



Published in Image Processing On Line on 2022-07-14.  
Submitted on 2021-10-15, accepted on 2022-05-12.  
ISSN 2105-1232 © 2022 IPOL & the authors CC-BY-NC-SA  
This article is available online with supplementary materials,  
software, datasets and online demo at  
<https://doi.org/10.5201/ipol.2022.385>

# Spectral Pre-Adaptation for Restoring Real-World Blurred Images using Standard Deconvolution Methods

Chaoqun Dong<sup>1,3</sup>, Filip Sroubek<sup>2</sup>, Javier Portilla<sup>3</sup>

<sup>1</sup> AII, Fraunhofer Singapore, Singapore ([chaoqun.dong@fraunhofer.sg](mailto:chaoqun.dong@fraunhofer.sg))

<sup>2</sup> UTIA, Czech Academy of Sciences ([sroubekf@utia.cas.cz](mailto:sroubekf@utia.cas.cz))

<sup>3</sup> Instituto de Óptica, Consejo Superior de Investigaciones Científicas, Spain ([javier.portilla@csic.es](mailto:javier.portilla@csic.es))

*Communicated by* Jérémy Anger      *Demo edited by* Jérémy Anger

## Abstract

Classical blur models are based on simplifying assumptions, namely shift-equivariance and circular boundary condition (CBC), that rarely hold in practice. Shift-equivariance means that a shift of the input produces the same shift of the output, which implies that blur is spatially invariant and image aliasing is not present. The CBC assumes that the image is rectangular and periodically repeating. Discrepancies between simplified models and real blurred observations cause strong artifacts in image restoration. The common remedy is to increase the model complexity and remove simplifying assumptions. However, this also brings extra computational complexity to the restoration task. We present spectral pre-adaptation (SPA) that pre-processes blurred images so they can be restored using fast standard deconvolution algorithms suitable for simplified models. The SPA serves as a connector between classical deconvolution methods and a variety of real observations involving blur. Experiments on simulated and real images show that standard deconvolution of SPA-interpolated images not only greatly reduces artifacts compared to direct deconvolution, but performs on a par with more complex restoration methods.

## Source Code

The reviewed source code and documentation for this algorithm are available from [the web page of this article](#)<sup>1</sup>.

**Keywords:** image restoration; non-circulant deconvolution; maximum likelihood interpolation; spectral models; missing samples; model discrepancies

<sup>1</sup><https://doi.org/10.5201/ipol.2022.385>

# 1 Introduction

## 1.1 Simple Blur Models in Restoration

A standard procedure for evaluating the performance of image restoration algorithms is to use a set of test images, plus a set of blur kernels and noise levels, and then apply simple convolution and add noise to simulate observations. The very same degradation model is used by the restoration algorithms. As this approach provides a ground truth, it allows for an objective performance comparison. Problems arise when we want to apply the algorithms to real captured blurred images. Because the simulated blur perfectly fits the blur assumed by the restoration algorithms, the referred performance tests provide no information about the robustness to different types and amounts of discrepancy between the model explaining a real blurred image and the observation model used for restoration. Such model discrepancies typically cause strong artifacts when restoring real images, considerably dropping the performance with respect to results of simulation tests.

Classical blur models used for image restoration assume a blur kernel that is constant across the image, and a sufficient sampling rate avoiding aliasing. These two conditions, jointly with that of spatially homogeneous noise statistics, translate into shift-invariant observation models. In many real situations, these assumptions are not acceptable approximations, and more sophisticated models and corresponding algorithms have been developed to tackle such scenarios. In particular, there is a significant amount of literature on shift-variant blur and its restoration; see e.g. [14, 2]. In contrast, the problem of restoring blurred aliased images, due in part to its greater difficulty, has been addressed rarely; see e.g. [6, 25]).

Another strong assumption, almost universally made by classical blur simulations, concerns the shape and boundary conditions of the blurred image support: Images are assumed to be discrete samples on a rectangular grid with CBC (also named periodic, providing continuity on a torus) for the blur. Whereas using rectangular arrays to represent images is simple and convenient for digital processing, the CBC seems, at first, a strange assumption<sup>2</sup>. The justification is purely numerical. The CBC allows us to calculate any linear shift-invariant transformation of a discrete signal in an efficient way, as the matrix representing this transformation is diagonalized by the discrete Fourier Transform (DFT). Even if modern, non-linear, restoration methods go far beyond a regularized kernel inversion in the Fourier domain, virtually all of them perform such operations as part of more complex, iterative, algorithms.

Summarizing, we write the *reference observation model* assumed by standard deconvolution algorithms as a discrete circular convolution with additive noise on a rectangular grid

$$z = x \circledast h + \omega. \quad (1)$$

## 1.2 Using a Spectral Model for Pre-processing the Observation

To address the above model discrepancy there are two possible strategies. The conventional one is to bring models closer to reality, i.e., to refine blur models used by the restoration algorithms, thus introducing extra complexity and degrees of freedom to provide a better fit for real observations. Improving the model is interesting from a scientific point of view, however its practical application has drawbacks, such as added computational cost and lower robustness due to the requirement of estimating more model parameters.

A less conventional venue to bridge the misfitting gap consists in *bringing reality closer to models*. This means transforming (pre-processing) the observation to enforce a good fit to the reference blur

---

<sup>2</sup>Note that the CBC is not just unrealistic but directly inconsistent in almost every real scenario. Two exceptions are full visual field (360 degrees) and periodic images (e.g. repeated tiles).

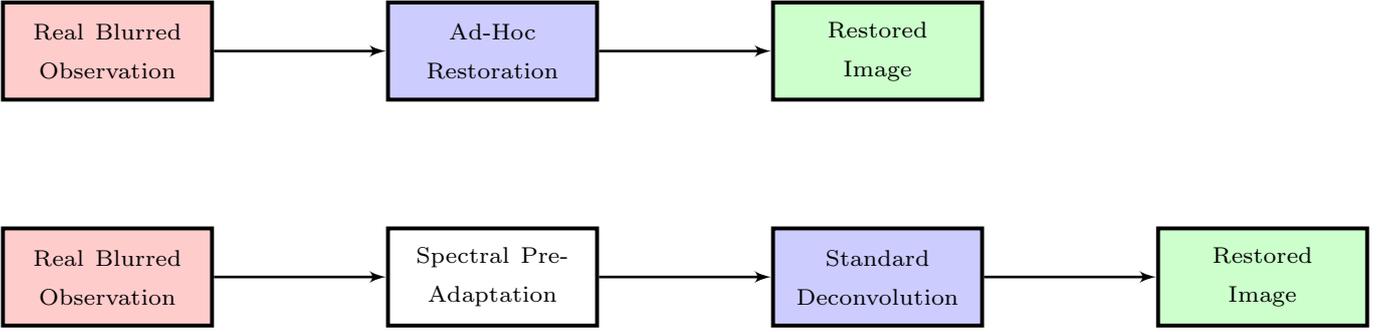


Figure 1: Two alternative approaches for restoring real blurred observations. Up: adding complexity to the observation model, which translates into an ad-hoc restoration algorithm. Bottom: include a versatile preprocessing (like SPA), and then using a standard deconvolution method, which assumes a simple blur observation model.

model in Equation (1); see Figure 1. Pre-processing the observation is a common practice when dealing with the boundary condition problem, e.g. by applying a smoothly decaying window [19] or edge extension strategies [11]. It is not usually applied to mitigate the effect of other model discrepancies, such as non-rectangular image support, aliasing, or slightly shift-variant blur kernels. Here we present a spectral pre-processing that partially compensates for such model violations. This makes it a powerful and versatile tool for deblurring real images using off-the-shelf, standard deconvolution algorithms.

Our pre-processing method, termed spectral pre-adaptation (SPA), requires that pixels violating the reference blur model are marked, and additional pixels are added to deal with the CBC and/or aliasing. Marked and added pixels are treated as unknowns that must be estimated. To estimate these pixels, we first obtain a power spectral density (PSD) model of the blurred image, which we compute from (i) a PSD model for uncorrupted images, (ii) blur kernel, and (iii) noise level (assuming noise is additive and decorrelated from the blurred image). Then we derive a linear, shift-variant, optimal solution, the most likely according to a multivariate Gaussian model given the image and degradation parameters. The solution is a rectangular, uniformly CBC-blurred image, consistent with the reference observation model of Equation (1) and thus a candidate for being successfully restored with a standard deconvolution algorithm.

As mentioned before, other methods exist that tackle the image restoration of real-world blurred images, which depart from the simple reference observation model described above. Among all methods recently proposed for dealing with model discrepancies in image restoration, we believe that two of them, one proposed by Almeida-Figueiredo [1] and one by Kotera et al. [10], deserve special attention, because of their applicability and good performance. The first one uses the idea from Reeves [17] to estimate a boundary extension of the restoration solution with the property of maximizing a smoothness criterion across boundaries under CBC. Almeida and Figueiredo extended the concept of using a mask for extending the image along the boundaries to using a more generic mask to estimate all pixels not obeying the CBC homogeneous blur model, because of occlusions, lost pixels and arbitrary image supports. They adapted the acclaimed ADMM restoration method to deal with unconstrained boundary conditions (UBC-ADMM). Kotera et al. went further by determining areas with model discrepancies automatically. They formulate the deconvolution problem as a Variational Bayesian inference similarly to [20] and decompose noise into Gaussian and non-Gaussian. The non-Gaussian part matches the areas where the convolutional model is violated. That method is showed to effectively deal not only with boundary pixels and occlusions, but also with clipping, such as dead under-exposed pixels, or over-exposed pixels (e.g. in highlights), both breaking the model’s linear response.

The two generic approaches in Figure 1 that deal with model discrepancies have, each, pros and cons. On the one hand, using a single-step restoration instead of two sequential steps, generally

achieves higher performance. On the other hand, encapsulating and compensating for model discrepancies *before* restoring the image is appealing, for it leaves open the choice of using any deconvolution method without requiring ad-hoc adaptations. In addition, the SPA computational cost, SPA being a linear estimate, is relatively modest when compared to non-linear restoration costs and even more when compared to non-standard blur methods. Therefore, the question of which generic option is better depends on the particular task, both in terms of performance and computational cost. Results presented here indicate that, at least in some cases, the SPA approach is superior in both terms and in other cases the performance is only moderately worse.

This paper unifies previous contents presented in a series of conference papers [15, 4, 5]. Besides providing a deeper, more comprehensive and thorough description of the involved models and methods (including open-source implementation), we have added an application of SPA to the case when there is a discrepancy between the assumed blur kernel and the real one. We have studied the regularizing effect of SPA pre-processing in two kernel mismatch scenarios (shift variant and inaccurate shift invariant blur), in the presence of aliasing. In addition, we have improved a previous method for computing the solution, which involves solving a large linear system of equations, by using a more efficient implementation based on the conjugate gradient method.

This paper is organized as follows. Section 2 describes our SPA method in detail including an explanation of the theory and description of the implementation. In Section 3, we demonstrate the usefulness of SPA via several applications. Visual and numerical results are shown and compared by experiments on both simulated and real images. Finally, in Section 4 we draw conclusions.

## 2 Models and Methods

We assume that a non-rectangular blurred observation  $y$  follows a model

$$y = (x * h + \omega)_{\downarrow Q}, \quad (2)$$

where  $x$  is the latent high-resolution image defined on an arbitrary non-rectangular region,  $*$  denotes valid convolution,  $h$  is the blur kernel inside the region,  $\omega$  is noise, and  $\downarrow Q$  denotes an optional downsampling factor along each direction, typically causing aliasing when  $Q > 1$  (see Section 3.2). We assume we have a good estimate of  $h$  in the original resolution of  $x$ , i.e. before downsampling. In real cases, however, we may not know exactly the kernel, or it may vary over the image. Here we consider a shift-invariant kernel model, but explore robustness of SPA pre-processing for non-exact and/or mildly shift-variant kernels (see Section 3.3). Since the blurred image support is non-rectangular and the complete estimated image is rectangular, we introduce a rectangular mask large enough to include  $y$ , which marks each pixel as either belonging to  $y$  or as unknown. The mask is assumed to be given. First, “unknown” pixels are those that do not follow the modeled blur. In addition, if the downsampling factor is present, we need to upsample the dropped samples and mark the new samples also as unknown. Finally, we also want to extend the rectangular support around  $y$  with stripes of several pixels wide along the edges, and mark them also as unknown. This is necessary to create a transition area for allowing the blur to obey the CBC; see Figure 3.

The SPA’s goal is to calculate all pixels marked as unknown, such that the resulting image, which we will denote as  $z$  and which includes  $y$ , is the most likely in the sense of a Gaussian image model following the reference observation model in Equation (1).

## 2.1 Observation and Image Models

The reference observation model in Equation (1) becomes the formation model for our interpolated blurred image  $z$ . It can be written in a matrix-vector form as

$$\mathbf{z} = \mathbf{H}_c \mathbf{x} + \mathbf{w}, \quad (3)$$

where  $\mathbf{H}_c$  is a square block-circulant matrix performing CBC convolution with the kernel  $h$ , and  $\mathbf{x}$ ,  $\mathbf{w}$  are the vectorized ideal image and noise, respectively. To complete our observation model, we implement the mask using a selection matrix  $\mathbf{S}$  ( $M \times N$ , with  $M < N$ ), which expresses the vectorized observation  $y$  as a subset of pixels of the extended image  $z$

$$\mathbf{y} = \mathbf{S}\mathbf{z}. \quad (4)$$

For defining a likelihood measurement, we have posed a standard stationary Gaussian model and additive white noise<sup>3</sup> with variance  $\sigma_w^2$

$$p(\mathbf{z}) = \frac{1}{(2\pi |\mathbf{C}_z|)^{1/2}} \exp\left(-\frac{1}{2} \mathbf{z}^T \mathbf{C}_z^{-1} \mathbf{z}\right), \quad (5)$$

where  $\mathbf{C}_z$  is the covariance matrix of  $\mathbf{z}$ , which is fixed for the observed image. The proposed interpolated solution is

$$\begin{aligned} \hat{\mathbf{z}} &= \arg \max_{\{\mathbf{z}, \mathbf{S}\mathbf{z}=\mathbf{y}\}} p(\mathbf{z}) \\ &= \arg \min_{\{\mathbf{z}, \mathbf{S}\mathbf{z}=\mathbf{y}\}} \mathbf{z}^T \mathbf{C}_z^{-1} \mathbf{z}. \end{aligned} \quad (6)$$

From Equation (3) follows that

$$\mathbf{C}_z = \mathbf{H}_c \mathbf{C}_x \mathbf{H}_c^T + \mathbf{C}_w,$$

where  $\mathbf{C}_x$  and  $\mathbf{C}_w$  are the covariance matrices of the uncorrupted image and noise, respectively. Here we assume a spatially homogeneous behaviour of signal and noise, and thus  $\mathbf{C}_x$  (and consequently also  $\mathbf{C}_z$ ) is a block-circulant Toeplitz matrix, that is diagonalizable by the DFT similarly as  $\mathbf{H}_c$ . If we assume white noise then  $\mathbf{C}_w = \sigma_w^2 \mathbf{I}$  and we can write

$$\begin{aligned} \mathbf{C}_z &= (\mathbf{F}\mathbf{D}_{H_c}\mathbf{F}^*)(\mathbf{F}\mathbf{D}_{P_x}\mathbf{F}^*)(\mathbf{F}\mathbf{D}_{H_c}^*\mathbf{F}) + \sigma_w^2 \mathbf{F}\mathbf{F}^* \\ &= \mathbf{F}(\mathbf{D}_{|H_c|^2} \mathbf{D}_{P_x} + \sigma_w^2 \mathbf{I})\mathbf{F}^* \\ &= \mathbf{F}\mathbf{D}_{P_z}\mathbf{F}^*, \end{aligned} \quad (7)$$

where  $\mathbf{F}^*$  and  $\mathbf{F}$  are the orthogonal matrices implementing the direct and inverse DFT, respectively, and  $\mathbf{D}$ 's are diagonal matrices with  $\mathbf{D}_{P_z} = \mathbf{D}_{|H_c|^2} \mathbf{D}_{P_x} + \sigma_w^2 \mathbf{I}$ . Using Equation (6) and inverting  $\mathbf{C}_z$  in Equation (7), the minus log-likelihood we want to minimize can be written in the Fourier domain as

$$-\log(p(\mathbf{z})) \propto \sum_{u,v} \frac{|Z(u,v)|^2}{P_Z(u,v)} + \text{const.}, \quad (8)$$

where  $Z(u,v)$  represents the DFT of  $z(n,m)$  and  $P_Z(u,v)$  is the corresponding Power Spectral Density (PSD), which appears (lexicographically reordered) in the coefficients of the diagonal matrix  $\mathbf{D}_{P_z}$

$$P_Z(u,v) = |H_c(u,v)|^2 P_X(u,v) + \sigma_w^2. \quad (9)$$

<sup>3</sup>For notation simplicity, and without loss of generality, we have dropped the mean in the following expressions, though it has been considered in the actual implementation, as shown in Algorithm 1.

Maximizing the likelihood of the interpolated image  $\mathbf{z}$  returns a likely observation of an image filtered by the blur kernel using a rectangular support under CBC conditions. As such, we ensure that the interpolated image will not overcome the noise level in the vicinity of the kernel zeros, thus avoiding strong artifacts in the restoration phase.

Note that our cost function is a more powerful and complete criterion to follow than mere smoothness as in [17, 11]. Besides forcing consistency to the blurring kernel spectrum, by including the  $P_X$  term, the model also provides a weight for each frequency corresponding to its relative abundance according to a natural image PSD model. In particular, for modeling  $P_X$ , we have chosen a separable stationary continuous Gauss-Markov process model [13], whose autocovariance is  $R(d_x, d_y) = \sigma_x^2 \rho^{|d_x|+|d_y|}$ , from which we obtain its PSD

$$P_X(u, v) = \frac{4\sigma_x^2 \log^2(\rho)}{(\log^2(\rho) + 4\pi^2 u^2)(\log^2(\rho) + 4\pi^2 v^2)}. \quad (10)$$

We have hand-optimized the parameters for optimizing the results in a small set of typical images, obtaining  $\sigma_x = 30$  and  $\rho = 0.65$ , the first adjusted to the image intensity range  $[0, 255]$ <sup>4</sup>. We have observed, nevertheless, a wide tolerance for the behavior of our algorithm with respect to the choice of these parameters. We have also used this model for doing Wiener restoration in the experiments.

## 2.2 Problem Formulation and Solution

Expressing Equation (8) back in a vector-matrix notation and including the consistency constraint on  $z$  from Equation (6) yields

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \|\mathbf{D}_{P_Z}^{-1/2} \mathbf{F}^* \mathbf{z}\|^2, \text{ s.t. } \mathbf{S} \mathbf{z} = \mathbf{y}. \quad (11)$$

To reformulate Equation (11) to an unconstrained optimization, we split the interpolated image  $\mathbf{z}$  into  $\mathbf{z}_i$  and  $\mathbf{z}_o$  by  $\mathbf{z}_P = \mathbf{P} \mathbf{z} = [\mathbf{z}_i \ \mathbf{z}_o]$ , where  $\mathbf{P}$  is a permutation matrix.  $\mathbf{z}_i$  and  $\mathbf{z}_o$  are two non-overlapping vectors representing the  $(N - M)$  interpolated pixels, and the  $M$  observed pixels, respectively. Similarly, the Fourier Transform matrix  $\mathbf{F}^*$  can also be split such that  $\mathbf{F}_P^* = \mathbf{F}^* \mathbf{P}^{-1} = [\mathbf{F}_i^* \ \mathbf{F}_o^*]$ , which performs a 2D DFT, but enforcing an explicit separation of observed and interpolated pixels. The interpolation constraint also implies  $\mathbf{z}_o = \mathbf{y}$ , and then we can rewrite the cost in Equation (11) as

$$\begin{aligned} \|\mathbf{D}_{P_Z}^{-1/2} \mathbf{F}^* \mathbf{z}\|^2 &= \|\mathbf{D}_{P_Z}^{-1/2} \mathbf{F}_P^* \mathbf{z}_P\|^2 \\ &= \|\mathbf{D}_{P_Z}^{-1/2} (\mathbf{F}_i^* \mathbf{z}_i + \mathbf{F}_o^* \mathbf{y})\|^2, \end{aligned} \quad (12)$$

which attains a minimum at

$$\hat{\mathbf{z}}_i = -(\mathbf{F}_i^* \mathbf{D}_{P_Z}^{-1} \mathbf{F}_i^*)^{-1} \mathbf{F}_i^* \mathbf{D}_{P_Z}^{-1} \mathbf{F}_o^* \mathbf{y}. \quad (13)$$

The previous solution is difficult to implement directly due to the involved permutations. For a practical implementation, we introduce the complementary  $N \times (N - M)$  selection matrix, which is the complement of  $\mathbf{S}$  and selects from  $z$  the non-observed pixels, i.e.  $\mathbf{z}_i = \mathbf{E} \mathbf{z}$ . The transposed matrices  $\mathbf{E}^T$  and  $\mathbf{S}^T$  perform expansions from the vectors corresponding to the selected pixels  $\mathbf{z}_i$  and  $\mathbf{z}_o = \mathbf{y}$  to the full domain of  $\mathbf{z}$  by filling the missing pixels with zeros. It is immediate to realize that  $\mathbf{F}_i^* \mathbf{z}_i = \mathbf{F}^* \mathbf{E}^T \mathbf{z}_i$  and  $\mathbf{F}_o^* \mathbf{y} = \mathbf{F}^* \mathbf{S}^T \mathbf{y}$ , which can be easily computed using fast operations: standard DFT and expansion with zeros. This gives us the solution for the interpolated pixels

$$\hat{\mathbf{z}}_i = -(\mathbf{E} \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{E}^T)^{-1} \mathbf{E} \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{S}^T \mathbf{y}. \quad (14)$$

<sup>4</sup>We can trivially adapt to other ranges by making  $\sigma_x \propto \max(x)$ .

And the final solution for the interpolated image  $z$  is

$$\hat{\mathbf{z}} = \mathbf{S}^T \mathbf{y} + \mathbf{E}^T \hat{\mathbf{z}}_i. \quad (15)$$

The solution in Equation (14) requires an inversion of a large shift-variant matrix. However in practice, we do not need to invert the matrix and instead apply the conjugate gradient method for solving the involved system of linear equations

$$\mathbf{A} \hat{\mathbf{z}}_i = \mathbf{b}, \quad (16)$$

where

$$\mathbf{A} = \mathbf{E} \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{E}^T \quad (17)$$

and

$$\mathbf{b} = -\mathbf{E} \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{S}^T \mathbf{y}. \quad (18)$$

Note that computing  $\mathbf{b}$  and applying  $\mathbf{A}$  involves only fast operations, namely: direct ( $\mathbf{F}^*$ ) and inverse ( $\mathbf{F}$ ) discrete Fourier transforms, point-wise division in the Fourier domain ( $\mathbf{D}_{P_Z}^{-1}$ ), masking ( $\mathbf{E}$  and  $\mathbf{S}$ ) and expansion with zeros ( $\mathbf{E}^T$  and  $\mathbf{S}^T$ ) in the spatial domain. The complete procedure is summarized in Algorithm 1.

## 2.3 Using SPA for Image Restoration

To use SPA pre-processing for image restoration, three steps are required: (i) prepare SPA input items (observed data, mask and parameters); (ii) perform SPA, and (iii) deconvolve the SPA output. Additionally, one may want to recombine a new image combining deblurred pixels with non-blurred pixels in the observation. Details are given in Algorithm 2. Figure 2 shows the flow chart of the algorithm.

### 2.3.1 Input Data and Parameters for SPA

First, for calculating the PSD of the ideal image  $\mathbf{x}$ , the variance  $\sigma_x^2$  is required. As a rule of thumb, we have experimentally validated that for typical images in an 8-bit range (from 0 to 255)  $\sigma_x^2 = 30$  is a suitable value. Second, as we mentioned before, we assume that the blurring kernel is known or has been previously estimated. If the kernel is unknown, one needs to apply assumptions, blur estimation methods or a combination of both with trial and error (e.g. if we know that the blur kernel comes from defocus we may approximate it by a disk, and just estimate its radius). In addition, the noise variance  $\sigma_w^2$  has to be estimated. This is usually easy using standard methods, such as median absolute value deviation (MAD) [8]. An essential component of the input is the mask defining which pixels are observed and which ones are to be interpolated. We discuss this in Section 2.3.3. One also needs to decide how much to extend the observed image. We discuss this in Section 2.3.2. Finally, the conjugate gradient used for solving the optimization problem of SPA requires another parameter, either the number of iterations (the one we used) or a relative tolerance criterion for the convergence.

### 2.3.2 Extending the Original Support

Here we analyze the problem of extending the support of the raw rectangular image  $y_{rec}$ , which includes the observed pixels  $y$  that we want to deblur. Besides extending the generally non-rectangular support of  $y$  to the rectangular support of  $y_{rec}$ , it is also necessary to extend the raw observed image  $y_{rec}$  with four stripes, one along each edge, to avoid the discontinuity across opposed boundaries (up-bottom, left-right) under our reference observation model with CBC. As we mentioned before,

**Algorithm 1:** Spectral pre-adaptation (SPA) pre-processing for image restoration

**input** : initial guess  $\mathbf{z}_{init}$ , kernel  $\mathbf{h}$ , *mask*, PSD of ideal image  $P_X(\sigma_x^2)$ , noise variance  $\sigma_\omega^2$   
**output**: maximum likelihood estimate  $\hat{\mathbf{z}}$

- step 1:** calculate  $\mathbf{b}$  // Equation (18)
- 1.1 calculate the mean of the observed pixels  $\mathbf{y}$ :  $\mathbf{sum}(\mathbf{z}_{init} \times \mathit{mask}) / \mathbf{sum}(\mathit{mask}) \rightarrow \mu$   
 $(\mathbf{y} = \mathbf{z}_{init}(\mathit{mask}))$
  - 1.2 subtract the mean from the observed pixels:  $(\mathbf{z}_{init} - \mu) \times \mathit{mask} \rightarrow \mathbf{y}_{ext}$
  - 1.3 Fourier Transform of blur kernel:  $FFT(h) \rightarrow \mathbf{H}$
  - 1.4 compute reference PSD for  $z$  using  $\mathbf{H}$ ,  $P_X(\sigma_x^2)$  and  $\sigma_\omega^2 \rightarrow \mathbf{D}_{P_Z}$  // Equation (7)
  - 1.5 Fourier Transform of the observed pixels:  $FFT(\mathbf{y}_{ext}) \rightarrow \mathbf{F}^* \mathbf{S}^T \mathbf{y}$
  - 1.6 Inverse Fourier Transform of  $(FFT(\mathbf{y}_{ext}) \times \mathbf{D}_{P_Z}^{-1}) \rightarrow q = \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{S}^T \mathbf{y}$
  - 1.7 select interpolated pixels as  $-q(\mathit{NOT} \mathit{mask}) \rightarrow \mathbf{b} = -\mathbf{E} \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{S}^T \mathbf{y}$
- step 2:** find the solution  $\hat{\mathbf{z}}_i$  using the conjugate gradient (CG) method // Equation (16)
- 2.1 set  $\mathbf{z}_i$  and  $\mathbf{z}$  according to the initial guess:  $\mathbf{z} = \mathbf{z}_{init} - \mu$ ,  $\mathbf{z}_i = \mathbf{z}(\mathit{NOT} \mathit{mask})$
  - 2.2 iteratively update  $\mathbf{z}_i$  until convergence, in each CG step, multiplication with  $\mathbf{A}$  is performed:
    - 2.2.1 Fourier Transf. of the extended interpolated pixels:  $FFT((\mathit{NOT} \mathit{mask}) \times \mathbf{z}) \rightarrow \mathbf{F}^* \mathbf{E}^T \mathbf{z}_i$
    - 2.2.2 Inverse Fourier Transf. of  $(FFT((\mathit{NOT} \mathit{mask}) \times \mathbf{z}) \times \mathbf{D}_{P_Z}^{-1}) \rightarrow r = \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{E}^T \mathbf{z}_i$
    - 2.2.3 select interpolated pixels as  $r(\mathit{NOT} \mathit{mask}) \rightarrow \mathbf{A} \mathbf{z}_i = \mathbf{E} \mathbf{F} \mathbf{D}_{P_Z}^{-1} \mathbf{F}^* \mathbf{E}^T \mathbf{z}_i$
- step 3:** construct the final solution // Equation (15)
- 3.1 combine observed and interpolated pixels:  $\mathbf{y}_{ext} + \mathit{expand}(\hat{\mathbf{z}}_i) + \mu \rightarrow \hat{\mathbf{z}}$

other authors (starting with [17]) have proposed to optimize a simple smoothness criterion for enforcing the so-extended image to be smooth under CBC. Note that, when dealing with blurred images we want to restore, it is crucial to keep smoothness across boundaries. However, *smoothness* (a particular, fixed, measurement of it), by itself, is arguably not the best criterion to optimize. As we have described, SPA not only provides *smooth* results under CBC, but also results that are likely in terms of the spectral behavior of typical images affected by the involved blur kernel. We have set the minimal width of each extended stripe as half the blur kernel size (rounded down), which is the minimal extension allowing to compute a non-CBC convolution with valid pixels. However, enlarging the estimated area even more may bring other benefits. For example, an interesting criterion is choosing the extension that minimizes the entropy rate under a Gaussian homogeneous model [9] of the extended image  $z$  according to the spectral model of Equation (9). This implies minimizing the average of the logarithm of the PSD. Here a key role is played by the spectrum of the blur kernel  $|H_c(u, v)|^2$ , whose low values produce a significant entropy decrease, especially when placed at spectral locations where  $P_X$  is large. The lowest possible values of  $|H_c|^2$  are zeros, and they can

---

**Algorithm 2:** Image Restoration using SPA
 

---

**input** : (partially) blurred observation on a rectangular support  $\mathbf{y}_{rec}$

**output**: restored image  $\hat{\mathbf{x}}$

**step 1:** prepare input parameters for SPA

- calculate PSD of the ideal image  $\rightarrow P_X(\sigma_x^2)$  // Equation (10)
- estimate/set the degradation parameters  $\rightarrow \mathbf{h}, \sigma_\omega^2$
- set boundary extension widths  $\rightarrow L_x^e, L_y^e$  // Section 2.3.2
- identify the mask  $\rightarrow mask$
- calculate a SPA initial guess:  $extrapol(\mathbf{y}, \mathbf{S}) \rightarrow \mathbf{z}_{init}$  // Section 2.3.4

**step 2:** estimate the extended blur image with SPA  $\rightarrow \hat{\mathbf{z}}$

**step 3:** deconvolution and recomposition

- deconvolution:  $\hat{\mathbf{z}} \rightarrow \hat{\mathbf{x}}_{rec}$
  - recomposition with the raw rectangular observation  $\mathbf{y}_{rec}$ :  $(\mathbf{y}_{rec}, \hat{\mathbf{x}}_{rec}, mask) \rightarrow \hat{\mathbf{x}}$
- 

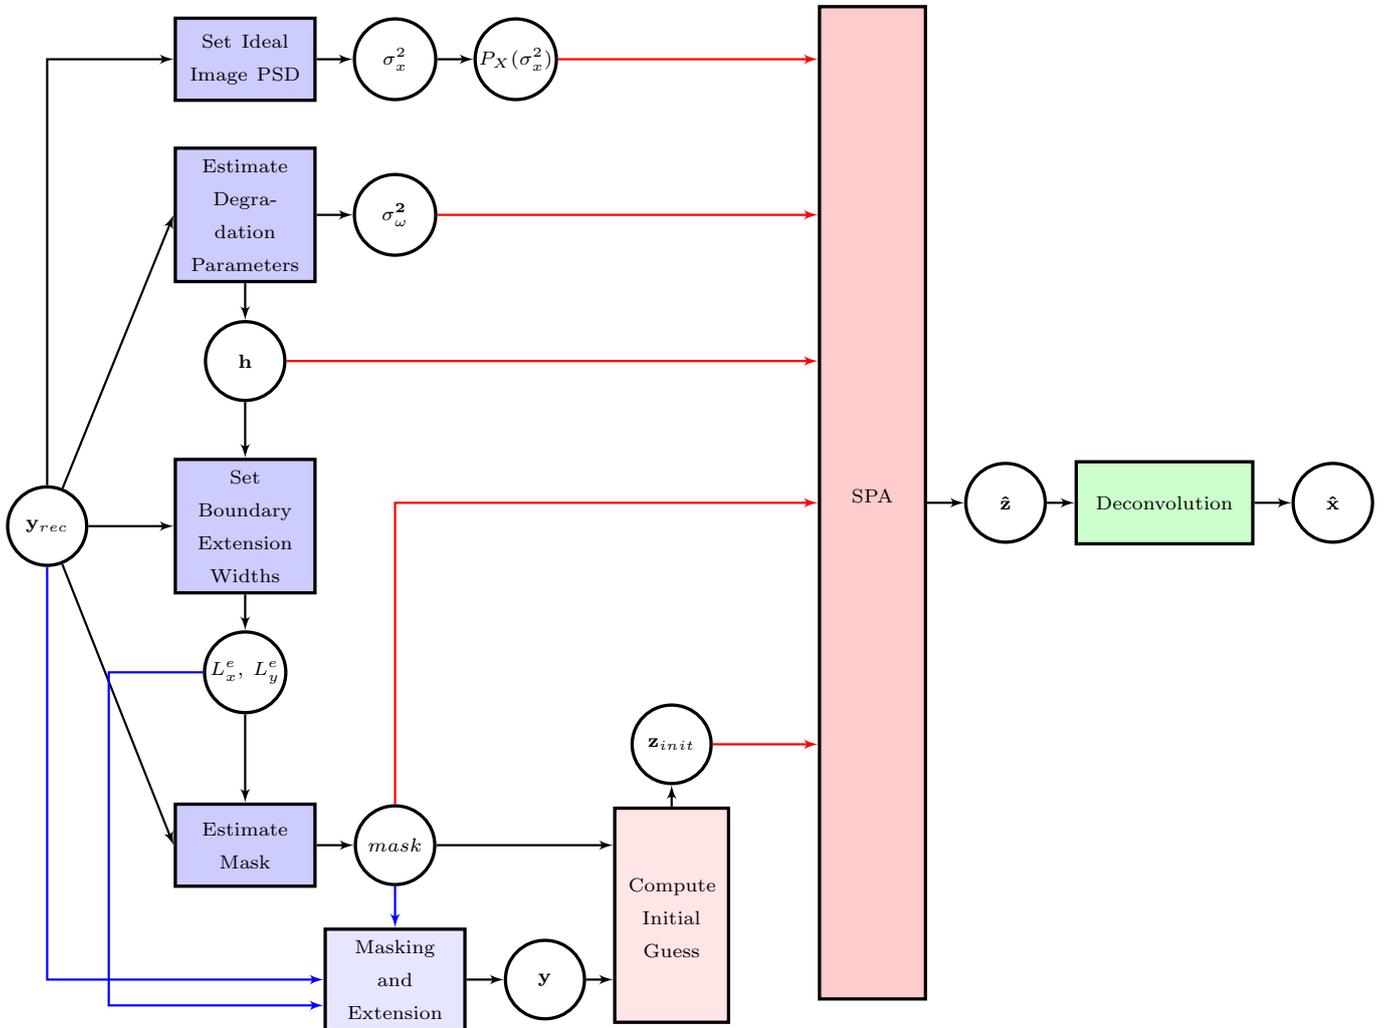


Figure 2: Flow chart of restoration algorithm with SPA.

be attained if the image size is chosen such that they fall exactly on sampled discrete frequencies.

For the case of  $xy$ -separable blur kernels, an independent adjustment of  $L_x$  and  $L_y$  is possible.

In particular, for the case of uniform blur kernels, or any convolved version of them, we can easily compute an extension that forces the zeros of the blur kernel to exactly fall on discrete samples in the Fourier domain. This can be done by extending the original image such that its new size becomes a multiple of the kernel size. Under this criterion, the width of each stripe is calculated by

$$L_x^e = (\lfloor N_x/n_x + k \rfloor n_x - N_x)/2, \quad (19)$$

where  $N_x$  and  $n_x$  are the dimensions of the original image and the kernel in the horizontal direction (same in vertical),  $\lfloor \cdot \rfloor$  is the floor rounding operator, and  $k$  is a non-negative integer, usually the smallest providing an integer solution to  $L_x^e$ . This makes the extended final image have dimensions  $(N_x + 2L_x^e, N_y + 2L_y^e)$ .

### 2.3.3 Setting up the Selection Mask

A crucial input for SPA is the selection mask  $s(n, m)$  (from which the matrix  $\mathbf{S}$  is derived), which divides pixels into two classes:  $s = 1$  for pixels in the observed uniformly blurred area compatible with the reference blur model of Equation (1), and  $s = 0$  for unknown pixels, which need to be substituted by their SPA interpolated values. Pixels in the extended areas along the rectangular boundaries (explained in Section 2.3.2) are labeled as  $s = 0$  since they are not observed and must be interpolated. For the other pixels, it is sometimes difficult to discern the boundaries of a uniformly blurred object that we want to restore. If the estimation of the mask is not straightforward, image matting methods such as [22] can be applied. As a general advice, it is worth noting that excluding a valid blurred pixel from the valid observation  $y$  does not cause a model violation, while the opposite (including a pixel that does not belong to the uniformly blurred area) does. Therefore, it is better to be conservative and not include pixels in the transition area<sup>5</sup>. For example, if a uniformly blurred object is observed over a sharp background, blurred pixels close to the object boundary are affected by the background and should be excluded from the mask.

Another interesting case is the construction of a mask for aliased observations. Although many aliasing scenarios are possible, in this article we have contemplated only the case where the aliased observation can be expressed as in Equation (2). For this case, we can construct the selection mask of the pixels to be interpolated by including  $Q - 1$  interleaved new samples between each pair of observed pixels. Figure 3 illustrates how to set the mask for a non-rectangular image patch suffering from aliasing with  $Q = 2$ . First it is framed into a rectangular patch, and then embedded into a larger image, which has extended boundaries to allow CBC and includes interleaved samples to address aliasing.

### 2.3.4 Computing an Initial Interpolation Guess

A good initial guess helps to speed up the convergence of the iterative step 2 in Algorithm 1. In order to obtain a starting point that is reasonably close to the final solution, we propose to use a method of our own: a shift-variant inter/extrapolation filtering using kernel  $h_e(n, m)$  with the mask  $s(n, m)$ . The inter/extrapolated image  $u_{int}(n, m)$  is computed by

$$u_{int}(n, m) = \frac{h_e(n, m) * (s(n, m)y(n, m))}{h_e(n, m) * s(n, m)}, \quad (20)$$

where the inter/extrapolation kernel  $h_e(n, m)$  is a decreasing function of its radius. Thus, the estimation is a weighted average of the known pixels, which gives more weight to the closer pixels and less weight to the farther ones. To avoid excessive blurring,  $h_e(n, m)$  must be carefully designed so

<sup>5</sup>A more powerful model is to consider non abrupt masks, as in [3].

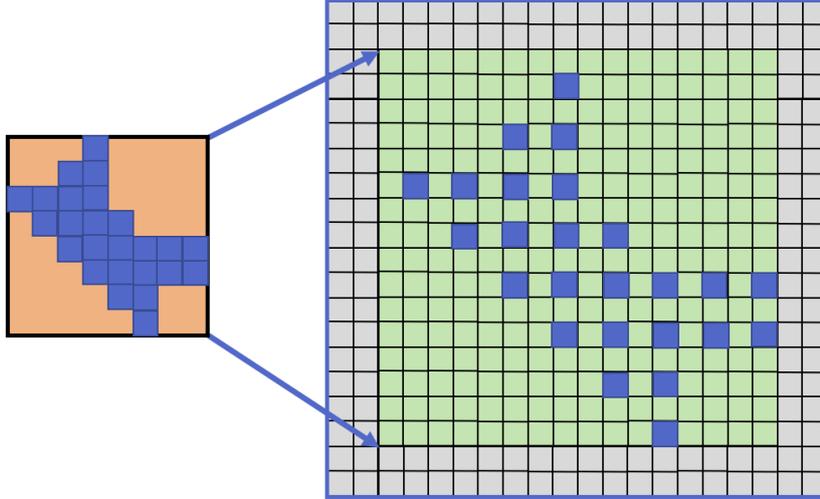


Figure 3: A graphical explanation of how a non-rectangular, non-CBC-complying, aliased blurred image (irregular dark blue area, on the left) is first framed (orange rectangle) and then reformatted for SPA interpolation, keeping the observed pixels (dark blue,  $s = 1$ ) and marking the rest (light color) for being inter/extrapolated ( $s = 0$ ). Here  $L_x^e = L_y^e = Q = 2$ .

that it behaves as a small filter at locations close to observed pixels, whereas acts like a large filter at locations far from the observed pixels. A good choice is a large radially symmetric filter, peaky at the center and rapidly decaying. Here we have chosen  $h_e(r) = r^{-n_e}$ , for  $r > 0$ , and  $h_e(0) = 0$  (as in any predictor filter, the central sample must vanish), with the power  $n_e = 7$  being chosen as the highest value not producing numerical instability. Using the above inter/extrapolation, our initial guess of  $z(n, m)$  is obtained as

$$z_{init}(n, m) = s(n, m)y(n, m) + (1 - s(n, m))u_{int}(n, m), \quad (21)$$

combining the interpolation with the observation in the same way as in Equation (15).

### 3 Applications and Experiments

SPA is a flexible method, applicable to many real-life imaging scenarios. To use it for different cases of blur model discrepancies, it only requires the mask indicating the pixels to interpolate, and setting the correct degradation parameters (blur kernel and noise). Here we give examples of using SPA to deal with circular boundary conditions, aliasing, kernel mismatch and arbitrary boundary shapes presented in blurred observations. This section describes four applications of SPA, three of them published in [15, 4, 5]. In all cases we have implemented the experiments using the latest versions of the SPA and  $L_2-r-L_0$  methods [16]<sup>6</sup> (the latter using  $N_{iter} = 10$  with its default parameters optimized for typical images, so no free parameter needs to be adjusted).

#### 3.1 Non-circulant-boundary Restoration

To use the DFT, image restoration techniques typically assume CBC for convolution. As this is an unnatural constraint for real-world blurred images, directly applying such methods for restoring real images usually results in strong artifacts.

<sup>6</sup>The latest version of the ©MATLAB source code for the  $L_2-r-L_0$  restoration method is available at: [https://www.researchgate.net/publication/325903515\\_L2rL0deblurMatlabToolbox2p1](https://www.researchgate.net/publication/325903515_L2rL0deblurMatlabToolbox2p1)

Among the many methods and recipes used for dealing with this problem, probably the most successful has been that of Reeves [17], initially applied to a Tikhonov-regularized blur inversion and later on integrated in non-linear restoration methods such as those in [18, 1, 10]. Reeves had the idea of extending the image along the boundaries in such a way that a classic smoothness measurement (namely, the Euclidean norm at the output of a Laplacian filter) was optimized across the periodic boundary transitions. A similar criterion was applied by [11], but aiming at producing a doubled-size CBC-tiling image including the three full mirror extensions of the original, as a pre-processing. Whereas this processing is effective at avoiding some artifacts during restoration, it is based on an elementary concept of smoothness. The smoothness criterion considers neither the particular blur kernel and noise level, nor the typical power spectrum of natural images. In contrast, SPA maximizes the overall likelihood (according to a still simple, but far richer, Gaussian model), on which both the particular shape of the blur kernel and the power spectral density of natural images have a strong impact. All this results in less artifacts during restoration, as we show here. Another successful method for non-circulant restoration is [12], which introduces a selection mask that removes invalid pixels affected by CBC in the estimation loop. This was later used by [10] in the context of automatic detection of pixels violating a blur model. The referred methods, except for [11], have in common that extra constraints are added to the restoration loop to deal with boundary conditions. In contrast, SPA and Liu’s method (besides other classical methods, like windowing), are pre-processing methods: they modify the observation so it can be restored using a standard deconvolution method.

**Images and degradation parameters.** For our experiments we have used three grayscale images all having size  $N_x = N_y = 256$  pixels for testing. Two of them are typical photographic images, *Cameraman* and *House*, and the other one is a highly textured image, *Straw*, having a lot of diagonal and obliquely oriented features. We have used 8 configurations of degradation parameters, consisting of 4 blur kernels and two noise levels (low and medium). For the 4 blur kernels, *PSF1* is  $h_{i,j} = (1 + i^2 + j^2)^{-1}$ , for  $i, j = -7 \dots 7$ , which is non-separable; *PSF2* is a  $9 \times 9$  uniform kernel; *PSF3* is a  $9 \times 1$  vertical motion blur; *PSF4* is a  $5 \times 7$  oblique kernel proportional to  $[00001111; 0012321; 0134310; 1232100; 1110000]$ . Details of each degradation pair (blur plus noise) are shown in Table 1.

**Simulation.** To obtain our simulated observation, we first performed a circular convolution and added noise. Then we kept the valid convolution area not affected by the CBC, resulting in an image with size  $M_x \times M_y$ , where  $M_x = N_x - 2n'_x$ ,  $M_y = N_y - 2n'_y$ , assuming odd-dimensions  $(2n'_x + 1) \times (2n'_y + 1)$  kernel sizes.

**Compared methods.** Our proposed method first uses SPA to estimate an extended blurry image and then applies a standard deconvolution method. The extended image has size  $(M_x + 2(n'_x + Le_x)) \times (M_y + 2(n'_y + Le_y))$ , which means our estimation goes beyond the original image support  $N_x \times N_y$ . Here we use  $Le_x = Le_y = 8$  and 50 iterations for SPA. We first compare SPA with mirror extension, which extends the image to the same size as in the case of SPA ( $N'_x = N_x + 2Le_x = M_x + 2n'_x + 2Le_x$ ), followed by edge tapering using ©MATLAB *edgetaper.m* (with a Gaussian kernel of  $\sigma = 6$ , hand-optimized value). We also compare to [11], as probably the most successful previous purely pre-processing method for avoiding boundary artifacts in deconvolution. To the best of our knowledge the only publicly available implementation of this method is the one from Sunghyun Cho<sup>7</sup>. In the

<sup>7</sup>We used [https://github.com/CoupeLibrary/handleoutlier/blob/master/code/wrap\\_boundary\\_liu.m](https://github.com/CoupeLibrary/handleoutlier/blob/master/code/wrap_boundary_liu.m) and passed a total size double of that of the observation. This simplifies the optimization problem to that of a membrane (minimizing the square of a Laplacian subject to boundary conditions), setting the parameters in [11] to  $\alpha = 1$ ,  $\lambda = 0$ .

results we have referred to this method as *WML*, as a mnemonic of “Wrapping Minimum Laplacian”. Finally, we also compare to an oracle, meant to provide an upper bound of the performance. This is calculated by directly deconvolving the simulation in its original  $N_x \times N_y$  support and then cropping the  $M_x \times M_y$  central area in the deconvolved result. Two deconvolution methods, *Wiener filter* ( $H_c(u, v)^* P_X(u, v) / P_Z(u, v)$ , using the spectral model of Equation (9) and (10)) and  $L_2$ - $r$ - $L_0$  [16], are used for comparison. Whereas the Wiener method performs well on images containing textures, the second one, based on sparsity, performs better on common real-world photographic images.

**Results and comparisons.** The performance of each method is measured by the Improvement in Signal-to-Noise Ratio (ISNR) calculated in the region not affected by CBC ( $M_x \times M_y = (N_x - 2n'_x) \times (N_y - 2n'_y)$ ). Figure 4 provides a comparison of the average performance of the two methods on the three images, for each degradation pair PSF+noise. It shows that our SPA is substantially more robust and it outperforms the two compared boundary extension methods (mirroring followed by edgetaper and [11]) for all cases, especially for degradation numbers #3, #5, #6 and #7, where the kernel has zeros in the low-medium frequencies or lacks horizontal/vertical symmetry. It is interesting to note that the only PSF, for which the SPA improvement is smaller, is *PSF1* that corresponds to a filter not having blurring zeros in low and medium frequencies. In general, the difference in gain is higher when the noise level is low, as low-noise implies higher amplification of the involved spectral components. In contrast, SPA significantly improves the results by correctly reproducing the blurring kernel’s zeros at their locations in the spectrum.

Figure 5 shows a visual comparison of mirroring plus edgetaper, WML [11], and SPA for *Straw* (a crop of the upper left corner without the pixels extended beyond the original image support,  $N_x \times N_y$ ) with degradation #7. Edgetaper, and WML to a lesser extent, produce strong artifacts along the image boundaries. SPA, in contrast, reduces drastically the artifacts and preserves the texture continuity.

Table 1 shows all the numerical results. The increase in performance using SPA is consistent for both deconvolution methods, as an indication of its universal applicability. We emphasize that not only average performance is much better using SPA, but also our method did not produce any catastrophic failures, like those of WML and edge tapering in *House* and *Straw* under the degradation #7 (oblique convolution kernel under low noise), using  $L_2$ -relaxed  $L_0$  restoration.

### 3.2 Aliasing-aware Restoration

Real-life imaging systems usually have a combination of blur and aliasing. The Point Spread Functions (PSFs) of real imaging devices are generally far from being ideal low-pass filters, which automatically gives rise to a certain amount of aliasing when going from a continuous to a discrete image representation. Also, the sensor’s detectors usually have an active area size comparable to the effective PSF size, and leave a significant proportion of their surface inactive. In real optical digital cameras, the combined effect of optics and pixel area integration in the sensor has been traditionally chosen to trade off the amount of aliasing against the spatial resolution. Usually a certain amount of aliasing in the digital images has been considered preferable to images that are alias-free but very blurry. In the last two decades, the spatial resolution of the sensor has grown faster than the improvement of the optical PSFs, which effectively reduces the amount of aliasing on the sensor. However, there are still medium resolution cameras having large detectors (e.g. thermal cameras), which are used for the cases where noise is a bigger concern than spatial resolution, for which the amount of aliasing is still considerable.

Another issue is that, when empirically measuring PSFs in a digital system affected by aliasing, the resulting PSF images are not shift-invariant anymore, and using them for restoration violates the shift-invariance implied in the conventional image degradation model, Equation (1). In addition,

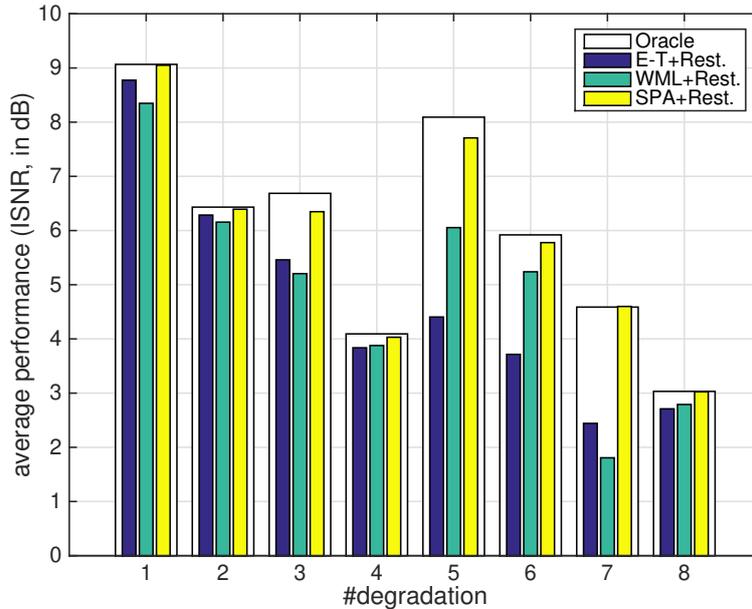


Figure 4: Average performance (using Wiener filter and  $L_2$ - $r$ - $L_0$ , and the three images of the set) of non-CBC image restoration, for three compared (edgetaper, WML [11], SPA) extension methods, plus the CBC-oracle, and 8 degradation (blur + noise) pairs. It is striking how close the SPA interpolation (yellow bars) gets to the ideal CBC-filtered image (oracle, drawn as a bounding box).

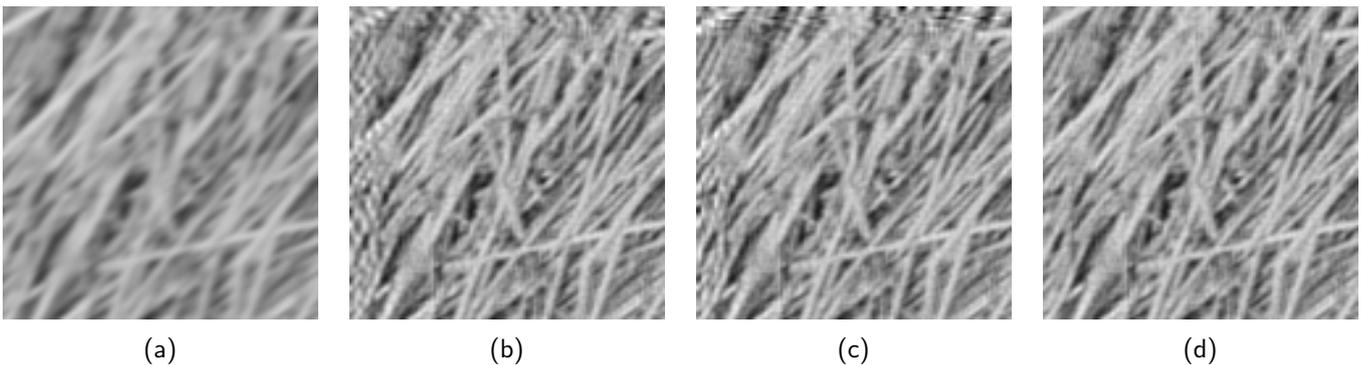


Figure 5: Visual results for CBC-handling in restoration. Upper left corner of *Straw*. (a) Simulated observation (degradation #7). (b) Deconvolution result using edgetaper and the Wiener filter. (c) Same as in (b), but using WML [11] for extending the image boundaries. (d) Same as in (b), but using SPA for extending the image boundaries.

albeit aliasing is typically not noticeable in blurred observations, it nevertheless results in serious artifacts after restoration as we show here.

Another interesting practical scenario combining blur and aliasing is the raw color channels in the camera color filter array (CFA) used in most consumer digital cameras. Camera lenses introduce blur, but not enough to avoid a certain amount of aliasing on each of these raw color mosaics. Again, preserving high spatial frequencies is usually considered preferable to completely removing the aliasing (causing artifacts in the demosaiced image).

Most often the aliasing problem in images is addressed by using anti-aliasing filters (which reduces both the aliasing and the spatial resolution) or making use of multiple observations (which is the base for most super-resolution algorithms), either as multiple captures, or as frames from a video sequence. Interpolation methods considering a single observation, such as [6] and [25], build sophisticated prior image models using self-similarity of natural images through scale and space to cope with

			Wiener				L2-r-L0			
			<i>ET+R</i>	<i>WML+R</i>	<i>SPA+R</i>	<i>Ora.</i>	<i>ET+R</i>	<i>WML+R</i>	<i>SPA+R</i>	<i>Ora.</i>
deg. #	<i>h</i>	$\sigma_w^2$	<i>HOUSE</i>							
#1	<i>PSF1</i>	0.25	7.41	6.95	<b>7.77</b>	<i>7.82</i>	9.23	8.15	<b>9.94</b>	<i>9.95</i>
#2	<i>PSF1</i>	2.00	5.01	4.94	<b>5.11</b>	<i>5.16</i>	7.82	7.17	<b>8.23</b>	<i>8.31</i>
#3	<i>PSF2</i>	0.31	5.75	5.46	<b>6.49</b>	<i>6.69</i>	8.29	7.88	<b>9.62</b>	<i>10.63</i>
#4	<i>PSF2</i>	4.00	4.04	4.00	<b>4.10</b>	<i>4.07</i>	6.97	7.14	<b>7.48</b>	<i>7.70</i>
#5	<i>PSF3</i>	1.00	2.40	5.78	<b>6.08</b>	<i>6.27</i>	3.26	8.03	<b>9.10</b>	<i>9.56</i>
#6	<i>PSF3</i>	4.00	1.80	4.30	<b>4.32</b>	<i>4.34</i>	3.02	7.22	<b>7.63</b>	<i>7.88</i>
#7	<i>PSF4</i>	0.25	2.73	2.28	<b>3.10</b>	<i>3.01</i>	5.13	2.77	<b>5.92</b>	<i>5.97</i>
#8	<i>PSF4</i>	4.00	1.97	2.02	<b>2.11</b>	<i>2.09</i>	3.96	3.99	<b>4.36</b>	<i>4.40</i>
deg. #	<i>h</i>	$\sigma_w^2$	<i>CAMERAMAN</i>							
#1	<i>PSF1</i>	0.25	8.67	8.63	<b>8.69</b>	<i>8.71</i>	10.90	10.64	<b>10.93</b>	<i>10.88</i>
#2	<i>PSF1</i>	2.00	<b>5.56</b>	5.55	5.53	<i>5.56</i>	<b>8.23</b>	8.10	<b>8.23</b>	<i>8.30</i>
#3	<i>PSF2</i>	0.31	5.58	5.65	<b>5.94</b>	<i>6.24</i>	8.43	7.86	<b>8.78</b>	<i>9.07</i>
#4	<i>PSF2</i>	4.00	<b>3.88</b>	3.85	3.87	<i>3.90</i>	5.48	5.51	<b>5.58</b>	<i>5.56</i>
#5	<i>PSF3</i>	1.00	5.47	5.41	<b>6.40</b>	<i>6.60</i>	7.30	7.10	<b>9.51</b>	<i>10.02</i>
#6	<i>PSF3</i>	4.00	4.03	4.13	<b>4.42</b>	<i>4.49</i>	6.28	6.61	<b>7.47</b>	<i>7.66</i>
#7	<i>PSF4</i>	0.25	3.77	3.73	<b>3.90</b>	<i>3.90</i>	6.44	5.68	<b>6.77</b>	<i>6.83</i>
#8	<i>PSF4</i>	4.00	2.32	2.35	<b>2.35</b>	<i>2.36</i>	3.56	3.59	<b>3.65</b>	<i>3.67</i>
deg. #	<i>h</i>	$\sigma_w^2$	<i>STRAW</i>							
#1	<i>PSF1</i>	0.25	8.97	8.87	<b>9.29</b>	<i>9.39</i>	8.22	7.70	<b>8.40</b>	<i>8.36</i>
#2	<i>PSF1</i>	2.00	5.85	5.91	<b>5.97</b>	<i>6.01</i>	6.16	6.00	<b>6.28</b>	<i>6.27</i>
#3	<i>PSF2</i>	0.31	3.50	3.65	<b>4.56</b>	<i>4.84</i>	3.54	3.05	<b>4.94</b>	<i>5.28</i>
#4	<i>PSF2</i>	4.00	2.27	2.37	<b>2.53</b>	<i>2.60</i>	2.22	2.28	<b>2.53</b>	<i>2.62</i>
#5	<i>PSF3</i>	1.00	4.95	6.18	<b>7.92</b>	<i>8.46</i>	4.71	4.64	<b>8.38</b>	<i>9.08</i>
#6	<i>PSF3</i>	4.00	4.21	5.24	<b>5.81</b>	<i>6.03</i>	4.19	4.79	<b>6.18</b>	<i>6.42</i>
#7	<i>PSF4</i>	0.25	<b>2.37</b>	<b>2.72</b>	<b>4.91</b>	<b>4.96</b>	-1.32	-2.07	<b>4.04</b>	<i>3.99</i>
#8	<i>PSF4</i>	4.00	2.62	2.94	<b>3.15</b>	<i>3.19</i>	2.19	2.21	<b>2.91</b>	<i>2.90</i>
average			4.39	4.70	<b>5.18</b>	<i>5.28</i>	5.59	5.67	<b>6.95</b>	<i>7.14</i>

Table 1: Performance comparison for non-CBC image restoration measured as ISNR in decibels. “E-T” refers to “edgetaper”, “WML” to “Wrap Minimum Laplacian” [11], “Ora” to the oracle (CBC-restoration result on CBC-complying simulation), and “SPA” to our SPA extension method. Results in Figure 5 are shown on yellow background.

aliasing. Unlike these approaches, SPA tackles aliasing in a single blurred observation by using spectral properties derived from the PSF, noise and power spectral density of images for interpolating the missing samples. As we will show in some examples, thanks to aliasing, it may achieve super-resolution to some extent, i.e., recover spectral contents beyond the Nyquist frequency of the observed image.

**Simulation procedure.** An important requirement to recover an approximation of the missing samples in the observed image is to be able to reliably estimate the blur kernel in high resolution (HR). As in previous applications, the observation  $y$  is modeled as a subset of pixels of an extended image  $z$  complying with the reference observation model in Equation (1). The simulation procedure first generates such image  $z$  with size  $N_x \times N_y$ , which is an alias-free CBC-blurred noisy image. Accordingly, our mask for the known pixels is a binary array in HR with ones at every other pixel location. In this set of experiments boundary pixels affected by CBC are neither masked nor added for interpolation, because we want to address the aliasing problem without being affected by boundary effects. As a consequence we set  $Le_x = Le_y = 0$  for SPA.

**Images and degradation parameters.** We have tested our method on two 8-bit grayscale images: *Barbara* and *Pirate*, with  $N_x = N_y = 512$ , from a public image database<sup>8</sup>. Three different PSFs are used as the HR kernel ( $h_{HR}$ ): *PSF1* is  $h_{i,j} = (1+i^2+j^2)^{-1}$ , for  $i, j = -7 \dots 7$ ; *PSF2* is a  $9 \times 9$  uniform kernel; and *PSF3* is a  $5 \times 7$  oblique kernel given by [00001111; 0012321; 0134310; 1232100; 1110000]. Each kernel is normalized to sum up to one. In addition, three different noise levels categorized as low ( $\sigma_\omega^2 = 0.09$ ), medium ( $\sigma_\omega^2 = 0.25$ ) and high ( $\sigma_\omega^2 = 2.25$ ) are added. Thus, in total nine degradation sets consisting of three blur kernel and three noise levels are used for experiments.

**Compared methods.** Likewise in non-CBC restoration, we have used two different methods, the classic *Wiener filter* and the  $L_2$ - $r$ - $L_0$  method [16]. We let SPA to run 50 iterations in all cases. The performance of the algorithm is measured as ISNR in decibels with respect to a double-size version of  $y$ , which is created by repeating each pixel three times. We have compared the results of three different restoration schemes with a baseline image:

- (1) *R+I*, (*Restoration+Interpolation*): directly deconvolving the observation ( $M_x \times M_y$ ) with a  $2 \times 2$  subsampled  $h_{HR}$  (always keeping the PSF's central sample), and then applying a bi-cubic spline interpolation to get the HR image ( $N_x \times N_y$ ).
- (2) *I+R* (*Interpolation+Restoration*): applying first a bi-cubic spline interpolation on the observation ( $M_x \times M_y$ ), and then deconvolving the HR image ( $N_x \times N_y$ ) with  $h_{HR}$ .
- (3) *SPA+R*, (*SPA+Restoration*): applying SPA interpolation on the observation ( $M_x \times M_y$ ) and then deconvolving the HR image ( $N_x \times N_y$ ) with  $h_{HR}$ .
- (4) *Oracle*: the baseline image obtained by deconvolving the alias-free image  $z$ . This sets an upper bound for the performance.

**Results and comparisons.** In order to understand the influence of different degradation sources on the results, we have plotted in Figure 6 the mean ISNR in dB of the compared methods for each simulated degradation. The mean values are calculated over both images and both deblurring algorithms. The plot shows that for the PSFs having zeros in the Fourier domain, namely *PSF2* (a uniform blur) and *PSF3* (an oblique blur), the non-aliasing-aware methods *R+I* and *I+R* produce poor results, especially for the low-noise scenarios. In contrast, the proposed *SPA+R* performs more evenly for all the tested degradation types, and is always significantly better than the other two methods.

Visual outcomes ( $256 \times 256$  close-ups of both images) are shown in Figure 7, with top row corresponding to *Barbara* with *PSF3* (oblique kernel) and  $\sigma_\omega^2 = 0.09$  (low noise level), and bottom row to *Pirate* with *PSF2* (uniform kernel) and  $\sigma_\omega^2 = 2.25$  (high noise level). As one can see, visually our *SPA+R* method outperforms the other two on both images by producing less artifacts, even in the case of *Barbara*, which suffers more from aliasing due to its large amount of high frequencies. By estimating the alias-free HR image, *SPA+R* is able to even recover some aliased details in the observation such as the strips on Barbara's shoulder. ISNRs for these two cases are highlighted in yellow in Table 2.

Table 2 presents the numerical results as ISNR in decibels, from which we have concluded that, on average, our *SPA+R* is 1.5 dB and 2.5 dB above *R+I* and *I+R*, respectively, and 2.5 dB below the *Oracle*. *Wiener filter* and  $L_2$ - $r$ - $L_0$  restoration methods behave similarly across all experiments. Roughly speaking, *SPA* method goes half-way between the standard (spline) interpolation and the ground truth.

<sup>8</sup>ImageProcessingPlace.com, [http://www.imageprocessingplace.com/root\\_files\\_V3/image\\_databases.htm](http://www.imageprocessingplace.com/root_files_V3/image_databases.htm)

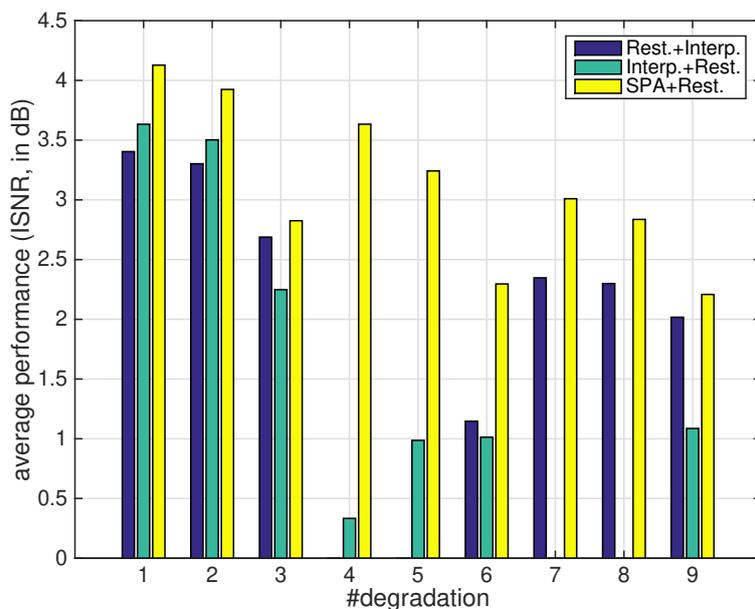


Figure 6: Aliasing-aware restoration results: Average performance (using Wiener filter and  $L_2$ - $r$ - $L_0$ , and two images, *Barbara* and *Pirate*) of the 3 compared restoration methods for the 9 degradation (blur + noise) experiments.



Figure 7: Results of aliasing-aware restoration. Top row: Crops of blurred and aliased simulated observations. Rest, from left to right: *R+I*, *I+R*, *SPA+R* (our result), *Oracle* using  $L_2$ - $r$ - $L_0$  [16] for restoration. *R+I* refers to first restoration, then spline interpolation. *I+R* the same, but in the reversed order. *SPA+R* refers to our method, spectral pre-adaptation plus image restoration. *Oracle* is the upper bound reference derived from direct deconvolution of the HR alias-free blurred image.

			Wiener				L2-r-L0			
			<i>R+I</i>	<i>I+R</i>	<i>SPA+R</i>	<i>Ora.</i>	<i>R+I</i>	<i>I+R</i>	<i>SPA+R</i>	<i>Ora.</i>
deg. #	<i>h</i>	$\sigma_w^2$	<i>BARBARA</i>							
#1	<i>PSF1</i>	0.09	2.23	2.63	<b>2.73</b>	<i>9.75</i>	2.27	2.61	<b>2.78</b>	<i>11.61</i>
#2	<i>PSF1</i>	0.25	2.14	2.50	<b>2.57</b>	<i>7.09</i>	2.23	2.54	<b>2.72</b>	<i>9.45</i>
#3	<i>PSF1</i>	2.25	1.63	1.48	<b>1.72</b>	<i>3.23</i>	1.96	1.80	<b>2.08</b>	<i>6.15</i>
#4	<i>PSF2</i>	0.09	-1.19	-0.60	<b>2.46</b>	<i>4.44</i>	-2.81	-1.03	<b>2.70</b>	<i>6.59</i>
#5	<i>PSF2</i>	0.25	-0.17	0.03	<b>2.13</b>	<i>3.43</i>	-0.80	-0.08	<b>2.35</b>	<i>5.15</i>
#6	<i>PSF2</i>	2.25	0.66	0.21	<b>1.36</b>	<i>1.97</i>	0.91	0.73	<b>1.72</b>	<i>2.68</i>
#7	<i>PSF3</i>	0.09	1.40	-1.72	<b>1.83</b>	<i>3.76</i>	<b>1.41</b>	<b>-4.99</b>	<b>1.92</b>	<b>5.61</b>
#8	<i>PSF3</i>	0.25	1.36	0.01	<b>1.68</b>	<i>2.84</i>	1.40	-2.83	<b>1.79</b>	<i>4.36</i>
#9	<i>PSF3</i>	2.25	1.13	0.90	<b>1.20</b>	<i>1.62</i>	1.29	0.34	<b>1.34</b>	<i>2.17</i>
deg. #	<i>h</i>	$\sigma_w^2$	<i>PIRATE</i>							
#1	<i>PSF1</i>	0.09	4.96	5.10	<b>6.11</b>	<i>9.37</i>	4.99	4.82	<b>6.19</b>	<i>10.16</i>
#2	<i>PSF1</i>	0.25	4.78	4.94	<b>5.64</b>	<i>7.54</i>	4.85	4.61	<b>5.86</b>	<i>8.65</i>
#3	<i>PSF1</i>	2.25	3.67	2.88	<b>3.64</b>	<i>4.80</i>	3.97	3.04	<b>4.42</b>	<i>6.69</i>
#4	<i>PSF2</i>	0.09	-1.63	1.49	<b>4.70</b>	<i>6.54</i>	-3.50	2.38	<b>5.39</b>	<i>8.63</i>
#5	<i>PSF2</i>	0.25	-0.18	1.82	<b>4.19</b>	<i>5.54</i>	-1.24	2.87	<b>4.94</b>	<i>7.47</i>
#6	<i>PSF2</i>	2.25	1.39	1.03	<b>2.77</b>	<i>3.61</i>	<b>1.71</b>	<b>2.35</b>	<b>3.71</b>	<b>5.23</b>
#7	<i>PSF3</i>	0.09	3.54	-3.53	<b>4.38</b>	<i>5.34</i>	3.57	-8.10	<b>4.74</b>	<i>7.56</i>
#8	<i>PSF3</i>	0.25	3.44	-0.63	<b>4.10</b>	<i>4.73</i>	3.50	-4.70	<b>4.54</b>	<i>6.59</i>
#9	<i>PSF3</i>	2.25	2.92	1.87	<b>3.10</b>	<i>3.70</i>	3.10	1.38	<b>3.73</b>	<i>4.93</i>
average			1.78	1.13	<b>3.13</b>	<i>4.96</i>	1.60	0.43	<b>3.50</b>	<i>6.65</i>

Table 2: Aliasing-aware restoration. Performance comparison, measured as ISNR in decibels. Results in Figure 7 are highlighted in yellow. See text for details.

### 3.3 Kernel Mismatch

Apart from real observations being aliased or not, we cannot expect to know the exact real blur kernel. Even for the cases in which we know the theoretical PSF (e.g. because we know every detail of the optical design), the real optics will not be exactly the same as the designed ones, due to tolerance errors in fabrication. Therefore, in reality, some extent of discrepancy between the real PSF (causing the blur in the observation) and the modeled PSF (the one assumed during deconvolution) is unavoidable. As expected, this blur kernel mismatch causes artifacts in the restoration result.

The obvious approach to deal with the inaccurate kernel estimation problem is to use more advanced kernel estimation tools. Goldstein and Fattal [7] explored the statistical irregularities in the power spectrum of blurred natural images. Xu and Jia [21] took a close look at the role of image edges in image deblurring, and proposed an efficient kernel estimation method based on a spatial prior and iterative support detection kernel refinement. Others, such as [23] and [24] obtain improvements through the power of deep neural networks. However, ultimately, it is hard to measure the achieved accuracy in the estimated kernel. Instead of pushing the model kernel closer to the actual one, SPA pushes the observation to be more likely under the assumption that the estimated blur kernel is correct. Of course, still, for this latter approach to be fruitful, there should not be significant gaps between the real PSF and the estimated one. Both approaches (bringing the model kernel closer to the actual one, and bringing the post-processed observation closer to the modeled kernel) are, rather than alternative, complementary: We can think of a sequential process, namely, first obtaining the best possible model kernel, subject to whatever constraints and model limitations, and, only then, applying SPA to bring the observation as close as possible to the assumed observation model, thus further reducing the restoration artifacts.

In our experiments, we use SPA to interpolate the unknown pixels in order to adapt the spectral properties of the modified observation, such that it becomes likely in terms of the modeled PSF. In

order to introduce the necessary degrees of freedom to make such a transformation, we start from an aliased observation, identical as the one described in the previous section. The SPA goal is, again, to interpolate those discarded pixels (3 out of 4, in  $2 \times 2$  neighborhoods), according to the observation model of Equation (2) with  $Q = 2$ . However, now, besides the observation being aliased, there is a mismatch between the real blur used to generate the observations and the modeled PSF assumed by the deconvolution method. Two sets of experiments are conducted for two different kinds of mismatches. In the first set, the real and modeled PSF are both spatially invariant, but have slightly different sizes. In the second, the real PSF slightly varies across the image, while the modeled PSF is kept fixed (shift-invariant), always according to the reference observation model of Equation (1). Since these are more challenging cases, with both aliasing and kernel mismatch, we increased the number of iterations for SPA to 100. We compared three restoration approaches:

- (1) Directly deconvolve an alias-free oracle simulation in HR, assuming the modeled PSF.
- (2) Use SPA to estimate a HR version of the aliased LR observation based on the modeled PSF, then deconvolve it assuming the modeled PSF.
- (3) Use a standard bi-cubic spline interpolation method to interpolate the aliased LR observation, then deconvolve the result assuming the modeled PSF.

For the deconvolution, we used the  $L_2$ - $r$ - $L_0$  method [16] as in previous experiments.

**Size mismatch.** We have used a uniform disk blur kernel here for simulation. We have simulated the kernel size mismatch by setting the diameter  $D$  of the real PSF to range from  $-40\%$  to  $+40\%$  of the modeled PSF diameter, in  $2\%$  steps. Figure 8 panel(a) shows the performance of each method with  $D_{modeled} = 9$ ,  $\sigma_\omega = 0.3$ , *Lena* image. One can see that, when there is both aliasing and mismatch (blue and black curves), SPA interpolation gives better restoration results compared to the standard spline method across all kernel size deviations. More importantly, even when there is no aliasing (red oracle curve) in the observation and the kernel size mismatch is significant (more than  $4\%$ ), directly deconvolving it with the modeled PSF gives worse results (up to 2 dB performance drop) than using SPA on the aliased observation. That is, when kernel size mismatch dominates, SPA can compensate its side effects even if aliasing is present at the same time, and this compensation has a larger positive effect than the loss of information caused by aliasing (at least in the examples tested by us).

**Spatially variant kernels.** For the spatially variant blur experiments, we have considered that the kernel diameter grows linearly with the radius from the center of the image, varying its diameter from  $(100 - p)\%$  of its nominal (modeled) value at the center of the image to  $(100 + p)\%$  at the center of the four image edges. We swept  $p$  from 0 to 40, in  $1\%$  incremental steps. We use the same diameter ( $D_{modeled} = 9$ ) for the modeled PSF and the same noise level ( $\sigma_\omega = 0.3$ ) with *Lena* image, as for the uniform size mismatch experiment. Figure 8 panel(b) shows the performance of each method. Similar as before, SPA-based restoration behaves very robustly against kernel mismatch (blue curve), as compared to a simple standard interpolation (black curve). Furthermore, it even outperforms alias-free oracle restoration when the kernel mismatch is large.

Figure 9 shows some visual results. The images in the first row are spatially-variant simulated observations, HR alias-free oracle (left), LR aliased after cubic spline interpolation (center), and LR aliased (right). The image on the left was downsampled by taking one out of four pixels (one in 2, along  $x$  and  $y$ , i.e.,  $Q = 2$ ), for simulating the aliased observation, on the right. Images in the second and third row, from left to right, are deconvolution results of the observation in HR, spline interpolated image and SPA interpolated image, all of them assuming the reference, shift-invariant,

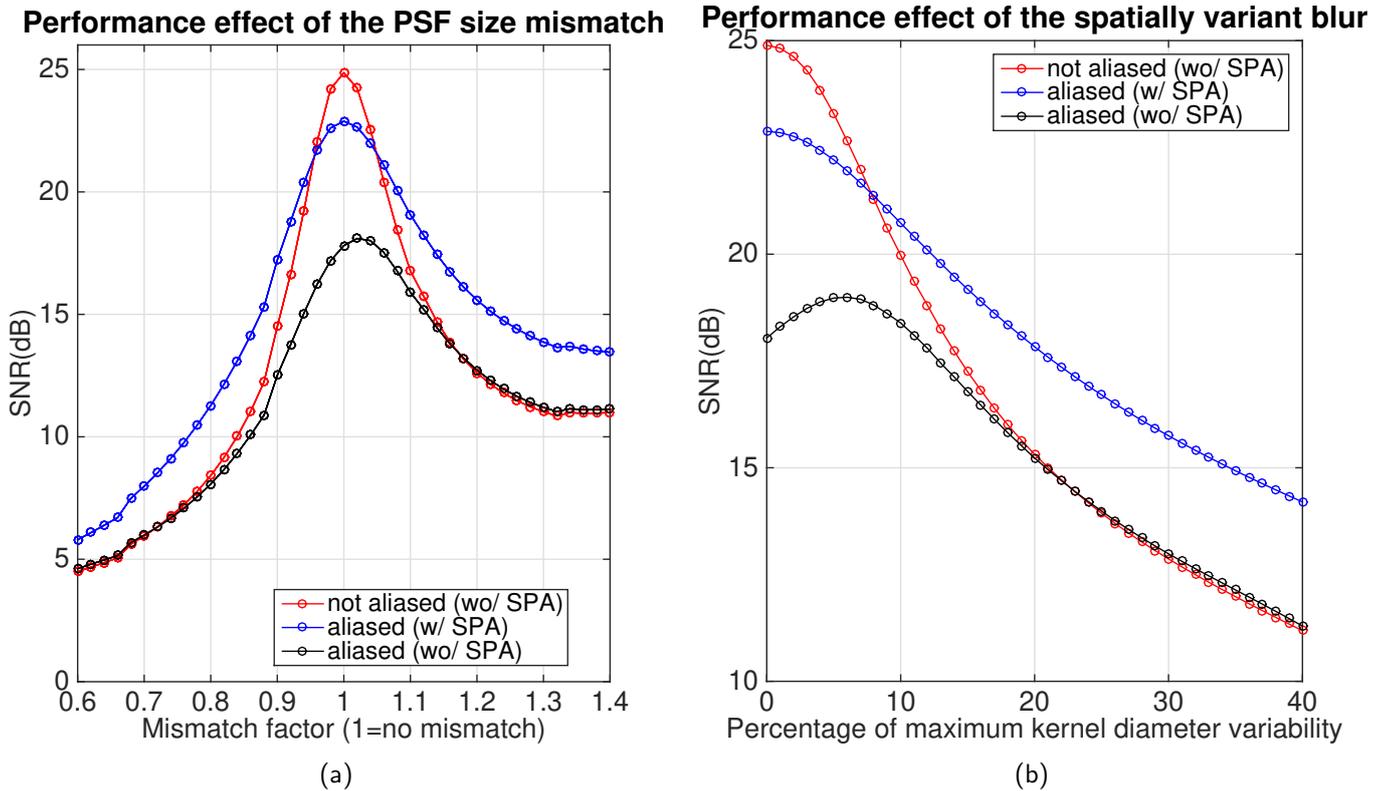


Figure 8: (a) Performance effect of PSF size mismatch. (b) Performance effect of PSF size variability.

PSF. Images in the second row are results at 5% variability. In this case, from Figure 8 panel(b) one can see that deconvolving directly on the HR simulation provides a higher SNR than on the SPA interpolation. Visually, it is true that the left image in the second row better preserves high frequency details than the SPA-processed one (right). However, SPA still gives a visually comparable result without increasing the artifacts. We believe this is a remarkable result, considering that, unlike for the alias-free oracle, the SPA-based restoration is dealing here both with aliasing and kernel mismatch. The third row shows results at the crossing point (8%) of the red and blue curves in Figure 8 panel (b), where the oracle and SPA-interpolated images give the same SNR. One can see that with larger kernel size variability, SPA produces relatively less artifacts than directly deconvolving the alias-free oracle. On the other hand, using a general purpose (cubic spline) interpolation applied to the aliased LR simulated observation (right) gives rise to strong artifacts.

### 3.4 Complex-shape Support Deconvolution

When blur changes abruptly from one image region to another, the image will have multiple distinct blurred areas of arbitrary shape. To restore such heterogeneously blurred images, we propose to frame these distinct uniformly blurred areas in rectangular regions and then apply SPA to estimate a rectangular shaped blurry image for each area. For each rectangular shaped sub-image, pixels not obeying the local blur model (single homogeneous blur, rectangular, CBC) are inter/extrapolated by SPA, thus enforcing them to follow the reference blurring model of this area. Then each estimated rectangular shaped blurred sub-image can be restored by any standard deconvolution algorithm. Finally, composing all masked restorations results with the original discarded pixels results in the final restored output.

Such two-step method can be very useful in practice. For example, an image with foreground in



Figure 9: Visual results of kernel mismatch: using a standard deconvolution method to restore an aliased spatially-variant blurred image after being interpolated using SPA (right column). We compare with direct deconvolution not using SPA: on an oracle, alias-free version (left column), and on the aliased observation after spline interpolation (central column). *First row*: Simulated observations for a 5% real kernel variation. *Second row*: Restoration results for a 5% real kernel variation. *Third row*: Restoration results for 8% real kernel variation. Note that SPA results not only have less artifacts compared to deconvolving the spline-interpolated LR image (central column), but also compared to deconvolving the alias-free oracle (left column).

focus and blurred background is considered as a two-area case, and can be restored with the help of SPA<sup>9</sup> as mentioned above. Although not shown here, this general model is applicable to more complex situations, such as having different motion blur in the scene, possibly in combination with out of focus elements, etc. In addition, this approach is powerful enough even for recovering missing blurred areas (*in-painting*) that have a size comparable to that of the blur kernel. We have compared the proposed two-step method with a successful unconstrained boundary conditions (UBC) variant of Alternating Direction Method of Multipliers (ADMM) restoration algorithm proposed in [1]. Experiments on real images have been conducted as well.

**Block in-painting.** For fair comparison with [1], we repeated the authors’ experiment with *Lena* image ( $256 \times 256$ ) and generated the blur kernel using their code<sup>10</sup>. The observation suffers from uniform blur with size  $19 \times 19$  and a very low noise level having a Blurred Signal-to-Noise Ratio (BSNR) of 60 dB. Also, pixels of three relatively large areas (eyes:  $19 \times 19$ , same size as the kernel; mouth:  $39 \times 39$ , almost twice of the kernel size) are discarded, together with the CBC-affected pixels along the boundaries (resulting in  $256 - 2 \times ((19 - 1)/2) = 238$  pixel size on both axes), as shown in Figure 10 panel (a). Since it is an ill-posed problem with extremely little noise, the number of iterations for SPA had to be increased to 12500. We needed 500 iterations in UBC-ADMM to guarantee convergence. The boundary extension was set to  $Le_x = Le_y = 5$  as follows from Equation (19) ( $(256 + 2 \times 5)/19 = 14$ ), and the original intensity range ( $[0, 205]$ ) was enforced in  $L_2-r-L_0$ .

Figure 10 panel (b) shows the SPA result ( $(256 + 2 \times 5) \times (256 + 2 \times 5) = 266 \times 266$ ), which gives an estimation of the blurred image in the missing blocks as well as in the extended boundary areas. After the deconvolution, the final in-painting results of SPA is presented in Figure 10 panel (c) along with the result of UBC-ADMM in panel (d) (both with size  $256 \times 256$ ). In this extreme case, and thanks to the non-linear regularization, UBC-ADMM clearly gives a better restoration than our method, both visually and in terms of SNR on the whole image (23.33dB vs. 22.52dB). However, the fact that our method successfully recovers the missing eyes of size comparable to the kernel indicates that a relatively simple spectral linear method like SPA suffices to regenerate incomplete blurred observations.



Figure 10: Recovering lost regions in a blurred image (a), using SPA interpolation (b) followed by a standard deconvolution (c), vs. using UBC-ADMM [1] (d). Note that the simulated observation (a) is smaller than the original size (and restored, (c) and (d)), which, in turn, is smaller than (b).

<sup>9</sup>Note that it is possible to use the same scheme for the reversed situation, where the foreground is blurred and the background is sharp. However, better results are obtained in that case using a more complex observation model [3].

<sup>10</sup>We thank M.S.C. Almeida and M.A.T. Figueiredo for making their code available.

**A sharp foreground on blurred background.** We tested our SPA-based restoration method on a simulated observation ( $N_x = N_y = 512$ ) composed of a foreground object in focus partially occluding a blurred background. We used a uniform disk of 9-pixel diameter for simulating the out-of-focus blur, and additive noise with  $\sigma_\omega = 0.3$ . SPA is applied to estimate the blurred image with  $L_{e_x} = L_{e_y} = 8$  and 100 iterations followed by the  $L_2$ - $r$ - $L_0$  deconvolution method [16] with default parameters. The result is compared with UBC-ADMM (FA-MD variant), with a hand-optimized parameter  $\lambda = 0.003$  and running 150 iterations to achieve practical convergence.

Visual results are shown in Figure 11. Panels (a) and (b) show the reference (ground-truth) image and the simulated observation, respectively. The result of UBC-ADMM is in panel (c). The result of SPA after deconvolution (using  $L_2$ - $r$ - $L_0$  [16]) is shown in panel (d). Stitching the deconvolved results with the observed sharp foreground object, we get the final restorations for both methods shown in panels (e) and (f). Both methods provide excellent results, yet our combined method better preserves texture and has less artifacts (see, for instance, the pirate’s chin, zoomed in panel (e)). This is a typical behaviour of  $L_2$ -relaxed  $L_0$ -sparse methods against TV-based methods. This leads to a significantly better Signal-to-Noise ratio (SNR) (measured on the restored background) of our method (24.96dB vs. 24.44dB). Running both algorithms on an Intel(R) Core(TM) i7-9700 CPU @ 3.00GHz using ©MATLAB code on R2019b version, UBC-ADMM takes ca. 25 s (in full silent mode), which is 5.7 times slower than the sum of SPA (ca. 0.8 s) plus  $L_2$ - $r$ - $L_0$  (ca. 3.6 s).

**Real images.** We have applied our method also to real images, in what we believe is the ultimate (and necessary) challenge for any method claiming to be applicable to real observations. The first testing photo, a crop of which is shown in Figure 12 panel (a), has been taken<sup>11</sup> by a professional digital camera<sup>12</sup>, indoors, with tripod and without flash. Although its sensor uses 12 bit pixels, it is mapped into a 16 bit intensity range ([0, 65535]). In this case, we have taken the mean of all three color channels to convert it into grayscale, which further reduced the noise. We eventually used a  $512 \times 512$  region of the photo for processing. The foreground (close up of a comb) can be easily segmented to obtain the mask, since it has a much darker color with respect to the background. For estimating the blur kernel, we first assumed a uniform disk, and then manually optimized the diameter by sweeping its value while deconvolving a different region of the original image containing only the background. The final estimated diameter is  $D \approx 29.5$ . For the noise, by applying [8], we have estimated  $\sigma_\omega \approx 4$ .

The second case is more challenging: restoring a picture taken with a mid-range mobile phone, indoors, indirect natural light, without tripod, and in JPEG format<sup>13</sup>. A  $512 \times 512$  crop of the picture is shown in Figure 13 left. For this case, we obtained the foreground object mask through high-pass filtering and morphological operations, and estimated  $D \approx 14.5$ ,  $\sigma_\omega \approx 1.25$  similarly as before. We processed the background image in gray-level, and after the restoration we re-scaled each color channel to match their means in the observation.

For the SPA method, we used  $L_x^e = L_y^e = 5$  and 100 iterations of the CG method for both pictures. For the  $L_2$ - $r$ - $L_0$  method, we have imposed the range [30000, 54000] for the first image and [120, 210] for the second one during the restoration loop. Results for both pictures are presented in Figure 12 panel (d) and Figure 13 right. One can see that some of the letters/words on the background text become readable after restoration.

<sup>11</sup>Many thanks to Isabel Portilla for kindly taking and providing us with this picture.

<sup>12</sup>Canon EOS 5D Mark II, 20 Mpix, CR2 to uint16 TIFF, focal length 50mm, exposure 1/2500 s., ISO 100, f/3.5.

<sup>13</sup>2Mpix camera from Sony Xperia D2203, focal length 2.78mm f/2.4, exposure time 1/250 s, JPEG compress: approx. 2.4 bits per RGB pixel.

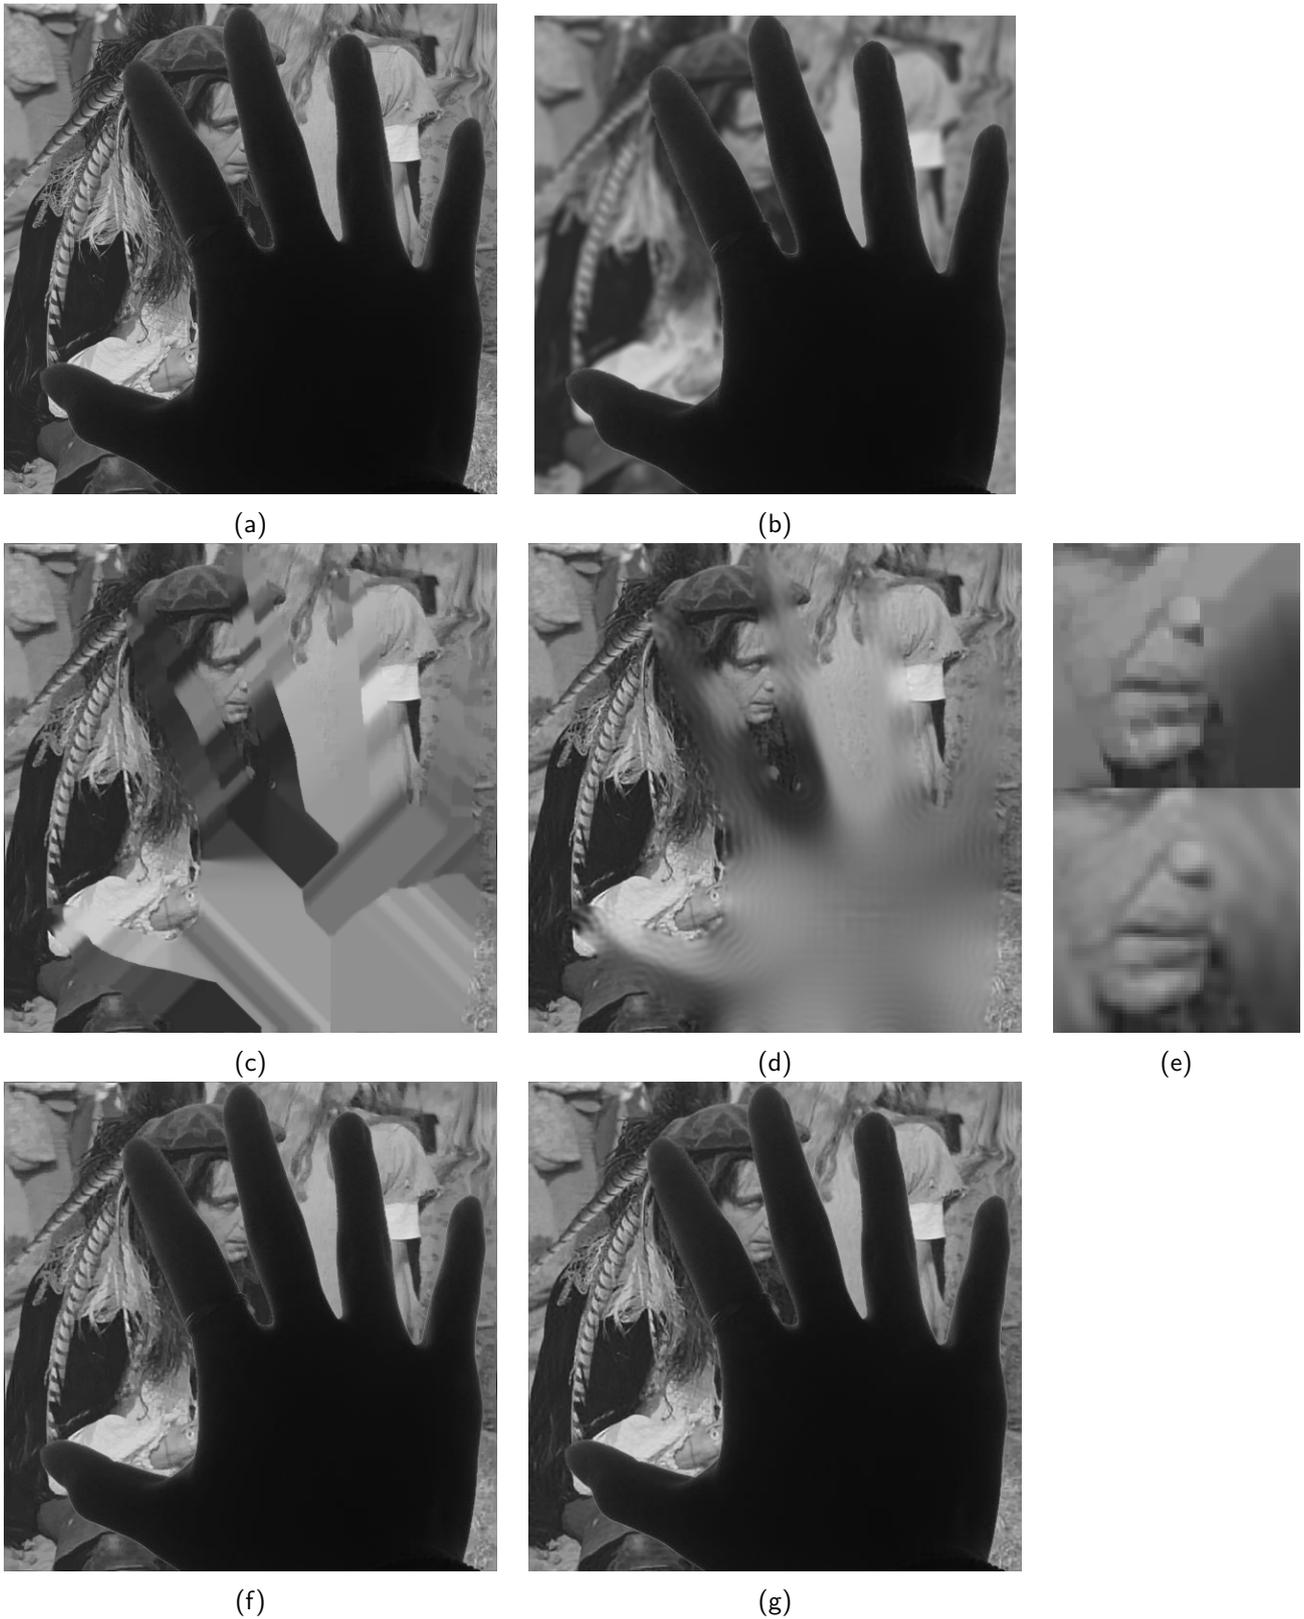


Figure 11: Visual results of complex-shape image restoration (realistic simulation). (a) Reference image (ground truth). (b) Simulated observation. (c) UBC-ADMM result. (d) Restored SPA result. (e) Detail of (c) (top) and (d) (bottom). (f) Final masked UBC-ADMM result. (g) Final masked SPA result.

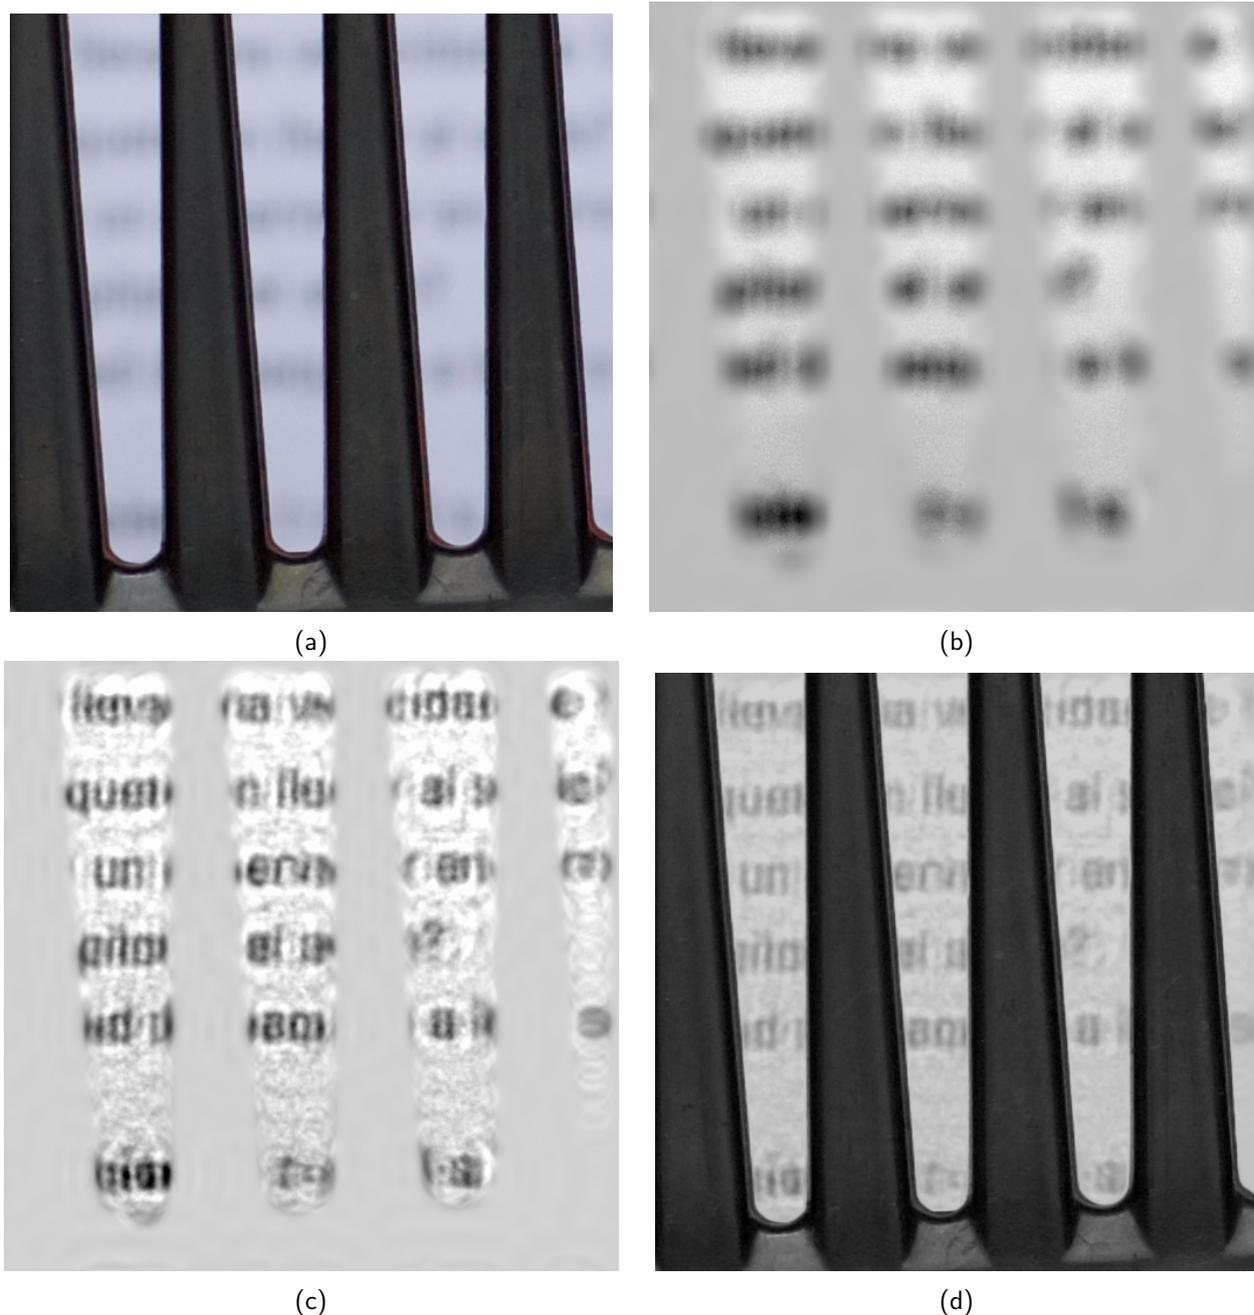


Figure 12: Restoring a high-quality photography. (a) Observation. (b) Background SPA result (scaled). (c) Restored background (scaled). (d) Final result.

## 4 Conclusions

We have proposed spectral pre-adaptation (SPA) as a new pre-processing method, which allows to apply any standard deconvolution method to observations with complex blur that goes beyond the standard convolution model. SPA uses a simple Gaussian spectral model and obtains a maximum likelihood estimate, given a subset of the observed pixels, of an hypothetical observation free from model discrepancies. SPA achieves this by substituting/adding interpolated pixels, through solving a large spatially-variant linear system of equations, by means of conjugate gradient optimization. By introducing this spectral pre-processing, we narrow the gap between the real observation and the blur model assumed in standard deconvolution methods. This makes restoration more robust



Figure 13: Restoring a JPEG-compressed medium quality photography. Left: observation. Right: restoration result using SPA.

against artifacts. After applying SPA, the user is free to choose any deconvolution method based on a standard blur model (shift-invariant, alias-free, based on a rectangular observation with circular boundary conditions), to which our pre-processing does not add significant computational overhead (note that state-of-the-art deconvolution methods are iterative, complex and non-linear, whereas SPA is linear and quite simple). Finally, usually one may want to post-process the deconvolution result, to recompose some areas of the image, by including previously discarded pixels (e.g. non-blurred pixels), and discarding some added ones (typically along the boundaries).

We have demonstrated the usefulness of the algorithm with several real-life applications by conducting experiments that deal with observations having non-circular boundary condition, aliasing, kernel mismatch, slightly shift-variant kernels and complex-shape blurred areas. Results indicate that SPA, when combined with efficient deconvolution methods, greatly mitigates artifacts with lower computational cost compared to state-of-the-art non-standard restoration methods. More importantly, for some experiments we have also tested SPA on real images and the results are as good as on simulations. Our restoration recovers heavily blurred texts in the background, becoming now partially readable. One experiment shows that this even applies to photos that are not taken with professional cameras but by inexpensive commercial cellphones, even subject to JPEG-compression. This fact suggests that the proposed algorithm has a large potential in its application to a broad set of real-life captured images.

We conclude that, when doing blurry image pre-processing prior to standard image deconvolution, preserving the observed image spectral behavior under CBC, and not merely its *smoothness* (as done traditionally) is a key factor for improving the restoration quality of real/realistic images. This is efficiently and robustly achieved by the spectral pre-adaptation method proposed here.

## Acknowledgment

This work has been funded by the Spanish Government grants FIS2016-75891-P, PID2020-118071GB-I00 and PID2020-113596GB-I00, the Czech Science Foundation grant GA20-27939S and the Praemium Academiae awarded by the Czech Academy of Sciences.

## Image Credits



Signal and Image Processing Institute, University of Southern California<sup>14</sup>



New York University Tandon School of Engineering<sup>15</sup>



Standard test image



Standard test image<sup>16</sup>

Other images shown in the paper belong to the authors.

## References

- [1] M.S.C. ALMEIDA AND M. FIGUEIREDO, *Deconvolving images with unknown boundaries using the alternating direction method of multipliers*, IEEE Transactions on Image Processing, 22 (2013), pp. 3074–3086. <https://doi.org/10.1109/TIP.2013.2258354>.
- [2] L. DENIS, E. THIÉBAUT, F. SOULEZ, J-M. BECKER, AND R. MOURYA, *Fast approximations of shift-variant blur*, International Journal of Computer Vision, 115 (2015), pp. 253–278. <https://doi.org/10.1007/s11263-015-0817-x>.
- [3] C. DONG, J. BARBERO, AND J. PORTILLA, *Nonconvex Bayesian restoration of blurred foreground images*, in IEEE International Conference on Image Processing (ICIP), 2019, pp. 754–758. <https://doi.org/10.1109/ICIP.2019.8803771>.
- [4] C. DONG AND J. PORTILLA, *Maximum likelihood interpolation for aliasing-aware image restoration*, in IEEE International Conference on Image Processing (ICIP), 2016, pp. 564–568. <https://doi.org/10.1109/ICIP.2016.7532420>.
- [5] —, *Spectral pre-adaptation for two-step arbitrary-shape-support image restoration*, in IEEE International Conference on Image Processing (ICIP), 2017, pp. 3515–3519. <https://doi.org/10.1109/ICIP.2017.8296936>.
- [6] D. GLASNER, S. BAGON, AND M. IRANI, *Super-resolution from a single image*, in IEEE International Conference on Computer Vision (ICCV), 2009. <https://doi.org/10.1109/ICCV.2009.5459271>.
- [7] A. GOLDSTEIN AND R. FATTAL, *Blur-kernel estimation from spectral irregularities*, in European Conference on Computer Vision (ECCV), Springer, 2012, pp. 622–635.
- [8] F.R. HAMPEL, *The influence curve and its role in robust estimation*, Journal of the American Statistical Association, 69 (1974), pp. 383–393. <https://www.jstor.org/2285666>.
- [9] S. IHARA, *Information Theory for Continuous Systems*, World Scientific, 1993. <https://doi.org/10.1142/1676>.
- [10] J. KOTERA, V. SMIDL, AND F. SROUBEK, *Blind deconvolution with model discrepancies*, IEEE Transactions on Image Processing, 26 (2017), pp. 2533–2544. <https://doi.org/10.1109/TIP.2017.2676981>.

<sup>14</sup><http://sipi.usc.edu/database/download.php?vol=textures&img=1.1.03>

<sup>15</sup><https://eeweb.engineering.nyu.edu/~yao/EL5123/SampleData.html>

<sup>16</sup>[http://www.imageprocessingplace.com/DIP-3E/dip3e\\_book\\_images\\_downloads.htm](http://www.imageprocessingplace.com/DIP-3E/dip3e_book_images_downloads.htm)

- [11] R. LIU AND J. JIA, *Reducing boundary artifacts in image deconvolution*, in IEEE International Conference on Image Processing (ICIP), 2008, pp. 505–508. <https://doi.org/10.1109/ICIP.2008.4711802>.
- [12] A. MATAKOS, S. RAMANI, AND J.A. FESSLER, *Image restoration using non-circulant shift-invariant system models*, in IEEE International Conference on Image Processing, 2012, pp. 3061–3064. <https://doi.org/10.1109/ICIP.2012.6467546>.
- [13] C.B. MEHR AND J.A. MCFADDEN, *Certain properties of Gaussian processes and their first-passage times*, Journal of the Royal Statistical Society. Series B (Methodological), 27 (1965), pp. 505–522. <https://www.jstor.org/stable/2345801>.
- [14] D. MIRAUT AND J. PORTILLA, *Efficient shift-variant image restoration using deformable filtering (Part I)*., EURASIP Journal of Advanced Signal Processing, 2012 (2012), p. 100. <https://doi.org/10.1186%2F1687-6180-2012-100>.
- [15] J. PORTILLA, *Maximum likelihood extension for non-circulant deconvolution*, in IEEE International Conference on Image Processing (ICIP), 2014, pp. 4276–4279. <https://doi.org/10.1109/ICIP.2014.7025868>.
- [16] J. PORTILLA, A. TRISTÁN-VEGA, AND I.W. SELESNICK, *Efficient and robust image restoration using multiple-feature L2-relaxed sparse analysis priors*, IEEE Transactions on Image Processing, 24 (2015), pp. 5046–5059. <https://doi.org/10.1109/TIP.2015.2478405>.
- [17] S.J. REEVES, *Fast image restoration without boundary artifacts*, IEEE Transactions on Image Processing, 14 (2005), pp. 1448–1453. <https://doi.org/10.1109/TIP.2005.854474>.
- [18] M. SOREL, *Removing boundary artifacts for real-time iterated shrinkage deconvolution*, IEEE Transactions on Image Processing, 21 (2012), pp. 2329–2334. <https://doi.org/10.1109/TIP.2011.2176344>.
- [19] K-C. TAN, H. LIM, AND B.T.G. TAN, *Windowing techniques for image restoration*, CVGIP: Graphical Models and Image Processing, 53 (1991), p. 491–500. [https://doi.org/10.1016/1049-9652\(91\)90033-G](https://doi.org/10.1016/1049-9652(91)90033-G).
- [20] D. TZIKAS, A. LIKAS, AND N.P. GALATSANOS, *Variational Bayesian Sparse Kernel-Based Blind Image Deconvolution With Student’s-t Priors*, IEEE Transactions on Image Processing, 18 (2009), pp. 753–764. <https://doi.org/10.1109/TIP.2008.2011757>.
- [21] L. XU AND J. JIA, *Two-phase kernel estimation for robust motion deblurring*, in European Conference on Computer Vision (ECCV), Springer, 2010, pp. 157–170.
- [22] N. XU, B. PRICE, S. COHEN, AND T. HUANG, *Deep image matting*, in IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 2970–2979. <https://doi.org/10.48550/arXiv.1703.03872>.
- [23] X. XU, J. PAN, Y. ZHANG, AND M. YANG, *Motion blur kernel estimation via deep learning*, IEEE Transactions on Image Processing, 27 (2017), pp. 194–205. <https://doi.org/10.1109/TIP.2017.2753658>.
- [24] M. YAMAC, B. ATAMAN, AND A. NAWAZ, *Kernelnet: A blind super-resolution kernel estimation network*, in IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 453–462. <https://doi.org/10.1109/CVPRW53098.2021.00056>.

- [25] K. ZHANG, X. GAO, D. TAO, AND X. LI, *Single image super-resolution with non-local means and steering kernel regression*, IEEE Transactions on Image Processing, 21 (2012), pp. 4544–4556. <https://doi.org/10.1109/TIP.2012.2208977>.