

Fast Space-Varying Convolution Using Matrix Source Coding with Applications to Camera Stray Light Reduction

Jianing Wei, *Member, IEEE*, Charles A. Bouman, *Fellow, IEEE*,
and Jan P. Allebach, *Fellow, IEEE*

Abstract—Many imaging applications require the implementation of space-varying convolution for accurate restoration and reconstruction of images. Here we use the term “space-varying convolution” to refer to linear operators whose impulse response has slow spatial variation. Moreover, these space-varying convolution operators are often dense, so direct implementation of the convolution operator is typically computationally impractical. One such example is the problem of stray light reduction in digital cameras, which requires the implementation of a dense space-varying deconvolution operator. However, other inverse problems, such as iterative tomographic reconstruction, can also depend on the implementation of dense space-varying convolution. While space-invariant convolution can be efficiently implemented with the Fast Fourier Transform (FFT), this approach does not work for space-varying operators. So direct convolution is often the only option for implementing space-varying convolution.

In this paper, we develop a general approach to the efficient implementation of space-varying convolution, and demonstrate its use in the application of stray light reduction. Our approach, which we call matrix source coding, is based on lossy source coding of the dense space-varying convolution matrix. Importantly, by coding the transformation matrix, we not only reduce the memory required to store it; we also dramatically reduce the computation required to implement matrix-vector products. Our algorithm is able to reduce computation by approximately factoring the dense space-varying convolution operator into a product of sparse transforms. Experimental results show that our method can dramatically reduce the computation required for stray light reduction while maintaining high accuracy.

Index Terms—Stray light, space-varying point spread function, image restoration, fast algorithm, inverse problem, digital photography.

I. INTRODUCTION

In many important imaging applications, it is useful to model the acquired data vector y as

$$y = Ax + w, \quad (1)$$

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

This research work was done when Jianing Wei was a Ph.D. student at the School of Electrical and Computer Engineering, Purdue University. Jianing Wei is currently with Google Inc., Mountain View, CA 94043, USA. Charles A. Bouman and Jan P. Allebach are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47906, USA. Email: weijianing@gmail.com, {bouman, allebach}@purdue.edu.

This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number 56541-CI, and the National Science Foundation under Contract CCR-0431024.

where x is the unknown image to be determined, A is a linear transformation matrix, and w is some additive noise that is independent of x with inverse covariance Λ_w . This simple linear model can capture many important inverse problems including image deblurring [1], tomographic reconstruction [2], and super-resolution [3].

There is a vast literature related to the effective solution of inverse problems with the form of Eq. (1). Solution to these inverse problems can be very difficult if the transformations A or $A^t\Lambda_w A$ are ill conditioned. This can and often does happen if the data is sparse, which will result in fewer rows than columns in A , or of low quality, which will result in small values in Λ_w . In these cases, many algorithms have been proposed to recover x from y . Classical approaches include algebraic reconstruction techniques (ART) [4]; simultaneous iterative reconstruction technique (SIRT) [5]; simultaneous algebraic reconstruction technique (SART) [6]; Lucy-Richardson algorithm [7]; Van Cittert’s method [1], [8]; approximate [9] or true [10] maximum likelihood estimation; regularized inversion [11]; and maximum a posteriori (MAP) estimate with different priors [12], [13], [14], [15], [16], [17]; Bayesian minimum mean squared error (MMSE) estimate with high-order Markov Random Field (MRF) prior [18]; and direct deconvolution followed by artifact removal using a machine learning approach [19]. A common objective in all these methods is to balance the goal of matching the observed (but noisy) data with the goal of producing a physically realistic image.

While many methods exist for addressing the inverse problem of Eq. (1), all these methods typically require the computation of matrix-vector products, such as Ax , or $A^t y$. In fact, iterative inversion methods typically require the repeated computation of the matrix-vector product $A^t Ax$ or $A^t \Lambda_w Ax$. If the matrix A is sparse, as is the case with a differential operator [20], or if it can be decomposed as a product of sparse transformations, as is the case with the wavelet [21] or fast Fourier transform (FFT) [22], then computation of Ax is fast and tractable. For example, if Ax is space-invariant convolution, then A can be decomposed as the product of two FFTs and a diagonal transform, resulting in a total computation for the evaluation of Ax that has an order of $O(P \log P)$ where P is the size of x .

However, in many important applications, A represents space-varying convolution where we define “space-varying convolution” to be a linear operator whose impulse response

varies slowly with position in the image. More specifically, when A represents 1D convolution, then it must have the property that

$$A_{q,p} = h(q - p) \quad (2)$$

for some function $h(\cdot)$ that represents the impulse response of the operator. Here $A_{q,p}$ represents the response at location q to a stimulus at location p . So we use the term “space-varying convolution” to mean that A can be written as

$$A_{q,p} = h(q - p, p), \quad (3)$$

where $h(\cdot, p)$ is a slowly varying function of p . The 2D case is similar, with q and p representing vector indices. If A represents space-varying convolution with a large point-spread function, then each direct evaluation of the matrix-vector product Ax or $A^t y$ can require enormous computation. Even when $A^t A$ is approximately Toeplitz, as is often the case in tomographic reconstruction [4], the matrix $A^t \Lambda_w A$ will not be Toeplitz when the noise is space varying.

The specific application considered in this paper is that of stray light reduction for digital photographs [23]. Stray light refers to the portion of the entering light flux that is misdirected to undesired locations in the image plane. It is sometimes referred to as lens flare or veiling glare [24]. It can be caused by imperfections in lens elements, or reflections between optical surfaces; but in the end, it tends to reduce the contrast of images by scattering light from bright regions of the image into dark regions. Since stray light is typically scattered across the imaging field, its associated point spread function (PSF) typically has large support and is space-varying [23]. With this in mind, an imaging system with stray light can be modeled using Eq. (1) with

$$A = (1 - \beta)I + \beta S, \quad (4)$$

where I is the identity matrix, S is a dense matrix describing the space-varying scattering of light across the imaging array¹, and β is a scalar constant representing the fraction of scattered light.

Typically the fraction of scattered light is small, and we know that the matrix A is very well conditioned, so in contrast to many inverse problems, stray light reduction is a well-posed inverse problem, and the solution can be computed using a formula such as that proposed by Jansson and Fralinger [25]

$$\hat{x} = (1 + \beta)y - \beta S y, \quad (5)$$

where y is the observed image, and \hat{x} is the estimate of the underlying image to be recovered.

However, even in this simple form, the computation of Eq. (5) is not practical for a modern digital camera, because the matrix S is both dense and lacks a Toeplitz structure. In typical stray light models, the PSF is space-varying with heavy tails, so the FFT cannot be directly used to speed up computation [23]. For example, if the image contains P pixels, the computation is $O(P^2)$; so for a 10^6 pixel image, it takes 10^6 multiplies to compute each output pixel, resulting in a total

of 10^{12} multiplies. This much computation requires hours of processing time, which makes accurate stray light reduction infeasible for implementation in low cost digital cameras.

There are several approaches in the literature to compute space-varying convolution. The overlap-add method [26], [27], [28] partitions the input space into isoplanatic patches and assumes that the PSF within each patch is spatially invariant. Then the FFT can be used to efficiently compute the result for each patch; and they are added together. The efficient filter flow approach [29] is also based on this method. Another approach is to approximate the space-varying convolution as the weighted linear combination of a small number of convolutions with fixed PSF [30], [31].

In this paper, we propose a novel algorithm for fast computation of space-varying convolution; and we demonstrate its effectiveness for the particular application of stray light reduction. This paper is an extension to our earlier conference paper [32], but it incorporates substantial new theoretical and experimental results. Our method, which we refer to as matrix source coding (MSC), is based on the use of lossy source coding techniques to compress the dense matrix S into a sparse form. Importantly, the effect of this compression is not only to reduce the memory required for storing S , but also to dramatically reduce the number of multiplies required to approximately compute the matrix-vector product Sy . Our approach works by first decorrelating the rows and columns of S , and then quantizing the resulting compacted matrix so that most of its entries become zero.

Our MSC method requires both on-line and off-line components. While the on-line component is fast, it requires the pre-computation of a sparse matrix decomposition in an off-line procedure. We introduce a method for efficiently implementing this off-line computation for a broad class of space-varying kernels whose wavelet coefficients are localized in space.

In order to assess the value of our approach, we apply it to the problem of stray light reduction for digital cameras. Using the stray light model of [23], we demonstrate that the MSC method can achieve an accuracy of 1% with approximately 7 multiplies per pixel at image resolution 1024×1024 , a dramatic reduction when compared to direct space-varying convolution. We also demonstrate the effectiveness of a practical pre-computation algorithm that can be performed in time proportional to the kernel size, and show that the MSC approach compares favorably to the overlap-add method.

In Section II we introduce the MSC method. The on-line algorithm for space-varying convolution is derived and presented in Section II-A, and the off-line pre-computation method is presented in Section II-B and Section IV. Section III presents the stray light application and model, and Section V follows with experimental results. Section VI concludes this paper.

II. MATRIX SOURCE CODING APPROACH TO SPACE-VARYING CONVOLUTION

In this section, we present the matrix source coding (MSC) approach for fast space-varying convolution. The MSC approach has two parts, which we will refer to as the on-line

¹Due to conservation of energy, we know that the columns of S sum to less than 1.

and off-line computations. The on-line algorithm computes the space-varying convolution using a product of sparse matrices. Section II-A presents the on-line algorithm along with its associated theory.

While the on-line computation is very fast, it depends on the result of off-line computation, which is to pre-compute a sparse version of the original space-varying convolution matrix. This computation is not input image-dependent, and only depends on the PSF of the imaging system. However, naive computation of this sparse matrix can be enormously computationally demanding. So in Section II-B, we introduce an efficient algorithm for accurate but approximate computation of the required sparse transformation for a broad class of problems.

A. Matrix Source Coding Theory

Our goal is to speed up the computation of general matrix-vector products with the form

$$z = Sy, \quad (6)$$

where y is the input, z is the output, and S is a dense $P \times P$ matrix without a Toeplitz structure.² Without loss of generality, we assume y is zero mean since it is always possible to subtract off y 's mean. If S is Toeplitz, then Sy can be efficiently computed using the fast Fourier transform (FFT) [33]. However, when S implements space-varying convolution, then it is not Toeplitz; and the FFT cannot be used directly.

Our strategy for speeding up computation of Sy is to use lossy transform coding of the rows and columns of S . The resulting coded matrix then becomes sparse after quantization, and this sparsity reduces both storage and computation. However, in order to perform lossy source coding of S , we must first determine the relevant distortion metric.

In general, the mean squared error (MSE) is not a good distortion metric for coding S because it does not account for any correlation in the vector y . In order to see this, let $[S]$ be a quantized version of S produced by lossy coding and decoding. Then the distortion introduced into the output z , which we denote as δz , is given by

$$\delta z \triangleq Sy - [S]y = \delta S y, \quad (7)$$

where $\delta S \triangleq S - [S]$ is the distortion introduced into S through lossy coding and δz is the resulting distortion introduced into z . Using the argument of [34], [35] and assuming that y is independent of δS , then the conditional expectation of the MSE in z given δS can be written as

$$\begin{aligned} \mathbb{E} [\|\delta z\|^2 | \delta S] &= \mathbb{E} [\text{trace}\{\delta z \delta z^t\} | \delta S] \\ &= \mathbb{E} [\text{trace}\{\delta S y y^t \delta S^t\} | \delta S] \\ &= \text{trace}\{\delta S \mathbb{E} [y y^t | \delta S] \delta S^t\} \\ &= \text{trace}\{\delta S R_y \delta S^t\}, \end{aligned} \quad (8)$$

where $\mathbb{E} [\|\delta z\|^2 | \delta S]$ means the expectation of MSE in z given δS , and $R_y = \mathbb{E}[y y^t] = \mathbb{E}[y y^t | \delta S]$. Notice that in the special

case when y is white, then we have that $R_y = I$, and

$$\mathbb{E} [\|\delta z\|^2 | \delta S] = \|\delta S\|^2. \quad (9)$$

In other words, when the input y is white, minimizing the squared error distortion of the matrix S (i.e., the Frobenius norm) is equivalent to minimizing the expected value of the squared error distortion for z . However, when the components of y are correlated, then minimum MSE coding of S can be quite far from minimizing the MSE in z .

So with this result in mind, our objective will be to find two transforms W_1 and T that simultaneously sparsify the matrix S and whiten the random vector y . More specifically, define the transformed matrix \tilde{S} as

$$\tilde{S} = W_1 S T^{-1}, \quad (10)$$

where W_1 is an orthonormal transform³ and T is an invertible transform. Then the output z can be computed as

$$z = W_1^{-1} \tilde{S} T y. \quad (11)$$

So if we temporarily ignore the computational cost of the transforms T and W_1 , then the computation of Eq. (11) is greatly reduced when \tilde{S} is made sparse. Furthermore, the sparsity can be increased by quantizing \tilde{S} . So if $[\tilde{S}]$ represents the quantized version of S , then the approximate output can be computed as

$$\hat{z} = W_1^{-1} [\tilde{S}] T y, \quad (12)$$

where \hat{z} is the approximation of z . If $\tilde{y} = T y$ is approximately white, then MMSE quantization of \tilde{S} will result in a MMSE estimate \hat{z} .

So this all begs the question of what transforms can best achieve this goal? In order to answer this question, we first adopt the following definition.

Definition 1: We say that (W_1, T) are optimal for the transform S and the zero-mean random input vector y , if the following three properties hold: a) W_1 is orthonormal; b) $\tilde{y} = T y$ is white so that $\mathbb{E}[\tilde{y} \tilde{y}^t] = I$; and c) the matrix $\tilde{S} = W_1 S T^{-1}$ is diagonal.

At this point, it is useful to make a few observations. First notice that if we can find (W_1, T) that produce a diagonal matrix \tilde{S} , then this is the most sparse matrix possible. In this case, the number of non-zero elements in \tilde{S} must equal the rank of S . Second, if T whitens the random vector y , then by Eq. (9) we know that minimum mean square error (MMSE) quantization of \tilde{S} will result in the MMSE representation of z . For these two reasons, we say that such a choice of (W_1, T) is optimal.

In fact, such an optimal choice of (W_1, T) always exists as shown in Appendix A. In particular, define the matrix

$$F = S E_y \Lambda_y^{1/2} \quad (13)$$

where $R_y = E_y \Lambda_y E_y^t$ is the eigen decomposition of the covariance of y , and let $F = U \Sigma V^t$ be the singular value

²In fact, all our results hold when S is a non-square $P_1 \times P_2$ matrix. However, we consider the special case of $P = P_1 = P_2$ here for notational simplicity. The extension to the more general case is direct.

³More generally, this transform can just be assumed to be orthogonal. But for notational simplicity we also assume it is normal, i.e. each row and column has a Euclidean norm of 1.

decomposition of F . Then the optimal choices of (W_1, T) are given by

$$T = V^t \Lambda_y^{-1/2} E_y^t, \quad (14)$$

and

$$W_1 = U^t. \quad (15)$$

Intuitively, the optimal transform W_1 decorrelates the rows of F and the optimal transform T^{-1} decorrelates the columns of S . When this is done, the transformed matrix of Eq. (10) is then diagonal with $\tilde{S} = \Sigma$. In fact, when the eigen values Λ_y and singular values Σ are unique, then these transforms must also be unique up to permutations and scaling by ± 1 (see Appendix A).

Unfortunately, the optimal choice of (W_1, T) is typically not practical. This is because to achieve this ideal sparsity, the transforms W_1 and T must typically be dense matrices. However, if the transforms T and W_1 are dense, then the application of these decorrelating transforms requires as much computation as the original multiplication by S . This would, of course, defeat the purpose of the sparsification of S .

So our practical objective is to find computationally efficient transformations, (W_1, T) , that approximately whiten y , and approximately decorrelate the rows of F and the columns of S . Intuitively, the transforms W_1 and T concentrate the energy in the rows and columns of the matrix S in much the same way that transform coding concentrates the energy in a signal.

To achieve this practical objective, we will use a wavelet transform to approximately decorrelate the vector y , the rows of F , and the columns of S . In our applications, y is an image which can be modeled as a stationary random process. In this case, it is well-known that the wavelet transform is a reasonable choice for the decorrelation of y because it is known to be an approximation to the Karhunen-Loeve transform for stationary sources [36], and it is commonly used as a decorrelating transform for various source coding applications [37].

Using this strategy, we form the matrix T by a wavelet transform followed by gain factors designed to normalize the variance of wavelet coefficients. This combination of decorrelation and scaling whitens or spheres the vector y . Specifically, we choose

$$T = \Lambda_w^{-1/2} W_2, \quad (16)$$

where W_2 is an orthonormal wavelet transform and $\Lambda_w^{-1/2}$ is a diagonal matrix of gain factors so that

$$\Lambda_w = \text{diag}(W_2 R_y W_2^t), \quad (17)$$

where $R_y = E[yy^t]$ is the covariance of y .

At the same time, notice that the transform $T^{-1} = W_2^{-1} \Lambda_w^{1/2}$ approximately decorrelates the columns of S when S represents space-varying convolution. This is true because if S is a slowly space-varying operator, then the covariance matrix of the columns given by $R_c = S^t S$ is approximately Toeplitz. Again, it is well-known that if covariance R_c is Toeplitz, then this corresponds to the covariance of a wide-sense stationary random process, and the wavelet transform is a good approximation to the diagonalizing Karhunen-Loeve transform [36].




Fig. 1. Block diagram of on-line computation of matrix source coding.

With this choice for T of Eq. (16), then the matrix F is given by

$$F = S W_2^t \Lambda_w^{1/2}. \quad (18)$$

Again, the matrix $R_r = F F^t$ that represents the covariance of the rows of F is approximately Toeplitz, so the rows of F can be approximately decorrelated with an orthonormal wavelet transform W_1 .

To summarize, Fig. 1 illustrates a block diagram of the on-line computation we will use for matrix source coding. In particular, the approximate output is computed as

$$\hat{z} = W_1^t [\tilde{S}] \Lambda_w^{-1/2} W_2 y, \quad (19)$$

where $[\tilde{S}]$ is a quantized version of the transformed matrix

$$\tilde{S} = W_1 S W_2^{-1} \Lambda_w^{1/2}. \quad (20)$$

First, the input data y is transformed and scaled. Then it is multiplied by a sparse matrix $[\tilde{S}]$. Then a second inverse wavelet transform is applied to compute the approximate result \hat{z} . The accuracy of the approximation is determined by the degree of quantization applied to \tilde{S} , with coarser quantization resulting in lower accuracy. So if little or no quantization is used, then the computation can be arbitrarily close to exact, but at the cost of less sparsity in \tilde{S} and therefore more computation and storage. In this way, matrix source coding allows for a continuous trade-off between accuracy and computation/storage in much the same way that signal source coding allows for a rate-distortion trade-off in applications such as image and video coding.

Finally, as a practical matter, the diagonal matrix Λ_w can be easily estimated from training images. In fact, we will assume that the gain factors for each subband are constant since the wavelet coefficients in the same subband typically have the same variance [1]. So to estimate Λ_w , we take the wavelet transform W_2 of the training images, compute the variance of the wavelet coefficients in each subband, and average over all images to obtain an estimate of the gain factors for each subband.

B. Efficient Off-Line Computation for Matrix Source Coding

In order to implement the fast space-varying convolution method described by Eq. (19), it is first necessary to compute the source coded matrix $[\tilde{S}]$. We will refer to the computation of $[\tilde{S}]$ as the off-line portion of the computation. Once the sparse matrix $[\tilde{S}]$ is available, then the space-varying convolution may be efficiently computed using on-line computation shown in Eq. (19).

However, even though the computation of $[\tilde{S}]$ is performed off-line, exact evaluation of this sparse matrix is still too large for practical problems. For example, if the image contains 16 million pixels (i.e., 2^{24} pixels) and each matrix entry is stored

as a 4-byte integer, then the temporary storage of S requires approximately 4×2^{48} bytes or equivalently one petabyte of storage. Of course, this much storage exceeds the capacity of most modern computers. So, the following section describes an efficient approach to compute $[\tilde{S}]$ which eliminates the need for any such temporary storage.

In our implementation, we use the Haar wavelet transform [21] for both W_1 and W_2 . The Haar wavelet transform is an orthonormal transform, so we have that

$$\begin{aligned}\tilde{S} &= W_1 S T^{-1} \\ &= W_1 S W_2^t \Lambda_w^{1/2}.\end{aligned}\quad (21)$$

Therefore, direct implementation of this off-line computation consists of a wavelet transform W_2 along the rows of S , scaling of the matrix entries, and a wavelet transform W_1 along the columns. Unfortunately, the computational complexity of these two operations is $O(P^2)$, where P is the number of pixels in the image, because the operation requires that each entry of S be touched. This computation can take a tremendous amount of time for high resolution images. Therefore, we aim at reducing the computational complexity to be linearly proportional to P and eliminating the need to store the entire matrix S . In order to achieve this goal, we will use a two stage quantization procedure combined with a recursive top-down algorithm for computing the Haar wavelet transform coefficients.

The first stage of this procedure is to zero out the entries in $SW_2^t \Lambda_w^{1/2}$ whose magnitudes are less than a specified threshold. We use $Q_t(SW_2^t \Lambda_w^{1/2})$ to describe this thresholding stage, where for a scalar x

$$Q_t(x) = \begin{cases} x & \text{for } |x| > t \\ 0 & \text{otherwise} \end{cases}.\quad (22)$$

The resulting matrix $Q_t(SW_2^t \Lambda_w^{1/2})$ is already quite sparse, thus there is no need to store the entire dense matrix. The value of the threshold t is chosen to achieve some target sparsity, $N \ll P$, where N is the average number of non-zero elements per row of the matrix $Q_t(SW_2^t \Lambda_w^{1/2})$.

The second stage of this procedure is to compute the wavelet transform W_1 over the columns of this already sparse matrix. Taking advantage of the sparsity resulting from the first stage, then the computation of the wavelet transform in the second stage is reduced from $O(P^2)$ to $O(NP)$. In summary, this two stage procedure is expressed mathematically as

$$[\tilde{S}] \approx [W_1 Q_t(SW_2^t \Lambda_w^{1/2})].\quad (23)$$

In practice, N should be chosen to be sufficiently large so that the accuracy-computation trade-off of the matrix-source coding is not impaired, but sufficiently small that the memory requirements for temporary storage are not excessive. In the experimental results section, we illustrate this trade-off.

However, there is still a problem as to how to efficiently compute $Q_t(SW_2^t \Lambda_w^{1/2})$. Since S is dense, direct evaluation of this expression requires $O(P^2)$ operations. However, after thresholding, most values in $Q_t(SW_2^t \Lambda_w^{1/2})$ are zero, so we can dramatically reduce computation by first identifying the regions where non-zero values are likely to occur. We refer

to these non-zero values as significant wavelet coefficients. As mentioned before, we use N to denote the number of significant wavelet coefficients per row. Once we can predict the location of the significant wavelet coefficients, we can use a recursive top-down approach to compute them without having to do the full wavelet transform for each row of S . We will discuss how to predict the location of the significant wavelet coefficients in constant time in Sec. IV.

Next, we describe our algorithm for computing the significant wavelet coefficients once their location is known. Our strategy is to first compute the approximation coefficients that are necessary to compute the significant wavelet coefficients using a recursive top-down approach. Then significant detail coefficients can be computed from the approximation coefficients at finer resolutions. Our approach is able to achieve $O(N)$ computational complexity for computing the significant Haar wavelet coefficients per row, which leads to a total complexity of $O(NP)$ for computing $Q_t(SW_2^t \Lambda_w^{1/2})$.

We use $f(q, k, i, j)$ to represent the approximation coefficient for row q of matrix S at location (i, j) and level k , where larger values of k correspond to coarser scales, and $f(q, 0, i, j)$ corresponds to the full resolution image for row q . We use $g_h(q, k, i, j)$, $g_v(q, k, i, j)$, and $g_d(q, k, i, j)$ to represent the detail coefficients in horizontal, vertical and diagonal bands, respectively. We can then compute $f(q, k, i, j)$ as the average of its corresponding 4 higher resolution neighbors:

$$\begin{aligned}f(q, k, i, j) &= \frac{1}{2}(f(q, k-1, 2i, 2j) + f(q, k-1, 2i+1, 2j) \\ &\quad + f(q, k-1, 2i, 2j+1) \\ &\quad + f(q, k-1, 2i+1, 2j+1)).\end{aligned}\quad (24)$$

Equation (24) specifies the parent/child relations of a quad-tree as shown in Figure 2. Our algorithm transverses this quad-tree in a pruned depth-first search; and at each step of the search, the corresponding wavelet approximation coefficient is computed. The search is pruned whenever the wavelet detail coefficients are determined to be ‘‘insignificant’’. In this case, the wavelet approximation coefficient is directly approximated by the local value of the actual image using the relationship

$$f(q, k, i, j) \approx 2^k f(q, 0, 2^k i, 2^k j).\quad (25)$$

When the search is not pruned, then the wavelet approximation coefficient is computed using Eq. (24), where the finer resolution coefficients are evaluated using recursion. Of course, we also terminate the search when $k = 0$ since this is the finest resolution at which the wavelet approximation coefficient can be evaluated exactly. Figure 2 provides a graphic illustration of how the search transverses from the coarsest to finest resolutions, and Fig. 3 provides a more precise pseudo-code specification of the algorithm. Note that the pseudo-code uses a function *SignificantCoef* that determines which detail coefficients are significant. Section IV will discuss how to efficiently evaluate the function *SignificantCoef*.

After obtaining the necessary approximation coefficients, we can compute the significant detail coefficients using the

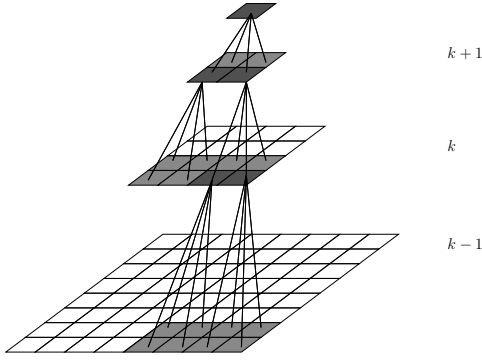


Fig. 2. An illustration of the quad-tree structure of Haar approximation coefficients. Dark gray squares are internal nodes of the pruned quad-tree. Corresponding detail coefficients are significant at these nodes, and we go on to the next level recursion. Light gray squares are leaf nodes of the pruned quad-tree. The termination condition of our recursive algorithm is met at these nodes; so the values are directly computed and returned. Empty squares are nodes pruned off from the quad-tree. Values at these nodes are not computed or touched.

following equations:

$$\begin{aligned}
 g_h(q, k, i, j) &= \frac{1}{2}(f(q, k-1, 2i, 2j) - f(q, k-1, 2i, 2j+1) \\
 &\quad + f(q, k-1, 2i+1, 2j) \\
 &\quad - f(q, k-1, 2i+1, 2j+1)), \\
 g_v(q, k, i, j) &= \frac{1}{2}(f(q, k-1, 2i, 2j) + f(q, k-1, 2i, 2j+1) \\
 &\quad - f(q, k-1, 2i+1, 2j) \\
 &\quad - f(q, k-1, 2i+1, 2j+1)), \\
 g_d(q, k, i, j) &= \frac{1}{2}(f(q, k-1, 2i, 2j) - f(q, k-1, 2i, 2j+1) \\
 &\quad - f(q, k-1, 2i+1, 2j) \\
 &\quad + f(q, k-1, 2i+1, 2j+1)).
 \end{aligned} \tag{26}$$

Notice that the complexity of our algorithm to compute $Q_t(SW_2^t \Lambda_w^{1/2})$ is $O(NP)$. To see that this is true, notice that the total complexity of the depth-first search is linearly proportional to the number of nodes in the tree; and the number of nodes in the tree is proportional to N . Moreover, the computation performed at each node has fixed computational complexity because we have assumed that the subroutine `SignificantCoef` of Fig. 3(c) has constant complexity. So putting this together, we know that the complexity of computing the wavelet transform for each row is $O(N)$. Thus, the total complexity of evaluating the wavelet transform for P rows must be order $O(NP)$.

The second stage of our matrix source coding is a wavelet transform W_1 along the columns of the sparse matrix $Q_t(SW_2^t \Lambda_w^{1/2})$. Storing the entire sparse matrix $Q_t(SW_2^t \Lambda_w^{1/2})$ prior to computing the wavelet transform W_1 may exceed the memory capacity of computers for high resolution images. So in practice, we use block wavelet transforms for W_1 to reduce memory usage. In this way, we can perform the transform block-by-block, and only need to keep $Q_t(SW_2^t \Lambda_w^{1/2})$ for each block prior to computing the wavelet transform W_1 . Once the transform W_1 for a block is complete, the values

```

Main routine:
K ← coarsest scale;
for all (i, j)
    f(q, K, i, j) ← FastApproxCoef(q, K, i, j, f);
end

```

(a)

```

float FastApproxCoef(q, k, i, j, f) {
    if (k==0) {
        f(q, 0, i, j) ← directly compute from PSF;
        return f(q, 0, i, j);
    }
    if ( SignificantCoef(q, k, i, j)==1 ) {
        f(q, k, i, j) = FastApproxCoef(q, k-1, 2i, 2j, f);
        f(q, k, i, j) += FastApproxCoef(q, k-1, 2i, 2j+1, f);
        f(q, k, i, j) += FastApproxCoef(q, k-1, 2i+1, 2j, f);
        f(q, k, i, j) += FastApproxCoef(q, k-1, 2i+1, 2j+1, f);
        f(q, k, i, j) ← f(q, k, i, j)/2;
    }
    else {
        f(q, k, i, j) ← 2k f(q, 0, 2ki, 2kj);
    }
    return f(q, k, i, j);
}

```

(b)

```

int SignificantCoef(q, k, i, j) {
    /* Define (iq, jq) to be the 2D index of row q. */
    iq = q / WIDTH;
    jq = q mod WIDTH;
    /* Let rk be a pre-computed value for level k. */
    if ((i -  $\frac{1}{2^k}i_q$ )2 + (j -  $\frac{1}{2^k}j_q$ )2 < rk2) {
        return 1;
    }
    else {
        return 0;
    }
}

```

(c)

Fig. 3. Pseudo code for fast computation of necessary Haar approximation coefficients. Part (a) is the main routine. Part (b) is a subroutine called by (a). Part (c) is a subroutine called by (b). The function `FastApproxCoef()` described in part (b) is a recursive function. The function `SignificantCoef()` described in part (c) determines whether any of the corresponding detail coefficients are significant. Note that values of r_k at level k in part (c) are pre-computed off-line. They represent radii of circles in different wavelet bands. Section IV discusses how to compute these values. After the main routine call, the necessary Haar approximation coefficients have been computed and stored in array f .

in $Q_t(SW_2^t \Lambda_w^{1/2})$ for this block may be discarded, and the process is repeated for each new block. This stage also has a total computational complexity of $O(NP)$. Thus the complete algorithm reduces the off-line computation of $[\tilde{S}]$ from $O(P^2)$ to $O(NP)$, making it feasible for implementation.

III. MODELS OF STRAY LIGHT REDUCTION

In this paper, we will use the example of stray light reduction to demonstrate the effectiveness of our matrix source coding algorithm. Stray light occurs in all optical imaging systems. It refers to the portion of the entering light flux that is misdirected to undesired locations in the image plane. It is sometimes referred to as lens flare or veiling glare [24]. Causes for this phenomenon include, but are not limited to: (1) Fresnel reflections from optical element surfaces; (2) scattering from surface imperfections on lens elements; (3) scattering from air bubbles in transparent glass or plastic lens elements; and (4)

	a (mm ⁻¹)	b (mm ⁻¹)	α	c (mm)
PSF parameter set 1	-1.65×10^{-4}	-8.35×10^{-4}	1.21	0.015
PSF parameter set 2	4×10^{-3}	0	1.15	0.02

TABLE I
TWO SETS OF PSF PARAMETERS

scattering from dust or other particles. Images contaminated by stray light suffer from reduced contrast and saturation.

We reduce the stray light by first formulating a PSF model of the imaging system and then inverting the model. This method is also used in previous work [23], [38], [39], [32]. The PSF model for a stray light contaminated optical imaging system is formulated as

$$y = ((1 - \beta)I + \beta S)x, \quad (27)$$

where I is an identity matrix, S accounts for stray light due to scattering and undesired reflections, and β represents the weight factor of the stray light. The entries of S are described by

$$S_{q,p} = s(i_q, j_q; i_p, j_p), \quad (28)$$

where (i_q, j_q) and (i_p, j_p) are the 2D indices corresponding to q and p respectively, and $s(i_q, j_q; i_p, j_p)$ is the value of the stray light PSF at (i_q, j_q) due to a point source located at (i_p, j_p) . We follow [23], [38], [39], [40] and model the stray light PSF as:

$$s(i_q, j_q; i_p, j_p) = \frac{1}{\gamma} \times \frac{1}{\left(1 + \frac{1}{i_p^2 + j_p^2} \left(\frac{(i_q i_p + j_q j_p - i_p^2 - j_p^2)^2}{(c + a(i_p^2 + j_p^2))^2} + \frac{(-i_q j_p + j_q i_p)^2}{(c + b(i_p^2 + j_p^2))^2} \right)\right)^\alpha}, \quad (29)$$

where γ is a normalizing constant that makes $\int_{i_q} \int_{j_q} s(i_q, j_q; i_p, j_p) di_q dj_q = 1$. An analytic expression for γ is $\gamma = \frac{\pi}{\alpha-1} (c + a(i_p^2 + j_p^2))(c + b(i_p^2 + j_p^2))$. Figure 4 shows example stray light PSFs due to different point source locations with the two sets of parameters in Table I. We can see that the PSF described by parameter set 2 is more space-varying than the PSF described by parameter set 1. The PSF parameters vary with different imaging systems. For a particular digital camera model and lens setting, we can use the apparatus and algorithm described in [23], [38], [39], [40] to estimate the specific stray light PSF parameters.

Once the model parameters are estimated, we follow the method in [41] to compute the restored image. The restoration formula is described in Eq. (5). Note that since our stray light PSF is global and space-varying, directly computing $z = Sy$ is quite expensive. For a 6 million pixel image, it requires 6 million multiplies per output pixel, and a total of 3.6×10^{13} multiplies. So our objective is to compress the transform matrix S using our matrix source coding technique described in Sec. II to reduce the on-line computation.

IV. PREDICTING LOCATIONS OF SIGNIFICANT WAVELET COEFFICIENTS

As discussed in Sec. II-B, we need a method to predict the location of the significant wavelet coefficients for each row of




Fig. 4. Examples of stray light PSFs with different parameters and due to different point source locations in the image. (a), (c), (e), and (g) show the stray light PSFs due to center, center right, bottom left, and top right point sources for PSF parameter set 1. (b), (d), (f), and (h) show the stray light PSFs for PSF parameter set 2. PSFs with parameter set 2 are more space-varying than with parameter set 1.

the matrix $SW_2^t \Lambda_w^{1/2}$. Notice that each row of S is an image, so the corresponding row of the matrix $SW_2^t \Lambda_w^{1/2}$ will be a scaled version of the wavelet transform of that image.

We use the stray light PSF as an example. Using stray light parameter set 1 described in Table I, we show the magnitude of a row of S in Fig. 5(a) for image resolution 256×256 . We show the magnitude of its corresponding eight-level Haar wavelet transform in Fig. 5(b). White pixels in Fig. 5(c) indicate the location of the largest (in magnitude) 1000 wavelet coefficients after being scaled by $\Lambda_w^{1/2}$. In order to predict the location of those 1000 significant wavelet coefficients, we use a circle, as depicted in Fig. 5(d), in each band to capture all significant wavelet coefficients. Then once we can predict the center location and radius of the circle in each band, we can locate the significant wavelet coefficients.

Given a row q of S , our prediction of the circle center

$(i_q^{(k)}, j_q^{(k)})$ in the k th level wavelet band is

$$\begin{aligned} i_q^{(k)} &= \frac{1}{2^k} i_q, \\ j_q^{(k)} &= \frac{1}{2^k} j_q, \end{aligned} \quad (30)$$

where $k = 1$ represents the finest resolution wavelet band. Then a wavelet coefficient in the k th level wavelet band is significant if its location (i, j) satisfies

$$\left(i - i_q^{(k)}\right)^2 + \left(j - j_q^{(k)}\right)^2 < r_k^2, \quad (31)$$

where r_k is the predicted radius of the circle in the k th level wavelet band.

We next need to predict the value of the radius, r_k , for each wavelet level k . In order to do that, we first need to decide the average number of significant wavelet coefficients we want to keep for each row of S . As mentioned before, we let this target number of significant wavelet coefficients per row be N . In the example shown in Fig. 5, $N = 1000$. We then take 49 rows of S with corresponding (i_q, j_q) 's uniformly distributed across the image. We take the eight-level Haar wavelet transform of those 49 row images, respectively. We then scale the resulting wavelet transforms using $\Lambda_w^{1/2}$, and keep only the largest N coefficients in each row. For each wavelet band, we find the maximum Euclidean distance between the location of the significant wavelet coefficients and the predicted center of the circle for each row, and then we take the maximum among all rows as the radius of the circle. After the center and radii of the circles are predicted, we consider all pixels falling inside the circles as significant wavelet coefficients.

V. EXPERIMENTAL RESULTS

In the following section, we evaluate the accuracy, computation, and memory trade-offs of the matrix source coding method using natural images, and we also compare it to alternative approaches such as the overlap-add method [26], [27], [28] using FFTs.

In all our experiments, we measure distortion using normalized root mean square error (NRMSE) computed as:

$$\text{NRMSE} = \sqrt{\frac{\|z - \hat{z}\|_2^2}{\|z\|_2^2}}. \quad (32)$$

All our on-line measurements of NRMSE are determined by averaging the NRMSE over a set of 10 natural test images.

We measure the on-line computation using the number of real multiplies per output pixel. The online computation includes a Haar wavelet transform W_2 , a scaling by $\Lambda_w^{-1/2}$, a multiplication with a sparse matrix $[\tilde{S}]$, and an inverse wavelet transform W_1^{-1} . The Haar wavelet transform W_2 and inverse wavelet transform W_1^{-1} can be implemented without multiplies. Scaling takes 1 multiply per output pixel. Multiplication with sparse matrix $[\tilde{S}]$ takes $||[\tilde{S}]||/P$ multiplies per output pixel, where $||[\tilde{S}]||$ represents the total number of non-zero entries in the quantized sparse matrix. So the total number of multiplies per output pixel is given by

$$\text{multiplies per output pixel} = 1 + ||[\tilde{S}]||/P. \quad (33)$$




Fig. 5. An example of a row of S displayed as an image together with its wavelet transform, and the locations of significant wavelet coefficients. (a) shows the amplitude map of a row image of S . (b) shows the absolute value of the scaled wavelet transform of the row image of S . (c) shows the locations of non-zero pixels after quantization of (b). These are the locations of significant wavelet coefficients. (d) zooms into one vertical band of (c), with a red circle to capture all significant wavelet coefficients in that band. Both (a) and (b) raise the image to the power $1/3$ in order to improve visibility of detail.

In many cases, computing the exact space-varying convolution result z over the whole image is prohibitively expensive. So we randomly select 1000 pixels in each image at which to compute the NRMSE. In addition, as described in Sec. II-A, we also use 30 natural images as training images to estimate the scaling matrix Λ_w . However, we only estimate a single scalar value for each subband. So a total of $3K + 1$ scalar values are estimated for Λ_w , where K is the number of levels of the wavelet transform W_2 .

Our experiments use two different parameter sets to model the stray-light PSF parameters. These two parameter sets are given in Table I and shown in Figure 4. Notice that parameter set 2 is more space-varying than parameter set 1.

In the following experiments, we first show the result if we only decorrelate the columns of S . Then we show the result if we decorrelate both the columns and the rows of S . We also compare our algorithm with the overlap-add method. Finally, we apply our matrix source coding algorithm to stray light reduction in natural images.

A. Results of Decorrelating Only the Columns of S

In this subsection, we only decorrelate the columns of S . So we approximate z as

$$\begin{aligned} \hat{z} &= [Q_t(SW_2^{-1}\Lambda_w^{1/2})]\Lambda_w^{-1/2}W_2y, \\ &\approx z. \end{aligned} \quad (34)$$

In this experiment, the image resolution is 256×256 . We use the eight-level Haar wavelet transform for W_2 . We choose to use stray light PSF parameter set 1 described in Table I in this experiment.

When we use our fast algorithm to compute the significant Haar wavelet coefficients for the rows of S , we can choose a number N as the target number of significant wavelet coefficients per row to compute before quantization. We conduct three experiments with $N = 1000$, $N = 2000$, and $N = 4000$, respectively. We then vary quantization intervals when quantizing $Q_t(SW_2^{-1}\Lambda_w^{1/2})$ in Eq. (34), and plot the NRMSE as a function of the number of multiplies per output pixel for these three choices in Fig. 6. From the plot, we can see that the more significant wavelet coefficients we compute, the better distortion-computation tradeoff we can achieve. However, more coefficients to compute also means more off-line computation. In addition, if we take a closer look at Fig. 6, we will find that $N = 4000$ produces little improvement compared to $N = 2000$. However, the off-line computation is approximately twice as great. So in the following experiments, unless otherwise specified, we choose $N = 2000$ for image resolution 256×256 . Figure 6 also shows that by only decorrelating the columns S , our matrix source coding algorithm produces 1.2% error with approximately 44 multiplies per output pixel.

B. Results of Decorrelating Both the Columns and Rows of S

In addition to decorrelating the columns, we also decorrelate the rows with wavelet transform W_1 . Then z is approximated

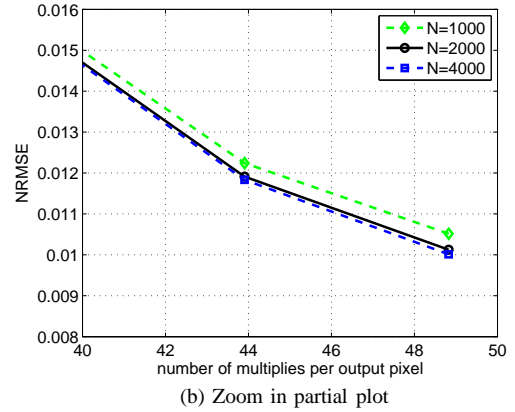
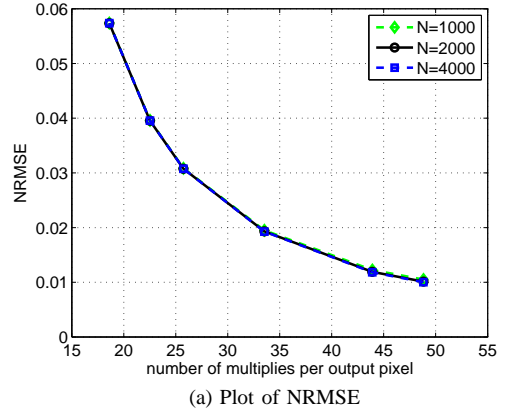


Fig. 6. Plot of NRMSE as a function of the number of multiplies per output pixel using our matrix source coding algorithm, but only decorrelating the columns of S . The green dashed line, the black solid line, and the blue dashed line represent the results when the target number of significant wavelet coefficients per row $N = 1000$, $N = 2000$, and $N = 4000$, respectively. We can see that for the same number of multiplies per output pixel, increasing N reduces NRMSE, although the difference is small.

as

$$\begin{aligned} \hat{z} &= W_1^{-1}[W_1Q_t(SW_2^t\Lambda_w^{1/2})]\Lambda_w^{-1/2}W_2y, \\ &\approx z. \end{aligned} \quad (35)$$

Figure 7 compares the distortion-computation trade-off for these two scenarios when we only decorrelate columns; and when we decorrelate both columns and rows. Both plots show the NRMSE versus the number of multiplies per output pixel. We can see that to achieve 1% error, decorrelating only columns requires approximately 48 multiplies per output pixel, whereas decorrelating both rows and columns requires approximately 14 multiplies per output pixel. Therefore, this experiment demonstrates the importance of decorrelating both rows and columns of the matrix S .

In addition, to better understand the computational savings, we use two different image resolutions of 256×256 and 1024×1024 in our experiments. In both cases, we set the target number of significant wavelet coefficients per row to be $N = 2000$, and we use the two sets of PSF parameters described in Table I. In the 1024×1024 case, we use a ten-level Haar wavelet transform for W_2 and a three-level block Haar wavelet transform for W_1 .

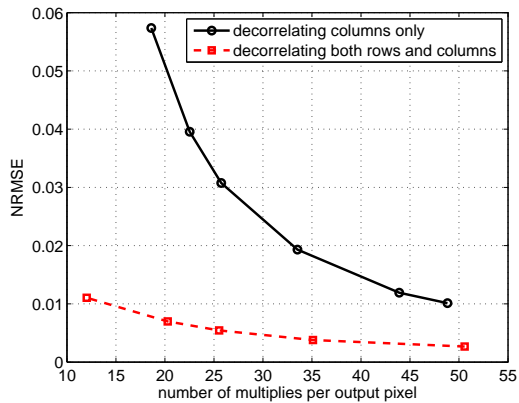



Fig. 7. Comparison of the distortion-computation tradeoff between decorrelating only columns and decorrelating both rows and columns. Notice that the decorrelation of both rows and columns produces a substantially more favorable tradeoff.

Figure 8 shows the tradeoff between distortion and computation for the two different resolutions. We can see that for PSF parameter set 1, at 1024×1024 resolution, with only 1% error, the matrix source coding algorithm reduces the number of multiplies per output pixel from $1048576 = 1024^2$ to only 7, which is a $149796 : 1$ reduction in computation. At 256×256 resolution, to achieve the same accuracy, we need approximately 14 multiplies per output pixel. So our algorithm achieves a higher reduction in computation for higher resolutions. In addition, we find that the error for the 1024×1024 case does not decrease much when the number of multiplies per pixel becomes large. That is because the error introduced in the first stage quantization $Q_t(SW_2^t \Lambda_w^{1/2})$ does not change for different quantization intervals in the second stage.

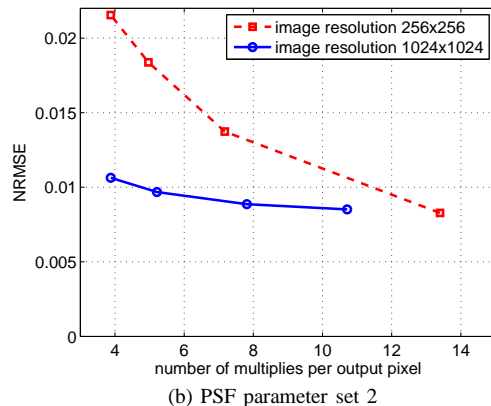
Figure 9 compares the effect of increasing the number of significant wavelet coefficients, N , for the two parameter sets and two image resolutions. For the 256×256 case, the distortion-computation tradeoff is almost the same for $N = 2000$ and $N = 4000$. This is because the error introduced in the first stage quantization is almost identical for $N = 2000$ and $N = 4000$, as shown in Fig. 6. However, for the 1024×1024 case, the increased value of N significantly reduces distortion.

C. Comparison with Overlap-add Method

A conventional way to compute space-varying convolution is to use the overlap-add method [27], [28]. This method first partitions the input image y into blocks, then computes the output for each block over a larger (overlapped) region, and finally adds the results together to approximate the final output image z . The reason why it can be fast is that when computing the output for each block, it assumes a spatially invariant PSF for that particular block, thus the fast Fourier transform (FFT) can be used to significantly reduce the computational complexity. Similar to our algorithm, this method also has a trade-off between speed and accuracy. More block partitions will produce higher accuracy, but requires more computation. The size of the output region for each block depends on the



(a) PSF parameter set 1



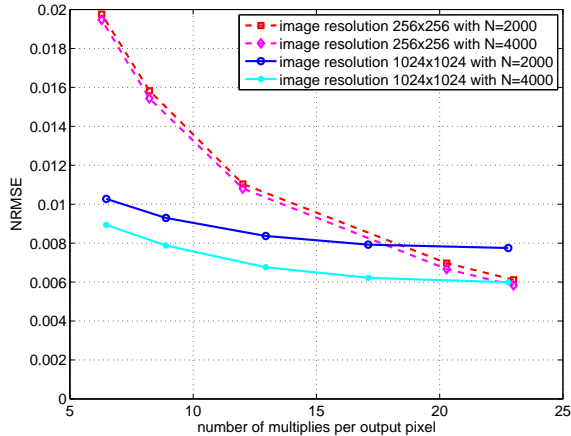
(b) PSF parameter set 2

Fig. 8. Comparison of the distortion-computation tradeoff of the matrix source coding algorithm using two different resolutions. The red dashed line shows the result for an image resolution of 256×256 , and the blue solid line shows the result for an image resolution of 1024×1024 . (a) and (b) show the results for PSF parameter set 1 and set 2, respectively.


size of the block and the spatial support of the PSF. If the block is of size $h \times h$, and the PSF has a spatial support of $(r+1) \times (r+1)$, then the output region for each block should be of size $(r+h) \times (r+h)$ [27]. Since the spatial support of stray light PSF is comparable to the size of the image, the output region for each block will be the entire image area.

Now we compare our algorithm with the overlap-add method [27] in terms of the amount of computation, storage requirement, and accuracy. The total amount of computation for the overlap-add method is $2P(M+1)\log_2(P) + 4MP$ real multiplies, where M is the number of blocks and P is the number of pixels in the image. Equivalently, the number of multiplies per pixel is $2(M+1)\log_2(P) + 4M$. To see this, notice that a discrete Fourier transform (DFT) of each of the input image blocks costs $2P\log_2(P)$ real multiplies using the conventional FFT [42]. Multiplication with the DFT of the PSF requires $4P$ real multiplies for each block. An inverse DFT takes $2P\log_2(P)$ real multiplies. So this results in a total of $2P(M+1)\log_2(P) + 4MP$ real multiplies.

In terms of storage or memory usage, the overlap-add method requires storing the Fourier transform of the PSF corresponding to each of the M blocks for a total of $2PM$ real-valued coefficients. So the overlap-add method requires storing $2M$ real values per output pixel. For our matrix source



(a) PSF parameter set 1



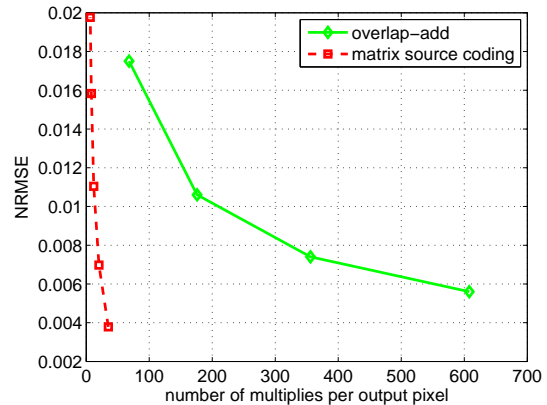
(b) PSF parameter set 2

Fig. 9. Comparison of distortion-computation tradeoff using different values of N , the target number of significant wavelet coefficients per row. Notice that a larger value of N is significant for the case of a larger image.


coding method, the storage requirement is $|\tilde{S}|/P$ real values per output pixel.

We perform experiments using the two sets of stray light PSF parameters given in Table I. We test on 10 natural images with resolution 256×256 . We use different numbers of blocks for the overlap-add method: $M = 1$, $M = 4$, $M = 9$, and $M = 16$. We then plot the average NRMSE versus the number of multiplies per output pixel in Fig. 10 to get a straightforward comparison with our matrix source coding method. We can see that our matrix source coding method outperforms the overlap-add method by a large margin in terms of the distortion-computation tradeoff. In addition, the distortion-computation tradeoff of our matrix source coding method does not change substantially for the two different sets of PSF parameters, whereas the distortion-computation tradeoff of the overlap-add method deteriorates significantly when the PSF becomes more space-varying with parameter set 2.

In Fig. 11, we also plot the average NRMSE versus the number of stored real values. For PSF parameter set 1, the overlap-add method does reasonably well in terms of its distortion-storage tradeoff. This is because it only must store M coefficients per pixel, where M is the number of blocks



(a) PSF parameter set 1



(b) PSF parameter set 2

Fig. 10. Plots of NRMSE versus the number of multiplies per output pixel for both the overlap-add method and matrix source coding algorithm over two different PSF parameter sets specified in Table I. The green solid line shows the result of the overlap-add method. The four data points shown correspond to $M = 1, 4, 9$, and 16 blocks, respectively. The red dashed line shows the result of matrix source coding algorithm. The PSF using parameter set 2 is more space-varying than the PSF using parameter set 1.

used in the overlap-add method. If the PSF is only slowly space varying, then M can be very small, and the storage requirements are modest. However, for the more rapidly varying PSF of parameter set 2, the matrix source coding method has a better distortion-storage tradeoff because M must be larger for the overlap-add method in this case. For example, with 5 real values stored per output pixel, the matrix source coding method produces an error of approximately 1.5%, whereas the overlap-add method produces an error of approximately 6%.

In order to demonstrate the results more intuitively, we use the PSF parameter set 2 provided in Table I, and compute the exact convolution of the stray light PSF with a sample input image shown in Fig. 12(a) at resolution 256×256 . The exact convolution result is shown in Fig. 12(b). We show the approximation results using the overlap-add method with $M = 1, 4$, and 16 blocks in Fig. 13 (a), (c), and (e) respectively. The absolute difference images between the approximation results and the ground truth in Fig. 12 (b) are shown in Fig. 13 (b), (d), and (f). The NRMSE is 9.89% for one block, 6.25% for four blocks, and 3.96% for sixteen blocks. We can see that the error is much reduced by using

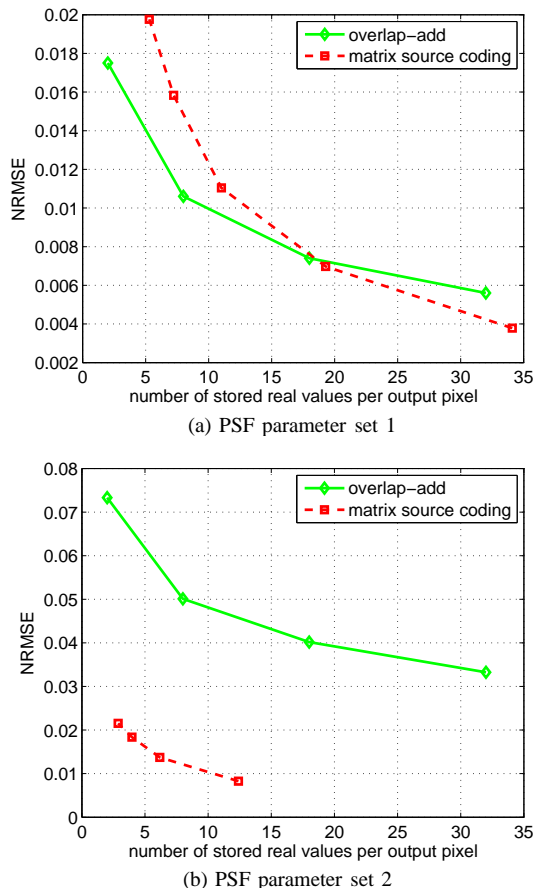


Fig. 11. Plots of NRMSE versus storage for both the overlap-add method and matrix source coding algorithm over two different PSF parameter sets specified in Table I. The green solid line shows the result of the overlap-add method. The red dashed line shows the result of matrix source coding algorithm. The PSF using parameter set 2 is more space-varying than the PSF using parameter set 1.




Fig. 12. A sample test image for space-varying convolution with stray light PSF. (a) shows an input image. (b) shows the exact result of the input image convolved with a space-varying stray light PSF with parameters set 2.

more blocks. We then compute the approximation using our matrix source coding algorithm with 15 multiplies per output pixel and show the result in Fig. 13(g). The difference image between our approximation and the ground truth is shown in Fig. 13(h). The NRMSE is 0.74%. This error is even less than that achieved by the overlap-add method with sixteen blocks, which needs 608 multiplies per output pixel.

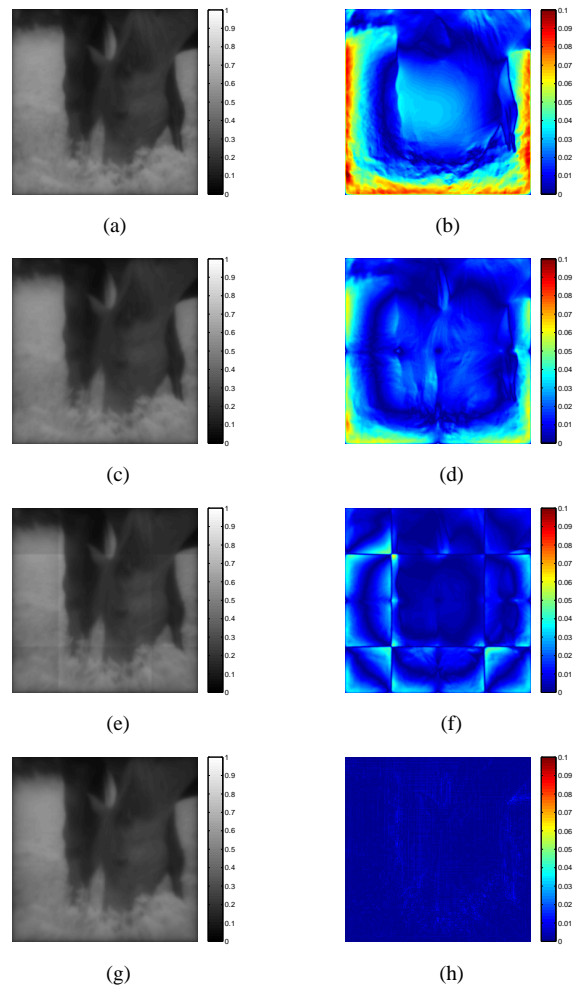
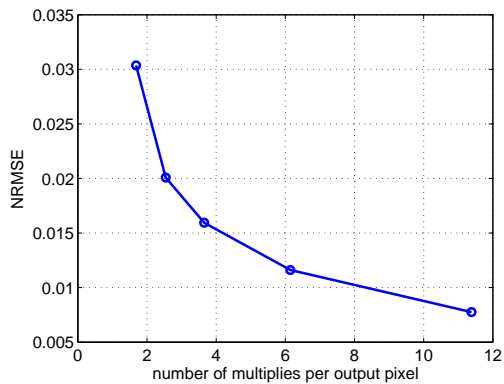


Fig. 13. Approximation to space-varying convolution with stray light PSF (parameter set 2) using the overlap-add method and our matrix source coding method. (a), (c), and (e) show the result using the overlap-add method with 1, 4, and 16 blocks, respectively. (b), (d), and (f) show the absolute difference images between ground truth (shown in Fig. 12(b)) and the approximation result in (a), (c), and (e), respectively. (g) shows the approximation result using matrix source coding method with 15 multiplies per output pixel, and Figure (h) shows the absolute difference image between (g) and the ground truth. The results (a), (c), (e), and (g) have NRMSE of 0.0989, 0.0625, 0.0396, and 0.0074, respectively, with corresponding number of multiplies per pixel 68, 176, 608, and 15.

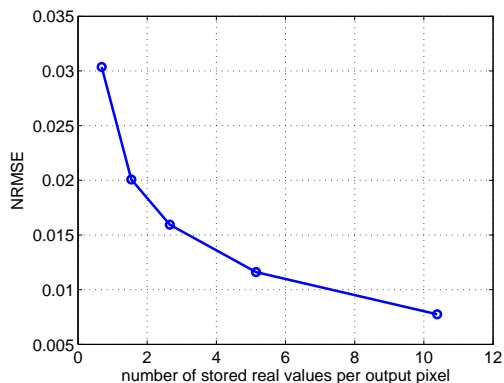
D. Matrix Source Coding Results for Stray Light Reduction

We use Eq. (5) to perform stray light reduction for images taken by an Olympus SP-510UZ camera⁴. The estimated stray light PSF parameters are $a = -1.65 \times 10^{-5} \text{ mm}^{-1}$, $b = -5.35 \times 10^{-5} \text{ mm}^{-1}$, $\alpha = 1.31$, $c = 1.76 \times 10^{-3} \text{ mm}$, $\beta = 0.3937$. In the stray light reduction formula described by Eq. (5), the space-varying convolution Sy is efficiently computed using our matrix source coding approach. The image resolution of this camera is 3088×2310 . In this experiment, we use a ten-level Haar wavelet transform for W_2 , and a three-level block Haar wavelet transform for W_1 , with block size 64×64 . We randomly select 1000 pixels on each of a new set of 10 testing images that were captured with this camera, and calculate the average NRMSE in computing Sy

⁴Olympus America Inc., Center Valley, PA 18034



(a)



(b)

Fig. 14. Plot of NRMSE for computing Sy versus the number of multiplies per output pixel and the number of stored real values per output pixel for a real camera stray light PSF for ten testing images with resolution 3088×2310 captured with this camera. (a) shows the plot of NRMSE versus the number of multiplies per output pixel. (b) shows the plot of NRMSE versus the number of stored real values per output pixel.

of Eq. (5). We vary the quantization intervals in the second stage quantization of our matrix source coding, and plot the average NRMSE versus the number of multiplies per output pixel and the number of stored real values per output pixel in Fig. 14.

We select a quantization level such that the online computation requires 6 multiplies per output pixel, and the corresponding storage requirement is 5 real values per output pixel. We then use our matrix source coding technique to compute Sy and apply Eq. (5) to perform stray light reduction. Figure 15 shows an example of the results of the stray light reduction. Figure 15(a) is the captured image. Figure 15(b) is the restored image. Figures 15(c-d) shows a comparison between the captured and restored versions for different parts of the image.

From this example, we can see that the stray light reduction algorithm increases the contrast and recovers lost details of the original scene.

VI. CONCLUSION

In this paper, we proposed a matrix source coding algorithm for efficiently computing space-varying convolution. Our approach reduced the computation from $O(P^2)$ to $O(P)$, where



(a) Captured image.



(b) Restored image.



(c) Captured \ restored.



(d) Captured \ restored.

Fig. 15. Example of stray light reduction on a 7 mega-pixel camera image: (a) shows a captured image, (b) shows the restored image, (c) and (d) show a comparison between captured and restored images for two parts of the image.

P is the number of pixels in the image. We presented both the on-line and the off-line parts of our algorithm, along with its associated theory. In the on-line part, we developed an algorithm to approximate a dense transform with wavelet transforms and a sparse matrix-vector product.

In the off-line part, we developed a two stage approach to make the sparsification of the dense transform matrix feasible. We also developed a fast algorithm for computing significant Haar wavelet coefficients of PSF images to accelerate our off-line computation. We use a top-down recursive approach, rather than the bottom-up approach used by a conventional filter-bank implementation. The value of this algorithm is to obtain the significant wavelet coefficients with $O(N)$ complexity, where N is the number of significant wavelet coefficients. By using this algorithm and our two stage approach, our off-line matrix source encoding process has a computational complexity of $O(NP)$.

Finally, we applied our matrix source coding algorithm to stray light reduction, where space-varying convolution with a PSF of large support is necessary. Our experimental results showed a trade-off between speed and accuracy. Our algorithm is able to achieve a 149796:1 reduction in computation with only 1% error at image resolution 1024×1024 (see Sec. V-B). The experimental results also demonstrated that our matrix source coding algorithm outperforms the overlap-add method by a large margin.

APPENDIX A

IDEAL DECORRELATING TRANSFORMS FOR MATRIX SOURCE CODING

Lemma 1: Let S be any $P \times P$ transformation, and let $y \in \mathbb{R}^P$ be a zero-mean random vector with covariance $R_y = E[y y^t]$, and eigen decomposition given by $R_y = E_y \Lambda_y E_y^t$. Then we define the matrix

$$F = S E_y \Lambda_y^{1/2}, \quad (36)$$

with a corresponding singular value decomposition given by

$$F = U \Sigma V^t, \quad (37)$$

where V is a square matrix. Then using Definition 1, the following choices of (W_1, T) are optimal for S and y ,

$$T = V^t \Lambda_y^{-1/2} E_y^t, \quad (38)$$

and

$$W_1 = U^t, \quad (39)$$

with $\tilde{S} = \Sigma$.

Furthermore, if the eigen values, Λ_y , and the singular values, Σ , are all non-degenerate (i.e., unique), then the transforms W_1 and T are unique up to permutation of rows/columns and multiplication by ± 1 .

The proof of the first part of the lemma only requires that we verify the properties of the Definition 1. To verify this, we

first show that \tilde{y} is whitened.

$$\begin{aligned} E[\tilde{y} \tilde{y}^t] &= T R_y T^t \\ &= V^t \Lambda_y^{-1/2} E_y^t R_y E_y \Lambda_y^{-1/2} V \\ &= V^t \Lambda_y^{-1/2} \Lambda_y \Lambda_y^{-1/2} V \\ &= V^t V = I. \end{aligned} \quad (40)$$

Next we show that $\tilde{S} = \Sigma$.

$$\begin{aligned} \tilde{S} &= W_1 S T^{-1} \\ &= U^t S E_y \Lambda_y^{1/2} V \\ &= U^t F V = \Sigma. \end{aligned} \quad (41)$$

So \tilde{S} is a diagonal matrix containing the singular values of the matrix $F = S E_y \Lambda_y^{1/2}$; and since it is diagonal, it has orthogonal rows and columns.

Now if the eigen values and singular values are all non-degenerate, then both the eigen decomposition and singular value decomposition are unique up to a permutation of the eigen vectors/singular vectors and multiplication of the eigen vectors/singular vectors by ± 1 . So this implies that W_1 is unique up to a permutation of its rows and multiplication of them by ± 1 , and T^{-1} is unique up to the same permutation/multiplication of its columns.


REFERENCES

- [1] A. Bovik, *Handbook of Image and Video Processing*. San Diego, CA: Academic Press, 2000.
- [2] G. T. Herman and A. Kuba, *Discrete Tomography: Foundations, Algorithms, and Applications*. New York, NY: Birkhauser Boston, 1999.
- [3] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 21–36, May 2003.
- [4] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2001.
- [5] P. Gilbert, "Iterative methods for the three-dimensional reconstruction of an object from projections," *Journal of Theoretical Biology*, vol. 36, no. 1, pp. 105–117, July 1972.
- [6] A. H. Anderson and A. C. Kak, "Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm," *Ultrasonic Imaging*, vol. 6, no. 1, pp. 81–94, January 1984.
- [7] W. H. Richardson, "Bayesian-based iterative method of image restoration," *Journal of the Optical Society of America*, vol. 62, no. 1, pp. 55–59, 1972.
- [8] P. H. V. Cittert, "Zum einfluss der spaltbreite auf die intensitätsverteilung in spektrallinien," *Z. Physik*, vol. 69, pp. 298–308, 1931.
- [9] H. M. Hudson and R. S. Larkin, "Accelerated image reconstruction using ordered subsets of projection data," *IEEE Transactions on Medical Imaging*, vol. 13, no. 4, pp. 601–609, 1994.
- [10] L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Transactions on Medical Imaging*, vol. 2, pp. 113–122, 1982.
- [11] A. N. Tychonoff and V. Y. Arsenin, *Solution of Ill-posed Problems*. Washington D. C.: V. H. Winston & Sons, January 1977.
- [12] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors," *IEEE Transactions on Medical Imaging*, vol. 8, pp. 194–202, 1989.
- [13] P. J. Green, "Bayesian reconstruction from emission tomography data using a modified EM algorithm," *IEEE Transactions on Medical Imaging*, vol. 9, pp. 84–93, 1990.
- [14] C. Bouman and K. Sauer, "A generalized gaussian image model for edge-preserving map estimation," *IEEE Transactions on Image Processing*, vol. 2, no. 3, pp. 296–310, July 1993.
- [15] A. Levin, R. Fergus, F. Durand, and W. Freeman, "Deconvolution using natural image priors," *Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory*, 2007.

- [16] D. Krishnan and R. Fergus, "Fast image deconvolution using hyper-Laplacian priors," *Advances in Neural Information Processing Systems*, pp. 1033–1041, 2009.
- [17] D. Zoran and Y. Weiss, "From learning models of natural image patches to whole image restoration," in *2011 IEEE International Conference on Computer Vision*, 2011, pp. 479–486.
- [18] U. Schmidt, K. Schelten, and S. Roth, "Bayesian deblurring with integrated noise estimation," in *2011 IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 2625–2632.
- [19] C. Schuler, H. Burger, S. Harmeling, and B. Scholkopf, "A machine learning approach for non-blind image deconvolution," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [20] C. Lanczos, *Linear Differential Operators*. Mineola, NY: Dover Publications, 1997.
- [21] S. Mallat, *A Wavelet Tour of Signal Processing*. San Diego, CA: Academic Press, 1998.
- [22] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, April 1965.
- [23] J. Wei, B. Bitlis, A. Bernstein, A. Silva, P. A. Jansson, and J. P. Allebach, "Stray light and shading reduction in digital photography – a new model and algorithm," in *Proceedings of the SPIE/IS&T Conference on Digital Photography IV*, vol. 6817, San Jose, CA, January 2008.
- [24] W. J. Smith, *Modern Optical Engineering : the Design of Optical Systems*. New York, NY: McGraw Hill, 2000.
- [25] P. A. Jansson and J. H. Fralinger, "Parallel processing network that corrects for light scattering in image scanners," *U.S. Patent 5,153,926*, 1992.
- [26] T. G. Stockham, "High-speed convolution and correlation," in *Proceedings of the April 26-28, 1966, Spring joint computer conference*, 1966, pp. 229–233.
- [27] J. G. Nagy and D. P. O'Leary, "Fast iterative image restoration with a spatially varying PSF," in *Proceedings of SPIE Conference on Advanced Signal Processing: Algorithms, Architectures, and Implementations VII*, vol. 3162, San Diego, CA, 1997, pp. 388–399.
- [28] J. Bardsley, S. Jefferies, J. Nagy, and R. Plemmons, "A computational method for the restoration of images with an unknown, spatially-varying blur," *Optics Express*, vol. 15, no. 5, pp. 1767–1782, 2006.
- [29] M. Hirsch, S. Sra, B. Scholkopf, and S. Harmeling, "Efficient filter flow for space-variant multiframe blind deconvolution," in *Proceedings of 2010 IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 607–614.
- [30] E. Gilad and J. von Hardenberg, "A fast algorithm for convolution integrals with space and time variant kernels," *Journal of Computational Physics*, vol. 216, pp. 326–336, 2006.
- [31] L. Denis, E. Thiebaut, and F. Soublez, "Fast model of space-variant blurring and its application to deconvolution in astronomy," in *2011 18th IEEE International Conference on Image Processing*, 2011, pp. 2817–2820.
- [32] J. Wei, G. Cao, C. A. Bouman, and J. P. Allebach, "Fast space-varying convolution and its application in stray light reduction," in *Proceedings of the SPIE/IS&T Conference on Computational Imaging VII*, vol. 7246, San Jose, CA, February 2009.
- [33] G. H. Golub and C. F. V. Loan, *Matrix Computations*. Baltimore, Maryland: The Johns Hopkins University Press, 1996.
- [34] G. Cao, C. A. Bouman, and K. J. Webb, "Fast and efficient stored matrix techniques for optical tomography," in *Proceedings of the 40th Asilomar Conference on Signals, Systems, and Computers*, October 2006.
- [35] —, "Results in non-iterative map reconstruction for optical tomography," in *Proceedings of the SPIE/IS&T Conference on Computational Imaging VI*, vol. 6814, San Jose, CA, January 2008.
- [36] D. Tretter and C. A. Bouman, "Optimal transforms for multispectral and multilayer image coding," *IEEE Transactions on Image Processing*, vol. 4, no. 3, pp. 296–308, 1995.
- [37] D. Taubman and M. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*. Norwell, MA: Kluwer Academic Publishers, 2002.
- [38] B. Bitlis, P. A. Jansson, and J. P. Allebach, "Parametric point spread function modeling and reduction of stray light effects in digital still cameras," in *Proceedings of the SPIE/IS&T Conference on Computational Imaging VI*, vol. 6498, San Jose, CA, January 2007.
- [39] P. A. Jansson, "Method, program, and apparatus for efficiently removing stray-flux effects by selected-ordinate image processing," *U.S. Patent 6,829,393*, 2004.
- [40] P. A. Jansson and R. P. Breault, "Correcting color-measurement error caused by stray light in image scanners," in *Proceedings of the Sixth*

Color Imaging Conference: Color Science, Systems, and Applications, Scottsdale, AZ, November 1998.

- [41] P. A. Jansson, *Deconvolution of Images and Spectra*. New York, NY: Academic Press, 1996.
- [42] P. S. R. Diniz, E. A. B. da Silva, and S. L. Netto, *Digital Signal Processing System Analysis and Design*. New York, NY: Cambridge University Press, 2002.



Jianing Wei received a B.S. degree from Shanghai Jiao Tong University in 2005, and a Ph.D. degree in Electrical Engineering from Purdue University in 2010. His Ph.D. work includes stray light reduction in digital cameras, algorithms for fast space-varying convolution, and functional MRI analysis. Between 2009 and 2013, he was with the U.S. Research Center of Sony Corporation, San Jose, CA, where he focused on computer vision and image processing algorithm development for digital cameras and televisions. Dr. Wei joined Google Inc., Mountain


View, CA, in 2013, where he is currently a software engineer working on imaging algorithm development for android cameras.



Charles A. Bouman received a B.S.E.E. degree from the University of Pennsylvania in 1981 and a MS degree from the University of California at Berkeley in 1982. From 1982 to 1985, he was a full staff member at MIT Lincoln Laboratory and in 1989 he received a Ph.D. in electrical engineering from Princeton University. He joined the faculty of Purdue University in 1989 where he is currently the Michael J. and Katherine R. Birck Professor of Electrical and Computer Engineering. He also holds a courtesy appointment in the School of Biomedical

Engineering and is co-director of Purdue's Magnetic Resonance Imaging Facility located in Purdue's Research Park.

Professor Bouman's research focuses on the use of statistical image models, multiscale techniques, and fast algorithms in applications including tomographic reconstruction, medical imaging, and document rendering and acquisition. Professor Bouman is a Fellow of the IEEE, a Fellow of the American Institute for Medical and Biological Engineering (AIMBE), a Fellow of the society for Imaging Science and Technology (IS&T), a Fellow of the SPIE professional society. He is also a recipient of IS&T's Raymond C. Bowman Award for outstanding contributions to digital imaging education and research, has been a Purdue University Faculty Scholar, received the College of Engineering Engagement/Service Award, and Team Award, and in 2014 he received the Electronic Imaging Scientist of the Year award. He was previously the Editor-in-Chief for the IEEE Transactions on Image Processing and a Distinguished Lecturer for the IEEE Signal Processing Society, and he is currently the Vice President of Technical Activities for IEEE Signal Processing Society. He has been an associate editor for the IEEE Transactions on Image Processing and the IEEE Transactions on Pattern Analysis and Machine Intelligence. He has also been Co-Chair of the 2006 SPIE/IS&T Symposium on Electronic Imaging, Co-Chair of the SPIE/IS&T conferences on Visual Communications and Image Processing 2000 (VCIP), a Vice President of Publications and a member of the Board of Directors for the IS&T Society, and he is the founder and Co-Chair of the SPIE/IS&T conference on Computational Imaging.



Jan P. Allebach is Hewlett-Packard Distinguished Professor of Electrical and Computer Engineering at Purdue University. His current research interests include image rendering, image quality, color imaging and color measurement, document aesthetics, and printer forensics. Allebach is a Fellow of the IEEE, the Society for Imaging Science and Technology (IS&T), and SPIE. He was named Electronic Imaging Scientist of the Year by IS&T and SPIE, and was named Honorary Member of IS&T, the highest award that IS&T bestows. He was the recipient of the 2013 IEEE Daniel E. Noble Award, and was elected to membership in the National Academy of Engineers in 2014, both on the basis of his work on digital halftoning.