathbf{x}) + \mu\|\mathbf{z} - \mathbf{x}\|_2^2, \tag{3}$$

which can be solved by performing iterations of the following sub-problems.

$$\mathbf{z}^t \leftarrow \min_{\mathbf{z}} \sum_p \|\mathbf{k}_p \otimes \mathbf{z}_p - \mathbf{y}_p\|_2^2 + \mu\|\mathbf{z} - \mathbf{x}^{t-1}\|_2^2, \tag{4}$$

$$\mathbf{x}^t \leftarrow \min_{\mathbf{x}} \mu\|\mathbf{z}^t - \mathbf{x}\|_2^2 + \lambda\Phi(\mathbf{x}). \tag{5}$$

Eq. (4) can be approximately solved in a closed form under circular boundary conditions as

$$\mathbf{z}_p^t = \mathcal{F}^{-1}\left(\frac{\overline{\mathcal{F}(\mathbf{k})}\mathcal{F}(\mathbf{y}_p) + \mu_p^k\mathcal{F}(\mathbf{x}_p^{t-1})}{\overline{\mathcal{F}(\mathbf{k})}\mathcal{F}(\mathbf{k}) + \mu_p^k}\right), \tag{6}$$

where $\mathcal{F}(\cdot)$ denotes the Fast Fourier Transform (FFT), $\mathcal{F}^{-1}(\cdot)$ denotes the inverse FFT, and $\overline{\mathcal{F}(\cdot)}$ denotes the complex conjugate of FFT. Notice that although we use a single μ notation in Eq. (6), we use different μ_p^t for different patches, channels and iterations when implementing the patch-wise deconvolution.

For Eq. (5), if we absorb μ_p^t into λ, it is a problem with $\mathbf{z}_p^t, \lambda_p^t$ as input and $\mathbf{x}^t$ as output. We treat $\Phi(\cdot)$ as a proximal operator and directly learn a deep projection network $\mathcal{P}$ similar to [2, 39, 9].

$$\mathbf{x}^t = \mathcal{P}(\mathbf{z}_p^t, \lambda_p^t). \tag{7}$$

Note that we have omitted the chop and assemble operations between $\mathbf{x}, \mathbf{z}$ and $\{\mathbf{x}_p\}, \{\mathbf{z}_p\}$

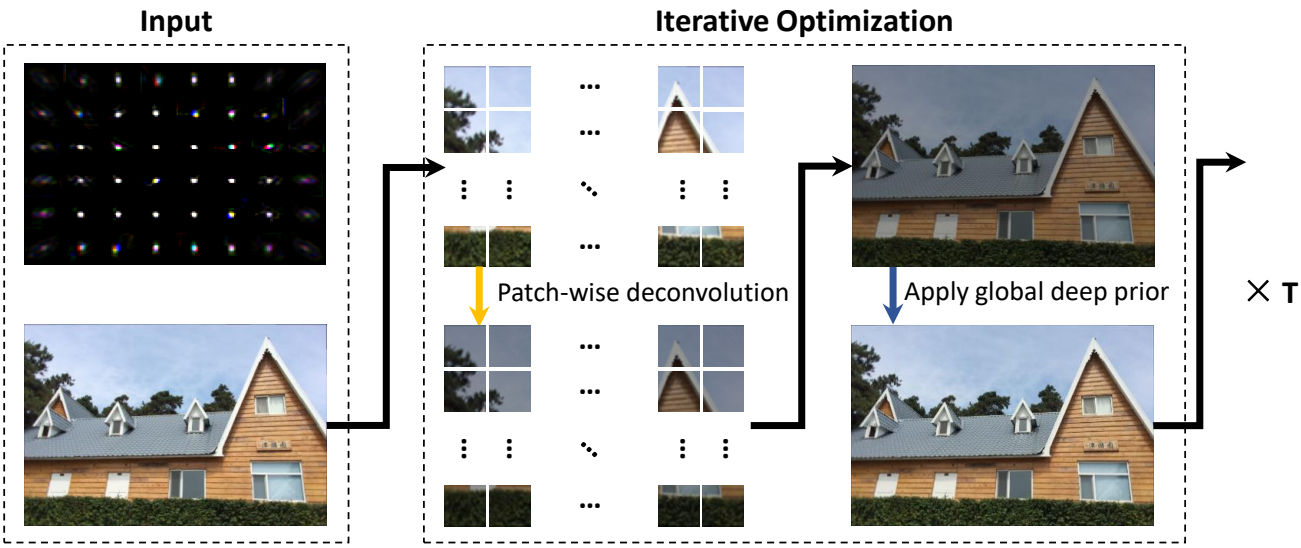

Figure 3: The flow chart of the proposed approach. The divided two parts—deconvolution and prior imposition, are illustrated with different colors.

3.3. Deep projector $\mathcal{P}$

The deep projector $\mathcal{P}$ is built mainly following the U-net [21], which is widely used for image-to-image translation problems. Here we replace the original plain convolution in original U-net with residual blocks for better convergence. Notice that as $\mathcal{P}$ is fully convolutional, it can be applied to images with an arbitrary size. Thus unlike the deconvolution stage which is done patch-wisely, we ensemble the patches together and update the whole image in this stage. The input of $\mathcal{P}$ is 6 channels(by assembling and concatenating $\mathbf{z}_p^t$ and λ_p^t) and output is 3 channels $\mathbf{x}$ with the same spatial dimension. The network structure is shown in Fig. 2. Ringing artifacts and blocking effects caused by the patch-wise deconvolution are effectively removed because information of neighboring patches will be used. The flow of the full forwarding process is illustrated in Fig. 3.

3.4. Pretrain and Adaption

End-to-end Training We use DIV2K and Flicker2K dataset [1], with 3450 high-definition high-resolution ($2K$) images in total, to train our model. We synthesize the training triplets (blurred image $\mathbf{y}$, aberration PSF $\mathbf{k}$ and clear image $\mathbf{x}$) as follows. The ground truth clear image $\mathbf{x}$ is drawn from the training dataset. The PSF $\mathbf{k}$ is generated by two ways: (i) We use Zemax software to simulate the PSF for various simple-lens designs. The kernel size of the PSF is 25×25 and the spatial resolution is set to 16×16 (in total, we have 256 different PSFs for each lens prescription). (ii) Anisotropic Gaussian kernels with different deviations are also used to increase the PSF diversity. The degenerated image $\mathbf{y}$ is generated via Zemax software by ray tracing (for PSF wth known prescriptions) or patch-wise convolution (for Gaussian kernels).

Algorithm 1: Framework of our proposed deconvolution method.

Input: $\mathbf{y}$, $\mathbf{k}$, max stage number T, $\{\mu_p^t\}$,$\{\lambda_p^t\}$.

Pre-processing;
1. Chop the input blur image into non-overlapping patches according to $\mathbf{k}$ and padding each patch with pixels from neighboring patches;
2. Pre-calculate $\{\mathcal{F}(\mathbf{k}_p)\}$, $\{\mathcal{F}(\mathbf{y}_p)\}$.
3. Initialize $\mathbf{x}^0$ as zero, thus the first stage update of $\mathbf{z}_p$ is the same as the Wiener filter.;

for $t = 1, ..., T$ **do**
1. Chop $\mathbf{x}^{t-1}$ into patches using same strategy in the pre-processing stage, and apply Eq. (6) to get $\{\mathbf{z}_p^t\}$.
2. Shave $\{\mathbf{z}_p^t\}$ and assemble them together to get $\mathbf{z}^t$. Then pass $\mathbf{z}^t$ and λ_t through the $\mathcal{P}$ to get $\mathbf{x}^t$.

end

Output: $\mathbf{x}^T$

Directly training on the full image with spatially-varying PSF is infeasible due to the GPU memory limits. Instead, for each training sample we randomly crop the image with 256×256 patch size, and also randomly crop a 2×2 patch from the whole 16×16 PSF map of a randomly drawn lens prescriptions. The blurred image thus has four regions with different PSFs. The network is trained by minimizing the ℓ_1

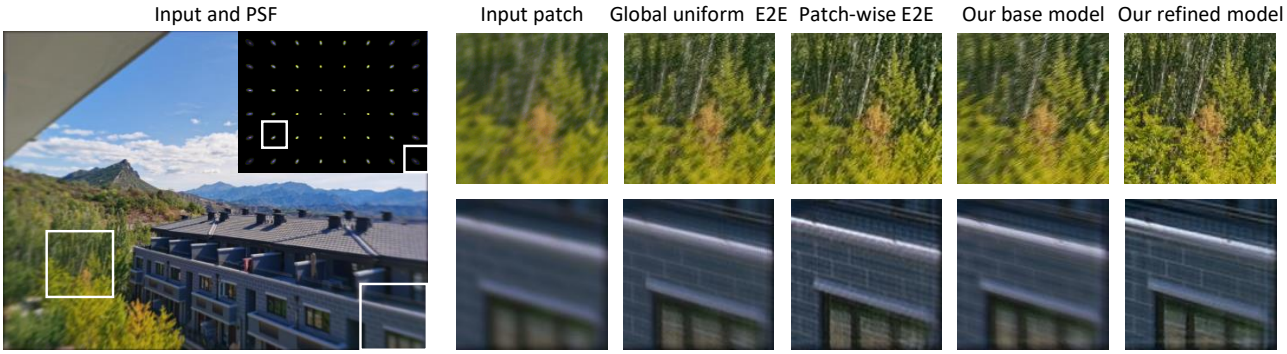

Figure 4: Comparison of visual results among different deep learning models, in correspondence to those in Table. 1.

reconstruction error. The batch-size is set to 3, resulting in a 10.6G memory cost on a single GPU during training. The μ, λ are set as free trainable parameters optimized together with the $\mathcal{P}$ network parameters using Adam optimizer [12] with the 0.001 learning rate. Altough we mentioned above that μ, λ are set different for each sub-region, in the pretrain stage, they are set as spatially uniform, only varies across different stages. It takes around 5 hours to train the model for 7K iterations using PyTorch implementation on a single 12G memory Nvidia RTX 2080 Ti GPU. One exemplar training curve is shown in Fig. 5.

Notice that for each training iteration, the network will see the assembled patches with different PSFs, thus the network needs to learn to remove the deconvolution artifacts but also the blocky artifacts. This implementation is crucial to the success of our methods.

Kernel-specific refine When a specific PSF is presented, we can refine the above pre-trained base model for performance boosting. We follow the same training protocol in the end-to-end training stage. Differently, instead of optimizing a spatially uniform μ, λ variable, we built a μ, λ map with each one corresponding to a specific local PSF. Experiments show that after lens specific adaption, the performance will raise significantly even with a very small number of training iterations, *e.g.*, 3 minutes for around 60 iterations. For the lenses with more severe aberrations in the marginal regions, we will use smaller patches in the outer regions than in the central ones.

4. Results

4.1. Quantitative evaluation

To demonstrate the efficiency and performance of our proposed scheme, we compare our model with three baseline lens aberration compensation methods based on deep learning: (i) globally uniform end-to-end (E2E) model by assuming the PSF being uniform and train an single E2E model applying on the full image, (ii) non-uniform patch-wise E2E model, in which we decompose the image into PSF-approximately-uniform patches and train one model for each image patch, and (iii) our pre-trained base model, *i.e.*, train a non-uniform model working for a set of diverse lenses. For efficiency evaluation, we compare the time for model training and inference for all the methods. We use DIV2K [1] validation set as the testing data and evaluate the restoration results in terms of peak-signal-to-noise-ratio

Table 1: Comparison of efficiency and performance of different strategies. For each comparison, best results are highlighted and second best one emphasized in bold font.

Approach	Training time	PSNR (dB)	SSIM
Global uniform E2E	$\sim$ 2 hrs	26.47	0.9106
Patch-wise E2E	$\sim$ 90 hrs	27.01	0.9110
Our base model	$\sim$ 5 hrs	26.27	0.9127
Ours refined model	$\sim$ **3 mins**	**27.36**	**0.9382**

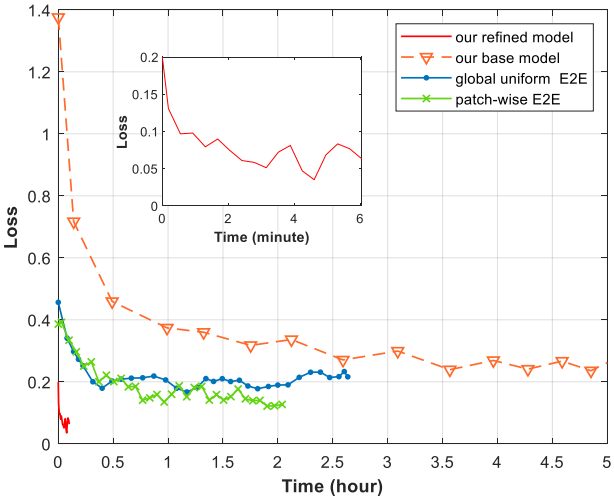

Figure 5: Training curves of four different learning strategies, in correspondence to those in Table 1. Here the curve for the patch-wise E2E is the average training time for a single patch.

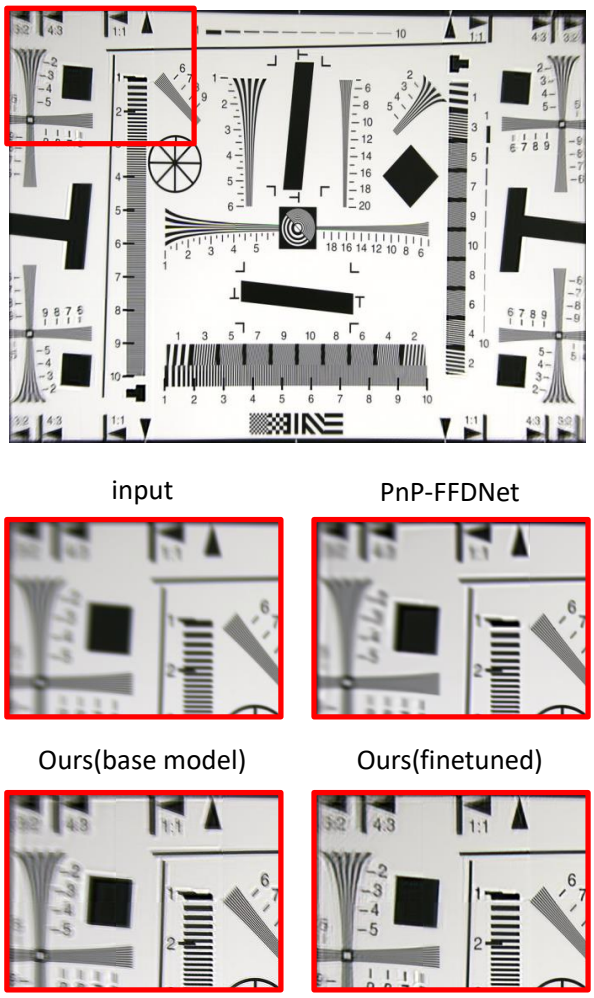

Figure 6: Comparison among different PnP schemes, with general denoiser (FFDNet), universal lens aberration compensation network (our base model) and lens specific compensation network (our refined model), incorporated in the same iterative framework.

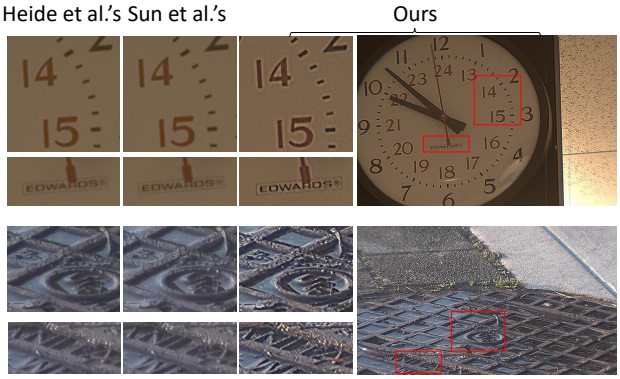

Figure 7: Comparison with one recently published method proposed by Sun et al. [29].

(PSNR) and SSIM [32].

For our base model, we use synthetic aberrant-sharp image pairs and PSFs of 20 different lenses with known Zemax prescriptions. For the other three settings, we use three simple lenses from Thorlabs Corp.: AC254-075-A-ML, LA4924-A, and LA5714. During the experiment, the synthetic image is set to 1024 $\times$ 768 pixels, and the patch size is set to be 128 $\times$ 128 pixels, under which the PSF can be roughly assumed uniform within a patch. The experiments are run on a computer with GTX TITAN 12G Memory GPU, Intel i9 CPU and 16G RAM.

The results are shown in Table 1 and one visual comparison is shown in Fig. 4, from which we can see that assuming the PSF to be globally uniform will cause low reconstruction quality, which again verifies that non-uniform compensation is crucial for lens aberration. Among the experiment settings applied to non-uniform deconvolution, our final results achieve comparable performance with the patch-wise E2E method, and superior performance to the base model before lens-specific refining.

Regarding time to train a model for a given lens, the patch-wise E2E model is exhausting since the model learning for the patches with severe degeneration is quite time consuming. and we usually need to learn tens of models for a mega-pixel image. For the globally uniform model, the required time is similar to that for learning a local E2E model. For the base model working generally for most lenses, around 5 hours are required for training, while for our lens specific adaption, one can easily build a high-quality aberration removal model on the top of the pre-trained base model within minutes, see the training curve in Fig. 5. In the inference stage, the deep learning based methods take similar time, around **2.4 seconds** for a **1024**$\times$**768** input image under our settings.

Overall, we show that non-uniform PSF awareness is a essential factor in solving the optical aberration correction problem using deep learning. Our universal pretrain plus lens-specific finetune strategy is also of both high efficiency and performance, making it a practical method.

4.2. Comparison to other plug-and-play schemes

Plug-and-play is an effective option for image restoration under non-uniform degeneration. Some recent researches treat the prior imposition as a denoising operation and use a generic deep image denoiser for different tasks. Therefore, one can also pre-train a denoising network and incorporate it into our iteration framework of Eq. 4 and Eq. 5 in a plug-and-play fashion. Besides, if we do not conduct lens-specific refinement, we can directly use the base model to any lens aberration model, as the plug-and-play method does.

To demonstrate the superiority of building aberration specific model and conducting lens-specific optimization,

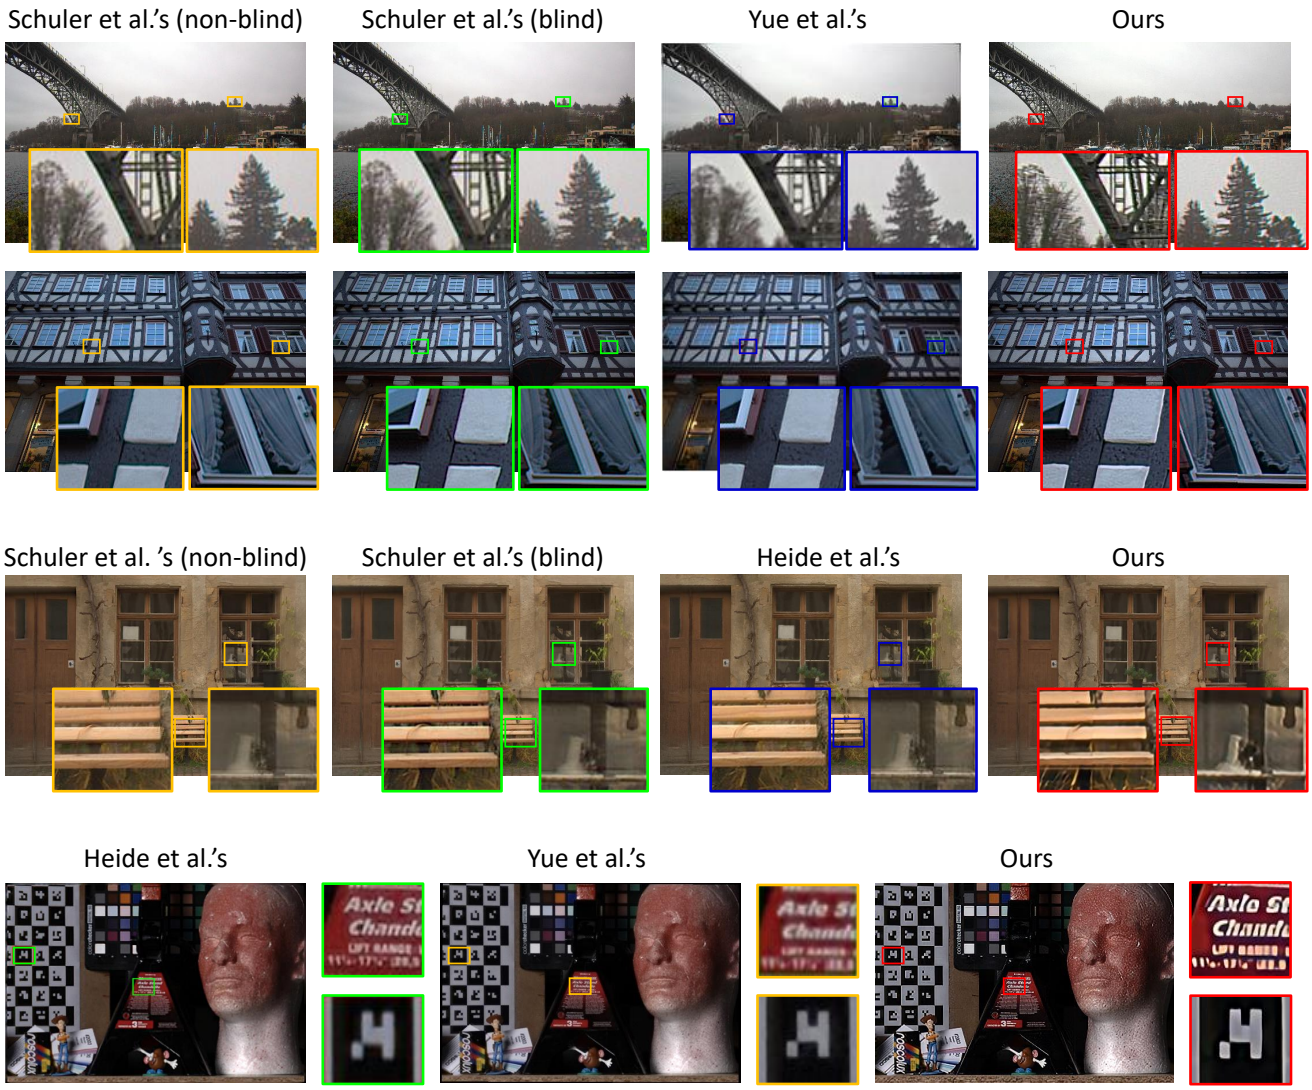

Figure 8: Comparison with state-of-the-art methods, including Schuler et al.'s [23][24], Yue et al.'s [36], and Heide et al.'s [7].

we conduct a comparison with above two plug-and-play schemes. For the generic denoiser we use the FFDNet [40]. We test different noise levels and find the one with best performance for the noise parameter setting in plug-n-play framework. The results are shown in Fig. 6, where we used the standard resolution chart for test and use the PSF of the same lens as in Fig. 1. We found that PnP-FFDNet produces more blurry results than our base model and the lens-specific model. This is reasonable because the deconvolution artifacts might be largely different from the Gaussian noise. The refined network performs better than the base model without lens-specific adaption. This validates the necessity of adaption.

4.3. Comparison to state-of-the-arts

We compare our method with other state-of-the-art aberration correction methods including Schuler et al.'s [24] [23], Yue et al.'s [36], Heide et al.'s [7], and Sun et al.'s [29] on real captured data. The results are presented in Figs. 7 and 8. As there is no ground-truth data, we only compare them qualitatively. It can be observed that: (i) our deep learning based methods have a superior or comparable performance over those generic prior optimization-based methods, which indicates that the deep model can represent the statistics of nature images quite better; (ii) blind compensation can achieve comparable performance with the non-blind ones, if PSF can be reliably estimated from the input degenerated image, and (iii) our approach performs better than (at least comparable to) state-of-the-

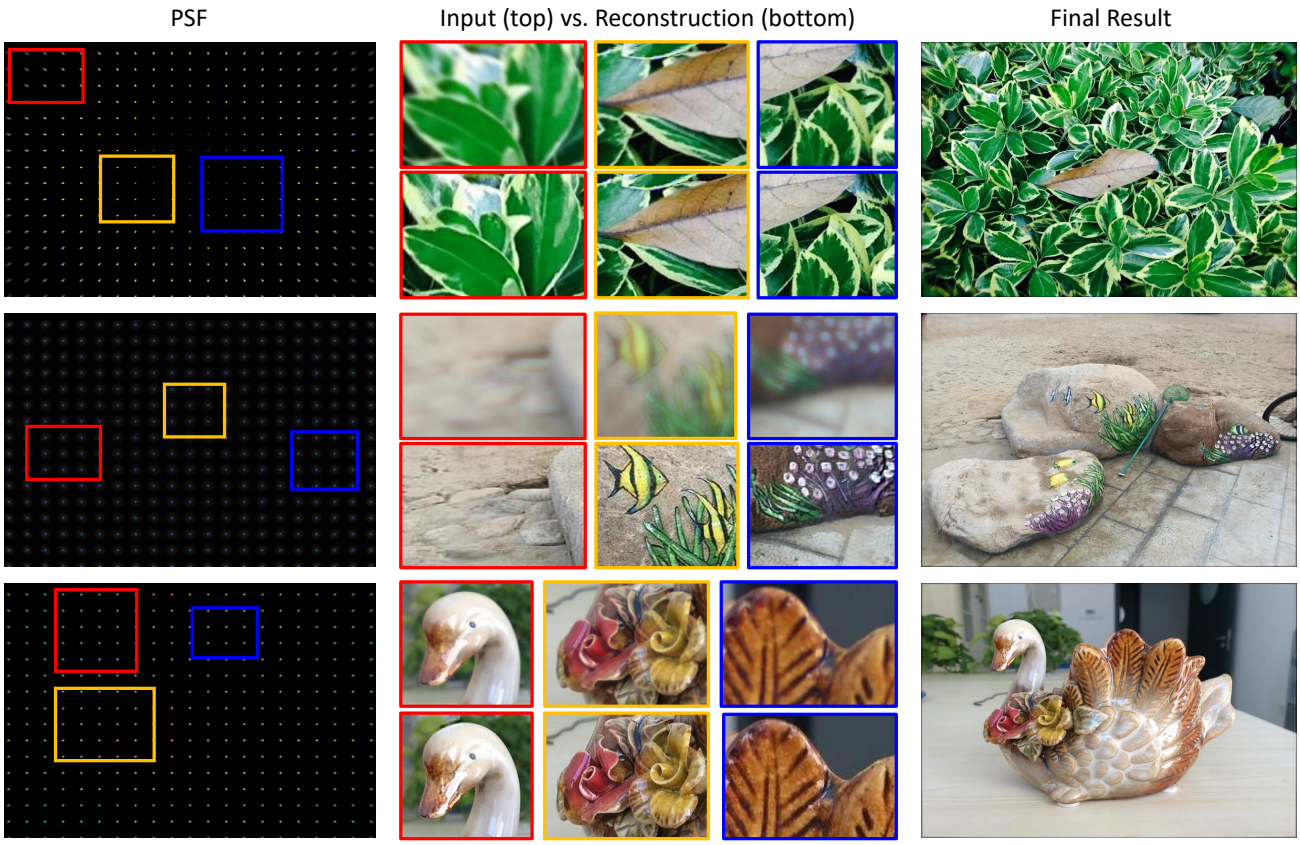

Figure 9: Results on real data captured with 3 different simple lenses. Left: estimated PSF. Middle: the comparison between the input and aberration-compensated version of some regions. Right: the holistic image of the final reconstruction. Please see more result in the supplemental material.

arts and is quite promising, considering its high flexibility and efficiency.

4.4. Results on real data

To test the performance of our approach on real data, we mount the simple lenses on a digital camera (HIKROBOT MV-CB013-20UC-C) to capture some photos, and conduct computational lens aberration. Here we demonstrate three cases with the same simple lenses as in Sec. 4.1. We use the algorithm by Sun et al. [28] to estimate the PSFs from images of several printed binary patterns and taking average to raise the performance, as shown in the leftmost column in Fig. 9. We then finetune our pretrained base model to each lens, and then feed the calibrated PSF together with the input degenerated image into our framework for final reconstruction, in the rightmost column in Fig. 9. We can see that the proposed approach can remove the blur effectively for different lenses, supporting our proposed method to be 'universal' and 'flexible'.

5. Summary and discussions

This paper reports a computational reconstruction approach using the end-to-end deep neural network to achieve high-quality imaging with a simple lens. Technically, the model is of high efficiency in both training and inference, and of high flexibility for adapting to different lenses.

Regarding the future work, we plan to build a training database from a customized lens set that covering tens of Zernike polynomial terms and learn a universal base model that can be adapted faster. Extending the model to explicitly incorporate cross-channel correlation and thus improving the final performance is another ongoing work.

Acknowledgements

This work is jointly funded by Ministry of Science and Technology of China(Grant No.2020AA0108202), National Natural Science Foundation of China(Grant No.61931012, 61627804 and 62088102), and Beijing Municipal Natural Science Foundation (Grant No.Z200021).

References

[1] E. Agustsson and R. Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, 2017. 4, 5

[2] J. H. R. Chang, C. Li, B. Póczos, B. V. K. Vijaya Kumar, and A. C. Sankaranarayanan. One network to solve them all — solving linear inverse problems using deep projection models. In *ICCV*, 2017. 3

[3] J. Dong, S. Roth, and B. Schiele. Deep wiener deconvolution: Wiener meets deep learning for image deblurring. In *NeurIPS*, 2020. 2

[4] T. Eboli, J. Sun, and J. Ponce. End-to-end interpretable learning of non-blind image deblurring. In *ECCV*, 2020. 3

[5] R. Fergus, B. Singh, A. Hertzmann, S. Roweis, and W. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graph.*, 25(3), 2006. 2

[6] A. Gupta, N. Joshi, C. Zitnick, M. Cohen, and B. Curless. Single image deblurring using motion density functions. In *ECCV*, 2010. 2

[7] F. Heide, M. Rouf, M. B. Hullin, B. Labitzke, W. Heidrich, and A. Kolb. High-quality computational imaging through simple lenses. *ACM Trans. Graph.*, 32(5), 2013. 2, 7

[8] M. Hirsch, C. J. Schuler, S. Harmeling, and B. Sch?lkopf. Fast removal of non-uniform camera shake. In *ICCV*, 2011. 2

[9] D. S. Jeon, S.-H. Baek, S. Yi, Q. Fu, X. Dun, W. Heidrich, and M. H. Kim. Compact snapshot hyperspectral imaging with diffracted rotation. *ACM Trans. Graph.*, 2019. 2, 3

[10] N. Joshi, R. Szeliski, and D. J. Kriegman. Psf estimation using sharp edge prediction. In *CVPR*, 2008. 2

[11] E. Kee, S. Paris, S. Chen, and J. Wang. Modeling and removing spatially-varying optical blur. In *ICCP*, 2011. 2

[12] D. P. Kingma and J. L. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5

[13] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-laplacian priors. In *NeurIPS*, 2009. 2

[14] J. Kruse, C. Rother, and U. Schmidt. Learning to push the limits of efficient fft-based image deconvolution. In *ICCV*, 2017. 2

[15] O. Kupyn, T. Martyniuk, J. Wu, and Z. Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *ICCV*, 2019. 2

[16] A. Levin, R. Fergus, F. Durand, and W. Freeman. Image and depth from a conventional camera with a coded aperture. *ACM Trans. Graph.*, 26(3), 2007. 2

[17] T. Michaeli and M. Irani. Blind deblurring using internal patch recurrence. In *ECCV*, 2014. 2

[18] X. Pan, X. Zhan, B. Dai, D. Lin, C. Loy, and P. Luo. Exploiting deep generative prior for versatile image restoration and manipulation. In *ECCV*, 2020. 3

[19] K. Rahbar and K. Faez. Blind correction of lens aberration using zernike moments. In *ICIP*, 2011. 2

[20] Y. Romano, M. Elad, and P. Milanfar. The little engine that could: Regularization by denoising(red). *SIAM J. Img. Sci.*, 10(4), 2017. 3

[21] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 4

[22] E. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin. Plug-and-play methods provably converge with properly trained denoisers. In *ICML*, 2019. 3

[23] C. Schuler, M. Hirsch, S. Harmeling, and B. Scholkopf. Non-stationary correction of optical aberrations. In *ICCV*, 2011. 2, 7

[24] C. Schuler, M. Hirsch, S. Harmeling, and B. Schölkopf. Blind correction of optical aberrations. In *ECCV*, 2012. 2, 7

[25] C. J. Schuler, H. C. Burger, S. Harmeling, and B. Schölkopf. A machine learning approach for non-blind image deconvolution. In *CVPR*, 2013. 2

[26] Y. Shih, B. Guenter, and N. Joshi. Image enhancement using calibrated lens simulations. In *ECCV*, 2012. 2

[27] H. Son and S. Lee. Fast non-blind deconvolution via regularized residual networks with long/short skip-connections. In *ICCP*, 2017. 2

[28] L. Sun, S. Cho, J. Wang, and J. Hays. Edge-based blur kernel estimation using patch priors. In *ICCP*, 2013. 2, 8

[29] T. Sun, Y. Peng, and W. Heidrich. Revisiting cross-channel information transfer for chromatic aberration correction. In *ICCV*, 2017. 2, 6, 7

[30] H. Tang and K. N. Kutulakos. What does an aberrated photo tell us about the lens and the scene? In *ICCP*, 2013. 2

[31] Y. Wang, J. Yang, W. Yin, and Y. Zhang. A new alternating minimization algorithm for total variation image reconstruction. *SIAM J. Img. Sci.*, 1(3), 2008. 2

[32] Z. Wang, A. C. Bovik, H. R. Hamid, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *TIP*, 13(4), 2004. 6

[33] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce. Non-uniform deblurring for shaken images. In *CVPR*, 2010. 2

[34] X. Yuan, Y. Liu, J. Suo, and Q. Dai. Plug-and-play algorithms for large-scale snapshot compressive imaging. In *CVPR*, 2020. 3

[35] T. Yue, J. Suo, X. Cao, and Q. Dai. Efficient method for high-quality removal of nonuniform blur in the wavelet domain. *IEEE TCSVT*, 27(9), 2017. 2

[36] T. Yue, J. Suo, J. Wang, X. Cao, and Q. Dai. Blind optical aberration correction by exploring geometric and visual priors. In *CVPR*, 2015. 2, 7

[37] H. Zhang, Y. Dai, H. Li, and P. Koniusz. Deep stacked hierarchical multi-patch network for image deblurring. In *CVPR*, 2019. 2

[38] J. Zhang, J. Pan, W. Lai, R. W. H. Lau, and M. Yang. Learning fully convolutional networks for iterative non-blind deconvolution. In *CVPR*, 2017. 2

[39] K. Zhang, L. V. Gool, and R. Timofte. Deep unfolding network for image super-resolution. In *CVPR*, 2020. 3

[40] K. Zhang, W. Zuo, and L. Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *TIP*, 27(9), 2018. 7