Discrete Multiscale Bayesian Image Reconstruction *†

Thomas Frese, Charles A. Bouman
Purdue University
Department of Electrical and Computer Engineering
1285 ECE Building
West Lafayette, IN 47907-1285, USA
{frese, bouman}@ecn.purdue.edu

Ken Sauer
University of Notre Dame
Department of Electrical Engineering
275 Fitzpatrick Hall
Notre Dame, IN 46556-5637, USA
sauer@nd.edu

Abstract

Statistical and discrete-valued methods can substantially improve reconstruction quality by incorporating prior information about both the imaging system and the object being imaged. A statistical method shown to perform well in the tomographic setting is Bayesian MAP estimation. However, computing the MAP estimate in the tomographic domain is a computationally involved optimization problem. Furthermore, discrete-valued MAP reconstruction requires accurate knowledge of the density or emission levels in the cross-section. In this paper we present an efficient multiscale algorithm for discrete-valued MAP reconstruction including estimation of the discrete levels. Experimental results indicate that the multiscale algorithm has improved convergence behavior over fixed scale reconstruction and is more robust with respect to local minima.

1. Introduction

Discrete reconstruction methods are based on the assumption that the object being imaged is composed of a discrete set of materials each with uniform properties. Thus, the problem of reconstruction reduces to one of determining the specific levels present in a reconstruction and then classifying each pixel to one of these discrete levels. Discrete reconstruction methods impose a very strong constraint on the reconstruction process, and therefore can substantially improve reconstruction quality.

Statistical methods have been shown to considerably improve reconstruction performance over conventional backprojection techniques under high noise conditions. Furthermore, statistical approaches easily incorporate special geometries such as limited or missing angle projection measurements.

A statistical method that has been shown to improve performance in many tomography problems is Bayesian maximum *a posteriori* (MAP) estimation [4, 5, 3]. Computation of the MAP estimate, however, poses a computationally involved optimization problem. In this work, we propose a technique for efficient computation of the discrete-valued MAP estimate directly in the projection domain. Using a Markov random field (MRF) prior, we adopt a pixel-wise update method known as iterative coordinate descent (ICD) [3]. ICD maximizes the MAP criterion by iteratively updating each pixel of the image. The discrete version of ICD used here essentially implements the iterated conditional modes (ICM) technique introduced by Besag [1]. However, while ICM was designed for image restoration tasks, the ICD algorithm is specifically designed for the tomographic reconstruction problem resulting in improved computational efficiency.

In addition to solving the MAP optimization problem, accurate estimation of the discrete emission rates or densities is essential for the reconstruction performance. In this paper, we present an efficient method to estimate the discrete levels concurrently with the reconstruction. In particular, we formulate the estimation of the discrete levels as a continuous-valued tomographic reconstruction problem and apply existing algorithms for continuous-valued reconstruction.

Finally, we extend our reconstruction method to a multiresolution algorithm. Multiresolution techniques achieve performance improvements in a variety of imaging problems [6] including continuous-valued tomographic reconstruction. The multiresolution algorithm presented here reconstructs the image in a coarse-to-fine fashion by initializing each resolution level with the interpolated reconstruction of the next coarser level. Our experimental results demonstrate that this multiscale algorithm is less prone to being trapped in local minima and in many cases, computa-

*This work was supported by the National Science Foundation under Grant MIP97-07763.

†appeared in *33rd Asilomar Conference on Signals, Systems, and Computers*, pp. 1687-1691, Pacific Grove, CA, November 1-4, 1998.

tionally more efficient than fixed resolution reconstruction.

2. Stochastic data models for tomography

In the following, we present the statistical framework for MAP reconstruction for both emission and transmission tomography. For the transmission case, define X as the N-dimensional vector of attenuation densities of the pixels in raster order. Let Y denote the vector of photon counts for all M projections at different angles and parallel offsets. Furthermore, let P_{ij} correspond to the length of intersection between the j^{th} pixel and the i^{th} projection. Then P is the matrix of elements P_{ij} and P_{i*} denotes the vector formed by its i^{th} row. Given these assumptions, the photon count Y_i, corresponding to projection i, is Poisson distributed with mean $y_T \exp(-P_{i*}x)$. The log-likelihood of the Y_i may then be written as

$$L(y|x) = \sum_{i=1}^{M} \Big(- y_T e^{-P_{i*}x} + y_i(\log y_T - P_{i*}x) - \log(y_i!) \Big). \quad (1)$$

For the emission case, we will use the same notation, but interpret x as the vector of pixel emission rates and Y as the observed photon counts. We define P_{ij} as the probability that an emission from pixel j is registered by the i^{th} detector pair. The photon counts Y are then Poisson distributed with parameter $P_{i*}x$ which yields the log-likelihood

$$L(y|x) = \sum_{i=1}^{M} \big(- P_{i*}x \ + y_i \log(P_{i*}x) - \log(y_i!) \big). \quad (2)$$

Both (1) and (2) have the common form $L(y|x) = -\sum_{i=1}^{M} f_i(P_{i*}x)$ where the $f_i(\cdot)$ are convex and differentiable. In the following, we will write all equations for the emission case; however, all methods apply analogously to the transmission case.

3 Discrete prior model

For discrete-valued tomographic reconstruction, we assume that each pixel has one of a fixed set $\mathcal{E}$ of known emission rates. We then apply a discrete Markov random field (MRF) prior model [1, 2]. The model encourages neighboring locations to have the same states or, in our case, emission rates. Define $t_1(x)$ as the number of horizontally and vertically neighboring pixel pairs with different emission rates in x, and $t_2(x)$ as the number of diagonally neighboring pixel pairs with different emission rates in x. The discrete density function for $x \in \mathcal{E}^N$ is then assumed to be of the form

$$\log p(x) = -(\beta_1 t_1(x) + \beta_2 t_2(x)) + \log(Z) \quad (3)$$

where Z is an unknown constant called the partition function. The regularization parameters β_1 and β_2 weight the influence of the prior in comparison to the likelihood term. Larger values of β_1 and β_2 assign higher cost to local pixel differences which results in smoother reconstructions. In the following, we will often write β for β_1 and assume $\beta_2 = \beta_1/\sqrt{2}$. Combining the prior (3) with the log-likelihood (2), we obtain the MAP optimization criterion

$$\hat{x} = \arg\max_{x \in \mathcal{E}^N} \{L(y|x) + \log p(x)\} \quad (4)$$

$$= \arg\max_{x \in \mathcal{E}^N} \Big\{ \sum_{i=1}^{M} (-P_{i*}x \ + y_i \log(P_{i*}x)) - (\beta_1 t_1(x) + \beta_2 t_2(x)) \Big\}. \quad (5)$$

4 Optimization

In order to compute the MAP reconstruction, we must perform the optimization of (5). A method that is well suited for the MAP optimization is a discrete version of iterative coordinate descent (ICD) [3]. The ICD method sequentially updates each pixel of the image. With each update, the current pixel is chosen to maximize the posterior probability (5). In comparison to ICM [1], the ICD algorithm takes advantage of the sparse structure of the forward projection matrix P to dramatically speed-up the optimization. Furthermore, ICD initializes the optimization with the convolution backprojection (CBP) instead of the ML initialization used by ICM.

Let $v_1(z, x_{\partial j})$ be the number of horizontal and vertical neighbors of x_j which do not have emission rate z, and $v_2(z, x_{\partial j})$ be the number of diagonal neighbors of x_j which do not have emission rate z. Then, the maximization of the MAP equation with respect to pixel x_j can be written as

$$x_j^{n+1} = \arg\min_z \{-L(y|X_j = z, X_k = x_k^n, k \neq j) + (\beta_1 v_1(z, x_{\partial j}) + \beta_2 v_2(z, x_{\partial j}))\} \quad (6)$$

where x^n is the image containing all previous pixel updates.

Computation of the log-likelihood $L(y|z, x^n)$ using (2) for each pixel update would lead to prohibitive computational complexity. This can be avoided by using only the log-likelihood difference

$$\Delta L(z) = \sum_{i \in \mathcal{I}_j} \big(- P_{ij}z + y_i \log(P_{i*}x^n + P_{ij}(z - x_j^n)) - y_i \log(P_{i*}x^n) \big) \quad (7)$$

where I_j is the set of projections i which intersect pixel x_j, i.e. $\mathcal{I}_j = \{i : P_{ij} \neq 0 \ , \ 1 \leq i \leq M\}$. Leaving out the terms which are constant with respect to z, the update

equation for x_j can then be written as

$$x_j^{n+1} = \arg\min_z \left\{ \sum_{i \in \mathcal{I}_j} (P_{ij} z - y_i \log(P_{i*} x^n + P_{ij}(z - x_j^n))) + (\beta_1 v_1(z, x_{\partial j}) + \beta_2 v_2(z, x_{\partial j})) \right\}. \quad (8)$$

Assuming a reasonably small set $\mathcal{E}$ of K fixed emission rates, the minimization can be carried out by trying all $z \in \mathcal{E}$ and selecting the one which minimizes (8). We store the M-dimensional state vector $S_i = P_{i*}x$ between iterations. After a pixel x_j is updated, S_i can be efficiently updated using

$$S_i^{n+1} = S_i^n + P_{ij}(x_j^{n+1} - x_j^n) \qquad \forall i \in \mathcal{I}_j. \quad (9)$$

In order to assess the computational complexity of the reconstruction, we first define M_0 as the average number of projections passing through a single pixel $M_0 = (1/N)\sum_{j=1}^{N} |\mathcal{I}_j|$. The complexity of a full-update of the reconstruction is then NKM_0. This is quite reasonable, considering that due to the sparsity of P, M_0 is typically small compared to M, i.e. $M_0 \ll M$.

5 Estimation of emission rates/densities

In practice, the set of densities or emission rates $\mathcal{E}$ corresponding to different regions in the cross-section may not be known. Since in the tomography problem the pixel likelihoods are not independent, even slight errors in the estimated emission rates may result in a large number of misclassified pixels. Thus, the accuracy of the emission rates is critical for the reconstruction. In the following, we show how ML estimation of the emission rates can be implemented by iteratively updating entire regions of pixels with equal emission rates [7].

Let $\theta_1 \ldots \theta_K$ denote the discrete emission rates so that $\mathcal{E} = \{\theta_1, \cdots, \theta_K\}$. Changing a single emission rate θ_k is equivalent to changing all pixels in the reconstruction that are classified to have emission rate θ_k. If we define a region as the collection of all pixels with the same emission rate, we obtain K different regions in the reconstruction. Analogously to the projection matrix P for individual pixels, we can now define a projection matrix Q for the regions such that Q_{ik} is the probability that an emission from the k^{th} region is registered by the i^{th} detector. In practice, Q can be efficiently computed by adding the contributions of all pixels in each region. This strategy allows for efficient updates of Q after pixel changes in the reconstruction updates. We can rewrite the i'th forward projection $P_{i*}x$ as follows

$$P_{i*}x = \sum_{j=1}^{N} P_{ij}x_j = \sum_{k=1}^{K} \left(\theta_k \sum_{\{j:x_j=\theta_k\}} P_{ij} \right) = \sum_{k=1}^{K} \theta_k Q_{ik} = Q_{i*}\theta \quad (10)$$

where

$$Q_{ik} = \sum_{\{j:x_j=\theta_k\}} P_{ij}. \quad (11)$$

Substituting (10) into (2) gives the likelihood for the emission case as

$$\log \mathcal{P}(Y = y|\theta) = \sum_{i=1}^{M} \left(-Q_{i*}\theta + y_i \log(Q_{i*}\theta) - \log(y_i!) \right). \quad (12)$$

This log-likelihood function is of the same form as (2), except that the discrete-valued N-component vector x is replaced by the continuous-valued K-component vector θ and the new projection matrix Q is of size $M \times K$. Thus, ML estimation of θ is equivalent to a continuous-valued tomographic ML reconstruction with K pixels.

Since θ is continuous-valued, the optimization of (12) is different from the discrete case. Again, we use an ICD optimization since it is easily implemented with constraints such as positivity of the emission rates. The ICD update equation of the θ_k is analogously to (8) given by

$$\theta_k^{n+1} = \arg\min_{v \geq 0} \left\{ \sum_{i \in \tilde{I}_k} (Q_{ik} v - y_i \log(Q_{i*}\theta^n + Q_{ik}(v - \theta_k^n))) \right\} \quad (13)$$

where $v \geq 0$ enforces the non-negativity of the emission rates and $\tilde{I}_k$ is defined as $\tilde{I}_k = \{i : Q_{ik} \neq 0\ ,\ 1 \leq i \leq M\}$. Since the cost function in (13) is well approximated by a quadratic, the optimization can be efficiently implemented using Newton minimization. The details of this algorithm are given in [3]. For efficient computation, we store and update the same state vector $S = Px = Q\theta$ as in the discrete MAP reconstruction (9). If θ_k^n is updated to θ_k^{n+1}, S can be updated as

$$S_i^{n+1} = S_i^n + Q_{ik}(\theta_k^{n+1} - \theta_k^n) \qquad \forall i \in \tilde{I}_k. \quad (14)$$

In order to assess computational complexity, we define H as the average number of Newton-iterations per class update (13) and use M as an upper bound for $|\tilde{I}_k|$. The complexity for a full update of θ is then on the order of KMH multiplies which is typically small in comparison to the complexity NKM_0 for a reconstruction update.

The iterations for the estimation of θ can be performed between full reconstruction updates. Each time a pixel changes during the reconstruction, the new Q matrix can be obtained as follows: If $x_j^n = \theta_k$ and $x_j^{n+1} = \theta_l$, then

$$\begin{aligned} Q'_{ik} &= Q_{ik} - P_{ij} \\ Q'_{il} &= Q_{il} + P_{ij} \end{aligned} \quad (15)$$

for all projections $i \in I_j$. This results in a computationally efficient algorithm since it avoids recomputing Q using (11) after each reconstruction update.

In order to apply the estimation of the emission rates as described above, it is necessary to obtain initial values for the estimates of θ_k. In practice, initial values for the θ_k can be extracted from the CBP reconstruction either manually or using clustering techniques.

6 Extension to multiscale model

We now extend the previous model to a multiresolution reconstruction algorithm. The multiscale algorithm performs image reconstruction at different resolutions, starting at coarse resolution and progressing to the desired finest resolution. The reconstruction at each resolution is initialized with the interpolated reconstruction at the next coarser level.

In comparison to the fixed scale method, the multiscale algorithm has improved convergence behavior. The coarse scale reconstructions are computationally inexpensive but already contain the large scale behavior of the solution. Thus, considerably fewer computationally expensive iterations are necessary at finer scales. Furthermore, the multiresolution algorithm is less prone to being trapped in local minima. This is especially true for the case where the discrete emission rates θ are estimated concurrently with the reconstruction.

To define the multiscale data model, let $x^{(n)}$ denote the reconstruction at resolution n, where $n = 0$ is the finest and $n = L - 1$ is the coarsest resolution. In order to calculate the log-likelihood function for level n, we simply compute a new projection matrix $P^{(n)}$ which incorporates the pixel size at level n. $P^{(n)}$ is of dimension $M \times 4^{-n}N$. The log-likelihood for the emission case is then given by

$$L^{(n)}(y|x^{(n)}) = \sum_{i=1}^{M} \Bigg(- P_{i*}^{(n)} x^{(n)} + y_i \log(P_{i*}^{(n)} x^{(n)}) - \log(y_i!) \Bigg). \quad (16)$$

This yields the MAP equation

$$\hat{x}^{(n)} = \arg\min_{x^{(n)}} \Big\{ - L^{(n)}(y|x^{(n)}) + \beta_1^{(n)} t_1(x^{(n)}) + \beta_2^{(n)} t_2(x^{(n)}) \Big\}. \quad (17)$$

Assuming that the spatial correlation is independent of the resolution, we choose the coarse resolution parameters $\beta_1^{(n)}$ and $\beta_2^{(n)}$ as outlined in [2] to be

$$\begin{aligned} \beta_1^{(n)} &= \beta_1 \\ \beta_2^{(n)} &= \beta_2. \end{aligned} \quad (18)$$

	Em.-Rate θ_1	Em.-Rate θ_2	Em.-Rate θ_3
Original Phantom	0.001	0.05	0.1
CBP Clustering	0.0005	0.0108	0.04
Fixed-Res. MAP	0.0007	0.0105	0.0632
Multi-Res. MAP	0.0010	0.0512	0.1028

Table 1. Original and estimated emission rates for synthetic cross-section. While the fixed scale algorithm gets trapped near the clustering initialization, the multiscale method estimates the emission rates quite accurately. All units are in mm^{-1}.

The parameters β_1 and β_2 can be chosen manually to achieve the amount of regularization desired. The L level multiresolution MAP reconstruction algorithm may then be summarized as follows:

1. Compute CBP, estimate initial emission rates θ.
2. Classify CBP pixels into discrete emission rates θ_k, decimate $(L-1)$-times to initialize $x^{(L-1)}$. Set $n = L - 1$.
3. Compute reconstruction $x^{(n)}$ using fixed scale MAP reconstruction. Include ML estimation of emission rates θ if necessary.
4. Initialize $x^{(n-1)}$ with pixel-replicated $x^{(n)}$.
5. if $n = 0$ stop. Otherwise set $n = n - 1$, goto 3.

7 Experimental Results

Reconstructions using the fixed and multiscale algorithms for a synthetic phantom are shown in Fig. 1. Fig. 1(a) shows the original cross-section of size 192 by 192 pixels containing pixels with three different emission rates. The projection data was calculated at 16 evenly spaced angles each with 192 parallel projections. The data samples where formed by Poisson random variables with the appropriate means.

Fig. 1(b) shows the CBP reconstruction using a generalized Hamming filter weighted by a Gaussian envelope. An unsupervised clustering routine applied to the CBP reconstruction identified three clusters with mean emission rates $\theta = [0.005, 0.0108, 0.04]$. The clustering result consists of two classes with very low emission rates corresponding to background pixels and only one class with higher emission rate corresponding to the discs in the foreground. The CBP thresholded at the midpoints between the classes determined by the clustering routine is shown in Fig. 1(c). In addition to the errors in the class estimates, the result contains noise and aliasing effects. Fig. 1(d) shows the fixed resolution MAP reconstruction using $\beta = 1.0$ where we assume that $\beta_1 = \beta$ and $\beta_2 = \beta/\sqrt{2}$. The reconstruction was initialized to the thresholded CBP reconstruction and the class estimates were initialized to the clustering result. While the reconstruction is less noisy than the thresholded CBP, the estimation of emission rates is trapped in a local minimum close to the initial values.

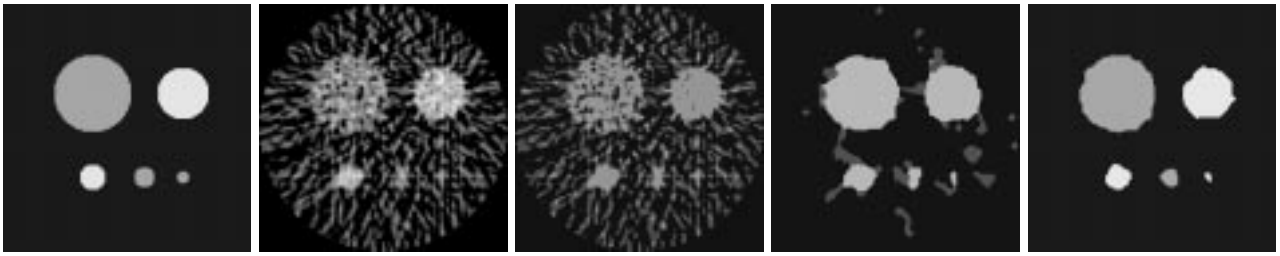

(a) Original phantom (b) CBP (c) Thresholded CBP (d) Fix-Resolution MAP $\beta = 1.0$ (e) Multiresolution MAP $\beta = 1.0$

Figure 1. Results for synthetic cross-section. Shown in (a) is the original cross-section. The continuous-valued CBP (b) contains considerable noise which is still present in the thresholded version (c), using the thresholds determined by unsupervised clustering. The fixed resolution algorithm (d), gets trapped in a local minimum resulting in class estimates close to the initialization. The multiresolution algorithm (e) estimates the classes correctly and achieves higher reconstruction performance.

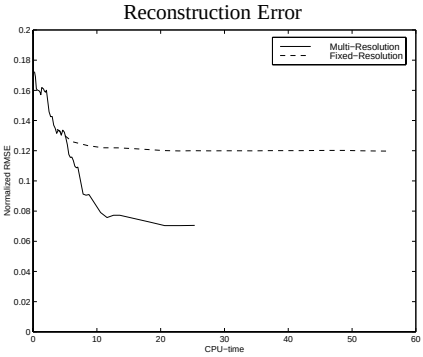

Figure 2. Normalized reconstruction error as a function of CPU-time for fixed- and multiresolution MAP reconstructions. The multiscale algorithm converges considerably faster than the fixed resolution method.

The reconstruction result using the multiscale algorithm with $L = 5$ resolution levels and $\beta = 1.0$ is shown in Fig. 1(e). The algorithm was initialized as in the fixed scale case. The estimated emission rates are very close to the true values as shown in Table 1. This results in correct classification of the 4 larger discs in the cross-section. Only the smallest disc is misclassified which is not surprising given the level of noise and the small size of the disc.

In addition to the superior reconstruction quality, the multiresolution method is faster than the fixed scale algorithm. Shown in Fig. 2 is the root mean square reconstruction error as a function of CPU-time. The multiscale algorithm converges considerably faster and achieves lower final reconstruction error than the fixed resolution method.

8 Conclusion

We have described an efficient multiresolution algorithm for discrete-valued tomographic MAP reconstruction. The algorithm includes an efficient method for estimating the discrete emission rates. The quality of the multiresolution reconstructions is significantly better than thresholded CBP reconstructions. In comparison to a fixed scale MAP reconstruction, the multiresolution method is less prone to local minima and converges faster.

References

[1] J. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society B*, 48(3):259–302, 1986.

[2] C. A. Bouman and B. Liu. Segmentation of textured images using a multiple resolution approach. In *Proc. of IEEE Int'l Conf. on Acoust., Speech and Sig. Proc.*, pages 1124–1127, New York, NY, April 11-14 1988.

[3] C. A. Bouman and K. Sauer. A unified approach to statistical tomography using coordinate descent optimization. *IEEE Trans. on Image Processing*, 5(3):480–492, March 1996.

[4] S. Geman and D. McClure. Bayesian image analysis: An application to single photon emission tomography. In *Proc. Statist. Comput. sect. Amer. Stat. Assoc.*, pages 12–18, Washington, DC, 1985.

[5] P. J. Green. Bayesian reconstruction from emission tomography data using a modified EM algorithm. *IEEE Trans. on Medical Imaging*, 9(1):84–93, March 1990.

[6] M. R. Luettgen, W. C. Karl, and A. S. Willsky. Efficient multiscale regularization with applications to the computation of optical flow. *IEEE Trans. on Image Processing*, 3(1), January 1994.

[7] K. Sauer and C. Bouman. Bayesian estimation of transmission tomograms using segmentation based optimization. *IEEE Trans. on Nuclear Science*, 39:1144–1152, 1992.