UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

Cosine Transform Preconditioners for High Resolution Image Reconstruction

Michael K. Ng
Raymond H. Chan
Tony F. Chan
Andy M. Yip

April 1999

CAM Report 99-17

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90095-1555

Cosine Transform Preconditioners for High Resolution Image Reconstruction*

Michael K. Ng† Raymond H. Chan‡ Tony F. Chan§ Andy M. Yip¶

Abstract

This paper studies the application of preconditioned conjugate gradient methods in high resolution image reconstruction problems. We consider reconstructing high resolution images from multiple undersampled, shifted, degraded frames with subpixel displacement errors. The resulting blurring matrices are spatially variant. The classical Tikhonov regularization and the Neumann boundary condition are used in the reconstruction process. The preconditioners are derived by taking the cosine transform approximation of the blurring matrices. We prove that when the L_2 or H_1 norm regularization functional is used, the spectra of the preconditioned normal systems are clustered around 1 for sufficiently small subpixel displacement errors. Conjugate gradient methods will hence converge very quickly when applied to solving these preconditioned normal equations. Numerical examples are given to illustrate the fast convergence.

1 Introduction

Due to hardware limitations, imaging systems often provide us with only multiple low resolution images. However, in many applications, a high resolution image is desired. For example, the resolution of the pictures of the ground taken from a satellite is relatively low and retrieving details on the ground becomes impossible. Increasing the image resolution by using digital signal processing technique [3, 11, 15, 17, 19, 20] is therefore of great interest.

*The authors would like to dedicate this work to Prof. Robert Plemmons in celebration of his 60th birthday.

†E-mail: mng@maths.hku.hk. Department of Mathematics, The University of Hong Kong, Pokfulam Road, Hong Kong. Research supported in part by HKU CRCG grant no. 10201939.

‡E-mail: rchan@math.cuhk.edu.hk. Department of Mathematics, The Chinese University of Hong Kong, Shatin, Hong Kong. Research supported in part by Hong Kong Research Grants Council Grant No. CUHK 4207/97P and CUHK DAG Grant No. 2060143.

§E-mail: chan@math.ucla.edu. Department of Mathematics, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90095-1555. Research supported by ONR under grant N00014-96-1-0277 and NSF under grant DMS 96-26755.

¶E-mail: mhyipa@hkusua.hku.hk. Department of Mathematics, The University of Hong Kong, Pokfulam Road, Hong Kong.

We consider the reconstruction of a high resolution image from multiple undersampled, shifted, degraded and noisy images. Multiple undersampled images are often obtained by using multiple identical image sensors shifted from each other by subpixel displacements. The reconstruction of high resolution images can be modeled as solving

$$\mathcal{H}f = g, \tag{1}$$

where g is the observed high resolution image formed from the low resolution images, f is the desired high resolution image and $\mathcal{H}$ is the reconstruction operator. If all the low resolution images are shifted from each others with exactly half-pixel displacements, $\mathcal{H}$ will be a spatially invariant operator. However, displacement errors present in practice, and the resulting operator $\mathcal{H}$ becomes spatially variant.

Since the systems are ill-conditioned and generally not positive definite, we solve them by using a minimization and regularization technique:

$$\min_f \left\{ \|\mathcal{H}f - g\|_2^2 + \alpha \mathcal{R}(f) \right\}. \tag{2}$$

Here $\mathcal{R}(f)$ is a functional which measures the regularity of f and the regularization parameter α is to control the degree of regularity of the solution. In this paper, we will use the L_2 and H_1 regularization functionals $\|f\|_2^2$ and $\|\mathcal{L}f\|_2^2$ where $\mathcal{L}$ is the first order differential operator.

Because of the blurring (convolution) process, the boundary values of g are not completely determined by the original image f inside the scene. They are also affected by the values of f outside the scene. Thus in solving f from (1), we need some assumptions on the values of f outside the scene. These assumptions are called boundary conditions. In [3], Bose and Boo used the traditional choice of imposing the zero boundary condition outside the scene, i.e., assuming a dark background outside the scene in the image reconstruction. However, when this assumption is not satisfied by the images, ringing effect will occur at the boundary of the reconstructed images. The problem is more severe if the images are reconstructed from a large sensor array since the number of pixel values of the image affected by the sensor array increases.

In this paper, we will use the Neumann boundary condition on the image, i.e., we assume that the scene immediately outside is a reflection of the original scene at the boundary. The Neumann boundary condition has been studied in image restoration [14, 1, 12] and in image compression [18, 13]. Our experimental results in [5] have shown that the Neumann image model gives better reconstructed high resolution images than that under the zero or periodic boundary conditions. In [5], we also proposed to use cosine transform preconditioners to precondition the resulting linear systems and preliminary numerical results have shown that these preconditioners are effective. The main aim of this paper is to analyze the convergence rate of these systems. We prove that when the L_2 or H_1 norm regularization functional is used, the spectra of the preconditioned systems are clustered around 1 for sufficiently small displacement errors.

The outline of the paper is as follows. In Section 2, we give a mathematical formulation of the problem. A brief introduction on the cosine transform preconditioners and the convergence analysis will be given in Section 3. In Section 4, numerical results are presented to demonstrate the effectiveness of the cosine transform preconditioners.

2 The Mathematical Model

We begin with a brief introduction of the mathematical model in high resolution image reconstruction. Details can be found in [3].

Consider a sensor array with $L_1 \times L_2$ sensors, each sensor has $N_1 \times N_2$ sensing elements (pixels) and the size of each sensing element is $T_1 \times T_1$. Our aim is to reconstruct an image of resolution $M_1 \times M_2$, where $M_1 = L_1 \times N_1$ and $M_2 = L_2 \times N_2$. To maintain the aspect ratio of the reconstructed image, we consider the case where $L_1 = L_2 = L$ only. For simplicity, we assume that L is an even number in the following discussion.

In order to have enough information to resolve the high resolution image, there are subpixel displacements between the sensors. In the ideal case, the sensors are shifted from each other by a value proportional to $T_1/L \times T_2/L$. However, in practice there can be small perturbations around these ideal subpixel locations due to imperfection of the mechanical imaging system. Thus, for $l_1, l_2 = 0, 1, \cdots, L-1$ with $(l_1, l_2) \neq (0, 0)$, the horizontal and vertical displacements $d^x_{l_1 l_2}$ and $d^y_{l_1 l_2}$ of the $[l_1, l_2]$-th sensor array with respect to the $[0, 0]$-th reference sensor array are given by

$$d^x_{l_1 l_2} = \frac{T_1}{L}(l_1 + \epsilon^x_{l_1 l_2}) \quad \text{and} \quad d^y_{l_1 l_2} = \frac{T_2}{L}(l_2 + \epsilon^y_{l_1 l_2}).$$

Here $\epsilon^x_{l_1 l_2}$ and $\epsilon^y_{l_1 l_2}$ denote respectively the normalized horizontal and vertical displacement errors.

We remark that the parameters $\epsilon^x_{l_1 l_2}$ and $\epsilon^y_{l_1 l_2}$ can be obtained by manufacturers during camera calibration. We assume that

$$|\epsilon^x_{l_1 l_2}| < \frac{1}{2} \quad \text{and} \quad |\epsilon^y_{l_1 l_2}| < \frac{1}{2}.$$

For if not, the low resolution images observed from two different sensor arrays will be overlapped so much that the reconstruction of the high resolution image is rendered impossible.

Let f be the original scene. Then the observed low resolution image $g_{l_1 l_2}$ for the (l_1, l_2)-th sensor is modeled by:

$$g_{l_1 l_2}[n_1, n_2] = \int_{T_2(n_2 - \frac{1}{2}) + d^y_{l_1 l_2}}^{T_2(n_2 + \frac{1}{2}) + d^y_{l_1 l_2}} \int_{T_1(n_1 - \frac{1}{2}) + d^x_{l_1 l_2}}^{T_1(n_1 + \frac{1}{2}) + d^x_{l_1 l_2}} f(x_1, x_2) dx_1 dx_2 + \eta_{l_1 l_2}[n_1, n_2], \tag{3}$$

for $n_1 = 1, \ldots, N_1$ and $n_2 = 1, \ldots, N_2$. Here $\eta_{l_1 l_2}$ is the noise corresponding to the (l_1, l_2)-th sensor. We intersperse the low resolution images to form an $M_1 \times M_2$ image by assigning

$$g[L(n_1 - 1) + l_1, L(n_2 - 1) + l_2] = g_{l_1 l_2}[n_1, n_2]. \tag{4}$$

Here g is an $M_1 \times M_2$ image and is called the *observed high resolution image*. Figure 1 shows the method of forming a 4×4 image g with a 2×2 sensor array where each g_{ij} has a 2×2 sensing elements, i.e. $L = 2$, $M_1 = M_2 = 4$, and $N_1 = N_2 = 2$.

Using a column by column ordering for g, we obtain $g = \mathcal{H}f + \eta$ where $\mathcal{H}$ is a spatially variant operator [3]. Since $\mathcal{H}$ is ill-conditioned due to the averaging of the pixel values in the

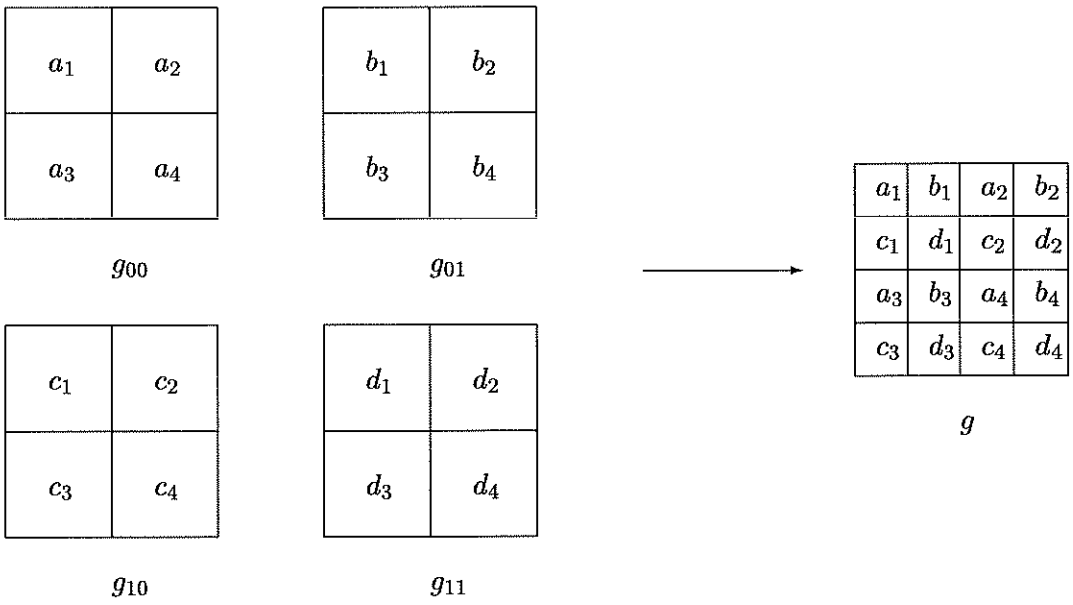

Figure 1: Construction of the observed high resolution image

image model in (3), the classical Tikhonov regularization is used and the minimization problem (2) is solved. In this paper, we use regularization functionals:

$$\mathcal{R}(f) = \|f\|_2^2 \quad \text{and} \quad \mathcal{R}(f) = \|\mathcal{L}f\|_2^2 \tag{5}$$

where $\mathcal{L}$ is the first order differential operator.

2.1 Image Boundary

The continuous image model in (3) can be discretized by rectangular rule and approximated by a discrete image model. Because of the blurring process (cf. (3)), the boundary values of g are also affected by the values of f outside the scene. Thus in solving f from (1), we need some assumptions on the values of f outside the scene. In [3], Bose and Boo imposed the zero boundary condition outside the scene, i.e., assuming a dark background outside the scene in the image reconstruction.

Let $\mathbf{g}$ and $\mathbf{f}$ be respectively the discretization of g and f using a column by column ordering. Under the zero boundary condition, the blurring matrix corresponding to the (l_1, l_2)-th sensor can be written as

$$\tilde{\mathbf{H}}_{l_1 l_2}(\epsilon) = \tilde{\mathbf{H}}^x_{l_1 l_2}(\epsilon) \otimes \tilde{\mathbf{H}}^y_{l_1 l_2}(\epsilon)$$

where $\tilde{\mathbf{H}}^x_{l_1l_2}(\epsilon)$ is an $M_1 \times M_1$ banded Toeplitz matrix with bandwidth $2L-1$

$$\tilde{\mathbf{H}}^x_{l_1l_2}(\epsilon) = \frac{1}{L}\begin{pmatrix} 1 & \cdots & 1 & h^{x+}_{l_1l_2} & & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \\ 1 & \ddots & \ddots & \ddots & \ddots & h^{x+}_{l_1l_2} \\ h^{x-}_{l_1l_2} & \ddots & \ddots & \ddots & \ddots & 1 \\ & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & & h^{x-}_{l_1l_2} & 1 & \cdots & 1 \end{pmatrix},$$

and

$$h^{x\pm}_{l_1l_2} = \frac{1}{2} \pm \epsilon^x_{l_1l_2}.$$

The $M_2 \times M_2$ banded blurring matrix $\tilde{\mathbf{H}}^y_{l_1l_2}(\epsilon)$ is defined similarly. We note that ringing effects will occur at the boundary of the reconstructed images if f is indeed not zero close to the boundary, see for instance Figure 3 in §4. The problem is more severe if the image is reconstructed from a large sensor array since the number of pixel values of the image affected by the sensor array increases.

In [5], we proposed to use the Neumann boundary condition on the image. It assumes that the scene immediately outside is a reflection of the original scene at the boundary. Our numerical results have shown that the Neumann boundary condition gives better reconstructed high resolution images than that by the zero or periodic boundary conditions. Under the Neumann boundary condition, the blurring matrices are still banded matrices with bandwidth $2L-1$, but the entries on the upper left part and the lower right part of the matrices are changed. The resulting matrices, denoted by $\mathbf{H}^x_{l_1l_2}(\epsilon)$ and $\mathbf{H}^y_{l_1l_2}(\epsilon)$, have a Toeplitz-plus-Hankel structure:

$$\mathbf{H}^x_{l_1l_2}(\epsilon) = \frac{1}{L}\begin{pmatrix} 1 & \cdots & 1 & h^{x+}_{l_1l_2} & & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \\ 1 & \ddots & \ddots & \ddots & \ddots & h^{x+}_{l_1l_2} \\ h^{x-}_{l_1l_2} & \ddots & \ddots & \ddots & \ddots & 1 \\ & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & & h^{x-}_{l_1l_2} & 1 & \cdots & 1 \end{pmatrix} + \frac{1}{L}\begin{pmatrix} 1 & \cdots & 1 & h^{x-}_{l_1l_2} & & 0 \\ \vdots & \iddots & \iddots & & & \\ 1 & \iddots & & & & h^{x+}_{l_1l_2} \\ h^{x-}_{l_1l_2} & & & & \iddots & 1 \\ & & & \iddots & \iddots & \vdots \\ 0 & & h^{x+}_{l_1l_2} & 1 & \cdots & 1 \end{pmatrix}, \tag{6}$$

and $\mathbf{H}^y_{l_1l_2}(\epsilon)$ is defined similarly. The blurring matrix corresponding to the (l_1, l_2)-th sensor under the Neumann boundary condition is given by

$$\mathbf{H}_{l_1l_2}(\epsilon) = \mathbf{H}^x_{l_1l_2}(\epsilon) \otimes \mathbf{H}^y_{l_1l_2}(\epsilon).$$

The blurring matrix for the whole sensor array is made up of blurring matrices from each sensor:

$$\mathbf{H}_L(\epsilon) = \sum_{l_1=0}^{L-1}\sum_{l_2=0}^{L-1} \mathbf{D}_{l_1l_2}\mathbf{H}_{l_1l_2}(\epsilon). \tag{7}$$

Here $\mathbf{D}_{l_1 l_2}$ are diagonal matrices with diagonal elements equal to 1 if the corresponding component of $\mathbf{g}$ comes from the (l_1, l_2)-th sensor and zero otherwise, see (4) or [3] for more details. With the Tikhonov regularization, our discretization problem becomes:

$$(\mathbf{H}_L(\epsilon)^t \mathbf{H}_L(\epsilon) + \alpha \mathbf{R})\mathbf{f} = \mathbf{H}_L(\epsilon)^t \mathbf{g} \tag{8}$$

where $\mathbf{R}$ is the discretization matrix corresponding to the regularization functional $\mathcal{R}(f)$ in (5).

3 Cosine Transform Based Preconditioners

The linear system (8) will be solved by using the preconditioned conjugate gradient method. In this section, we construct the cosine transform preconditioner of $\mathbf{H}_L(\epsilon)$ which exploits the banded and block structures of the matrix.

Let $\mathbf{C}_n$ be the $n \times n$ discrete cosine transform matrix, i.e., the (i, j)-th entry of $\mathbf{C}_n$ is given by

$$\sqrt{\frac{2 - \delta_{i1}}{n}} \cos\left(\frac{(i-1)(2j-1)\pi}{2n}\right), \qquad 1 \le i, j \le n,$$

where δ_{ij} is the Kronecker delta. Note that the matrix-vector product $\mathbf{C}_n \mathbf{z}$ can be computed in $O(n \log n)$ operations for any vector $\mathbf{z}$, see [13, pp. 59–60]. For an $m \times m$ block matrix $\mathbf{B}$ with the size of each block equals to $n \times n$, the cosine transform preconditioner $c(\mathbf{B})$ of $\mathbf{B}$ is defined to be the matrix $(\mathbf{C}_m \otimes \mathbf{C}_n)\Lambda(\mathbf{C}_m \otimes \mathbf{C}_n)$ that minimizes

$$||(\mathbf{C}_m \otimes \mathbf{C}_n)\Lambda(\mathbf{C}_m \otimes \mathbf{C}_n) - \mathbf{B}||_F$$

in Frobenius norm, see [7]. Here Λ is any diagonal matrix. Clearly, the cost of computing $c(\mathbf{B})^{-1}\mathbf{y}$ for any vector $\mathbf{y}$ is $O(mn \log mn)$ operations. For banded matrices in (7), which has $(2L-1)^2$ non-zero diagonals and is of size $M_1 M_2 \times M_1 M_2$, the cost of constructing $c(\mathbf{H}_L(\epsilon))$ is of $O(L^2 M_1 M_2)$ operations only, see [6].

3.1 Spatially Invariant Case

When there are no subpixel displacement errors, i.e., when all $\epsilon^x_{l_1,l_2} = \epsilon^y_{l_1,l_2} = 0$, the matrices $\mathbf{H}^x_{l_1 l_2}(0)$ and also $\mathbf{H}^y_{l_1 l_2}(0)$ are the same for all l_1 and l_2. We will denote them simply by $\mathbf{H}^x_L$ and $\mathbf{H}^y_L$. We claim that in this case, the blurring matrix $\mathbf{H}_L \equiv \mathbf{H}_L(0) = \mathbf{H}^x_L \otimes \mathbf{H}^y_L$ can always be diagonalized by the discrete cosine transform matrix.

We begin with $L = 2$. The blurring matrix $\mathbf{H}_2 = \mathbf{H}^x_2 \otimes \mathbf{H}^y_2$, where $\mathbf{H}^x_2$ is an $M_1 \times M_1$ tridiagonal matrix given by

$$\mathbf{H}^x_2 = \frac{1}{2}\begin{pmatrix} \frac{3}{2} & \frac{1}{2} & & & \\ \frac{1}{2} & 1 & \frac{1}{2} & & \\ & \ddots & \ddots & \ddots & \\ & & \frac{1}{2} & 1 & \frac{1}{2} \\ & & & \frac{1}{2} & \frac{3}{2} \end{pmatrix}$$

and $\mathbf{H}_2^y$ is an $M_2 \times M_2$ matrix with the same structure. It is easy to see that in this case, the matrices $\mathbf{H}_2^x$ and $\mathbf{H}_2^y$ can be diagonalized by $\mathbf{C}_{M_1}$ and $\mathbf{C}_{M_2}$ respectively. Thus $\mathbf{H}_2$ can be diagonalized by $\mathbf{C}_{M_1} \otimes \mathbf{C}_{M_2}$.

Next we observe that the blurring matrix is ill-conditioned.

Lemma 1 *Under the Neumann boundary condition, if there are no subpixel displacement errors, i.e., $h_{l_1 l_2}^{x-} = h_{l_1 l_2}^{x+} = 1/2$, then the $M_1 \times M_1$ matrix $\mathbf{H}_2^x$ can be diagonalized by the discrete cosine transform matrix and its eigenvalues are given by*

$$\lambda_j(\mathbf{H}_2^x) = \cos^2\left(\frac{(j-1)\pi}{2M_1}\right), \quad 1 \le j \le M_1. \tag{9}$$

In particular, the condition number $\kappa(\mathbf{H}_2^x)$ of the matrix $\mathbf{H}_2^x$ satisfies

$$\kappa(\mathbf{H}_2^x) \ge O(M_1^2). \tag{10}$$

Proof: The formula for the eigenvalues can be derived easily using the definition of the discrete cosine transform matrix. Since $\lambda_{\max}(\mathbf{H}_2^x) = 1$ and

$$\lambda_{\min}(\mathbf{H}_2^x) = \cos^2\left(\frac{(M_1-1)\pi}{2M_1}\right) \le \sin^2\left(\frac{\pi}{M_1}\right) \le \frac{\pi^2}{M_1^2},$$

the estimate of the condition number is then given by (10). □

It follows from Lemma 1 that the condition number of the matrix $\mathbf{H}_2 (= \mathbf{H}_2^x \otimes \mathbf{H}_2^y)$ is of $O(M_1^2 M_2^2)$. The matrix is very ill-conditioned. For $L > 2$, we have the following theorem.

Theorem 1 *Under the Neumann boundary condition, if there are no subpixel displacement errors, i.e., $h_{l_1 l_2}^{x-} = h_{l_1 l_2}^{x+} = 1/2$, then the matrix $\mathbf{H}_L^x$ can be diagonalized by the discrete cosine transform matrix and its eigenvalues are given by*

$$\lambda_i(\mathbf{H}_L^x) = \frac{4}{L}\cos^2\left(\frac{(i-1)\pi}{2M_1}\right) p_L\left(\frac{(i-1)\pi}{M_1}\right), \quad 1 \le i \le M_1, \tag{11}$$

where

$$p_L\left(\frac{(i-1)\pi}{M_1}\right) = \begin{cases} \displaystyle\sum_{j=1}^{L/4} \cos\left(\frac{(i-1)(2j-1)\pi}{M_1}\right), & L = 4k \text{ for some positive integer } k, \\ \displaystyle\frac{1}{2} + \sum_{j=1}^{(L-2)/4} \cos\left(\frac{(i-1)2j\pi}{M_1}\right), & \text{otherwise.} \end{cases} \tag{12}$$

Proof: We first establish a relationship between the matrices $\mathbf{H}_L^x$ and $\mathbf{H}_2^x$. From (6), for $L > 2$, we have

$$\mathbf{H}_L^x = \begin{cases} \frac{2}{L}\sum_{j=1}^{L/4} \mathbf{S}_{2j-1}\mathbf{H}_2^x, & L = 4k \text{ for some positive integer } k, \\ \frac{2}{L}\sum_{j=0}^{(L-2)/4} \mathbf{S}_{2j}\mathbf{H}_2^x, & \text{otherwise,} \end{cases} \tag{13}$$

where $\mathbf{S}_0$ is the $M_1 \times M_1$ identity matrix and

$$\mathbf{S}_k = \text{ Toeplitz}(\mathbf{e}_{k+1}) + \text{ Hankel}(\mathbf{e}_k), \quad 1 \le k \le M_1 - 1.$$

Here Toeplitz($\mathbf{e}_k$) is the $M_1 \times M_1$ symmetric Toeplitz matrix with the k-th unit vector $\mathbf{e}_k$ as the first column, and Hankel($\mathbf{e}_k$) is the $M_1 \times M_1$ Hankel matrix with $\mathbf{e}_k$ as the first column and $\mathbf{e}_k$ in the reverse order as the last column.

We remark that the Toeplitz part in $\mathbf{S}_k$ can be interpreted as the decomposition of the discrete blurring function $[0.5, 1, \ldots, 1, \ldots, 1, 0.5]$ into the sum of the elementary discrete blurring function $[0.5, 1, 0.5]$ with different shifts. For example, for $L = 4$, we have

$$[\tfrac{1}{2}, 1, 1, 1, \tfrac{1}{2}] = [\tfrac{1}{2}, 1, \tfrac{1}{2}, 0, 0] + [0, 0, \tfrac{1}{2}, 1, \tfrac{1}{2}],$$

where the two terms on the right together gives the Toeplitz part in $\mathbf{S}_1$. For $L = 6$, we have

$$[\tfrac{1}{2}, 1, 1, 1, 1, 1, \tfrac{1}{2}] = [\tfrac{1}{2}, 1, \tfrac{1}{2}, 0, 0, 0, 0] + [0, 0, \tfrac{1}{2}, 1, \tfrac{1}{2}, 0, 0] + [0, 0, 0, 0, \tfrac{1}{2}, 1, \tfrac{1}{2}],$$

where the first and the third terms on the right together gives the Toeplitz part in $\mathbf{S}_2$ while the middle term gives the Toeplitz part of $\mathbf{S}_0$. Because we are considering the Neumann boundary condition, entries outside the blurring matrix $\mathbf{H}_L^x$ are flipped into the matrix (cf. (6)). This is done by means of the Hankel part of $\mathbf{S}_k$. Thus the resulting shift matrices are given by $\mathbf{S}_k$ and we obtain (13).

Since $\{\mathbf{S}_k\}_{k=0}^{M_1-1}$ is exactly a basis for the space containing all matrices that can be diagonalized by $\mathbf{C}_{M_1}$, see [2], it follows that the matrix $\mathbf{H}_L^x$ can be diagonalized by the discrete cosine transform matrix. We also note that the eigenvalues of $\mathbf{S}_k$ are given by

$$\lambda_i(\mathbf{S}_k) = 2\cos\left(\frac{(i-1)k\pi}{M_1}\right), \quad 1 \le i \le M_1,$$

see for instance [2]. Using (13) and (9), the eigenvalues of $\mathbf{H}_L^x$ are given in (11). □

Theorem 1 states that the matrices $\mathbf{H}_L (= \mathbf{H}_L^x \otimes \mathbf{H}_L^y)$ are also very ill-conditioned and their condition numbers are at least of order $M_1^2 M_2^2$ (cf. (11)). We remark that some of these matrices may even be singular. For instance, when $L = 4$ and $M_1 = M_2 = 64$, $\lambda_{33}(\mathbf{H}_4^x) = 0$. Thus a regularization procedure such as (8) should be imposed to obtain a reasonable estimate for the original image in the high resolution reconstruction.

In this paper, we consider the L_2 and H_1 norm regularization functionals in (8). Correspondingly, we are required to solve the following linear systems:

$$(\mathbf{H}_L^t\mathbf{H}_L + \alpha\mathbf{I})\mathbf{f} = \mathbf{H}_L^t\mathbf{g} \quad \text{or} \quad (\mathbf{H}_L^t\mathbf{H}_L + \alpha\mathbf{L}^t\mathbf{L})\mathbf{f} = \mathbf{H}_L^t\mathbf{g}, \tag{14}$$

where $\alpha > 0$, $\mathbf{I}$ is the identity matrix and $\mathbf{L}^t\mathbf{L}$ is the discrete Laplacian matrix with the Neumann boundary condition. We note that $\mathbf{L}^t\mathbf{L}$ can be diagonalized by the discrete cosine transform matrix, see for instance [4]. Thus if we use the Neumann boundary condition for both the blurring matrix $\mathbf{H}_L$ and the regularization operator $\mathbf{L}^t\mathbf{L}$, then the coefficient matrix in (14) can be diagonalized by the discrete cosine transform matrix and hence its inversion can be done in three 2-dimensional fast cosine transforms (one for finding the eigenvalues of the coefficient matrix, two for transforming the right hand side and the solution vector, see [16] for instance). Thus the total cost of solving the system is of $O(M_1M_2 \log M_1M_2)$ operations.

We remark that for the zero boundary condition, discrete sine transform matrices can diagonalize Toeplitz matrices with at most 3 bands (e.g., $\tilde{\mathbf{H}}_2$) but not dense Toeplitz matrices in general (e.g., $\tilde{\mathbf{H}}_4$), see [8] for instance. Therefore, in general we have to solve large block-Toeplitz-Toeplitz-block systems. The fastest direct Toeplitz solvers require $O(M_1^2M_2^2)$ operations, see [10]. The systems can also be solved by the preconditioned conjugate gradient method with some suitable preconditioners, see [7]. We note however that the cost *per iteration* is at least four 2-dimensional fast Fourier transforms. Thus we see that the cost of using the Neumann boundary condition is significantly lower than that of using the zero boundary condition.

3.2 Spatially Variant Case

When there are subpixel displacement errors, the blurring matrix $\mathbf{H}_L(\epsilon)$ has the same banded structure as that of $\mathbf{H}_L$, but with some entries slightly perturbed. It is a near block-Toeplitz-Toeplitz-block matrix but it can no longer be diagonalized by the cosine transform matrix. Therefore we solve the corresponding linear system by the preconditioned conjugate gradient method. We will use the cosine transform preconditioner $c(\mathbf{H}_L(\epsilon))$ of $\mathbf{H}_L(\epsilon)$ as the preconditioner.

Below we study the convergence rate of the preconditioned conjugate method for solving the linear systems

$$[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha\mathbf{I}]^{-1}[\mathbf{H}_L(\epsilon)^t\mathbf{H}_L(\epsilon) + \alpha\mathbf{I}]\mathbf{f} = \mathbf{H}_L(\epsilon)^t\mathbf{g} \tag{15}$$

and

$$[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha\mathbf{L}^t\mathbf{L}]^{-1}[\mathbf{H}_L(\epsilon)^t\mathbf{H}_L(\epsilon) + \alpha\mathbf{L}^t\mathbf{L}]\mathbf{f} = \mathbf{H}_L(\epsilon)^t\mathbf{g}, \tag{16}$$

where α is a positive constant. We prove that the spectra of the preconditioned normal systems are clustered around 1 for sufficiently small subpixel displacement errors. Hence when the conjugate gradient method is applied to solving the preconditioned systems (15) and (16), we expect fast convergence. Our numerical results in §4 show that the cosine transform preconditioners can indeed speed up the convergence of the method. We begin the proof with the following lemma.

Lemma 2 *Let* $\epsilon^* = \max_{0\le l_1,l_2\le L-1}\{\epsilon^x_{l_1l_2}, \epsilon^y_{l_1l_2}\}$. *Then for all* M_1 *and* M_2, *we have*

$$\|\mathbf{H}_L(\epsilon) - \mathbf{H}_L\|_2 \le 4\epsilon^* \quad \text{and} \quad \|c(\mathbf{H}_L(\epsilon)) - \mathbf{H}_L\|_2 \le 4\epsilon^*. \tag{17}$$

Proof: From (7), each row or column of $\mathbf{H}_L(\epsilon)$ and $\mathbf{H}_L$ differ in at most $4L$ entries and each entry is bounded by ϵ^*/L. It follows that

$$\|\mathbf{H}_L(\epsilon) - \mathbf{H}_L\|_\infty \le 4\epsilon^* \quad \text{and} \quad \|\mathbf{H}_L(\epsilon) - \mathbf{H}_L\|_1 \le 4\epsilon^*.$$

Hence the first inequality in (17) follows by using $\|\cdot\|_2 \le \sqrt{\|\cdot\|_1\|\cdot\|_\infty}$. For the second inequality, we first note that by Theorem 1, $c(\mathbf{H}_L) = \mathbf{H}_L$. Hence we have

$$\|c(\mathbf{H}_L(\epsilon)) - \mathbf{H}_L\|_2 = \|c(\mathbf{H}_L(\epsilon) - \mathbf{H}_L)\|_2 \le \|\mathbf{H}_L(\epsilon) - \mathbf{H}_L\|_2,$$

where the last inequality follows from $\|\cdot\|_2 \le \|c(\cdot)\|_2$, see [7]. □

Lemma 3 *Let* $\epsilon^* = \max_{0\le l_1,l_2\le L-1}\{\epsilon^x_{l_1l_2}, \epsilon^y_{l_1l_2}\}$. *Then*

$$\|\mathbf{H}_L(\epsilon)^t\mathbf{H}_L(\epsilon) - c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon))\|_2 < d_L(\epsilon^*)$$

where $d_L(\cdot)$ *is a function independent of* M_1 *and* M_2 *and* $\lim_{\epsilon^*\to 0} d_L(\epsilon^*) = 0$.

Proof: We note that

$$\begin{aligned} & \|\mathbf{H}_L(\epsilon)^t\mathbf{H}_L(\epsilon) - c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon))\|_2 \\ \le\ & \|\mathbf{H}_L(\epsilon)^t[\mathbf{H}_L(\epsilon) - c(\mathbf{H}_L(\epsilon))]\|_2 + \|[\mathbf{H}_L(\epsilon)^t - c(\mathbf{H}_L(\epsilon))^t]c(\mathbf{H}_L(\epsilon))\|_2. \end{aligned}$$

By Theorem 1, $\|H_L\|_2$ is bounded above by a constant independent of M_1 and M_2. Hence by (17), $\|\mathbf{H}_L(\epsilon)^t\|_2$ and $\|c(\mathbf{H}_L(\epsilon))\|_2$ are also bounded above by some constants independent of M_1 and M_2. Moreover, by (17) again, $\|\mathbf{H}_L(\epsilon) - c(\mathbf{H}_L(\epsilon))\|_2$ and $\|\mathbf{H}_L(\epsilon)^t - c(\mathbf{H}_L(\epsilon))^t\|_2$ are less than $8\epsilon^*$. The result therefore follows. □

Using the above lemmas, we can analyze the convergence rate of the preconditioned systems (15) and (16).

Theorem 2 *Let* $\epsilon^* = \max_{0\le l_1,l_2\le L-1}\{\epsilon^x_{l_1l_2}, \epsilon^y_{l_1l_2}\}$. *If* ϵ^* *is sufficiently small, then the spectra of the preconditioned matrices*

$$[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha\mathbf{I}]^{-1}[\mathbf{H}_L(\epsilon)^t\mathbf{H}_L(\epsilon) + \alpha\mathbf{I}]$$

are clustered around 1 and their smallest eigenvalues are bounded away from 0 by a positive constant independent of M_1 *and* M_2.

Proof: We just note that

$$\|[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{I}]^{-1}\|_2 \leq \frac{1}{\alpha}$$

and

$$\begin{aligned} & [c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{I}]^{-1}[\mathbf{H}_L(\epsilon)^t \mathbf{H}_L(\epsilon) + \alpha \mathbf{I}] \\ = \; & \mathbf{I} + [c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{I}]^{-1}[\mathbf{H}_L(\epsilon)^t \mathbf{H}_L(\epsilon) - c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon))]. \end{aligned}$$

Hence the result follows by applying Lemma 3. □

Theorem 3 *Let* $\epsilon^* = \max_{0 \leq l_1, l_2 \leq L-1}\{\epsilon^x_{l_1 l_2}, \epsilon^y_{l_1 l_2}\}$. *If* ϵ^* *is sufficiently small, then the spectra of the preconditioned matrices*

$$[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{L}^t \mathbf{L}]^{-1}[\mathbf{H}_L(\epsilon)^t \mathbf{H}_L(\epsilon) + \alpha \mathbf{L}^t \mathbf{L}]$$

are clustered around 1 and their smallest eigenvalues are uniformly bounded away from 0 by a positive constant independent of M_1 *and* M_2.

Proof: Since

$$\begin{aligned} & [c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{L}^t \mathbf{L}]^{-1}[\mathbf{H}_L(\epsilon)^t \mathbf{H}_L(\epsilon) + \alpha \mathbf{L}^t \mathbf{L}] \\ = \; & \mathbf{I} + [c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{L}^t \mathbf{L}]^{-1}[\mathbf{H}_L(\epsilon)^t \mathbf{H}_L(\epsilon) - c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon))], \end{aligned}$$

it suffices to show that $\|[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{L}^t \mathbf{L}]^{-1}\|_2$ is bounded above by a constant independent of M_1 and M_2. Since $\lambda_{\min}(A) + \lambda_{\min}(B) \leq \lambda_{\min}(A + B)$ for any Hermitian matrices A and B (see [9, Corollary 8.1.3, p.411]), we have

$$\begin{aligned} & \|[c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) + \alpha \mathbf{L}^t \mathbf{L}]^{-1}\|_2 \\ \leq \; & \frac{1}{\lambda_{\min}(\mathbf{H}_L^t \mathbf{H}_L + \alpha \mathbf{L}^t \mathbf{L}) - \|c(\mathbf{H}_L(\epsilon))^t c(\mathbf{H}_L(\epsilon)) - \mathbf{H}_L^t \mathbf{H}_L\|_2}. \end{aligned} \tag{18}$$

Because the matrix $\mathbf{H}_L^t \mathbf{H}_L + \alpha \mathbf{L}^t \mathbf{L}$ can be diagonalized by the 2-dimensional discrete cosine transform matrix, we can estimate the smallest eigenvalue of this matrix. We first note that

$$\lambda_{(i-1)M_2+j}(\mathbf{L}^t \mathbf{L}) = 4 \sin^2\left(\frac{(i-1)\pi}{2M_1}\right) + 4 \sin^2\left(\frac{(j-1)\pi}{2M_2}\right), \tag{19}$$

for $1 \leq i \leq M_1$ and $1 \leq j \leq M_2$, see [4]. By using (11) and the fact that $\mathbf{H}_L = \mathbf{H}_L^x \otimes \mathbf{H}_L^y$, we obtain

$$\lambda_{(i-1)M_2+j}(\mathbf{H}_L^t \mathbf{H}_L) = \left(\frac{4}{L}\right)^4 \cos^4\left(\frac{(i-1)\pi}{2M_1}\right) \cos^4\left(\frac{(j-1)\pi}{2M_2}\right) p_L^2\left(\frac{(i-1)\pi}{M_1}\right) p_L^2\left(\frac{(j-1)\pi}{M_2}\right), \tag{20}$$

for $1 \le i \le M_1$ and $1 \le j \le M_2$, where $p_L(\cdot)$ is defined in (12).

Clearly the function $\sin^2(x/2)$ is zero at $x = 0$ and positive in $(0, \pi]$, whereas the function $\cos^4(x/2)p_L^2(x) \ge 1/4$ at $x = 0$ and is nonnegative in $[0, \pi]$. Thus we see that the function

$$4\alpha \sin^2\left(\frac{x}{2}\right) + 4\alpha \sin^2\left(\frac{y}{2}\right) + \left(\frac{4}{L}\right)^4 \cos^4\left(\frac{x}{2}\right) \cos^4\left(\frac{y}{2}\right) p_L^2(x) p_L^2(y)$$

is positive for all x and y in $[0, \pi]$. It follows from (19) and (20) that the matrix $\mathbf{H}_L^t \mathbf{H}_L + \alpha \mathbf{L}^t \mathbf{L}$ is positive definite and its smallest eigenvalue is bounded away from 0 by a positive constant independent of M_1 and M_2. In view of Lemma 3, the right hand side of (18) is therefore bounded by a positive constant independent of M_1 and M_2 for sufficiently small ϵ^*. □

Thus we conclude that the preconditioned conjugate gradient method applied to (15) and (16) with $\alpha > 0$ will converge superlinearly for sufficiently small displacement errors, see for instance [7]. Since $\mathbf{H}_L(\epsilon)$ has only $(2L-1)^2$ non-zero diagonals, the matrix-vector product $\mathbf{H}_L(\epsilon)\mathbf{x}$ can be done in $O(L^2 M_1 M_2)$. Thus the cost per each iteration is $O(M_1 M_2 \log M_1 M_2 + L^2 M_1 M_2)$ operations, see [9, p.529]. Hence the total cost for finding the high resolution image vector is of $O(M_1 M_2 \log M_1 M_2 + L^2 M_1 M_2)$ operations.

4 Numerical Examples

In this section, we illustrate the effectiveness of using cosine transform preconditioners for solving high resolution image reconstruction problems. The original image is shown in Figure 2 (left). The conjugate gradient method is employed to solving the preconditioned systems (15) and (16). The stopping criteria is $||\mathbf{r}^{(j)}||_2/||\mathbf{r}^{(0)}||_2 < 10^{-6}$, where $\mathbf{r}^{(j)}$ is the normal equations residual after j iterations. In the tests, the parameters $\epsilon^x_{l_1 l_2}$ and $\epsilon^y_{l_1 l_2}$ are random values chosen between $-1/2$ and $1/2$. A Gaussian white noise with signal-to-noise ratio of 30dB is added to the low resolution images.

Tables 1–2 show the numbers of iterations required for convergence for $L = 2$ and 4 respectively. In the tables, "cos", "cir" or "no" signify that the cosine transform preconditioner, the level-2 circulant preconditioner [7] or no preconditioner is used respectively. We see from the tables that the cosine transform preconditioner converges much faster than the circulant preconditioners for different M and α, where $M(= M_1 = M_2)$ is the size of the reconstructed image and α is the regularization parameter. Also the convergence rate is independent of M for fixed α as predicted by Theorems 2 and 3.

Next we show the 256×256 reconstructed images from four 128×128 low resolution images, i.e., a 2×2 sensor array is used. One of the low resolution images is shown in Figure 2 (middle). The observed high resolution image $\mathbf{g}$ is shown in Figure 2 (right). We tried the Neumann, zero and periodic boundary conditions to reconstruct the high resolution images. Figure 3 shows the reconstructed images. The optimal regularization parameter α is chosen such that it minimizes the relative error of the reconstructed image $\mathbf{f}_r(\alpha)$ to the original image $\mathbf{f}$, i.e., it minimizes $\|\mathbf{f} - \mathbf{f}_r(\alpha)\|_2/\|\mathbf{f}\|_2$. By comparing the figures in Figure 3, it is clear that the trees in the image

α	1×10^{-2}			1×10^{-3}			1×10^{-4}		
M	cos	cir	no	cos	cir	no	cos	cir	no
32	8	27	48	12	58	127	20	83	325
64	8	27	48	11	64	130	19	125	347
128	8	27	48	11	68	129	17	173	345
256	8	27	48	10	68	129	17	181	348

Table 1a. Number of iterations for $L = 2$ where the L_2 norm regularization is used.

α	1×10^{-2}			1×10^{-3}			1×10^{-4}		
M	cos	cir	no	cos	cir	no	cos	cir	no
32	7	16	26	9	38	68	13	70	178
64	7	16	26	9	36	69	13	88	180
128	7	16	26	9	38	69	13	99	180
256	6	16	26	8	38	69	13	99	180

Table 1b: Number of iterations for $L = 2$ where the H_1 norm regularization is used.

α	1×10^{-2}			1×10^{-3}			1×10^{-4}		
M	cos	cir	no	cos	cir	no	cos	cir	no
32	7	33	45	10	67	111	16	145	256
64	6	34	47	10	84	123	16	180	314
128	6	32	47	10	96	125	15	237	323
256	6	32	47	9	92	125	15	262	323

Table 2a: Number of iterations for $L = 4$ where the L_2 norm regularization is used.

α	1×10^{-2}			1×10^{-3}			1×10^{-4}		
M	cos	cir	no	cos	cir	no	cos	cir	no
32	5	23	33	8	46	72	12	86	159
64	5	23	33	8	63	83	12	127	182
128	5	23	34	7	65	87	11	155	204
256	5	22	34	7	63	86	11	178	216

Table 2b: Number of iterations for $L = 4$ where the H_1 norm regularization is used.