

FAST RECONSTRUCTION ALGORITHMS FOR OPTICAL TOMOGRAPHY USING SPARSE MATRIX REPRESENTATIONS

Guangzhi Cao, Charles A. Bouman, and Kevin J. Webb

School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47907 USA

ABSTRACT

This paper introduces a novel method for reconstructing optical tomography images using pre-computed transforms. Our approach is to pre-compute and store the inverse matrix required for MAP reconstruction using lossy source coding techniques. We show how lossy source coding techniques can be used to store the large and non-sparse matrix by applying a wavelet transform in the image space and appropriate orthonormal transforms in the sensor space. Lossy coding dramatically reduces the the number of non-zero coefficients, thereby proportionately reducing both the required storage and computation time. However, if the number of sensor measurements is large, the storage and computation of the orthonormal transforms can become prohibitive. For this purpose, we introduce a general method for approximating any orthonormal transform by a series of sparse binary transforms. This sparse matrix transform technique is then used together with lossy coding to result in a fast reconstruction algorithm for optical tomography. Simulations indicate that the technique can dramatically reduce the storage and computation requirements in reconstruction by exploiting redundancy in the transformed matrices.

Index Terms— Optical tomography, sparse matrices, image coding

1. INTRODUCTION

The objective of optical tomography is to reconstruct tissue cross-sections from measurements of diffusely scattered light [1]. This is an ill-posed and highly nonlinear inverse problem, but it is of great potential value since optical spectroscopy methods offer the potential for much greater specificity in clinical and research applications. Model based methods are ideally suited for optical tomography, but they result in very computationally expensive optimization problems. Numerical optimization methods such as multigrid optimization [2] and adjoint differentiation [3], can reduce computation, but

This work was supported by the National Science Foundation under Contract CCR-0431024.

the computational challenge remains a potential barrier to the use of optical tomography in real-time clinical applications.

In this paper we present a novel non-iterative reconstruction approach for optical tomography with the potential to be much faster than iterative reconstruction methods. Our approach is based on pre-computation and storage of the inverse operator using lossy image coding techniques. We show how lossy source coding techniques can be used to store the large and non-sparse matrix by applying appropriate orthonormal transforms in the sensor space and a wavelet transform in the image space. Importantly, these orthonormal transforms are chosen to minimize the mean squared error in the reconstructed image by decorrelating both the sensor inputs and the columns of the matrix. Using typical imaging geometries for optical tomography, we have shown [4] that lossy coding of this transformed matrix can dramatically reduce the the number of non-zero matrix coefficients, thereby proportionally reducing both the required storage and computation time.

However, if the number of sensor measurements is large, the storage and computation of the required orthonormal transform can become prohibitive. For this purpose, we present a general method for approximating any orthonormal transform by a series of sparse binary rotations. This sparse matrix transform technique is then used together with the lossy coding to result in a fast reconstruction algorithm for optical tomography.

2. OPTIMAL LOSSY COMPRESSION FOR THE INVERSE SYSTEM MATRIX

The reconstruction problem in optical tomography can be linearized and formed as an inversion

$$\hat{x} = Hy \quad (1)$$

where $\hat{x} \in R^N$ denotes the reconstructed 3D image of the tissue's optical parameters, $y \in R^M$ is the surface measurement, and H is the inverse system matrix which can depend on an unknown bulk parameter vector θ of the material. For notational simplicity, we will suppress the dependence on θ . Our

approach to reconstruction will be to pre-compute H , compress it using lossy compression methods, and then decompress H and compute the required matrix-vector product on-the-fly to reconstruct the image \hat{x} . In practice, this may mean that a series of H matrices will be compressed for differing values of θ . Notice that the columns of H are actually 3D images corresponding to the reconstruction of different sensor measurements, so these columns are well suited for compression by conventional lossy image coding methods. However, it is critical to perform this lossy compression with the appropriate distortion metric on H , so that the mean-squared error is minimized in the reconstructed image. In order to obtain this metric, we must account for the distribution of the measurement y .

Let R_y be the auto-correlation matrix of y defined by

$$R_y \triangleq E[yy^T] = E\Lambda_y E^T \quad (2)$$

where E and Λ_y are the matrices of eigenvectors and eigenvalues of R_y , respectively. Define the transformed inverse system and data as

$$\tilde{H} = HE\Lambda_y^{\frac{1}{2}} \quad (3)$$

$$\tilde{y} = \Lambda_y^{-\frac{1}{2}} E^T y. \quad (4)$$

Notice that $E[\tilde{y}\tilde{y}^T] = I$, and $\hat{x} = \tilde{H}\tilde{y}$. If we further define δH as the quantization error in H and $\delta\hat{x}$ as the resulting quantization error in \hat{x} , then we have that $\delta\hat{x} = \delta\tilde{H}\tilde{y}$. From this we can obtain an expression for the conditional MSE in \hat{x} given δH . Assume the measurement y is independent of the quantization error δH . Then we have the fact that

$$E[\|\delta\hat{x}\|^2 | \delta H] = \|\delta\tilde{H}\|^2 = \text{trace}\{\delta H^T \delta H R_y\} \quad (5)$$

where $\|\delta\tilde{H}\|$ is the Frobenius norm [4]. This means that if we minimize $\|\delta\tilde{H}\|^2$ we obtain a reconstructed image $\hat{x} + \delta\hat{x}$ with minimum distortion. Therefore, \tilde{H} is the appropriate representation for the storage and coding of the inverse matrix H .

Next, we decorrelate along rows of H with Karhunen-Loeve (KL) transform and decorrelate along columns with wavelet transform [5]. These transforms can be expressed as

$$\tilde{H} = W\tilde{H}\Phi \quad (6)$$

where Φ consists of the eigenvectors of $\tilde{H}^T\tilde{H}$, and W represents 3D wavelet transform. Since both the KL and wavelet transform are orthonormal, we know $\|\delta\tilde{H}\|^2 = \|\delta\tilde{H}\|^2$. Accordingly, define the transform data as

$$\tilde{y} = \Phi^T \tilde{y} = \Phi^T \Lambda_y^{-\frac{1}{2}} E^T y. \quad (7)$$

Then \hat{x} can be expressed as:

$$\hat{x} = W^{-1}\tilde{H}\tilde{y} \quad (8)$$




Fig. 1. Illustration of inverse system matrix compression. \tilde{H} is a sparse representation of H through KL and wavelet transforms. The shaded region represents the effective non-zero entries in \tilde{H} .

After obtaining the sparse representation \tilde{H} , the SPIHT algorithm [6] is used to quantize and encode each column of the matrix with the same quantization table.

Figure 1 illustrates compression of H , where \tilde{H} is the sparse representation of H obtained through KL and wavelet transforms. Due to sparsity of \tilde{H} , evaluation of $\tilde{H}\tilde{y}$ requires many fewer multiplications than evaluation of $H y$. Therefore, lossy compression both reduces the data that must be stored, and reduces computation required to reconstruct the image.

So the final algorithm for reconstruction is given by

$$\hat{x} = W^{-1}[\tilde{H}]T y, \quad (9)$$

where $[\tilde{H}]$ denotes the quantized version of \tilde{H} and T is defined by

$$T = \Phi^T \Lambda_y^{-\frac{1}{2}} E^T. \quad (10)$$

If the number of sensor measurements, M , is small, then the computation and storage are dominated by the number of nonzero elements in $[\tilde{H}]$.

However, if the number of measurements, M , is large, then the computation and storage may be dominated by multiplication and storage of the matrix T . The next section addresses how to efficiently compute and store the transform T .

3. SPARSE MATRIX TRANSFORM

In this section, we present a sparse matrix transform (SMT) method which allows for efficient storage and multiplication with T . We do this by approximating the orthonormal transforms in T as a product of sparse binary orthonormal rotations. Since the distortion-rate performance is primarily dominated by the whitening process as shown in Figure 3, we illustrate this method by computing an approximation to the orthonormal matrix E from equation (2).

We approximate E as

$$E \approx \prod_{k=0}^{K-1} A_k \quad (11)$$

where i_k and j_k are two distinct indices chosen so that $i_k < j_k$, and A_k is a sparse matrix of the form

$$[A(i_k, j_k, \theta_k)]_{i,j} = \begin{cases} 1 & i = j \neq \{i_k \text{ or } j_k\} \\ \cos(\theta_k) & \text{if } i = j = \{i_k \text{ or } j_k\} \\ -\sin(\theta_k) & \text{if } i = i_k \text{ and } j = j_k \\ \sin(\theta_k) & \text{if } i = j_k \text{ and } j = i_k \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

Each matrix A_k can be viewed as a rotation in the coordinate pair (i_k, j_k) . Our strategy is to use K (order $M \log_2 M$) of these sparse rotations to approximate the orthonormal transform E . Since multiplication by A_k only requires 3 multiplies, an approximate sparse factorization of E can dramatically reduce complexity for large matrices E . In practice, we represent the matrix A_k in the computer by two shorts for i_k and j_k , and two doubles for $\cos(\theta_k)$ and $\sin(\theta_k)$.

We compute the matrices A_k in a recursive procedure starting with $k = 0$. To do this, apply the following recursion

$$\begin{aligned} C &\leftarrow R_y \\ \text{For } k = 0 : K - 1 \{ \\ &F_{i,j} \leftarrow \frac{C_{i,j}^2}{C_{i,i}C_{j,j}} \\ &(i_k, j_k) \leftarrow \arg \max_{i < j} \{F_{i,j}\} \\ &\theta_k \leftarrow \arctan(2 * C_{i_k, j_k} / (C_{i_k, i_k} - C_{j_k, j_k})) / 2 \\ &A_k \leftarrow A(i_k, j_k, \theta_k) \\ &C \leftarrow A_k^T C A_k \\ \} \end{aligned}$$

Intuitively, each new sparse rotation A_k , decorrelates the indices i_k and j_k , and it can be shown that the indices i_k and j_k are selected so that $\prod_{i=0}^{M-1} C_{ii}$ is minimized after the update. Accordingly, Λ_y in (3) and (4) is replaced by the diagonal entries of the resulting C . The same factorization procedure can also be used to approximate Φ in (6) using its sparse matrix transform.

4. NUMERICAL RESULTS

Here we give some numerical results of fluorescence optical diffusion tomography (FODT) [7] from reflectance measurements based on the proposed algorithms. The measurement geometry is shown in Fig. 2(a) where a 6 cm \times 6 cm probe scans on top of the semi-infinite medium. The probe contains 4 sources and 117 detectors as shown in Fig. 2(b). The background optical values are set to $\mu_{a_x, m} = 0.02 \text{ cm}^{-1}$, $D_{x, m} = 0.03 \text{ cm}$ and $\eta\mu_{a_f} = 0 \text{ cm}^{-1}$, which approximates the optical property of tissue. The measurements were generated with a spherical heterogeneity of radius 0.5 cm present 2 cm below the center of the probe. The optical values of the heterogeneity are $\mu_{a_x} = 0.12 \text{ cm}^{-1}$, $\mu_{a_m} = 0.02 \text{ cm}^{-1}$, $D_x = D_m = 0.03 \text{ cm}^{-1}$ and $\eta\mu_{a_f} = 0.05 \text{ cm}^{-1}$. The size

of the computation domain is $8 \times 8 \times 4 \text{ cm}^3$ with a resolution 0.25 cm. Additive noise was introduced based on the shot noise model of [8] giving an average SNR of 35.6 dB. The inverse matrix was computed column by column using an iterative optimization method [4]. We modeled the covariance of y by $R_y = AA^T$ where A is the linearized forward model for the FODT system. The wavelet transform was constructed with the popular biorthogonal 9/7 tap filters and using a symmetric boundary extension. We used QccPack library [9] to implement the 3D wavelet SPIHT coding.

Figure 3 plots the normalized root mean squared error (NRMSE) of the reconstructed image versus the number of bits per matrix entry in H . Each plot corresponds to a different measurement transform. The best distortion-rate performance is obtained by using the exact KL transform for both the whitening of y and the decorrelation of \tilde{H} . However, simply whitening y yields nearly the same distortion-rate performance, while whitening with the SMT transform results in a small but significant increase in bit rate. The SMT uses $8M \log_2 M$ rotations. Notice that using no transform, results in a large increase in bit rate.

Table 1 shows the effect of compression with the KLT and SMT methods. The KLT method uses both whitening and decorrelation with a single stored transform, and the SMT method uses just whitening with the sparse transform. The bit rate for both methods was adjusted to achieve a distortion of approximately 10%, and this resulted in a SPIHT compression ratio of 192:1 for the KLT and 122:1 for the SMT. The total storage includes both the storage of H and the storage of the required transform. Notice that the SMT reduces storage by a factor of approximate 1.7 for this problem.

Fig. 4 shows the reconstruction results at the depth of 2 cm with the different compression methods. The normalized root mean squared error (NRMSE) in the two compression cases are both approximate 10%.

5. CONCLUSION

We introduce a method for compressing the transforms required to reconstruct images for FODT data. Our method incorporates a sparse-matrix transform formed by a series of rotation matrices to both reduce storage and computation. We show that the method has promise on a simulation of a practical imaging geometry.

6. REFERENCES

- [1] S.R. Arridge, "Optical tomography in medical imaging," *Inverse Problems*, vol. 15, pp. R41–R93, 1999.
- [2] S. Oh, A.B. Milstein, C.A. Bouman, and K.J. Webb, "A general framework for nonlinear multigrid inversion," *IEEE Trans. on Image Processing*, vol. 14, no. 1, pp. 125–140, Jan. 2005.
- [3] S. Ahn, A. Chaudhari, F. Darvas, C.A. Bouman, and R. Leahy, "Fast image reconstruction methods for fully 3d multispectral




Fig. 2. (a) Measurement system geometry; (b) Probe configuration: \times - source; o - detector

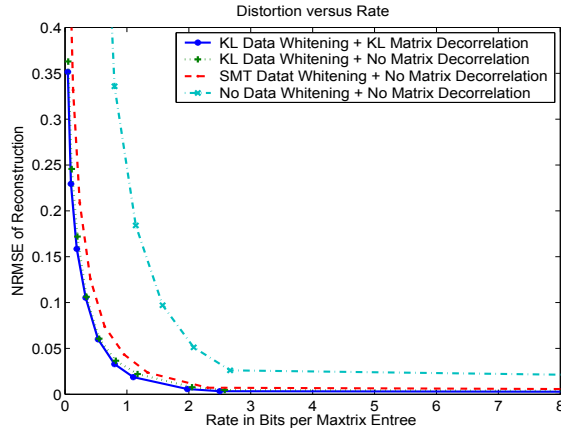


Fig. 3. Distortion versus rate rate for compression using various measurement transforms. Notice that KL whitening works nearly as well as the optimal transform, and that SMT whitening close to KL whitening.

	Computation		Storage	
	Order	Seconds	Order	Mbytes
Conjugate Grad.	MNI	70	MN	66.1
Uncompress	NM	0.25	NM	66.1
Compress with KLT	$\frac{NM}{c_1} + M^2$	0.015	$\frac{NM}{c_1} + M^2$	2.02
Compress with SMT	$\frac{NM}{c_2} + K$	0.019	$\frac{NM}{c_2} + K$	1.18

Table 1. Comparison of conjugate gradient inversion, direct multiplication by the uncompressed inverse, and multiplication by the compressed inverse using KLT and SMT. Results use number of voxels $N = 33 \times 33 \times 17$, number of measurements $M = 468$, number of sparse rotations $K = 8M \log_2 M$, and SPIHT compression ratios of $c_1 = 192 : 1$ and $c_2 = 122 : 1$ for KLT and SMT whitening respectively.




Fig. 4. Reconstruction results using different compression methods (a) original image, (b) uncompressed reconstruction, and reconstruction results using (c) KLT transform for whitening (10.2% distortion with SPIHT compression ratio of $c_1 = 192 : 1$), and (d) SMT transform for whitening (10.5% distortion with SPIHT compression ratio of $c_1 = 122 : 1$).

optical bioluminescence tomography,” in *IEEE International Symposium on Biomedical Imaging*, April 12-15 2007.

[4] 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*, Oct 29-Nov 1 2006.

[5] D. Tretter and C.A. Bouman, “Optimal transforms for multi-spectral and multilayer image coding,” *IEEE Trans. on Image Processing*, vol. 4, no. 3, pp. 296–308, March 1995.

[6] A. Said and W.A. Pearlman, “A new, fast, and efficient image codec based on set partitioning in hierarchical trees,” *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243–250, June 1996.

[7] A.B. Milstein, S. Oh, K.J. Webb, C.A. Bouman, Q. Zhang, D.A. Boas, and R.P. Millane, “Fluorescence optical diffusion tomography,” *Applied Optics*, vol. 42, no. 16, pp. 3081–3094, June 2003.

[8] J.C. Ye, K.J. Webb, C.A. Bouman, and R.P. Millane, “Optical diffusion tomography using iterative coordinate descent optimization in a Bayesian framework,” *J. Optical Society America A*, vol. 16, no. 10, pp. 2400–2412, October 1999.

[9] J.E. Fowler, “Qccpack: An open-source software library for quantization, compression, and coding,” in *DCC '00: Proceedings of the Conference on Data Compression*, Washington, DC, USA, 2000, p. 554, IEEE Computer Society.