Least Square Based Convolution Mask for Image Restoration

Shuozhong Wang and Xinpeng Zhang
School of Communication and Information Engineering
Shanghai University, P. R. China

Abstract: A small convolution mask for image restoration is obtained by passing a test pattern through the degradation filter. The generated test matrix is then pseudo-inversed to give a least square operator. The operator can be decomposed into a bank of singular value sub-filters. Regularization based on truncated SVD and an optimization procedure is introduced to suppress high order SV components. Only small sized matrices are involved in calculation, making the technique suitable for software and hardware implementation.

1. Introduction

Many image restoration approaches based on singular value decomposition involve operations on very large matrices, and thus require huge computation capacity, both in storage space and execution speed. Therefore efforts have been made to find efficient algorithms. Fish et al.[1] proposed a scanning SVD approach in which SVD is performed over a scanning window that is much smaller than the image size. In order to alleviate blocking effects, only the center portion of this window is retained for the restored output. This necessitates overlapping of the scanning windows. To design a small-sized FIR filter in solving an early real-time restoration problem, Schutten and Vermeij[2] first formed their filter in the frequency domain. The coefficients are then inverse-transformed, and truncated heuristically. Reichenbach and Park[3] proposed a restoration method using a small convolution kernel that is a constrained Wiener filter based upon a compact support. Wiener filtering resolves the ill-posed difficulty by incorporating power spectra of both image and additive noise into the equation, and is most conveniently implemented in the frequency domain. In [3], after a small kernel is obtained, image restoration is performed in the space domain by convolution. Under the constraint of a spatially limited support, the performance of Reichenbach's filter is optimal in the MMSE sense under the finite-support constraint, and quite close to the unconstrained Wiener filter. The method, like the unconstrained Wiener filter, requires knowledge of power spectra of both the image prior to degradation, and the noise. If this knowledge is not available, the power spectra have to be estimated from the observed image, resulting in a somewhat degraded performance.

In this paper, we introduce a small-sized image restoration operator in a form of truncated SVD filter-bank. An algebraic solution is sought to give an algorithm that is easy to implement in the space domain. The method uses a regularized pseudoinverse technique to generate a convolution kernel. Unlike conventional approaches, the SVD operation is performed not on a huge matrix representing the blurring filter, but rather, on a much smaller data matrix. In order to combat noise, regularization is introduced, achieved by decomposing the restoration operator into a bank of SVD sub-filters, and omitting some high-order filters responsible for noise amplification. An appropriate truncating threshold is obtained based on analyses of errors in the restored image. The proposed method does not rely on the detailed image and noise power spectra, but only on an estimate of the noise variance. The performance of the TSVD filter-bank operator is close to the Wiener filter, and better if the image power spectrum has to be estimated from the degraded observation in the Wiener filtering.

2. Methodology

Image degradation is commonly described by the following expression[4]:

$$\mathbf{y} = \mathbf{H} * \mathbf{x} + \mathbf{n} \tag{1}$$

where $*$ is a convolution operator, **x** and **y** are the ideal and observed images sized $N_1N_2\times1$ and $M_1M_2\times1$, respectively; and **n,** sized $M_1M_2\times1$, is an independent additive noise with a zero mean. **H** is a linear and shift-invariant degradation operator that may be expressed as an $M_1M_2\times N_1N_2$ block Toeplitz matrix. The objective of image restoration is to find an estimate of the ideal image, $\hat{\mathbf{x}}$, given the observation **y** and the degradation operator **H**. The estimate should resemble the original ideal image as closely as possible.

Image degradation can sometime be characterized with a compact lowpass convolution mask, also denoted **H** as in the following discussion. The aim of this work is to derive a deconvolution, or restoration, mask that “undoes” the degradation.

First consider a one-dimensional case. Pass a binary step series of length N, $\mathbf{f} = [0, 0, ..., 0, 1, 1, ..., 1]^{\mathrm{T}}$, through a degradation filter **h** of length L to generate a blurred output, **g**, that is the convolution of **f** and **h**:

$$g_m = \sum_{l=1}^{L} f_{m+l-1}\, h_{L-l+1}, \quad m = 1, 2, \cdots, M;\, M = N - L + 1 \tag{2}$$

Assume that **b** of length K (usually $K \ge L$) is the restoration operator to be sought. Convolving **g** with **b** produces an estimate of the test pattern **f**, as can be expressed in the following way:

$$\begin{bmatrix} \hat{f}_1 \\ \hat{f}_2 \\ \vdots \\ \hat{f}_J \end{bmatrix} = \begin{bmatrix} g_1 & g_2 & \cdots & g_K \\ g_2 & g_3 & \cdots & g_{K+1} \\ \vdots & \vdots & \ddots & \vdots \\ g_J & g_{J+1} & \cdots & g_{J+K-1} \end{bmatrix} \begin{bmatrix} b_K \\ b_{K-1} \\ \vdots \\ b_1 \end{bmatrix} \tag{3}$$

where $J=N-K-L+2$. The minimal allowed value of J is $K+L-1$, with $N_{\min}=2K+2L-3$. Matrix **G** may be termed the *degraded test pattern* (DTP). Since the original test pattern is known to be the step series, a least-square solution to the restoration operator is obtained:

$$\mathbf{b} = \mathbf{G}^{+}\mathbf{f} \tag{4}$$

where $\mathbf{G}^{+}$ is the Moore-Penrose pseudo inversion of **G**.

Generalization to non-separable 2D cases is straightforward. A small rectangle with four alternate black and white quadrants, is chosen as the test image represented by a $J_1\times J_2$ matrix **F**, or equivalently, by a column-stacked vector **f**. The degraded output vector is obtained by convolving **F** with an $L_1\times L_2$ degradation mask **H**:

$$g_{m_1,m_2} = \sum_{l_1=1}^{L_1}\sum_{l_2=1}^{L_2} f_{m_1+l_1-1,m_2+l_2-1}\, h_{L_1-l_1+1,L_2-l_2+1} \qquad m_1 = 1, 2, \ldots, M_1, m_2 = 1, 2, \ldots, M_2 \tag{5}$$

where $M_i = N_i - L_i + 1$, $i = 1, 2$. With the elements appropriately arranged, a DTP matrix, **G**, is obtained:

$$\begin{aligned} &G_{j,k} = g_{j_1+k_1-1,\, j_2+k_2-1} \\ &j = (j_1 - 1)J_2 + j_2,\ j_i = 1,2,\ldots,J_i,\ j = 1,2,\ldots,J_1J_2 \\ &k = (k_1 - 1)K_2 + k_2,\ k_i = 1,2,\ldots,K_i, k = 1,2,\ldots,K_1K_2 \end{aligned} \tag{6}$$

Thus an expression formally identical to (4) can be used to produce the required LS operator **b**, which is then rearranged to give a $K_1\times K_2$ mask **B**. **b** is a $K_1K_2\times1$ vector:

$$\mathbf{b} = [\mathbf{b}_{K_1} \quad \mathbf{b}_{K_1-1} \quad \cdots \quad \mathbf{b}_{k_1} \quad \cdots \quad \mathbf{b}_1]^T \tag{7}$$

where the sub-vectors

$$\mathbf{b}_{k_1} = [b_{k_1,K_2} \quad b_{k_1,K_2-1} \quad \cdots \quad b_{k_1,k_2} \quad \cdots \quad b_{k_1,1}]^T \tag{8}$$

3. Regularization

The operator in (2) performs poorly in the presence of noise. Experiments show that performance deteriorates rapidly as SNR decreases. The restored image is overwhelmed by noise when SNR ≤ 70dB. This can be remedied by a regularizing procedure based on SVD of **G**.

Rewrite (2) in the SVD form:

$$\mathbf{b} = \sum_{r=1}^{R} \frac{1}{\lambda_r} (\mathbf{v}_r \mathbf{u}_r^{\mathrm{T}}) \mathbf{f} = \sum_{r=1}^{R} \frac{1}{\lambda_r} \mathbf{b}_r \tag{9}$$

where R is the rank of **G**; $\mathbf{v}_r$ and $\mathbf{u}_r$ are the r-th characteristic vectors of $\mathbf{G}^{\mathrm{T}}\mathbf{G}$ and $\mathbf{G}\mathbf{G}^{\mathrm{T}}$, respectively; λ_rs are the singular values of **G** in a descending order, and $\mathbf{b}_r$s are the singular-value components of **b**.

In (9), $\mathbf{b}_1$ is a low-pass operator, while others are band-pass operators. The center frequencies rise with r. When acting on a degraded image, the individual sub-filters extract different frequency contents from the data. The outputs are then multiplied by $1/\lambda_r$, and added together to form the restored image. Noise is amplified by the large values of $1/\lambda_r$ in the high frequency region. In this region the data contain relatively less energy of the image. This is shown in Fig.1 where the SVD spectra of the pure image data (solid), noise (dotted), and the noisy image (dashed) are plotted. A 64×64 block from a test image Lena was used in the computation. The additive noise was Gaussian with a zero mean, and independent of the image. The signal-to-noise ratio was 30dB. A 9×9 restoration kernel was assumed so that **G** has 81 singular values. It is observed that the signal power is basically concentrated in the region of small SV indices corresponding to large singular values and low frequencies. It quickly decays towards the high end, while the decay of noise power is much slower. Considerable amount of noise power is distributed in the high SV index region, significantly exceeding the signal power.

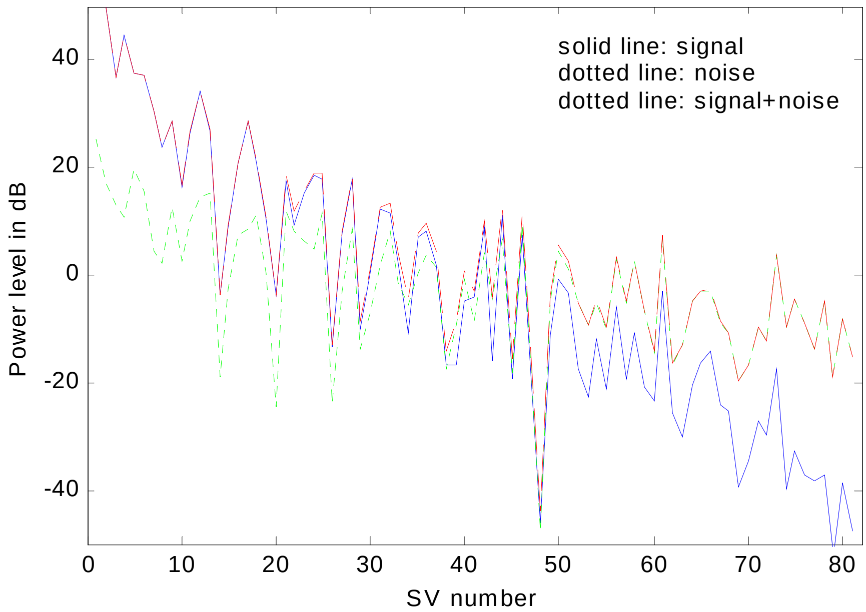

Fig.1 SV spectra of signal and noise components

In order to prevent the restored image from being corrupted by noise amplification, a truncated SVD technique can be used to drop some of the sub-filters corresponding to the smallest singular values. Analytical or empirical methods for choosing a truncating threshold in conventional SVD restoration approaches have been discussed in the literature[1,5,6]. But some was found to be inconsistent in different applications[7]. Most of these methods are not directly applicable to the present filter-bank formulation. In this paper, the truncating threshold is chosen from an observation of the energy contained in the restored image data and the error.

Assume that only ρ filters of the lowest orders are used, resulting in a partially restored image:

$$\hat{\mathbf{x}}^{(\rho)} = \mathbf{B}^{(\rho)}\mathbf{H}\mathbf{x} + \mathbf{B}^{(\rho)}\mathbf{n} \tag{10}$$

where $\mathbf{B}^{(\rho)}$ is a block Toeplitz matrix obtained from

$$\mathbf{b}^{(\rho)} = \sum_{r=1}^{\rho} \mathbf{b}_r / \lambda_r \tag{11}$$

Let the partially restored image undergo the same degradation as the observed image described in (1). If a good value of ρ is chosen, the energy contained in the degraded result should be close to that in the observation $\mathbf{y}$. A criterion for a good SVD truncation is thus obtainable by minimizing

$$\begin{aligned} d(\rho) &= \mathrm{E}\{[(\mathbf{H}\hat{\mathbf{x}}^{(\rho)} + \mathbf{n}) - \mathbf{y}]^T \ [(\mathbf{H}\hat{\mathbf{x}}^{(\rho)} + \mathbf{n}) - \mathbf{y}]\} \\ &= \mathrm{E}\{[\mathbf{H}\hat{\mathbf{x}}^{(\rho)} - \mathbf{y}]^T \ [\mathbf{H}\hat{\mathbf{x}}^{(\rho)} - \mathbf{y}]\} + 2\mathrm{E}\{\mathbf{n}^T \mathbf{B}^{(\rho)T} \mathbf{H}^T \mathbf{n}\} - \sigma_n^2 \end{aligned} \tag{7}$$

where σ_n^2 is the noise variance. The first expectation represents energy contained in the difference between observation and the degraded result of restored image less noise. The second expectation can most conveniently be evaluated in the frequency domain:

$$\mathrm{E}\{\mathbf{n}^T \mathbf{B}^{(\rho)T} \mathbf{H}^T \mathbf{n}\} = \sigma_n^2 \sum_{u,v} |B^{(\rho)}(u,v)||H(u,v)| = \sigma_n^2 \ \beta^{(\rho)} \tag{8}$$

$B^{(\rho)}(u, v)$ and H(u, v) are discrete Fourier transform of the $K\times K$ operator $\mathbf{B}^{(\rho)}$ and the $L\times L$ operator $\mathbf{H}$ respectively, padded with 0s to match the image size. The summation is carried out over the entire range of the discrete spatial frequencies, u and v. Therefore, (7) becomes

$$d(\rho) = \mathrm{E}\{[\mathbf{H}\hat{\mathbf{x}}^{(\rho)} - \mathbf{y}]^T [\mathbf{H}\hat{\mathbf{x}}^{(\rho)} - \mathbf{y}]\} + [2\beta^{(\rho)} - 1]\sigma_n^2 \tag{9}$$

If restoration were perfect and without noise, both terms in (9) would have vanished to make d = 0. Fig.2 shows an example in which the test image Lena is degraded by a 5×5 mask and AWGN at SNR=30dB. A 9×9 support for the restoration operator is used, with a total of 81 SVD filters in the SVD restoration filter bank. Only the first ρ filters are used to process the degraded image. The values of $d(\rho)$ (solid), RMS error in the restored image (dashed), and P_i /100 where P_i is power of the restored image (dotted) are shown against ρ, the number of sub-filters used. All three curves show stepping patterns that are strongly inter-related, indicating that certain filters have more impact to the performance than others. Recall that the insignificant SV components have been dropped in producing the spectral curves. When ρ = 28~34, both d and RMSE show a minimum, therefore, a good choice of SVD truncation is obtained. As more high-order filters are included, both the error and the power in restored image rise steeply, showing the devastating effect of the noise.

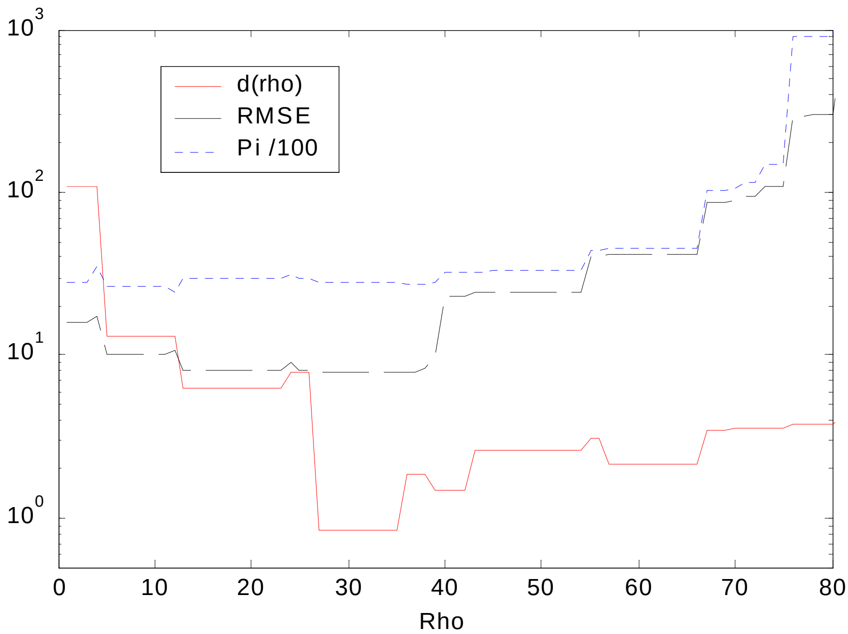

Fig.2 Performance of TSVD filter bank

Since the restored image power is attainable from the filter output, it can be used as an indicator for an appropriate choice of ρ although the minimum in the output power may not strictly coincide with that in the RMS error. In some cases, local minima may occur, necessitating human intervention based on visual inspection to choose a satisfactory result from multiple candidates. From a practical point of view, nevertheless, visual assistance is useful, as the RMSE criterion does not always agree with perceptual judgements.

As in any other regularization approaches, knowledge of the noise variance is important in obtaining a good tradeoff between the fidelity in the restored image and the smoothness of the solution. Methods for estimating the noise variance can be found in the literature such as [1], [4], [8], and [9].

4. Experimental Results

In Fig.3, the original, degraded and restored versions of test image Lena are shown. The degradation was done with a 5×5 Gaussian mask and AWGN at SNR=25dB. Restoration was accomplished using a 9×9 LS mask with TSVD and ρ =33. RMSE in the restored image was 7.32 gray levels, or 2.86% with respect to the full gray scale. A similar experiment was performed on Baboon (not shown here), with RMSE of 19.5 gray levels (7.6% with respect to the full gray scale).

Fig.4 shows a performance comparison between the proposed LS mask and a Wiener filter. The original Lena was degraded with the same 5×5 mask and AWGN. RMS errors in the degraded image ranged from 40 (SNR ≥ 50dB) up to 75 (SNR = 0dB). The image was restored with a 9×9 TSVD filter-bank operator and a Wiener filter, respectively. For the Wiener filter, two cases have been considered. The image power spectrum was computed either from the ideal image giving a minimal RMS error, or from the degraded image with an inferior performance. It is observed that the proposed method has a performance fairly close to the Wiener filter of the first case. The performance is better than that of the second case in which P_x was unknown.

As Baboon contains more high-frequency contents, quality of the restored image is some what inferior to that of Lena. Different sizes of the restoration operator have been tested. Generally speaking, the resultant RMS errors do not change substantially over a fairly wide range of the K value, e.g., K = 5~13 when L = 5. A further increase of K does not provide discernible improvement in terms of root mean square error and, on the contrary, even causes a rise in the error. A larger K also requires more computation and produces a wider erroneous border in the restored image. The situation of course depends on the specific image and signal-to-noise ratio. The choice of K within the said range, however, does have an impact on the perceptive quality of the restored image. Experiments show that an operator equal to, or moderately larger than the degradation filter, e.g., K = 5~9 for L = 5, usually gives a satisfactory result. If too small an operator is used, the sharpness of the restored image will be sacrificed, while a large operator tends to cause more noise and some ringing effects.

Fig.3 Original, degraded and restored images.

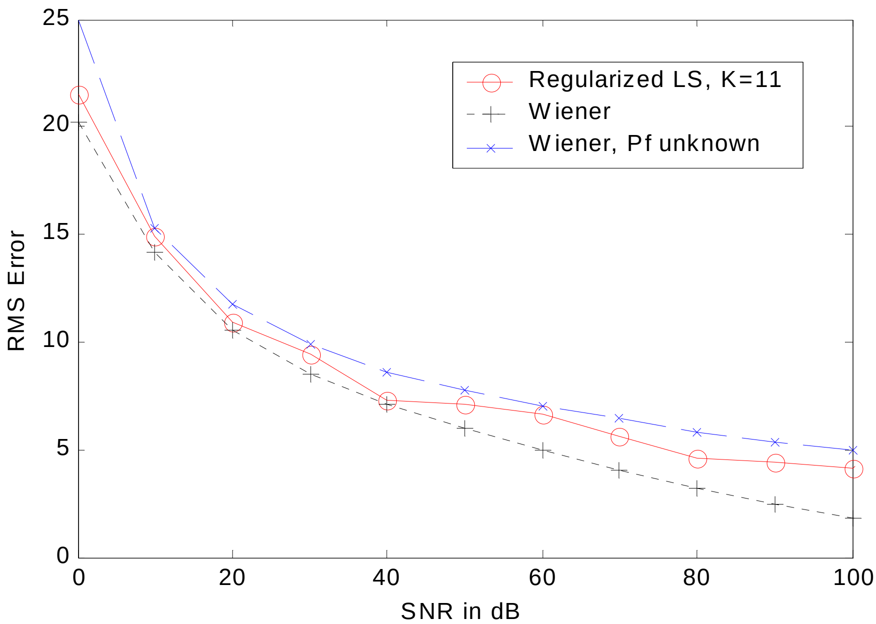

Fig.4 Comparison between LS mask and Wiener filter

5. Remarks

The proposed restoration method provides a small convolution mask, and has a performance approaching the Wiener filter without knowing the power spectrum of the ideal image. A TSVD-based regularizing procedure is used to deal with noise problems. Unlike the methods described in the previous papers[3,4], however, only small sized matrices are involved in the computation. The size of matrices under consideration is only related to L and K, the sizes of the degradation and restoration operators respectively, and independent of the image size.

The SVD expansion of the operator makes it possible to partially separate noise contributions from useful information buried in the observed data. The operator is viewed as a bank of sub-filters with pass-bands ranging from DC to high frequencies, extracting different contents over the entire frequency domain. Outputs from the sub-filters form a spectrum of the restored image components. Regularization is realized by throwing away some sub-filters that mainly contribute to noise amplification. A truncating threshold is obtained by minimizing energy contained in the difference between the observation and a "blurred image estimate".

Having obtained the operator, image restoration is implemented in the space domain by convolution. Image convolution with a small operator can be made very efficient, whether by software or hardware, therefore is of great interest in many applications.

With only knowledge of noise variance, performance of the TSVD filter-bank operator is comparable to the Wiener filter that is optimal in the MMSE sense. Apart from the noise variance, the Wiener filter requires the image power spectrum. Errors of the TSVD filter-bank are generally greater than Wiener with a perfectly known power spectrum of the original image. However, when the Wiener filter uses an inaccurate power spectrum estimated from the degraded observation, the proposed method has an advantage.

References

[1] D. A. Fish, J. Grochmalicki, and E. P. Pike, Scanning singular-value-decomposition method for restoration of images with space-variant blur, *J. Opt. Soc. Am. A*, 13, 464–469, 1996

[2] R. W. Schutten, and G. F. Vermeij, The approximation of image blur restoration filter by finite impulse response, *IEEE Trans. Pattern Anal. Mach. Intell.*, PAMI-2, 176–180, 1980

[3] S. E. Reichenbach, et al., Small Convolution Kernels for High-Fidelity Image Restoration, *IEEE Trans. Signal Processing*, 39, 2263–2274, 1991

[4] H. Andrews, et al., *Digital Image Restoration*, Prentice Hall, 1977

[5] J. Shen, H. Erdogan, and E. S. Ebbini, "An optimal operator design technique for coded excitation ultrasound imaging system," in *Proc. IEEE Ultrasonics Symposium*, vol.3, pp.1777–1781, 1994

[6] A. Sono, "Optimally regularized inverse of singular value decomposition and application to signal extrapolation," *Signal Processing*, vol.30, pp.163–176, 1993

[7] I. F. Gorodnitsky, and B. D. Rao, "Analysis of error produced by truncated SVD and Tikhnov regularization methods," in *Conference Record of the 28th Asilomar Conference on Signal, Systems and Computers*, vol.1, pp.25–29, 1994

[8] K. Rushforth, et al., Least-Square Reconstruction of Objects with Missing High-Frequency Components, *J. Opt. Soc. Am*, 72 (1), 24–35, 1982

[9] J. Maeda, et al., Restoration of Band-Limited Images by an Iterative Regularized Pseudoinverse Method, *J. Opt. Soc. Am.*, A1(1), 28–34, 1984