

Optimal Filters from Calibration Data for Image Deconvolution with Data Acquisition Error

Published in Springer's *Journal of Mathematical Imaging and Vision* 2011

Julianne Chung · Matthias Chung · Dianne P. O'Leary

Received: date / Accepted: date

Abstract Data acquisition errors due to dead pixels or other hardware defects can cause undesirable artifacts in imaging applications. Compensating for these defects typically requires knowledge such as a defective pixel map, which can be difficult or costly to obtain and which is not necessarily static. However, recent calibration data is readily available in many applications. In this paper, we compute optimal filters for image deconvolution with denoising using only this calibration data, by minimizing the empirical Bayes risk. We derive a bound on how the reconstruction changes as the number of dead pixels grows. We show that our approach is able to reconstruct missing information better than standard filtering approaches and is robust even in the presence of a large number of defects and to defects that arise after calibration.

Keywords inverse problems · deconvolution · dead pixels · spectral filtering · missing data artifacts · optimal filter · defective pixel · hot pixel · camera artifact · CCD defect · data acquisition error · MRI artifact · CT scan

artifact · empirical risk · Bayes risk · Bayesian risk · machine learning.

1 Introduction

Missing data artifacts are common in imaging systems such as MRI [19], CT scans [1], and charge-coupled-device (CCD) detectors. For definiteness, we discuss CCD cameras, but our techniques apply more generally. Common CCD defects include *dead pixels* or *hot pixels*, whose response is independent of light input; *dead columns*, where an entire column of pixels are dead [20]; and *pixel traps*, where a defective pixel causes loss of information in all pixels above it in a column [14, 22].

Defects in CCDs become more prevalent as detectors age. If an accurate map of defective pixels is available, as in Figure 1, digital correction techniques such as image inpainting or interpolation can be used to pre-process the data and “fill-in” the missing information [2, 3, 5]. However, obtaining an accurate map may require tedious visual inspection of calibration data. In applications where images are processed without manual intervention, artifacts in the reconstructed images may be missed by the scientists who use the results.

In this paper, we assume that no defective pixel map is available. However, we do rely on recent calibration data, which can be gathered in many applications. As in our previous work [7], we use this calibration data to compute optimal filters for image deconvolution. Our approach uses an empirical Bayes risk minimization framework to learn the missing information. We adopt the assumption, common in the machine learning literature, that the calibration data are samples drawn from

J. Chung
Department of Mathematics
University of Texas at Arlington
Arlington, TX, 76019 USA
E-mail: chungj@uta.edu

M. Chung
Department of Mathematics
Texas State University
San Marcos, TX, 78666 USA
E-mail: mc85@txstate.edu

D. P. O'Leary
Department of Computer Science and UMIACS
University of Maryland
College Park, MD, 20742 USA
E-mail: oleary@cs.umd.edu

a space of images according to some probability distribution. We describe two approaches for deconvolution that can be efficiently implemented within the empirical risk framework. One is image deblurring via spectral filtering and another is a combined image deblurring and denoising approach. We show that our approach is robust even when a large number of pixels are defective or few training images are available.

In Section 2 we describe the basic problem of image deconvolution and introduce our approach to defining optimal filters for deblurring and denoising. In Section 3 we assess the effect of missing pixels on the reconstruction, deriving a bound on how the reconstruction changes as the number of dead pixels grows. Numerical results are shown in Section 4, demonstrating superior performance of our method relative to standard algorithms such as Tikhonov regularization and inpainting. In these experiments, our algorithms provide reconstructions with small error even when 50% of the pixels are defective or when additional pixels become defective after calibration. Conclusions are presented in Section 5.

2 Problem Formulation

The goal of image deconvolution is to reconstruct a clear image from one that has been blurred and obscured by noise. Mathematically, the observed blurred image $\mathbf{b} \in \mathbb{R}^n$ can be written as

$$\mathbf{b} = \mathbf{A}\boldsymbol{\xi} + \boldsymbol{\delta}, \quad (1)$$

where $\boldsymbol{\xi} \in \mathbb{R}^n$ is the unknown true image, $\mathbf{A} \in \mathbb{R}^{n \times n}$ models the blurring process, and $\boldsymbol{\delta} \in \mathbb{R}^n$ is additive noise. The blurring matrix \mathbf{A} is defined by the point spread function (PSF), which describes how a point source is blurred [15]. In this paper, we assume that the PSF is known.

The goal of image deconvolution is to reconstruct an approximation of $\boldsymbol{\xi}$, given \mathbf{b} and \mathbf{A} . The problem is ill-posed, meaning small errors in the data may lead to large errors in the solution. When n is small enough, or when the blur is spatially invariant and boundary conditions permit a fast spectral decomposition of \mathbf{A} [15], spectral filtering is often used to regularize the problem and compute stable solutions.

2.1 Spectral Filtering

Let $\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ be the singular value decomposition of \mathbf{A} , where $\boldsymbol{\Sigma}$ is a diagonal matrix containing




Fig. 1: Sample dead pixel map. At a white (dead) pixel, the observed image value is zero.

the singular values, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0$, and orthogonal matrices \mathbf{U} and \mathbf{V} contain the left and right singular vectors \mathbf{u}_i and \mathbf{v}_i , $i = 1, 2, \dots, n$, respectively. The filtered solution can be written as

$$\begin{aligned} \mathbf{x}_A(\boldsymbol{\phi}_A) &= \sum_{i=1}^n (\boldsymbol{\phi}_A)_i \frac{\mathbf{u}_i^\top \mathbf{b}}{\sigma_i} \mathbf{v}_i, \\ &= \mathbf{V} \text{diag}(\mathbf{U}^\top \mathbf{b}) \boldsymbol{\Sigma}^{-1} \boldsymbol{\phi}_A, \end{aligned} \quad (2)$$

where $\boldsymbol{\phi}_A$ is the vector of spectral filter factors $(\boldsymbol{\phi}_A)_i$. The choice of filter factors is crucial to the reconstruction. Standard filters typically use a particular functional representation for the filter factors. For example, Tikhonov filter factors can be written as $(\boldsymbol{\phi}_A)_i = \sigma_i^2 / (\sigma_i^2 + \alpha^2)$, where α is a regularization parameter, chosen to make the filter factors corresponding to large singular values close to 1 and filter factors corresponding to small singular values close to 0. In contrast, building on our work in [7], our approach allows each filter factor to be selected independently, without assuming a particular functional form, so we treat each of the n values $(\boldsymbol{\phi}_A)_i$ as a separate parameter.¹

Ideally, we would use image-dependent filter factors that minimize the error between the true image and the reconstructed image. However, the true image is not known in practice, so we choose to learn good filter factors from a set of calibration or training images [10], using a Bayes formulation. As in [7], we assume that the operator \mathbf{A} is deterministic, not subject to noise, that the true images (both calibration and test) are chosen according to

¹ The extension to a matrix \mathbf{A} that has more rows than columns, or that has rank less than n , is straightforward. In that case, replace every $\boldsymbol{\Sigma}^{-1}$ by the pseudo-inverse $\boldsymbol{\Sigma}^\dagger$ and every summation from $i = 1$ to n by a summation $i = 1$ to $\text{rank}(\mathbf{A})$, so that the number of filter factors is reduced to $\text{rank}(\mathbf{A})$.

a probability distribution with finite second moments, and that we have a probability distribution for the noise with zero mean and finite second moments. Let $\boldsymbol{\xi}^{(k)}$, $k = 1, 2, \dots, N$, be a set of training (calibration) images, with corresponding observed (noisy) blurred images $\mathbf{b}^{(k)}$. Then we compute the *optimal p -norm error filter* [7] as the solution to the *empirical Bayes risk minimization problem*

$$\hat{\phi}_A = \arg \min_{\phi_A} \frac{1}{pnN} \sum_{k=1}^N \|\mathbf{x}_A^{(k)}(\phi_A) - \boldsymbol{\xi}^{(k)}\|_p^p, \quad (3)$$

where $\|\cdot\|_p$ is the p -norm and $\mathbf{x}_A^{(k)}$ is the filtered solution (2) corresponding to $\mathbf{b}^{(k)}$. For $p = 2$, Equation (3) reduces to solving a linear problem, since the n parameters $(\phi_A)_i$ appear linearly in (2), but for all other p , we solve this optimization problem using a Gauss-Newton optimization method with Armijo line search [17].

2.2 A Combined Approach: Deblurring and Denoising

It has been observed in the image processing literature [16, 18] that spectral filtering followed by image denoising is an effective approach for the deconvolution of images. The basic idea is to first apply a spectral filter to deblur the image and then use a second low-pass filter in some other frequency-related basis to denoise the resulting image. Finding the right balance of deblurring and denoising is key to the success of this approach. In this section, we use calibration data to determine the optimal combination of deblurring and denoising, by employing a variable projection approach within an empirical Bayes risk framework.

Let $\mathbf{W} \in \mathbb{R}^{n \times n}$ represent any orthogonal discrete 2-D frequency transform. For example, \mathbf{W} and its inverse operation, \mathbf{W}^\top , could be defined using a wavelet, Fourier, or cosine transform.

Then the filtered solution can be written as

$$\mathbf{x}_{AW}(\phi_A, \phi_W) = \mathbf{W}^\top \text{diag}(\phi_W) \mathbf{W} \mathbf{x}_A(\phi_A), \quad (4)$$

where ϕ_A is a spectral filter for deblurring, ϕ_W is a frequency filter for denoising, and $\mathbf{x}_A(\phi_A)$ is defined in (2). The optimal filters $\hat{\phi}_A, \hat{\phi}_W$ are computed as solutions to the joint optimization problem

$$\arg \min_{\phi_A, \phi_W} f(\phi_A, \phi_W), \quad (5)$$

where

$$f(\phi_A, \phi_W) = \frac{1}{pnN} \sum_{k=1}^N \|\mathbf{x}_{AW}^{(k)}(\phi_A, \phi_W) - \boldsymbol{\xi}^{(k)}\|_p^p, \quad (6)$$

and $\mathbf{x}_{AW}^{(k)}$ is the filtered solution (4) corresponding to $\mathbf{b}^{(k)}$.

Computing solutions to (5) can be numerically difficult due to nonlinearities. However, for the case $p = 2$, we can exploit the separability of the problem and use a variable projection approach [12, 13]. That is, we solve the reduced optimization problem

$$\min_{\phi_A} f(\phi_A, \hat{\phi}_W(\phi_A)), \quad (7)$$

where

$$\hat{\phi}_W(\phi_A) = \arg \min_{\phi_W} f(\phi_A, \phi_W). \quad (8)$$

Notice that for fixed ϕ_A , the filtered solution

$$\mathbf{x}_{AW}(\phi_A, \phi_W) = \mathbf{W}^\top \text{diag}(\mathbf{W} \mathbf{x}_F(\phi_A)) \phi_W,$$

is linear in ϕ_W . Thus, optimization problem (8) for $p = 2$ has a closed form expression for the solution. Problem (7) can then be solved by Newton's method, or a variant. Every time a function value is needed, (8) is solved, and derivatives, if needed, are computed using the chain rule. Using variable projection reduces the number of variables in the optimization problem and tends to reduce the number of local minimizers, which can make it easier to find a global minimizer.

Furthermore, if Tikhonov regularization is used for deblurring, then ϕ_A is determined by a single parameter α , so that (7) reduces to a one-parameter optimization problem.

3 Effect of Missing Pixels on the Reconstruction

In this section, we compute error bounds for reconstructions when missing pixels are present in the data. We first compare the reconstruction for a blurred image with missing pixels to that of the blurred image without missing pixels. We show that, for $p = 2$, the mean squared difference in reconstructions can be bounded linearly in the number of missing pixels.

The filtered solution corresponding to no missing pixels can be represented as $\mathbf{x}_A = \mathbf{V} \boldsymbol{\Phi}_A \boldsymbol{\Sigma}^{-1} \mathbf{U}^\top \mathbf{b}$, where \mathbf{b} is the blurred image in (1) and $\boldsymbol{\Phi}_A$ is a

diagonal matrix with filter factors $(\phi_A)_i$ on the diagonal. Let \mathcal{M} be the set of indices corresponding to missing pixels in the image, and let \mathbf{M} be the identity matrix with ones replaced by zeros for missing pixels. Then the observed image with missing pixels can be written as

$$\tilde{\mathbf{b}} = \mathbf{M}\mathbf{b},$$

and the corresponding filtered solution is given by

$$\tilde{\mathbf{x}}_A = \mathbf{V}\Phi_A\Sigma^{-1}\mathbf{U}^\top\tilde{\mathbf{b}}. \quad (9)$$

Assume that the observed image has been normalized so that $0 \leq b_i \leq 1$. Then, using the fact that

$$\mathbf{M} = \mathbf{I} - \sum_{i \in \mathcal{M}} \mathbf{e}_i \mathbf{e}_i^\top,$$

where \mathbf{e}_i is the i -th column of the identity matrix, we get a bound for the mean squared difference in reconstructions,

$$\begin{aligned} \|\mathbf{x}_A - \tilde{\mathbf{x}}_A\|_2^2 &= \|\mathbf{V}\Phi_A\Sigma^{-1}\mathbf{U}^\top \left(\sum_{i \in \mathcal{M}} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{b}\|_2^2 \\ &\leq \gamma_A^2 \sum_{i \in \mathcal{M}} b_i^2 \\ &\leq \gamma_A^2 |\mathcal{M}|, \end{aligned}$$

where

$$\gamma_A^2 = \max_i \left(\frac{(\phi_A)_i}{\sigma_i} \right)^2.$$

Similarly, for the combined approach, we have

$$\tilde{\mathbf{x}}_{AW} = \mathbf{W}^\top \Phi_W \mathbf{W} \tilde{\mathbf{x}}_A. \quad (10)$$

and

$$\begin{aligned} \|\mathbf{x}_{AW} - \tilde{\mathbf{x}}_{AW}\|_2^2 &= \|\mathbf{W}^\top \Phi_W \mathbf{W} \mathbf{V} \Phi_A \Sigma^{-1} \mathbf{U}^\top \left(\sum_{i \in \mathcal{M}} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{b}\|_2^2 \\ &\leq \|\Phi_W \mathbf{W} \mathbf{V} \Phi_A \Sigma^{-1}\|_2^2 \cdot \|\mathbf{U}^\top \left(\sum_{i \in \mathcal{M}} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{b}\|_2^2 \\ &\leq \|\Phi_W\|_2^2 \cdot \|\Phi_A \Sigma^{-1}\|_2^2 \sum_{i \in \mathcal{M}} b_i^2 \\ &\leq \gamma_{AW}^2 |\mathcal{M}|, \end{aligned}$$

where

$$\gamma_{AW}^2 = \gamma_A^2 \cdot \max_i ((\phi_W)_i)^2.$$

We have shown that the squared difference can be bounded linearly in the number of missing pixels. Using this bound and the fact that

$$\begin{aligned} \|\boldsymbol{\xi} - \mathbf{x}_A\|_2 &= \|\boldsymbol{\xi} - \mathbf{V}\Phi_A\Sigma^{-1}\mathbf{U}^\top\mathbf{b}\|_2 \\ &= \|(\Phi_A - \mathbf{I})\mathbf{V}^\top\boldsymbol{\xi} + \Phi_A\Sigma^{-1}\mathbf{U}^\top\boldsymbol{\delta}\|_2 \\ &\leq \|(\Phi_A - \mathbf{I})\mathbf{V}^\top\boldsymbol{\xi}\|_2 + \gamma_A\|\boldsymbol{\delta}\|_2, \end{aligned}$$

the total reconstruction error can be bounded as

$$\begin{aligned} \|\boldsymbol{\xi} - \tilde{\mathbf{x}}_A\|_2 &\leq \|\boldsymbol{\xi} - \mathbf{x}_A\|_2 + \|\mathbf{x}_A - \tilde{\mathbf{x}}_A\|_2 \\ &\leq \|\mathbf{e}_A\|_2 + \gamma_A(\|\boldsymbol{\delta}\|_2 + |\mathcal{M}|^{1/2}), \end{aligned}$$

where $\|\mathbf{e}_A\| = \|(\Phi_A - \mathbf{I})\mathbf{V}^\top\boldsymbol{\xi}\|$ is the regularization error, obtained by using regularized inverse matrix $\mathbf{V}\Phi_A\Sigma^{-1}\mathbf{U}^\top$. The second term corresponds to perturbation errors. Similarly, we have for the AW filters,

$$\begin{aligned} \|\boldsymbol{\xi} - \tilde{\mathbf{x}}_{AW}\|_2 &\leq \|\boldsymbol{\xi} - \mathbf{x}_{AW}\|_2 + \|\mathbf{x}_{AW} - \tilde{\mathbf{x}}_{AW}\|_2 \\ &\leq \|\mathbf{e}_{AW}\|_2 + \gamma_{AW}(\|\boldsymbol{\delta}\|_2 + |\mathcal{M}|^{1/2}), \end{aligned}$$

where $\mathbf{e}_{AW} = (\mathbf{W}^\top \Phi_W \mathbf{W} \mathbf{V} \Phi_A \mathbf{V}^\top - \mathbf{I})\boldsymbol{\xi}$. These results are summarized in the following theorem.

Theorem 1 *Assume $\mathbf{A}, \boldsymbol{\xi}$, and $\tilde{\mathbf{b}}$ are given. Let $\tilde{\mathbf{x}}_A$ and $\tilde{\mathbf{x}}_{AW}$ be filtered solutions defined in Equations (9) and (10) respectively. Then the reconstruction errors are bounded in terms of the square root of the number of missing pixels $|\mathcal{M}|$, i.e.,*

$$\|\boldsymbol{\xi} - \tilde{\mathbf{x}}_A\|_2 \leq \|\mathbf{e}_A\|_2 + \gamma_A(\|\boldsymbol{\delta}\|_2 + |\mathcal{M}|^{1/2})$$

and

$$\|\boldsymbol{\xi} - \tilde{\mathbf{x}}_{AW}\|_2 \leq \|\mathbf{e}_{AW}\|_2 + \gamma_{AW}(\|\boldsymbol{\delta}\|_2 + |\mathcal{M}|^{1/2}).$$

In this section, we have computed error bounds for the reconstructions, when missing pixels are present in the data. Previous results for inpainting [4] do not include the blurring operator and only consider errors in the inpainting region.

4 Numerical Results

In this section, numerical results illustrate the performance of the optimal filters for image deconvolution, in the presence of missing data artifacts. Experiments were conducted to address the following questions: First, how do the optimal filters compare to standard methods for deconvolution (Section 4.1)? Second, how do our learning methods perform as the number of dead pixels increases, and how do they compare to standard inpainting methods (Section 4.2)? Third, how sensitive are the optimal filters to additional dead pixels, which have not been learned by the empirical Bayes framework (Section 4.3)? Lastly, how many training images are needed to get reasonable results (Section 4.4)?

4.1 Comparing Optimal Filters to Tikhonov Filters

Satellite images from NASA were used in the experiments. Each image was 256×256 . Forty training images were obtained by randomly rotating, translating, and zooming 8 NASA images 5 times each. We defined \mathbf{A} using a spatially-invariant symmetric Gaussian point spread function with mean $\mathbf{0}$ and covariance matrix equal to the identity ². More precisely, a blurred image $\mathbf{b}^{(k)}$ was produced from a true image $\boldsymbol{\xi}^{(k)}$ by

$$\mathbf{b}^{(k)} = \mathbf{M}(\mathbf{A}\boldsymbol{\xi}^{(k)} + \boldsymbol{\delta}^{(k)}),$$

where the noise $\boldsymbol{\delta}^{(k)}$ is an independent and identically distributed sample from a normal distribution with mean zero and covariance a multiple of the identity matrix, with the multiple uniformly sampled from $[0.05, 0.1]$. Other NASA satellite images were processed in the same way for validation.

In this experiment, training data is used to compute optimal filters from Equation (3) with $p = 2$ and $p = 10$. We refer to these as the *A2* and *A10* filters respectively. The *A2* filter corresponds to a widely used measure of the error, namely, the mean squared error. It has been shown that the optimal *A2* filter numerically approximates the Wiener filter [7]. The 10-norm was selected (as a proxy for the infinity or max norm) so that large errors, such those arising at dead pixel locations, are strongly penalized and small errors, such as background noise, are weakly penalized. The solution changes slowly for larger values of p , but finding the optimal parameters is a more difficult optimization problem.

We also computed the optimal filters from Equation (5) with $p = 2$, and we refer to them as *AW2* filters. Matlab’s `fminbnd` routine was used for the 1D optimization, and \mathbf{W} represented a single level discrete 2-D Haar wavelet decomposition.

The missing pixel map consisted of 1310 randomly distributed dead pixels (accounting for approximately 2% of the image), a dead column (column 156), and a pixel trap (in column 129). Figure 1 illustrates the dead pixel map used in our experiment, and the top row of Figure 2 shows the two validation images.

For each of the validation images, we compare our optimal filters to the Tik-GCV filter, which is

² Any type of blurring operator \mathbf{A} (spatially variant or invariant) may be considered, with the only restriction that the singular value decomposition should be available.

the Tikhonov filter obtained by using the generalized-cross-validation (GCV) method [11] to choose the regularization parameter α . We also compare to the Tik-MSE filter, where α is chosen to minimize the mean square error. Tik-MSE is not a practical method, since it requires knowledge of the true validation images, but it indicates the best performance that Tikhonov filtering could obtain. It is important to note that optimal filters are computed off-line, while the Tikhonov parameter is typically re-estimated for each image.

To illustrate the performance of the filters in the presence of dead pixels, we show in Figure 2 the absolute error images in inverted colormap (white corresponding to zero). All of the optimal filter approaches were able to compensate for the dead pixels better than the Tikhonov approaches. Dark spots and vertical lines in the error images for Tikhonov indicate large errors at the dead pixel locations.

Figure 3 shows parts of columns 129 and 156 in the restored images. The plots clearly show that Tikhonov-based methods have difficulty reconstructing dead pixels, while the optimal filters do a surprisingly good job. In particular, the *AW2* filter seems to reconstruct the missing information very accurately.

In order to investigate how different our computed filters are from standard ones, we display them as images. Figures 4(a)–(c) display the Tik-MSE filter (for the second validation image), the *A2* filter and the *A10* filter. The filter factors are arranged in a two-dimensional format, using the fact that the Gaussian blurring function is separable to provide a natural layout from largest singular value in the top left corner to smallest in the bottom right, with decreasing singular values as we go down each column or across each row. The color of the pixel for a particular singular value σ_i indicates the magnitude of its filter factor $(\phi_A)_i$. These filters are visibly different. For the *AW2* filter, $\widehat{\phi}_A$ is similar to the filter in Figure 4(a) (image not shown). The corresponding filter $\widehat{\phi}_W$ is computed for the wavelet basis and is shown in Figure 4(d). The filter is arranged such that the top left block corresponds to the scaling coefficients, and the other blocks correspond to the horizontal, vertical, and diagonal orientations [9].

We notice that the dead column and pixel trap appear in the filter factors for the scaling coefficients and the vertical orientation. A reconstruction that provides a good approximation of dead columns and pixel traps can be obtained using




Fig. 2: Validation images (top row), and absolute error images (inverted colormap: white corresponds to zero) for reconstructions using optimal filters (A2, A10, and AW2) and Tikhonov filters.

only these filter factors (with all others set to one). However, obtaining a good reconstruction that also compensates for random missing pixels requires the full denoising filter.

4.2 Sensitivity to the Number of Dead Pixels

In the next experiment, we investigate the sensitivity of the reconstructions to the number of dead pixels in the observation. We consider randomly positioned dead pixels with 1% to 50% of the image missing. Here, we do not include pixel traps or dead columns.

We again use 40 training images, but we evaluate the filters using 90 different validation images. These validation images were obtained by taking 9 different satellite images and performing 10 random rotation, translation, and zoom transformations of each, and then applying blurring followed by the dead pixel mask. For each validation image, we record the MSE between the true and reconstructed image for the A2 and the Tik-GCV filters.

In this experiment, we also compare the performance of the optimal filters to methods using image inpainting prior to deblurring. A variety of inpainting techniques have been proposed [5, 6, 8], and in this paper, we consider two of these methods, one that does not require a dead pixel map and a second approach that does. In the first method, for each of the validation images, we compute a reconstruction using a 3×3 median filter for image inpainting followed by Tik-GCV for deblurring. In the second method, we compute a reconstruction using total variation (TV) for inpainting followed by Tik-GCV for deblurring. For TV inpainting, software from [8] was used, and the correct dead pixel map was applied.

In Figure 5(a) we provide box plots for the reconstruction errors from the A2 filter, corresponding to different percentages of missing pixels. The box shows the 25th and 75th quartiles, with a circle at the median. The whiskers extend to extreme values, and outliers are plotted individually. The general trend is that, as expected, the error increases with increasing numbers of dead pixels. Similarly, box plots for reconstruction errors from the AW2 filter for different percentages of missing pixels is presented in Figure 5(b). Figure 5(c) presents median reconstruction errors for all five reconstruction approaches on the same scale.

It is remarkable to see that reconstructions with optimal filters result in smaller median errors than

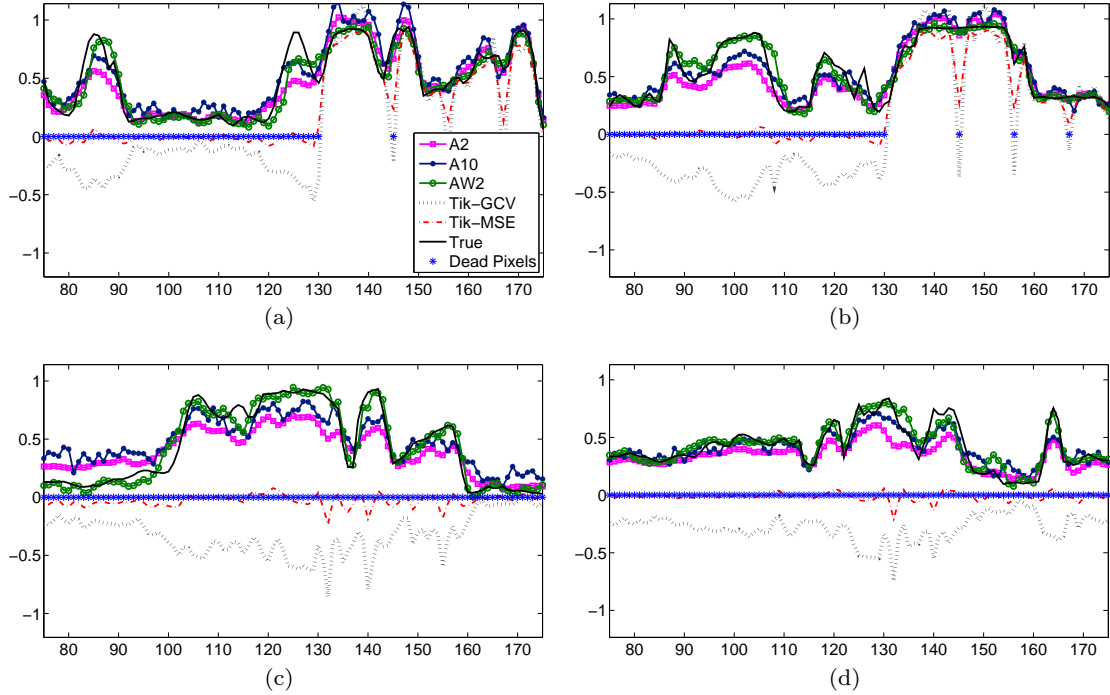


Fig. 3: Detail of reconstructions for the two validation images (left and right). Plots (a) and (b) show pixels 75 to 175 in column 129, containing a pixel trap and random dead pixels. Plots (c) and (d) show the same pixels in column 156, a dead column. The horizontal axis denotes the row of the pixel within the column. The solid curve plots the true pixel values, and the blue stars identify dead pixels.






Fig. 4: (a) Tikhonov filter (for the second validation image) with regularization parameter computed using the MSE criterion. (b) A2 filter. (c) A10 filter. (d) Denoising filter in the wavelet domain for the AW2 filter.

total variation inpainting followed by Tik-GCV. Total variation inpainting requires an accurate dead pixel map, while optimal filters just require training data, which is often available. Furthermore, since optimal filters can be computed off-line, computing reconstructions for optimal filters can be

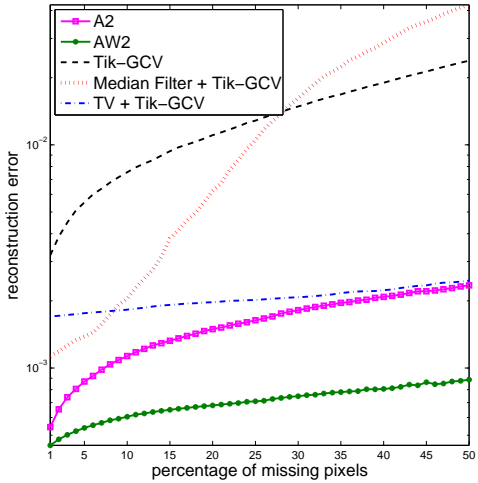
significantly faster than inpainting followed by deconvolution.



(a)



(b)



(c)

Fig. 5: Figure (a) and (b) contain box plots for reconstruction error for varying percentages of missing pixels for A2 and AW2 filters respectively. (c) Median reconstruction errors for A2, AW2, and Tik-GCV filters at varying percentages of missing pixels, as well as median reconstruction errors for 3×3 median filter followed by Tik-GCV and total variation inpainting followed by Tik-GCV. The solid lines with squares and dots in (c) correspond to median values in (a) and (b) respectively.

4.3 Sensitivity to Unlearned Dead Pixels

Next we consider the sensitivity of the optimal filters to unlearned missing pixels. That is, we compute the A2 and AW2 filters for 40 training im-

ages with a dead pixel map of 10 percent randomly generated missing pixels, and we tested these optimal filters on 90 validation images with additional dead pixels. More specifically, we considered an additional 0–39 percent of missing pixels in the validation images. For each additional percentage, we record the MSE between the true and reconstructed image for each of the 90 validation images and present the median of the reconstruction errors in Figure 6. Since an additional percentage of missing pixels of 0 corresponds to using 10 percent missing pixels for both the training and validation data, the data points at 0 on the x-axis in Figure 6 are the same as the data points at 10 on the x-axis in Figure 5(c).

Figure 6 shows that the optimal filters are still useful when there are additional dead pixels which are not present in the training data. Furthermore, the desired level of accuracy in reconstruction can inform the user when it may be necessary to update the training data and compute new optimal filters. For example, if the desired median accuracy is 0.01 (see dashed line), then optimal filters should be re-computed after an additional 34 percent of missing pixels for A2 and an additional 30 percent for AW2. An interesting phenomena is that for up to an additional 5 percent of missing pixels, the AW2 filter produces smaller median errors than the A2 filter.

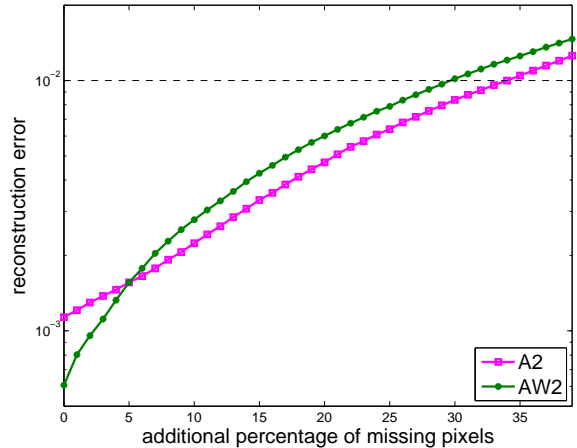


Fig. 6: Median reconstruction errors for A2 and AW2 filters for increasing percentages of unlearned missing pixels in the validation data.

4.4 Sensitivity to the Number of Training Images

In the next experiment, we investigate the sensitivity of the reconstructions to the number of training images. We compute the A2 and AW2 filters on 1 to 60 training images, and test the filters on a set of 50 validation images. The problem setup is described in Section 4.1, and the missing pixel map in Figure 1 is used. The median reconstruction errors for the validation images are found in Figure 7. This plot is often referred to as the Pareto curve.

The expected trend is that for an increasing number of training images, we get better reconstructions on average for the validation data [21, Chap. 3]. Both the A2 and AW2 filters display this trend. For the AW2 filter, additional improvement in reconstruction error is minimal after approximately 16 training images. However, for the A2 filter, 6 images may be sufficient.




Fig. 7: Pareto curve of median reconstruction errors for A2 and AW2 filters.

5 Conclusion

We have described an approach that uses calibration data to determine optimal filters for image deconvolution, and we have shown that the resulting reconstructions have high quality, even when the number of defective pixels is large or changing. Training or calibration data are available in a variety of applications, making the method widely (but not universally) applicable. A key advantage of this approach is that reconstructions can be done efficiently off-line and without manual intervention or knowledge of a dead pixel map.

Acknowledgement

We are grateful to the referees for helpful suggestions. The work of JC was partially supported by NSF grant DMS 0902322. The work of DPO was partially supported by NSF grant DMS 1016266. The satellite images were obtained from NASA's website: www.nasa.gov.

References

1. Barrett, J.F., Keat, N.: Artifacts in CT: Recognition and avoidance. *RadioGraphics* **24**, 1679–1691 (2004). DOI:10.1148/rg.246045065
2. Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: *SIGGRAPH '00 Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Technique*, pp. 417–424. ACM Press/Addison-Wesley, New York (2000). Doi:10.1145/344779.344972
3. Bushberg, J.T., Seibert, J.A., Leidholdt Jr., E.M., Boone, J.M.: *The Essential Physics of Medical Imaging*, 2nd Edition. Lippincott Williams & Wilkins, Philadelphia (2002)
4. Chan, T., Kang, S.: Error analysis for image inpainting. *Journal of Mathematical Imaging and Vision* **26**(1-2), 85–103 (2006). Doi:10.1007/s10851-006-6865-7
5. Chan, T., Shen, J.: *Image Processing and Analysis: Variational, PDE, Wavelet, and Stochastic Methods*. Society for Industrial and Applied Mathematics, Philadelphia (2005)
6. Chan, T., Yip, A., Park, F.: Simultaneous total variation image inpainting and blind deconvolution. *International Journal of Imaging Systems and Technology* **15**(1), 92–102 (2005). Doi:10.1002/ima.20041
7. Chung, J., Chung, M., O'Leary, D.P.: Designing optimal spectral filters for inverse problems. *SIAM J. Sci. Comput.* **33**(6), 3132–3252 (2011). Doi:10.1137/100812938
8. Dahl, J., Hansen, P., Jensen, S., Jensen, T.: Algorithms and software for total variation image reconstruction via first-order methods. *Numerical Algorithms* **53**(1), 67–92 (2010). Doi:10.1007/s11075-009-9310-3
9. Daubechies, I.: *Ten Lectures on wavelets*, CBMS-NSF Regional Conference Series in Applied Mathematics, vol. 61. SIAM, Philadelphia (1992)
10. De Vito, E., Rosasco, L., Caponnetto, A., Giovannini, U.D., Odone, F.: Learning from examples as an inverse problem. *Journal of Machine Learning Research* **6**, 883–904 (2005)
11. Golub, G., Heath, M., Wahba, G.: Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* **21**(2), 215–223 (1979)
12. Golub, G., Pereyra, V.: The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate. *SIAM Journal on Numerical Analysis* **10**(2), 413–432 (1973)
13. Golub, G.H., Pereyra, V.: Separable nonlinear least squares: the variable projection method and its applications. *Inverse Problems* **19**(2), R1–R26 (2003). Doi:10.1088/0266-5611/19/2/201

14. Hainaut, O.: Basic image processing (1996). URL <http://www.eso.org/~ohainaut/ccd>
15. Hansen, P.C., Nagy, J., O’Leary, D.P.: Deblurring Images: Matrices, Spectra, and Filtering. Society for Industrial and Applied Mathematics, Philadelphia (2006)
16. Neelamani, R., Choi, H., Baraniuk, R.: ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Transactions on Signal Processing* **52**(2), 418–433 (2004)
17. Nocedal, J., Wright, S.: Numerical Optimization. Springer, New York (1999)
18. Nowak, R., Thul, M.: Wavelet-vaguelette restoration in photon-limited imaging. In: Acoustics, Speech and Signal Processing, 1998. Proceedings of the 1998 IEEE International Conference on, vol. 5, pp. 2869–2872. IEEE (1998). Doi:10.1109/ICASSP.1998.678124
19. Pusey, E., Lufkin, R., Brown, R., Solomon, M., Stark, D., Tarr, R., Hanafee, W.: Magnetic resonance imaging artifacts: Mechanism and clinical significance. *RadioGraphics* **6**(5), 891–911 (1986)
20. Raven, C.: Numerical removal of ring artifacts in microtomography. *Review of Scientific Instruments* **69**(8), 2978–2980 (1998). Doi:10.1063/1.1149043
21. Vapnik, V.N.: Statistical Learning Theory. Wiley, New York (1998)
22. Whitmore, B., Wiggs, M.: Charge transfer traps in the WFPC2. In: A. Koratkar, C. Leitherer (eds.) Calibrating HST: Post Servicing Mission, STScI Calibration Workshop 1, Conf. Proceedings. Space Telescope Science Institute, Baltimore, Maryland (1995). <http://www.stsci.edu/hst/wfpc2/documents/isr/9503.1.html>