

Polyblur: Removing Mild Blur by Polynomial Reblurring

Mauricio Delbracio [✉], *Member, IEEE*, Ignacio Garcia-Dorado, Sungjoon Choi, Damien Kelly, and Peyman Milanfar, *Fellow, IEEE*

Abstract—We present a highly efficient blind restoration method to remove mild blur in natural images. Contrary to the mainstream, we focus on removing slight blur that is often present, damaging image quality and commonly generated by small out-of-focus, lens blur, or slight camera motion. The proposed algorithm first estimates image blur and then compensates for it by combining multiple applications of the estimated blur in a principled way. To estimate blur we introduce a simple yet robust algorithm based on empirical observations about the distribution of the gradient in sharp natural images. Our experiments show that, in the context of mild blur, the proposed method outperforms traditional and modern blind deblurring methods and runs in a fraction of the time. Our method can be used to blindly correct blur before applying off-the-shelf deep super-resolution methods leading to superior results than other highly complex and computationally demanding techniques. The proposed method estimates and removes mild blur from a 12MP image on a modern mobile phone in a fraction of a second.

Index Terms—Efficient image deblurring, mild blur estimation, mobile imaging.

I. INTRODUCTION

IMAGE sharpness is undoubtedly one of the most relevant attributes defining the visual quality of a photograph. Blur is caused by numerous factors such as the camera’s focus not being correctly adjusted, objects appearing at different depths, or when relative motion between the camera and the scene occurs during exposure. Even in perfect conditions, there are unavoidable physical limitations that introduce blur. Light diffraction due to the finite lens aperture, integration of the light in the sensor and other possible lens aberrations introduce blur leading to a loss of details. Additionally, other components of the image processing pipeline itself, particularly demosaicing and denoising, can introduce blur.

Removing blur from images is a longstanding problem in image processing and computational photography spanning more than 50 years [1]–[3]. Progress has been clear and sustained. From image enhancement algorithms [4], [5], blind and non-blind deconvolution methods [6]–[9] where sophisticated priors

Manuscript received March 9, 2021; revised June 4, 2021; accepted July 12, 2021. Date of publication July 30, 2021; date of current version August 12, 2021. (Corresponding author: Mauricio Delbracio.)

The authors are with Google Research, Mountain View, CA 94043 USA (e-mail: mdelbra@google.com; ignacioid@google.com; sungjoonc@google.com; damienkelly@google.com; milanfar@google.com).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2021.3100998>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2021.3100998



Fig. 1. Mild blur, as shown in these examples, can be efficiently (42 ms/MP on a mid-range smartphone CPU) removed by combining multiple applications of the estimated blur. Readers are encouraged to zoom-in for better visualization.

are combined with optimization schemes, to very recent years with the incipient application of deep neural models [10]–[16].

Many of these methods are capable of processing significantly degraded images revealing previously unseen image details. The gain in quality is in many cases extraordinary, but impractical. These methods make extensive use of prior information (learned or modeled) producing images that are often unrealistic. This is mainly because the inverse problem they tackle is seriously ill-posed.

In this work, we detach ourselves from the current trend and focus on the particular case where the blur in the image




Fig. 2. The deblurred image is generated by combining multiple re-applications of the estimated blur (polyblur).

is small. As we show in our experimental results, for this case, the vast majority of existing methods generate notorious image artifacts in addition to requiring great computational power. This is mainly because most of the algorithms were not specifically designed for this very common and practical use case. We introduce a highly efficient blind image restoration method that removes mild blur in natural images. The proposed algorithm first estimates image blur and then compensates for it by combining multiple applications of the estimated blur in a principled way (Figure 2). The method is inspired by the fact that the inverse of an operator that is close to the identity (e.g., mild blur), can be well approximated by means of a low-degree polynomial on the operator itself. We design polynomial filters that use the estimated blur as a base and approximate the inverse without neglecting that image noise can get amplified.

The removal of blur commonly leads to the introduction of halos (oversharpening) mostly noticeable near image edges. To address this, we present a mathematical characterization of halos and propose a blending mechanism to render an artifact-free final image. This step, which is also highly efficient, is important to achieve consistent high quality.

Experiments with both real and synthetic data show that, in the context of mild blur in natural images, our proposed method outperforms traditional and modern blind deblurring methods and runs in a fraction of the time. The simplicity of the polynomial filter, together with the choice of the Gaussian function as blur the model, enables a highly efficient implementation. The method runs at interactive rates so it can be used on mobile devices. Polyblur can estimate and remove mild blur on a 12MP image on a modern mobile phone in a fraction of a second. Additionally, we show that Polyblur can be used to blindly correct blur on an image before applying an off-the-shelf deep super-resolution method. Learning-based single image super-resolution algorithms (SISR) are usually trained on images where the degradation operator has been modeled in an over-simplistic way (typically by bilinear or bicubic blur). Polyblur in combination with a standard deep SISR network leads to superior results than other highly complex and computationally demanding methods.

Contributions. We introduce a novel method to estimate and remove mild blur (very common in mobile photography) that is (i) highly efficient and simple, (ii) theoretically sound, and (iii) produces competitive or better results while being orders of magnitude faster. Additionally, our method can blindly correct blur before applying off-the-shelf deep super-resolution methods leading to superior results vs. highly complex and computationally demanding techniques. Our algorithm’s components are built in a principled way: from explicit, empirically verified

assumptions in the blur estimation, to mathematical models in the polynomial deblurring and halo removal. Each component is designed to maximize efficiency (12MP image processed in milliseconds on a mobile CPU).

The remainder of the paper is organized as follows. In Section II we discuss the closely related work, while in Section III we present the polynomial approximation of the blur inverse and the adopted polynomial deblurring family. Section IV introduces the anisotropic Gaussian blur model and the necessary tools that allow an efficient estimation of the parameters. Section V summarizes the main components of the proposed blind estimation and removal algorithm, and experimental results on real and simulated data are shown in Section VI. Conclusions are finally summarized in Section VII.

II. RELATED WORK

Image blur is generally modeled as a linear operator acting on a sharp latent image. If the blur is shift-invariant then the blurring operation amounts to a convolution with a blur kernel. This implies,

$$v = k * u + n, \quad (1)$$

where v is the captured image, u the underlying sharp image, k the unknown blur kernel and n is additive noise. In what follows we review the different families of methods that remove image blur.

Blind deconvolution and variational optimization. The typical approach to image deblurring is by formulating the problem as one of blind deconvolution where the goal is to recover the sharp image u without knowing the blur kernel k [2], [3], [17]. Most blind deconvolution methods run on two steps: a blur kernel k is first estimated and then a non-blind deconvolution technique is applied [6], [18]–[22]. Fergus *et al.* [6] is one of the best examples of the blind deconvolution variational family that seeks to combine image priors, assumptions on the blurring operator, and optimization frameworks, to estimate both the blurring kernel and the sharp image [8], [9], [22]–[27]. Estimating the blur kernel is easier than jointly estimating the kernel and the sharp image together. Levin *et al.* [7], [28] show that it is better to first solve a maximum a posteriori estimation of the kernel than the latent image and the kernel simultaneously. Notwithstanding, even in non-blind deblurring, the significant attenuation of the image spectrum by the blur and model imperfections, lead to an ill-posed inverse problem [29], [30].

Sharpening methods. Sharpening methods aim to reduce mild blur and increase overall image contrast by locally modifying the image. Unsharp masking, arguably the most popular sharpening

algorithm is very sensitive to noise and generally leads to over-sharpening artifacts [31], [32]. Zhu and Milanfar [33] propose an adaptive sharpening method that uses the local structure and a local sharpness measure to address noise reduction and sharpening simultaneously. Bilateral filtering [34], and its many variants/adaptations [4], [5], [35], [36] can be used to enhance the local contrast and high frequency details of an image. The overall idea is to proceed similarly to that of unsharp masking: the input image is filtered and then a proportion of the residual image is added back to the original input. This procedure boosts high-frequencies that are removed by the adaptive filter. Since sharpening methods do not explicitly estimate image blur, their deblurring performance is limited. In fact, sharpening algorithms are strictly local operators and do not have access to additional information from outside the local neighborhood. Whereas deblurring algorithms like ours use the estimation of blur, which in general is global or done on a larger neighborhood, to remove the blur.

Deblurring meets deep-learning. In the past five years, with the popularization of deep convolutional networks several image and video deblurring methods have been introduced [10]–[14], [16], [37]–[42]. Most of these methods target strong motion blur and are in general trained with large datasets with realistically synthesized image blur. Efficient deblurring using deep models is a very challenging task [43] particularly on mobile devices [44]. In fact, one major difficulty for training deep deblurring models is the challenge of collecting sharp and blurry paired image data needed for supervised training. To bypass this costly process, different techniques have been employed, from directly simulating motion blur kernels [37], [40], or using high frame rate videos and averaging successive frames mimicking a longer capture exposure [13], [38]. Most methods are trained to predict the sharp image using the blurry image as the only input and target minimizing some pixel or perceptual loss. In [45], authors present a video deblurring framework that jointly estimates the optical flow and the latent sharp frames, enforcing a reblurring consistency model that takes into account the estimated optical flow and the sharp images. Other techniques have exploited the use of adversarial training as a loss function [11], [12], [15], [39] and also to generate synthetic realistic data [15].

Blind deconvolution methods do a remarkable job when the image is seriously damaged and manage to enhance very low quality images. However, their performance in the particular case where blur is mild is limited since they introduce, in general, noticeable artifacts and have a significant computational cost. At the other extreme, adaptive sharpening methods do a fine job boosting image contrast when the input image has very little blur. Since they do not incorporate any explicit measure of the blur, their deblurring capabilities are restricted. Deep learning based methods require high computational cost, making them impractical in many contexts. This work focuses on the problem of estimating and removing slight blur. As we show in the experimental part, in this setting, the proposed method produces superior results to blind deconvolution techniques and other popular adaptive sharpening algorithms, and similar results vs. advanced deep networks while being significantly more efficient.

III. DEBLURRING WITH POLYNOMIAL FILTERS

A. Approximating the Blur Inverse

Let us first assume that we have an estimation of the blur kernel k . To recover the image u we need to solve an inverse (deconvolution) problem. If the blur kernel is very large, the inversion is significantly ill-posed and some form of image prior will be required. In this work, we assume that the image blur is mild (see examples on Figure 1) so, as we will show in the experimental section, there is no need to incorporate any sophisticated image prior. Instead we proceed by building a linear restoration filter constructed from the estimated blur. An interesting observation [46]–[48] is that by carefully combining different iterations of the blur operator on the blurry image we can approximate the inverse of the blur.

One way to illustrate this is as follows.

Lemma 1: Let K be the convolution operator with the blur kernel k and I the identity operator, then if $\|I - K\| < 1$ under some matrix norm, K^{-1} the inverse of K exists and,

$$K^{-1} = \sum_{i=0}^{\infty} (I - K)^i. \quad (2)$$

The proof is a direct consequence of the convergence of the geometric series. If we assume circular boundary conditions for the convolution, then the eigenvectors of the matrix K are the Fourier modes of k . This implies that $(I - K)^i = F^H (I - D)^i F$, where F is the Fourier basis, F^H is the Hermitian transpose of F , and D is a diagonal matrix with the eigenvalues of K .

In particular, the series converges if the kernel Fourier coefficients $\hat{k}(\zeta)$ satisfy $|\hat{k}(\zeta) - 1| < 1$. In the case of blur filters that conserve the overall luminosity, a reasonable hypothesis is that $k(\mathbf{x}) \geq 0$, and $\int k(\mathbf{x}) d\mathbf{x} = 1$. This implies that $|\hat{k}(\zeta)| \leq 1$ which is not enough to guarantee convergence.

Polynomial filters. According to Lemma 1 we can approximate the inverse with a polynomial filter using the blur kernel as a base. For instance, if we truncate the power series and keep up to order 3, the polynomial approximate inverse of K is,

$$K_3^{\text{inv}} = 4I - 6K + 4K^2 - K^3. \quad (3)$$

This motivates the use of more general polynomials

$$p(K) = \sum_{i=0}^d a_i K^i, \quad (4)$$

where the order d and coefficients (a_0, \dots, a_d) can be designed to amplify or attenuate differently depending on how the blur is affecting a particular component.

Symmetric filters with non-negative Fourier coefficients. Let us first assume that the blur filter $k(\mathbf{x})$ is symmetric $k(\mathbf{x}) = k(-\mathbf{x})$, and has non-negative Fourier coefficients. We will later show how we can generalize the approach to any filter. In this setting, $\hat{k}(\zeta)$ the Fourier coefficients of the filter k are in $[0,1]$. The Fourier coefficients of the polynomial filter are related with the ones from the base filter through the same polynomial,




Fig. 3. Polynomial filters can be analyzed directly on the *Real* line. We present examples of the proposed $p_{3,\alpha,b}$ family, the third order polynomial approximation of the inverse given by Eq. (3) (`poly-inverse3`). Following Eq. (5) a polynomial filter can be analyzed in the Fourier domain by looking into how $x \in [0, 1]$ is mapped by the polynomial. This reflects how the differently attenuated Fourier components get amplified or attenuated by the polynomial filter. The left figure can be thought of as essentially the magnitude frequency response of the proposed filters.

that is,

$$\widehat{p}(k)(\zeta) = \sum_{i=0}^d a_i \widehat{k}(\zeta). \quad (5)$$

A polynomial filter can be analyzed in the Fourier domain by looking into how the interval $I = [0, 1]$ is mapped by the polynomial. This reflects how the differently attenuated Fourier components get amplified or attenuated by the polynomial filter. In Figure 3 we show a plot of different example polynomials $p(x)$ when applied to a signal $x \in I$.

If we apply the polynomial filter to an image v that has been affected by the same base blur k as the one used to build the polynomial filter, we get

$$p(K)v = p(K)Ku + p(K)n, \quad (6)$$

where it becomes evident that there is a trade-off between inverting the blur (i.e., $p(K)K \approx Id$.) and avoiding noise amplification, i.e., $p(K)n \approx 0$. Our goal is to design polynomials that (i) try to invert the effect of blur in the frequencies that have not been significantly attenuated, and (ii) avoid over-amplifying frequencies affected by the blur. The right plot in Figure 3 shows $p(x)x$, which can be interpreted as the combined effect of applying the polynomial on a signal that has been attenuated by the same blur.

B. Designing Deblurring Polynomial Filters

We want to build a polynomial filter $p(K)K \approx Id$, when K is close to the identity, or in terms of the polynomial, $p(x) \approx 1/x$ if $x \approx 1$. Note that a Taylor expansion of $1/x$ at $x = 1$ of degree d , leads to the polynomial filter from truncating (2) to degree d . This polynomial does a good job approximating the inverse, but significantly amplifies noise, particularly in the most attenuated components. Instead, we propose to approximate the inverse but have better control of how noise is amplified. Let us assume we design a polynomial of degree d , i.e., we need to define the $d + 1$ coefficients. We will approximate the inverse by forcing the polynomial to have equal derivatives at the ones of the inverse




Fig. 4. Effect of using different polynomials from the $p_{3,\alpha,b}$ family. Increasing α boosts mid-frequencies (first row), increasing b leads to more sharper but slightly noisier results. Readers are encouraged to zoom-in for better visualization.

function at $x = 1$. This is done up to order $d - 2$,

$$p^{(i)}(x = 1) = (-1)^i i!, \quad (7)$$

for $i = 1, \dots, d - 2$. We also have the additional constraint that $p(x = 1) = 1$ (no change in luminosity). The remaining two degrees of freedom are left as *design parameters*. We can control how the mid-frequencies get amplified by controlling $p^{(d-1)}(x = 1) = \alpha$, and also, how noise is amplified at the frequencies that are completely attenuated by the blur, by controlling $p(x = 0) = b$. This system of $d + 1$ linear equations leads to a (closed-form) family of polynomials, $p_{d,\alpha,b}$. The value of α and b should vary in a range close to the one given by the truncated power series, i.e., $\alpha = (-1)^d d!$ and $b = d + 1$.

The polynomial computational cost is proportional to the order, so the lower, the better. We choose order three since it is the lowest order to control both mid-frequencies boosting (α), and noise amplification (b). The impact of using a low-order polynomial is that in a single application of the filter blur may still remain. However, the remaining blur may be removed by repeated applications of the filter as we show in the experimental results. The polynomial filter family of degree $d = 3$ is

$$p_{3,\alpha,b}(x) = (\alpha/2 - b + 2)x^3 + (3b - \alpha - 6)x^2 + (5 - 3b + \alpha/2)x + b. \quad (8)$$

In Figure 3 we show different polynomials and how their shape is affected by the design parameters α and b . All the results presented in this article are with polynomials of this third order family. As an illustration, we show in Figure 4 the behavior using different coefficients in the polynomial deblurring.

C. Generalization to Any Blur Filter

Kernels with negative or complex Fourier coefficients (as the one shown in Fig. 5 a) cannot be directly deblurred using our Polynomial filtering (Fig. 5 b). A way around this restriction is to apply a correction filter $c_k(x)$ to the input image so that total blur kernel $h = c_k * k$ has non-negative Fourier coefficients and the solution to the deblurring problem remains the same,

$$v = u * k \implies v_{\text{eq}} = v * c_k = u * k * c_k = u * k_{\text{eq}}. \quad (9)$$




Fig. 5. Kernels with negative or complex Fourier coefficients (as the one shown in blue) need to be compensated before applying our polynomial deblurring. In blue the input blur kernel k , in orange each respective correction filter c_k , while in red corrected kernels $h = c_k * k$ used in the Polynomial deblurring (b-d).

Filtering the image with the flipped kernel, i.e., $c_k(\mathbf{x}) = k(-\mathbf{x})$ will lead to $h(\mathbf{x})$ having real non-negative Fourier coefficients $\hat{h}(\zeta) = |\hat{k}(\zeta)|^2$. However, this correction filter introduces additional blur to the image, making the deblurring problem more ill-posed than the original one (Fig. 5 c). An alternative correction filter is to compensate for the *phase* but without introducing any additional attenuation of the spectra. This can be done by the pure phase filter, $\hat{c}_k(\zeta) = \bar{\hat{k}}(\zeta)/|\hat{k}(\zeta)|$, where $\bar{\hat{k}}(\zeta)$ denotes the complex conjugate of $\hat{k}(\zeta)$. c_k is a pure phase filter, i.e., it has a constant Fourier magnitude of one at all frequencies.

In practice, given v we can estimate k , and compute c_k . Then, we compute v_{eq} and k_{eq} . The solution u , to $v_{\text{eq}} = u * k_{\text{eq}}$ is the same as in the original problem but k_{eq} has non-negative Fourier coefficients allowing us to apply our polynomial deblurring. Note that in the Gaussian blur case no correction is needed. As can be shown in the example in Figure 5 d, phase correction leads to a deblurred artifact-free image.

IV. PARAMETRIC MILD BLUR MODEL AND ESTIMATION

Since we target mild-blur removal, there is not much gain in having a very fine model of the blur. We thus propose to model the blur kernel with an anisotropic Gaussian function, specified with three parameters: σ_0 , the standard deviation of the main axis, $\rho = \sigma_1/\sigma_0$, the ratio between the the principal axis and the orthogonal one standard deviations, and θ , the angle between the major axis and the horizontal. Thus, the Gaussian blur kernel at pixel (x, y) is:

$$k(x, y) = Z \exp(- (a_0 x^2 + 2a_1 xy + a_2 y^2)), \quad (10)$$

where $a_0 = \frac{\cos(\theta)^2}{2\sigma_0^2} + \frac{\sin(\theta)^2}{2\rho^2\sigma_0^2}$, $a_1 = \frac{\sin(2\theta)}{4\sigma_0^2}(\frac{1}{\rho^2} - 1)$, $a_2 = \frac{\sin(\theta)^2}{2\sigma_0^2} + \frac{\cos(\theta)^2}{2\rho^2\sigma_0^2}$ and Z is a normalization constant so the area of the kernel is one. Although this is a rough model for arbitrary blur in images, if the blur is small, either directional or isotropic, the anisotropic Gaussian parameterization is reasonable. Figure 6(left) shows examples of Gaussian blur kernels.

A. From Natural Image Model to Blur Estimation

Estimating a blur kernel given only an image is a challenging ill-posed problem. It is necessary to make some assumptions, for instance by using an image prior (learned from data or




Fig. 6. Gaussian blur model and examples (left). The middle plot shows the relationship between the computed gradient feature f_θ and σ_0 on simulated images (each point represents an image), and the calibrated model (Eq. (13), red curve). The right plot shows the error on the estimation of the blur angle θ , as a function of ρ . Angle estimation error is very low for anisotropic (directional) kernels (low ρ value). The color in middle and right panels represents the density of points, low density in blue, high density in yellow.

from a statistical or any kind of mathematical model). In this work, we focus on efficiency and require that the estimation of the blur is very fast. Thus, variational models, for instance those typically involved in blind-deconvolution approaches are prohibitively expensive. Instead, we will rely on the following rough observations about the image gradient distribution.

Assumption 1: In a sharp image, the maximum value of the image gradient in any direction is mostly constant and roughly independent of the image.

Assumption 2: If a sharp image is affected by Gaussian blur, the blur level in the direction of the principal axis will be linearly related to the inverse of the maximum image gradient in the principal directions.

The validity of these assumptions is discussed in Appendix A. Based on these assumptions we proceed as follows. We first estimate the maximum magnitude of the image derivative at all possible orientations and then take the minimum value among them. This leads to

$$f_\theta = \min_{\psi \in [0, \pi)} f_\psi = \min_{\psi \in [0, \pi)} \max_{\mathbf{x}} |\nabla_\psi v(\mathbf{x})|, \quad (11)$$

where $v(\mathbf{x})$ is the input image, and $\nabla_\psi f(\mathbf{x}) = \nabla v(\mathbf{x}) \cdot (\cos \psi, \sin \psi)$, is the directional derivative of v at direction ψ . Then, the blur kernel parameters are

$$\theta = \arg \min_{\psi \in [0, \pi)} f_\psi, \quad \sigma_0 = \frac{c}{f_\theta} \quad \text{and} \quad \sigma_1 = \frac{c}{f_{\theta_\perp}}, \quad (12)$$

where c is a coefficient controlling the assumed linear relation between the gradient feature and the level of blur.

Observation. The estimation of the image gradient at any given image pixel introduces additional blur (since some sort of interpolation is needed). This will be in addition to the image blur that the image already has. Let us assume the blur estimation introduces an isotropic Gaussian blur of strength σ_b , then the total blur of the image will be approximately $\sigma_{0_T}^2 = \sigma_b^2 + \sigma_0^2$ (due to the semigroup property of Gaussian function). This leads to

$$\sigma_0 = \sqrt{\frac{c^2}{f_\theta^2} - \sigma_b^2}, \quad \sigma_1 = \sqrt{\frac{c^2}{f_{\theta_\perp}^2} - \sigma_b^2}, \quad (13)$$

where c and σ_b are two parameters to be calibrated.

B. Model Calibration.

To calibrate the parameters c and σ_b we proceed as follows. Given a set of sharp high quality images, we simulate $K = 1000$ random Gaussian blurry images, by randomly sampling the blur space and the image set. The Gaussian blur kernels are generated by sampling random values for $\sigma_0 \in [0.3, 4]$ and $\rho \in [0.15, 1]$. Additive Gaussian white noise of standard deviation 1% is added to each simulated blurry image.

For each of the blurry images we compute the gradient features and the maximum and minimum values according to (11). The parameters c and σ_b are estimated by minimizing the mean absolute error. Figure 6 presents the calibration results. Additional details are included in Appendix A.

V. POLYBLUR IMPLEMENTATION DETAILS

In this section we describe a Polyblur implementation to allow us to achieve the performance to deblur a 12MP image on a modern mobile platform in a fraction of a second. The core parts of the algorithm are implemented using the language *Halide* [49]. Halide separates the algorithm implementation from its execution schedule (i.e., loop nesting, parallelization, etc) enabling us to more easily fuse pipeline stages for locality and to make use of SIMD and thread parallelism. Note that despite of this choice, most part of the algorithm can run in parallel and a parallel implementation could take full advantage of the parallel processing power of modern GPUs.

Blur estimation. We follow Sec. IV-A to estimate the blur. As the first step we normalize the image using quantiles ($q = 0.0001$ and $1 - q$) to be robust to outliers. From the image gradient we compute the directional derivative at $m = 6$ angles uniformly covering $[0, \pi)$. The maximum direction of the magnitude at each angle is found and among the m maximum values, we find the minimum value and angle (f_0, θ_0) through bicubic interpolation. Using (13) we compute σ_0 and σ_1 . More details about the blur estimation, as long as a pseudo-code for the blur estimation step, are included in Appendix B.

Deblurring step. The deblurring filter has a closed form given by the estimated blur and the polynomial in (8). In the case of using a third order polynomial, the restoration filter support is roughly three times the one of the estimated base blur. Large blur kernel convolutions can be efficiently computed in Fourier domain. If the Gaussian blur is separable, the convolution can be efficiently computed in-place. In this case, we do not compute the polynomial filter and directly apply and accumulate repeated applications of the Gaussian blur as follows:

$$v_0 = a_d v, \quad v_i = k * v_{i-1} + a_{d-i} v, \quad \text{for } i = 1, \dots, d, \quad (14)$$

where v is the input blurry image, k the estimated blur kernel, and a_0, \dots, a_d the deblurring polynomial coefficients, and v_d the deblurred output image. Non-separable Gaussian filtering can also be efficiently computed by two 1D Gaussian filters in non-orthogonal axis. This can further optimize our procedure.

Halo detection and removal. Halos can be generated due to misestimation of blur or more generally due to model mismatch. Halos appear at pixels where the blurry image and restored image have opposite gradients (gradient reversal). Let $v(\mathbf{x})$ be




Fig. 7. Halo Removal (HR). Top row shows a crop for the input image, the deblurred one, and the final merge with the halo removal step. In the bottom we show a profile (and a zoom-in on the right) of the insensitive values orthogonal to the balloon boundary (orange segment). The gradient reversal has been eliminated.

the blurry image and $\bar{v}(\mathbf{x})$ the deblurred one. Pixels with gradient reversal are those where $M(\mathbf{x}) = -\nabla v(\mathbf{x}) \cdot \nabla \bar{v}(\mathbf{x})$, is positive.

Let us compute a new image $\bar{v}_z(\mathbf{x})$ formed as a per pixel convex combination (blending) of the input image $v(\mathbf{x})$ and the deblurred one $\bar{v}(\mathbf{x})$ using a weights $z(\mathbf{x}) \in [0, 1]$,

$$v_z(\mathbf{x}) = z(\mathbf{x})v(\mathbf{x}) + (1 - z(\mathbf{x}))\bar{v}(\mathbf{x}). \quad (15)$$

We want to avoid halos in the final image $v_z(\mathbf{x})$, but keep as much as possible of the deblurred one $\bar{v}(\mathbf{x})$. If $z(\mathbf{x})$ does not change too fast, then

$$\nabla v_z(\mathbf{x}) = z(\mathbf{x})\nabla v(\mathbf{x}) + (1 - z(\mathbf{x}))\nabla \bar{v}(\mathbf{x}). \quad (16)$$

To avoid halos on the final image we require $\nabla v(\mathbf{x}) \cdot \nabla v_z(\mathbf{x}) > 0$. Then, we get that for pixels where $M(\mathbf{x}) > 0$, having

$$z(\mathbf{x}) \leq \frac{M(\mathbf{x})}{\|\nabla v(\mathbf{x})\|^2 + M(\mathbf{x})},$$

leads to a new image that does not have any gradient reversal introduced in the deblurring step. To keep as much as possible of the deblurred image, $z(\mathbf{x})$ should be as small as possible. Then the final image is generated by adopting

$$z(\mathbf{x}) = \max\left(\frac{M(\mathbf{x})}{\|\nabla v(\mathbf{x})\|^2 + M(\mathbf{x})}, 0\right), \quad (17)$$

in the convex combination given by (15). Figure 7 shows an example of a filtered image and the halo correction.

Polyblur iterated. To remove some remaining image blur, we can re-apply Polyblur. This implies re-estimating the image blur on the previous deblurred image, and applying a new polynomial deblurring. In the next section we present some results when iterating polyblur multiple times.

VI. EXPERIMENTS

We carried out a series of experiments to evaluate Polyblur and compare against other blind deblurring [12], [14], [50]–[53] and sharpening methods [33], [54]. The comparison is threefold: we compare traditional metrics such as PSNR, SSIM and perceptual ones such as LPIPS [55]; we do a visual inspection of artifacts (qualitative); and we report processing times. We also present results on images *in the wild* having mild blur. Finally, we show how Polyblur can be used to remove blur before applying an off-the-shelf image superresolution deep method. Table I shows

TABLE I

TIMES IN MS RUN ON DESKTOP (MOBILE). DESKTOP: INTEL XEON HASWELL 2.3 GHZ, MOBILE: SNAPDRAGON 855

Algorithm phase	12MP	3MP	1MP
1. Blur estimation	302 (357)	88 (95)	23 (25)
2a. Polyblur (separable)	66 (194)	16 (82)	9 (35)
2b. Polyblur (Fourier)	266 (472)	35 (212)	16(134)
3. Halo removal	14 (70)	61 (20)	3 (7)

TABLE II

MILD DEBLURRING ON THE SYNTHETIC DATASET GENERATED FROM DIV2K. LOWER LPIPS VALUES IMPLY BETTER QUALITY. POLYBLUR- N IT INDICATES APPLYING POLYBLUR N TIMES, AND *noHR* INDICATES POLYSHARP WITHOUT THE HALO-REMOVAL STEP. CONVDEB [51]* DID NOT PRODUCE REASONABLE RESULTS IN 32/100 IMAGES (AVERAGE VALUES ARE COMPUTED IN THE REST). TIMES ARE IN MS (FOR A 1MP IMAGE), AND WE PRESENT THEM HERE AS AN INDICATIVE REFERENCE. † INDICATES GPU

Method	PSNR	MS-SSIM	LPIPS	Time
SRN [14]	27.58	0.948	0.304	75-152†
SparseDeb [53]	27.48	0.940	0.296	> 5000
DeblGANv2 [12]	29.79	0.963	0.213	350†
DeblGANv2m [12]	29.53	0.958	0.209	60†
Spectral [50]	27.14	0.950	0.230	> 5000
Deblur-10 [52]	27.31	0.945	0.263	> 5000
GLAS [33]	28.15	0.957	0.284	> 5000
GuidFilt [54]	26.98	0.934	0.282	80
ConvDeb [51]*	27.18	0.941	0.233	90
Polyblur-1it (noHR)	30.35	0.968	0.238	39
Polyblur-1it	30.38	0.968	0.236	42
Polyblur-2it	30.06	0.964	0.255	84
Polyblur-3it	29.44	0.955	0.269	126
Blurry-noisy	29.31	0.956	0.245	-

average execution times for each step in the Polyblur algorithm. Our method can process a 12MP image on a modern mobile platform in 600 ms.

Comparison on simulated blur. We generated a mild-blur dataset by artificially blurring sharp images from the DIV2K validation dataset. Mild blur is simulated by applying a random Gaussian kernel of different sizes, shapes, and orientations ($\sigma_0 \sim \mathcal{U}[0.3, 4]$). Additive white Gaussian noise of standard deviation 1% is added on top. More details are included in the supplementary material. Polyblur produces the best results in terms of PSNR and MS-SSIM while being significantly faster than most of the other deblurring methods (Table II). As shown in Figure 8, Polyblur leads to naturally pleasant images, while most of the compared methods introduce artifacts. The visual quality of the output results is similar to highly complex methods such as DeblurGAN-v2 [12]. Our CPU implementation of Polyblur is still 50% faster than the highly optimized mobile version of DeblurGANv2 that runs on GPU.

Re-applying Polyblur produces slightly worse quantitative results on the simulated Gaussian blur dataset (Table II) but in general leads to better qualitative results on real images where blur might not be perfectly Gaussian (e.g., Figures 23–27 in the supplemental material).

Dealing with noise and compression artifacts. If computational resources are available, a prefiltering step separating



Fig. 8. Comparison of Polyblur with other deblurring and adaptive sharpening methods on synthetically blurred DIV2K dataset.

noise-like structure from the rest can be applied. If the input image is very noisy or has compression artifacts, this prefiltering step will keep Polyblur from amplifying artifacts present in the image. Since the residual image can be added back at the end, this step does not need to be carried out by a state-of-the-art denoiser. Fig. 9 shows an example of Polyblur on a noisy and compressed image. We evaluated two alternatives, the Pull-Push denoiser [56] and the Domain Transform [57] edge-preserving filter. In both cases we achieve the desired deblurring, and a pleasant result.

Results on images in the wild. In Figure 1 we present a selection of results of Polyblur applied to some images in the wild. As shown, Polyblur manages to remove mild blur, as the one present in most images, without introducing any new



Fig. 9. Pre-filtering. When the input image presents noise and other artifacts, Polyblur can amplify artifacts. As a pre-step, we may apply any filter that separates high-frequency texture and other details [56], [57], apply Polyblur ($\alpha = 6$, $b = 1$), and finally add back the residual.

TABLE III
4X SUPER-RESOLUTION ON DIV2KRRK DATASET [60]. POLYBLUR-NIT INDICATES APPLYING POLYBLUR N TIMES. PSNR AND SSIM VALUES ARE DIRECTLY TAKEN FROM [60], [61]. IN RED FONT THE BEST METHOD, IN BLUE THE SECOND BEST

Method	PSNR	SSIM
Bicubic	25.354	0.6775
EDSR [59]	25.628	0.6971
KernelGAN [60]	26.810	0.7316
CorrectionFilter[61] + EDSR	26.204	0.736
Polyblur-1it + EDSR	26.966	0.7396
Polyblur-2it + EDSR	27.115	0.7452
Polyblur-3it + EDSR	26.955	0.7368

artifacts. Please refer to the supplemental material for comparison with other state-of-the-art deblurring and sharpening methods.

Deblurring before super-resolution. Single image super-resolution has seen remarkable progress in the last few years mostly due to the use of deep models trained on image datasets. The common practice [58] is to simulate low-resolution training images by simply applying a bicubic downsampling operator. Unfortunately, inference performance suffers significantly if images do not tightly follow the training distribution. We evaluate the performance of applying Polyblur as a pre-step before using an off-the-shelf deep network for doing $4\times$ image upscaling. We trained from scratch an EDSR [59] network with 32 layers and 64 filters using DIV2K training dataset. We compare against KernelGAN [60] and the Correction filter introduced in [61]. Figure 10 shows examples of applying Polyblur before upscaling the image. Table III shows a quantitative comparison of our approach and the competitors. Polyblur produces the best quantitative and qualitative results.

VII. CONCLUSION

We introduced a highly efficient algorithm for estimating and removing mild blur that is ubiquitously present in many captured images (especially on handheld devices). A parametric blur kernel is first estimated using simple gradient features that encode the direction and intensity of the blur. The estimated blur is then removed by combining multiple applications of the blur




Fig. 10. Image $4\times$ super-resolution using an off-the-shelf deep SISR model with unknown PSF (EDSR [59]). More results are given in the supplementary material.

kernel in a well founded way that allows us to approximate the inverse while controlling the noise amplification. Our method successfully handles blurs that are reasonably well captured by an anisotropic Gaussian of standard deviations 0.3–3.0. This includes small unidirectional motion blur (e.g., shown in Fig. 1 middle panel), lens blur, and slight defocus blur that are very common in mobile photography. Since we approximate the inverse with a low-order polynomial, the algorithm fails gracefully without introducing jarring artifacts. The most severe failure case is due to kernel mismatch, in which case the algorithm either fails to remove all blur in the image, or over-sharpens the image. The halo removal is designed to handle the latter case (real example in Figure 6). The current blur estimation is global: one blur kernel per image. If there is non-uniform (shift-variant) blur, the algorithm removes the blur common to all regions, acting conservatively. As future work, we would like to explore how to extend our method to address stronger and non-uniform blurs while maintaining the high efficiency of the algorithm. Our experiments show that, in the context of mild blur, the algorithm produces similar or better results than other state-of-the-art image deblurring methods while being significantly faster. The whole deblurring process runs in a fraction of a second on a 12MP image on a modern mobile platform. Our method can be used to blindly correct blur before applying an off-the-shelf deep super-resolution model leading to superior results than other computationally demanding techniques.

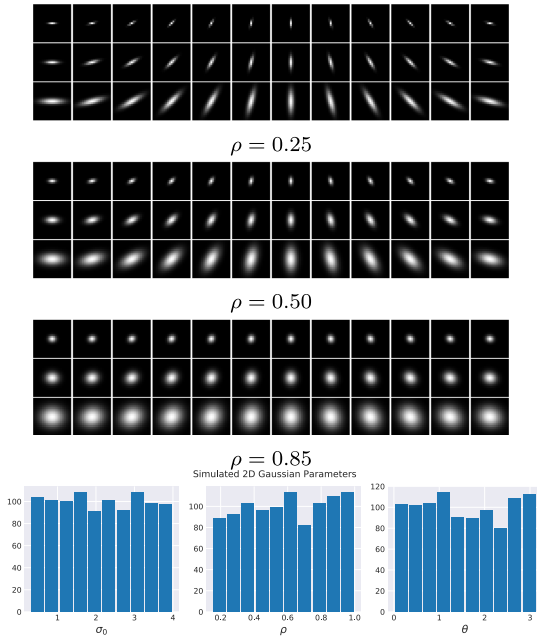


Fig. 11. Gaussian Blur kernel examples and distribution of simulated parameters.




Fig. 12. Gaussian blur model calibration. Error in the estimation of σ_0 and σ_1 (and ρ).

APPENDIX A

BLUR MODEL VALIDATION AND CALIBRATION

The Gaussian blur estimation presented in Section IV is based on computing gradient features. To calibrate c and σ_b , we proceed as follow. Given a set of 50 sharp high quality images, we simulate $K = 1000$ random Gaussian blurry images, by randomly sampling the blur space and the image set. The Gaussian blur kernels are generated by sampling random values for $\sigma_0 \in [0.3, 4]$ and $\rho \in [0.15, 1]$. Additive Gaussian white noise of standard deviation 1% is added to each simulated blurry image.

Examples of simulated Gaussian blur kernels are shown in Figure 11. For each of the blurry images we compute the gradient features according to (11). The parameters c and σ_b are estimated by minimizing the mean absolute error (MAE). The calibrated parameters are $c = 89.8$ and $\sigma_b = 0.764$. Note that the values of c and σ_b are implementation dependent (e.g., the finite difference scheme used to compute image gradient).

Figure 12-left shows the relation between the inverse of the estimated gradient feature (i.e., $1/f_\theta$) and the simulated blur kernel σ_0 value. Each of the blur points represents one simulated image. As we can see, sharp images that have very low blur

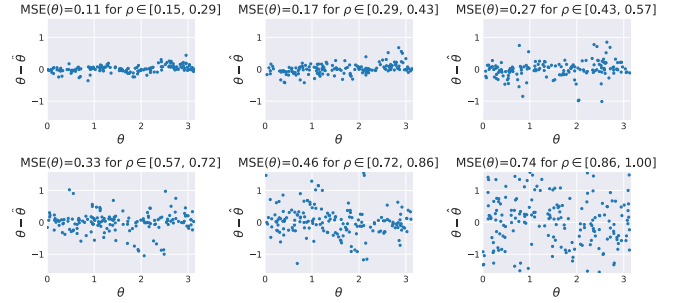


Fig. 13. Blur model calibration. Error in estimated parameters. Calibrated parameters are $c = 89.8$, $\sigma_b = 0.764$.

values ($\sigma_0 \ll 1$) have very similar feature values. The same analysis holds for the gradient feature at the orthogonal direction, and its relation to $\sigma_1 = \rho\sigma_0$ (Figure 12-middle). This validates our *Assumption 1*.

Additionally, Figure 12 shows the almost linear relation (with some spread) between the inverse of the gradient feature $1/f_\theta$ and the blur level σ_0 . This is exactly *Assumption 2*.

Our model contemplates for slight blur that the gradient operator may have introduced when computing the gradient features (σ_b). This is further analyzed at the end of this appendix for a synthetic image.

In Figure 12-right we show a plot of the real ρ value of the simulated blur kernel and the estimated one. Although there are some outliers, the estimation is in general close to the real value ($\text{MSE}(\rho) = 0.121$).

Estimation of blur direction θ . The estimation of the blur direction is done by computing the angle θ with minimum gradient feature value in Eq. (11). Figure 13 shows the error on the estimation of the angle for each simulated blur. As we can see, the error in the angle is quite low for blur kernels highly directional ($\rho < 0.5$). For large values of ρ (e.g., $\rho \in [0.75, 1.0]$), the kernels are almost isotropic, and the estimation is inaccurate. Nevertheless, being almost isotropic, the kernel shape is not affected by the angle value in this case.

Error metrics on estimated kernel values. The ultimate goal in blur estimation is to estimate a blur that is close to the real one. In Figure 14 we present the distribution of two different error metrics that evaluate the distance between the estimated kernel and the simulated one directly on the kernel space. The first metric is the kernel similarity, which is the normalized cross correlation:

$$\text{ksim}(k, \hat{k}) = \frac{1}{\|k\| \|\hat{k}\|} \sum_i k_i \hat{k}_i. \quad (18)$$

When both kernels are equal the kernel similarity is 1. The second metric we compute is the ℓ_1 norm between estimated kernel values. To give an idea of the range of both metrics, we also present the error distribution of the real kernel and an average kernel.¹ This shows the estimations are accurate.

Concentric circles: A synthetic example. We generated a

¹The average kernel is defined as the isotropic kernel having standard deviation the central value of the simulated range $[0.35, 4.0]$.

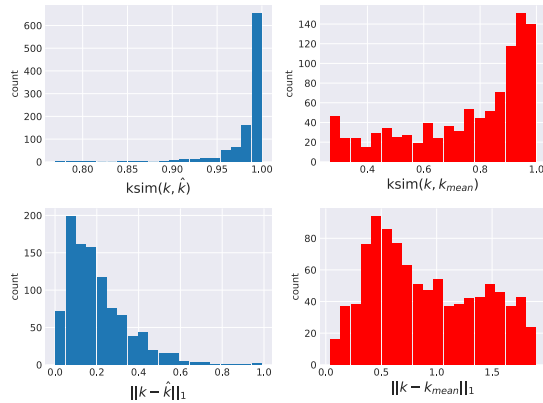
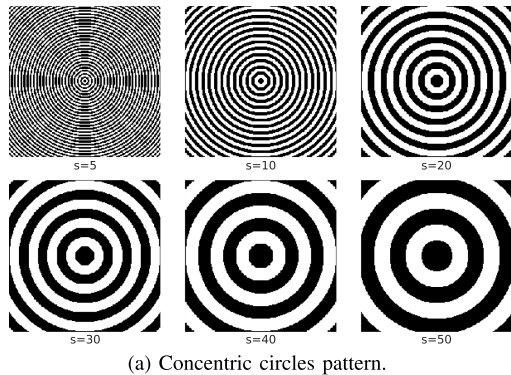



Fig. 14. Kernel error metrics on the simulated data. Top-row shows the histogram of the kernel similarity between estimated kernels and the respective ground-truth one (left), and between the ground-truth kernels and a fixed isotropic Gaussian kernel (denoted by k_{mean}) having standard deviation the mid-value on the simulated range (right). Bottom-row shows the histogram of the ℓ_1 difference between the estimated kernels and the respective ground-truth one (left), and between the ground-truth kernels and k_{mean} .



(a) Concentric circles pattern.



(b) Gradient feature f_θ as a function of simulated blur level σ_0 .

Fig. 15. Blur model on concentric circles (synthetic image).

family of synthetic images with concentric circles at different distances (controlled by a parameter $s \in [5, 50]$), see Figure 15 a and blurred them with isotropic Gaussian blur of intensity ($\sigma_0 \in [0.3, 5]$). In Figure 15 b we show the computed gradient feature (f_θ) for the different blurry images (different σ_0) on each simulated pattern ($s = 5$ to $s = 50$). We calibrated the blur model (i.e., estimated c and σ_b values) using the $s = 50$ synthetic pattern image. The model is highly accurate when the circular rings are separated enough so the image can be locally considered a step-edge ($s \geq 20$). However, when the concentric circles are very close (e.g., $s = 5$), the model is not accurate and the blur estimation is biased.

Algorithm 1: Gaussian Blur Estimation.

input : image $u, q, n_{\text{angles}}, \sigma_{\text{max}}, \sigma_{\text{min}}, \rho_{\text{max}}, \rho_{\text{min}}$
output: Blur parameters: σ_0, ρ, θ .

// 1. Compute gradient features;

$n = \text{Normalize}(u, q)$;

$u_x, u_y = \text{ComputeGradient}(n, \sigma_{\text{max}})$;

for $i \in [0, n_{\text{angles}}]$ **do**

$\Psi_i = i \cdot \pi / n_{\text{angles}}$;

$u_{\Psi_i} = u_x \cdot \cos(\Psi_i) - u_y \cdot \sin(\Psi_i)$;

$f_{\Psi_i} = \text{Max}(|u_{\Psi_i}|)$;

end

$I_\Psi = \text{Interpolate}(\{f_\Psi\})$;

$f_\theta, \theta_0 = \text{Min}(I_\Psi)$;

$f_{\theta_\perp}, \theta_{0_\perp} = \text{Ortho}(f_\theta, \theta_0)$;

// 2. Compute and clamp Gaussian;

$\sigma_0, \sigma_1 = \text{Regression}(f_\theta, f_{\theta_\perp})$ // From Eq. (13);

$\text{Clamp}(\sigma_0, \sigma_1, \sigma_{\text{max}}, \sigma_{\text{min}})$;

$\rho = \sigma_0 / \sigma_1$;

$\text{Clamp}(\rho, \rho_{\text{min}}, \rho_{\text{max}})$;

For comparison purposes we also fit a purely linear model that directly maps the inverse of the gradient feature to the sigma value. The linear model (which is the proposed model with $\sigma_b = 0$) is very close to the proposed model except in low sigma values. This is because the proposed model takes into account the (very little) blur σ_b that the gradient operator introduces, leading to a more precise estimation.

APPENDIX B

BLUR ESTIMATION IMPLEMENTATION DETAILS

In Algorithm 1 we present the pseudo-code for our estimation of Gaussian blur. We follow Section 4 to estimate the blur. As the first step the input image is normalized using quantiles ($q = 0.0001$ and $1 - q$) to be robust to outliers. From the image gradient (u_x, u_y) we compute the directional derivative at n_{angles} (usually $n_{\text{angles}} = 6$) uniformly covering $[0, \pi)$. The maximum direction of the magnitude for each sample angle is found. Among the n_{angles} maximum values, we find the minimum value and angle (f_θ, θ_0) through bicubic interpolation. Using Eq. 13 we compute σ_0 and σ_1 (and compute $\rho = \sigma_1 / \sigma_0$).

The gradient features can be efficiently computed in parallel not just between different angles but also between each different pixel. Computing the maximum can be represented as a gather operation that can be optimized using shared memory and tiling. Within our pipeline this represents around 40 – 53% of the computation, this comes from the fact that we compute a maximum magnitude n_{angles} times.

REFERENCES

- [1] L. S. Kovácsznay and H. M. Joseph, "Image processing," in *Proc. IRE*, vol. 43, no. 5, pp. 560–570, 1955.
- [2] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *IEEE Signal Process. Mag.*, vol. 13, no. 3, pp. 43–64, May 1996.

- [3] W.-S. Lai, J.-B. Huang, Z. Hu, N. Ahuja, and M.-H. Yang, "A comparative study for single image blind deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1701–1709.
- [4] K. He, J. Sun, and X. Tang, "Guided image filtering," in *Euro. Conf. Comput. Vis.*, Springer, 2010, pp. 1–14.
- [5] H. Talebi and P. Milanfar, "Fast multilayer laplacian enhancement," *IEEE Trans. Comput. Imag.*, vol. 2, no. 4, pp. 496–509, Dec. 2016.
- [6] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman, "Removing camera shake from a single photograph," in *Proc. ACM SIGGRAPH*, 2006, pp. 787–794.
- [7] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding and evaluating blind deconvolution algorithms," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 1964–1971.
- [8] D. Perrone and P. Favaro, "Total variation blind deconvolution: The devil is in the details," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 2909–2916.
- [9] Q. Shan, J. Jia, and A. Agarwala, "High-quality motion deblurring from a single image," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 1–10, 2008.
- [10] H. Gao, X. Tao, X. Shen, and J. Jia, "Dynamic scene deblurring with parameter selective sharing and nested skip connections," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3848–3856.
- [11] O. Kupyn, V. Budzan, M. Mykhailych, D. Mishkin, and J. Matas, "Deblurgan: Blind motion deblurring using conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8183–8192.
- [12] O. Kupyn, T. Martyniuk, J. Wu, and Z. Wang, "Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 8878–8887.
- [13] S. Su, M. Delbracio, J. Wang, G. Sapiro, W. Heidrich, and O. Wang, "Deep video deblurring for hand-held cameras," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1279–1288.
- [14] X. Tao, H. Gao, X. Shen, J. Wang, and J. Jia, "Scale-recurrent network for deep image deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8174–8182.
- [15] K. Zhang *et al.*, "Deblurring by realistic blurring," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2737–2746.
- [16] P. Wieschollek, M. Hirsch, B. Scholkopf, and H. Lensch, "Learning blind motion deblurring," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 231–240.
- [17] T. F. Chan and C.-K. Wong, "Total variation blind deconvolution," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 370–375, Mar. 1998.
- [18] S. Cho and S. Lee, "Fast motion deblurring," in *Proc. ACM SIGGRAPH Asia*, 2009, pp. 1–8.
- [19] J. Pan, Z. Hu, Z. Su, and M.-H. Yang, "Deblurring text images via l_0 -regularized intensity and gradient prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 2901–2908.
- [20] J. Pan, D. Sun, H. Pfister, and M.-H. Yang, "Deblurring images via dark channel prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 10, pp. 2315–2328, Oct. 2018.
- [21] L. Xu and J. Jia, "Two-phase kernel estimation for robust motion deblurring," in *Proc. Euro. Conf. Comput. Vis.*, Springer, 2010, pp. 157–170.
- [22] L. Xu, S. Zheng, and J. Jia, "Unnatural l_0 sparse representation for natural image deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 1107–1114.
- [23] J.-F. Cai, H. Ji, C. Liu, and Z. Shen, "Blind motion deblurring from a single image using sparse approximation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 104–111.
- [24] L. Chen, F. Fang, T. Wang, and G. Zhang, "Blind image deblurring with local maximum gradient prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1742–1750.
- [25] M. Jin, S. Roth, and P. Favaro, "Normalized blind deconvolution," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 668–684.
- [26] D. Krishnan, T. Tay, and R. Fergus, "Blind deconvolution using a normalized sparsity measure," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2011, pp. 233–240.
- [27] T. Michaeli and M. Irani, "Blind deblurring using internal patch recurrence," in *Proc. Euro. Conf. Comput. Vis.*, Springer, 2014, pp. 783–798.
- [28] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Efficient Marginal Likelihood Optimization in Blind Deconvolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2011, pp. 2657–2664.
- [29] J. Anger, G. Facciolo, and M. Delbracio, "Modeling realistic degradations in non-blind deconvolution," in *Proc. 25th IEEE Int. Conf. Image Process.*, 2018, pp. 978–982.
- [30] D. Krishnan and R. Fergus, "Fast image deconvolution using hyperlaplacian priors," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 1033–1041.
- [31] S. H. Kim and J. P. Allebach, "Optimal unsharp mask for image sharpening and noise removal," *J. Electron. Imag.*, vol. 14, no. 2, p. 0 23005, 2005.
- [32] A. Polesel, G. Ramponi, and V. J. Mathews, "Image enhancement via adaptive unsharp masking," *IEEE Trans. Image Process.*, vol. 9, no. 3, pp. 505–510, Mar. 2000.
- [33] X. Zhu and P. Milanfar, "Restoration for weakly blurred and strongly noisy images," in *Proc. IEEE Workshop Appl. Comput. Vis.*, 2011, pp. 103–109.
- [34] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. IEEE 6th Int. Conf. Comput. Vis.*, 1998, pp. 839–846.
- [35] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2005, vol. 2, 2005, pp. 60–65.
- [36] P. Milanfar, "A tour of modern image filtering: New insights and methods, both practical and theoretical," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 106–128, Jan. 2013.
- [37] A. Chakrabarti, "A neural approach to blind motion deblurring," in *Proc. Euro. Conf. Comput. Vis.*, Springer, 2016, pp. 221–235.
- [38] S. Nah, T. Hyun Kim, and K. Mu Lee, "Deep multi-scale convolutional neural network for dynamic scene deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3883–3891.
- [39] S. Ramakrishnan, S. Pachori, A. Gangopadhyay, and S. Raman, "Deep generative filter for motion deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, 2017, pp. 2993–3000.
- [40] J. Sun, W. Cao, Z. Xu, and J. Ponce, "Learning a convolutional neural network for non-uniform motion blur removal," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 769–777.
- [41] P. Wieschollek, B. Scholkopf, H. P. Lensch, and M. Hirsch, "End-to-end learning for image burst deblurring," in *Proc. Asian Conf. Comput. Vis.*, Springer, 2016, pp. 35–51.
- [42] W. Ren, J. Zhang, J. Pan, S. Liu, J. Ren, J. Du, X. Cao and M.H. Yang, "Deblurring Dynamic Scenes via Spatially Varying Recurrent Neural Networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, 2021. doi: [10.1109/TPAMI.2021.3061604](https://doi.org/10.1109/TPAMI.2021.3061604).
- [43] S. Nah, S. Son, R. Timofte, and K. Mu Lee, "Ntire 2020 challenge on image and video deblurring," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 416–417.
- [44] C.-M. Chiang *et al.*, "Deploying image deblurring across mobile devices: A perspective of quality and latency," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 502–503.
- [45] H. Chen, J. Gu, O. Gallo, M.-Y. Liu, A. Veeraraghavan, and J. Kautz, "Reblur2deblur: Deblurring videos via self-supervised learning," in *Proc. IEEE Int. Conf. Comput. Photography*, 2018, pp. 1–9.
- [46] J. Kaiser and R. Hamming, "Sharpening the response of a symmetric nonrecursive filter by multiple use of the same filter," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 25, no. 5, pp. 415–422, Oct. 1977.
- [47] P. Milanfar, "Rendition: Reclaiming what a black box takes away," *SIAM J. Imag. Sci.*, vol. 11, no. 4, pp. 2722–2756, 2018.
- [48] X. Tao, C. Zhou, X. Shen, J. Wang, and J. Jia, "Zero-order reverse filtering," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 222–230.
- [49] J. Ragan-Kelley, A. Adams, S. Paris, M. Levoy, S. Amarasinghe, and F. Durand, "Decoupling algorithms from schedules for easy optimization of image processing pipelines," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 1–12, 2012.
- [50] A. Goldstein and R. Fattal, "Blur-kernel estimation from spectral irregularities," in *Proc. Eur. Conf. Comput. Vis.*, Springer, 2012, pp. 622–635.
- [51] M. S. Hosseini and K. N. Plataniotis, "Convolutional deblurring for natural imaging," *IEEE Trans. Image Process.*, vol. 29, pp. 250–264, 2019.
- [52] J. Pan, Z. Hu, Z. Su, and M.-H. Yang, " l_0 -Regularized intensity and gradient prior for deblurring text images and beyond," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 2, pp. 342–355, Feb. 2017.
- [53] H. Zhang, D. Wipf, and Y. Zhang, "Multi-image blind deblurring using a coupled adaptive sparse prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 1051–1058.
- [54] K. He, J. Sun, and X. Tang, "Guided image filtering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 6, pp. 1397–1409, Jun. 2013.
- [55] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 586–595.
- [56] J. R. Isidoro and P. Milanfar, "Pull-push non-local means with guided and burst filtering capabilities," in *Denosing Photographic Images Video*. Springer, 2018, pp. 267–294.
- [57] E. S. Gastal and M. M. Oliveira, "Domain transform for edge-aware image and video processing," in *Proc. ACM SIGGRAPH*, 2011, pp. 1–12.

- [58] W. Yang, X. Zhang, Y. Tian, W. Wang, J.-H. Xue, and Q. Liao, "Deep learning for single image super-resolution: A brief review," *IEEE Trans. Multimedia*, vol. 21, no. 12, pp. 3106–3121, Dec. 2019.
- [59] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, "Enhanced deep residual networks for single image super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. workshops*, 2017, pp. 136–144.
- [60] S. Bell-Kligler, A. Shocher, and M. Irani, "Blind super-resolution kernel estimation using an internal-gan," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 284–293.
- [61] S. A. Hussein, T. Tirer, and R. Giryes, "Correction filter for single image super-resolution: Robustifying off-the-shelf deep super-resolvers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 1428–1437.



Damien Kelly received the B.A./B.A.I. degree in computer & electronic engineering from Trinity College Dublin, Dublin, Ireland, in 2005 and the Ph.D. degree in 2010. Since then, he has been a Research Fellow with Media Processing Group, Trinity College Dublin and with Green Parrot Pictures Ltd. Developing software tools for video enhancement. In 2011, he joined Google and has worked in the Chrome Media, and YouTube Video Infrastructure teams on VR audio, audio or video transcoding, and quality enhancement. In 2018, he joined the Computational Imaging Team with Google Research and is currently working on image or video super-resolution and enhancement.



Mauricio Delbracio (Member, IEEE) received the B.Sc. degree in electrical engineering from Universidad de la República, Montevideo, Uruguay, in 2006, and the M.Sc. and Ph.D. degrees in applied mathematics from École Normale Supérieure de Cachan, France, in 2009 and 2013, respectively. He is currently a Research Scientist with Google Research. Before joining Google in 2019, he was an Assistant Professor with the Department of Electrical Engineering, Universidad de la República. From 2013 to 2016, he was a Postdoctoral Researcher with the Department

of Electrical and Computer Engineering, Duke University, Durham, NC, USA. His current research interests include algorithms, data analysis and applications of machine learning to image and signal processing. In 2016, he was awarded the Early Career Prize from the Society for Industrial and Applied Mathematics (SIAM) Activity Group on Imaging Science in 2016 for his important contributions to image processing.



Ignacio Garcia-Dorado received the M.S. degree in electrical engineering from Universidad Politécnica de Madrid, Madrid, Spain, the M.S. degree in computer engineering from Lund University, Lund, Sweden, and the M.S. and Ph.D. degrees in computer science from Purdue University, West Lafayette, IN, USA, in 2008, 2014, and 2015, respectively. From 2008 to 2010, he was a Computer Engineer with ESA, Noordwijk, The Netherlands. From January to May 2010, he was a Research Assistant with McGill University, Montreal, QC, Canada. During this time,

he was a Research Intern with NVidia during the Summer of 2013 and as a Research Assistant with the University of California, Berkeley, Berkeley, CA, USA, during the Summer of 2014. After his Ph.D. defense in October 2015, he moved to Mountain View to work with Computational Photography Team with Google Research. After this, he was awarded the Fulbright Scholarship to initiate Ph.D. studies.



Sungjoon Choi received the B.S. degree in computer science and engineering from Seoul National University, Seoul, South Korea, in 2007, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 2012 and 2015, respectively. From 2008 to 2010, he was with Microsoft, Redmond, WA, USA. In 2015, he joined Google, to work on spatial perception in computer vision where his current research interests include computational photography, image processing, and machine learning. He was the recipient of the Samsung Scholarship

for graduate studies. He was also a finalist of the ACM-ICPC World Finals in 2002 and 2006.



Peyman Milanfar (Fellow, IEEE) received the undergraduate education in electrical engineering and mathematics from the University of California, Berkeley, Berkeley, CA, USA, and the M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA. He is currently a Principal Scientist Director at Google Research, where he leads the Computational Imaging Team. Prior to this, he was a Professor of electrical engineering with the University of California, Santa Cruz, from 1999 to 2014. From 2012 to

2014, he was an Associate Dean of research with the School of Engineering. From 2012 to 2014, he was on leave with Google-x, where he helped develop the imaging pipeline for Google Glass. Most recently, his team at Google developed the digital zoom pipeline for the Pixel phones, which includes the multi-frame super-resolution (Super Res Zoom) pipeline, and the RAISR upscaling algorithm. In addition, the Night Sight mode on Pixel 3 uses Super Res Zoom technology to merge images (whether you zoom or not) for vivid shots in low light, including astrophotography. He holds 15 patents, several of which are commercially licensed. He founded MotionDSP, which was acquired by Cubic Inc. (NYSE:CUB). He is a Keynote Speaker at numerous technical conferences, including the Picture Coding Symposium (PCS), SIAM Imaging Sciences, SPIE, and the International Conference on Multimedia (ICME). Along with his students, he has won several Best Paper Awards from the IEEE Signal Processing Society. He is a Distinguished Lecturer of the IEEE Signal Processing Society.