Preface

The Preface for the book will be written by the editors after they have received all the individual contributions.

Contents

Chapter 1

Multiscale Bayesian Methods for Discrete Tomography

Thomas Frese[1]
Charles A. Bouman[2]
Ken Sauer[3]

ABSTRACT Statistical methods of discrete tomographic reconstruction pose new problems both in stochastic modeling to define an optimal reconstruction, and in optimization to find that reconstruction. Multiscale models have succeeded in improving representation of structure of varying scale in imagery, a chronic problem for common Markov random fields. This chapter shows that associated multiscale methods of optimization also avoid local minima of the log a posteriori *probability better than single-resolution techniques. These methods are applied here to both segmentation/reconstruction of the unknown cross-sections, and estimation of unknown parameters represented by the discrete levels.*

1.1 Introduction

The reconstruction of images from projections is important in a variety of problems including tasks in medical imaging and non-destructive testing. Perhaps, the reconstruction technique most frequently used in commercial applications is convolution backprojection (CBP) [1]. While CBP works well for reconstruction problems with a complete set of projections having high signal-to-noise ratio (SNR), special cases benefit from alternative algorithms which can better model the imaging geometry and measurement process. These cases arise, for example, in low dosage medical imaging [2], non-destructive testing of materials with widely varying densities [3] and applications with limited angle projections [4] or hollow projections [5]. In such cases, statistical and discrete-valued methods can substantially im-

[1]Purdue University, Department of Electrical and Computer Engineering, West Lafayette, IN 47907-1285, USA. E-mail: frese@ecn.purdue.edu
[2]Purdue University, Department of Electrical and Computer Engineering, West Lafayette, IN 47907-1285, USA. E-mail: bouman@ecn.purdue.edu
[3]University of Notre Dame, Department of Electrical Engineering, 275 Fitzpatrick Hall, Notre Dame, IN 46556-5637, USA. E-mail: sauer@nd.edu

prove the reconstruction quality by incorporating important prior information about both the imaging system and the object being imaged. 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. Therefore, an ideal reconstruction should only contain pixel or voxel values from a corresponding set of discrete levels. In this case, 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.

Early methods for discrete-valued reconstruction focused on reconstructions of binary arrays from only the horizontal and vertical projections [6]. The deterministic projections were treated as a system of linear equations. Attention was particularly paid to the ambiguity of reconstructions which was formulated in the context of switching components [6, 7]. Algorithms for unambiguous reconstruction were developed by assuming object constraints such as connectedness in 2-D [8] or convexity in 3-D [9]. In addition, these concepts were extended to four or more projection angles including the analysis of the ambiguity problem [7, 10]. However, all of these techniques assume deterministic projection measurements and do not perform optimally under high noise conditions.

A second approach to discrete-valued reconstruction detects parameterized objects directly in the projection domain. This strategy is applicable when the objective is to detect specific objects or regions such as tumors in medical imaging or material defects in non-destructive testing. Rossi and Willsky [11] introduced this approach by performing maximum likelihood (ML) estimation of the location of a single object in the imaging plane. This concept was later extended to a three-dimensional parameterization supporting multiple objects per plane [12]. Here, constrained objects in 3-D are formed as a combination of basic cylinders whose parameters are estimated as part of the reconstruction. A review of object parameterization methods as well as a new algorithm for the approximate reconstruction of compact objects modeled by polyhedral shapes is given in the chapter by A. Mohammad-Djafari and C. Soussen in this book. Parameterized object reconstruction methods are specifically designed for low SNR conditions. However, they rely on *a priori* knowledge about shape characteristics of the objects in the cross-section. These methods are therefore not applicable in cases where such information is unavailable or the objects in the cross-section cannot easily be parameterized.

In this work, we focus on discrete-valued reconstruction from noisy projections using statistical methods. Statistical methods model the random nature of the physical data collection process, then seek the solution that best matches the probabilistic behavior of the data. Consequently, statistical methods can improve performance considerably in cases of low SNR.

Statistical approaches also easily incorporate special geometries such as limited or missing angle projection measurements. Common statistical techniques incorporate implicit information about desired characteristics of the reconstruction without explicit modeling of objects in the cross-section.

A statistical method that is well suited for tomographic reconstruction is Bayesian maximum *a posteriori* (MAP) estimation. Bayesian methods in general have been shown to improve performance in many emission and transmission tomography problems [13, 14, 15, 16] as well as in image restoration tasks [17, 18]. Bayesian MAP estimation reconstructs the image as a tradeoff between matching the projection data and regularizing the solution by a prior probability distribution. The regularization imposed by the prior reflects assumed characteristics of feasible reconstructions. Due to this regularization, the MAP estimation problem is well-posed and avoids the high noise sensitivity frequently encountered in maximum likelihood (ML) estimation.

Priors for Bayesian reconstruction methods are often chosen to impose smoothness constraints on the reconstruction to eliminate high frequency noise. A prior model that has generally proven to be useful in the tomographic setting is the Markov random field (MRF) image model [13, 19, 14, 15, 20]. The chapter by M. T. Chan, G. T. Herman and E. Levitan in this book presents a new MRF for modeling image prior distributions and describes methods for estimating the MRF's parameters. These parameter estimation techniques allow the model to be adapted to the specific characteristics of an ensemble of images. Importantly, the Bayesian estimate is computed directly from the convolution backprojection reconstruction, rather than from the original projection data. This approach has the advantage of reducing computation.

In contrast, we use a simple discrete MRF model [21, 22, 33], and instead focus on the computational difficulties resulting from direct Bayesian reconstruction from the tomographic data. In order to solve this optimization problem, we employ multiscale algorithms to both reduce computation and improve convergence to the global minimum. In addition, we introduce a method for estimating the densities of the discrete regions as part of the reconstruction process. This is important because precise knowledge of these discrete densities is required for accurate Bayesian reconstruction.

The MAP reconstruction itself can be formulated as an optimization problem which can be solved using a number of different techniques. The expectation-maximization (EM) algorithm, suitable for ML reconstruction [23], has been adopted for MAP estimation with Gaussian priors [24, 25, 26, 27]. Extensions of these models to more general MRF priors were proposed in [14, 28, 15, 29]. However, application of the EM algorithm for MAP estimation is difficult and usually suffers from slow convergence.

Instead of using EM techniques, we focus on the direct optimization of the MAP equation. We adopt a pixel-wise update method known as iterative coordinate descent (ICD) [16, 30] which 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 [22]. However, while ICM was designed for image restoration tasks, the ICD algorithm is specifically designed for the tomographic reconstruction problem resulting in dramatically improved computational efficiency.

In addition to solving the optimization problem, the discrete-valued reconstruction requires knowledge of the density or emission rate levels in the cross-section. In practice, exact information about these discrete levels is often unavailable. In such cases it is desirable to estimate the densities or emission rates as part of the reconstruction. In Section 1.5, we discuss our method to estimate these discrete classes concurrently with the reconstruction process. We show how the class estimation can be formulated as a continuous-valued tomographic reconstruction problem with the number of points equal to the number of classes.

Finally, we extend our reconstruction method to a multiresolution algorithm. Multiresolution techniques achieve performance improvements in a variety of imaging problems [31, 32] including image segmentation [33, 34] and continuous-valued tomographic reconstruction [35]. Multiresolution algorithms reconstruct the image at different resolutions, typically progressing from coarse to fine scale. The coarse scale solutions serve as initialization or prior information for reconstructions at finer scales. Due to the improved initialization and the higher SNR at coarse scale, multiresolution algorithms are typically more robust with respect to local minima. In addition, local pixel interactions at coarse scale are equivalent to large scale interactions at fine scale. This combined with the low computational complexity at coarse scale makes multiscale algorithms very efficient.

The multiscale algorithm presented here is a straightforward extension of our fixed scale algorithm. The reconstruction is performed in a coarse-to-fine fashion by initializing each resolution level with the interpolated reconstruction of the next coarser level. The reconstructions at each level are computed using the fixed resolution method. Our experimental results demonstrate that this multiscale algorithm is less prone to being trapped in local minima and in many cases, computationally more efficient than the fixed resolution version.

1.2 Stochastic data models for tomography

In this section, we will develop the statistical framework for MAP reconstruction in computed tomography. Our framework is applicable to both transmission and emission measurements and supports general imaging geometries. The models presented here are based on the exact Poisson statistics of the photon measurements. Computationally more efficient but

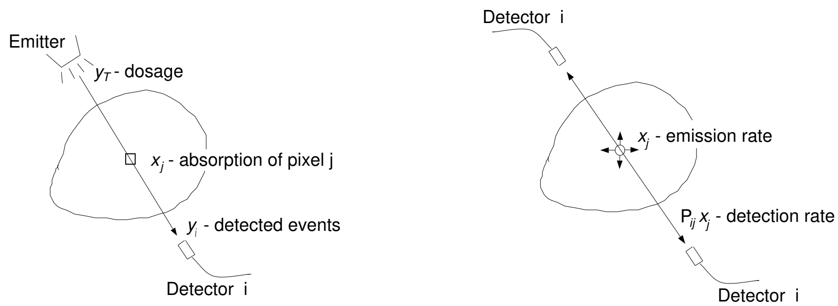

(a) Transmission setup (b) PET emission setup

FIGURE 1.1. Physical setup for transmission and positron emission tomography (PET). In transmission tomography (a), photons are induced into the cross-section. After attenuation by the absorptive material, detectors measure the remaining photon rates. In positron emission tomography (b), the cross-section contains a radioisotope and is surrounded by a ring of detectors. If two detectors register photons at the same time, a pixel emission is assumed to have occurred on the projection line between them.

approximate models can be obtained by using a Taylor expansion of the likelihood function [16, 30].

In transmission tomography, the objective is to measure photon attenuation for different projections through a cross-section of absorptive material. An illustration of the physical setup is shown in Fig. 1.1(a). The cross-section is surrounded by a ring of photon emitters and detectors. The emitters induce a calibrated photon rate y_T directed along certain angular and parallel projections. After attenuation by the absorptive material, the photon rates are measured by the detector opposite the respective emitter. The photon rates measured by the detectors are not direct measurements of attenuation. Instead they are noisy photon counts which can be modeled as Poisson-distributed random variables.

In order to write the probability density for the Poisson measurements, 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 distribution of the Y_i may then be written as

$$\mathcal{P}(Y = y|x) = \prod_{i=1}^{M} \frac{\exp(-y_T e^{-P_{i*}x})(y_T e^{-P_{i*}x})^{y_i}}{y_i!}. \tag{1.1}$$

We use upper case letters for random variables and lower case letters for particular realizations. Taking the logarithm of (1.1), we obtain the log-likelihood

$$\text{(transmission)} \quad L(y|x) = \log \mathcal{P}(Y = y|x) = \sum_{i=1}^{M} \left(-y_T e^{-P_{i*}x} + y_i(\log y_T - P_{i*}x) - \log(y_i!) \right). \tag{1.2}$$

In emission tomography, no dosage is induced into the cross-section. Instead, the image plane contains some photon emitting material. A physical setup for the specific example of positron emission tomography (PET) is shown in Fig. 1.1(b). In this case, the cross-section contains a radioactive isotope. Recombination of positrons in the radioisotope results in emission of gamma rays in two opposite directions. These gamma rays are detected by a ring of detectors around the cross-section. If two detectors register photons at the same time, this is counted as an emission on the projection line between them. In the following, we will develop the statistical model for the general emission case.

The objective in emission tomography is to reconstruct the emission rates of all pixels in the image plane. Again, the photon detections can be modeled as Poisson distributed random variables. In order to emphasize the similarity to the transmission problem, we will use the same notation, but interpret x as the vector of emission rates for all N pixels 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 distribution

$$\mathcal{P}(Y = y|x) = \prod_{i=1}^{M} \frac{\exp(-P_{i*}x)(P_{i*}x)^{y_i}}{y_i!}. \tag{1.3}$$

The log-likelihood is therefore given by

$$\text{(emission)} \quad L(y|x) = \log \mathcal{P}(Y = y|x) = \sum_{i=1}^{M} \left(-P_{i*}x + y_i \log(P_{i*}x) - \log(y_i!) \right). \tag{1.4}$$

The log-likelihood functions for both the transmission and the emission case have the form

$$\log \mathcal{P}(Y = y|x) = -\sum_{i=1}^{M} f_i(P_{i*}x) \tag{1.5}$$

where $f_i(\cdot)$ are convex and differentiable functions. This common form will lead to similar methods of solving these two problems. In the following, we will write all equations for the emission case; however, all methods apply analogously to the transmission case.

For the emission problem, maximum likelihood (ML) estimation of x from y yields the optimization problem

$$\hat{x}_{ML} = \arg\min_x \sum_{i=1}^{M} \left(P_{i*}x \; - y_i \log(P_{i*}x)\right) \ . \tag{1.6}$$

For low signal-to-noise-ratio medical imaging problems, the ML estimate has well documented shortcomings [36, 37, 38]. Noise and sampling limitations can produce high frequency noise in the ML reconstruction that is not present in the original cross-section. It is therefore desirable to regularize tomographic inversion by some means. Maximum *a posteriori* probability (MAP) estimation addresses this problem by treating the original image as a random field, X, with prior distribution, $p(x)$. Again, we use a lower case x to denote a particular realization of the random vector X. The prior distribution regularizes the optimization problem so that a unique solution always exists [39]. The logarithm of the *a posteriori* distribution of X given Y may be computed using Bayes' formula.

$$\begin{aligned} L_p(x|y) &\stackrel{\triangle}{=} \log \mathcal{P}(X = x | Y = y) \\ &= L(y|x) + \log p(x) - \log \mathcal{P}(Y = y) \end{aligned} \tag{1.7}$$

The maximum *a posteriori* (MAP) estimate is then the value of $\hat{x}$ which maximizes the *a posteriori* density given the observations y

$$\begin{aligned} \hat{x} &= \arg\max_x L_p(x|y) \\ &= \arg\max_x \{L(y|x) + \log p(x)\} \\ &= \arg\max_x \{L(y,x)\} \, . \end{aligned} \tag{1.8}$$

The last equation indicates that the MAP estimate also maximizes the log of the joint distribution, $L(y,x) = \log \mathcal{P}(X = x, Y = y)$.

The MAP estimate has been shown to substantially improve performance in many image reconstruction and estimation problems. While computation of the exact MAP estimate is computationally intractable for the discrete problem, approximate solutions can be obtained with reasonable complexity as outlined in the next section. We will treat only the MAP estimation problem, since the ML estimate is the special case of a constant prior distribution.

1.3 Markov random field prior models

While the likelihood term $L(y|x)$ in the MAP equation (1.8) is determined by the physics of the data collection process, the prior distribution is selected by the experimenter to model desired characteristics of typical reconstructions. Most commonly, the prior models are chosen to reflect the high correlation of adjacent pixels. A model that has proven particularly useful is the Markov random field (MRF) [40]. Similar to a Markov chain in one dimension, the 2D MRF limits pixel interactions to a local neighborhood of pixels. This localization allows for efficient optimization of the MAP equation.

In order to write the equations for the MRF prior, we first need to define the concept of a neighborhood. If i denotes a single pixel location, we will denote its neighborhood by ∂i. This neighborhood can consist of any set of pixels $\{k : k \neq i\}$ which satisfies the symmetry property that $i \in \partial k \Rightarrow k \in \partial i$. Given this definition, a MRF is a random field which has the property

$$p(x_i | x_j, j \neq i) = p(x_i | x_{\partial i}). \tag{1.9}$$

In other words, the prior conditional probability of a pixel value depends only on a local neighborhood of pixels. Under some weak technical conditions, a random field is a MRF if and only if it has a probability distribution corresponding to a Gibbs distribution [40, 41]. This result, which is known as the Hammersley-Clifford theorem, may be used to express the likelihood function $\log p(x)$. While the theory of MRF's is quite extensive [42, 43, 44], we will restrict ourselves to a simple model based on at most an 8-point neighborhood.

Since we are interested in discrete-valued tomographic reconstruction, we assume that each pixel takes on a value from a set $\mathcal{E}$ of K discrete emission rates. We then apply a discrete MRF prior model that is frequently used in segmentation problems [21, 22, 33]. The model encourages neighboring locations to have the same states or, in our case, emission rates. To define the model, we must first define two simple functions, $t_1(x)$ and $t_2(x)$. $t_1(x)$ is the number of horizontally and vertically neighboring pixel pairs with different emission rates in x, and $t_2(x)$ is 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

$$\log p(x) = -(\beta_1 t_1(x) + \beta_2 t_2(x)) + \log(Z) \tag{1.10}$$

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 will result in a smoother reconstruction. Based on the geometry of the 8-point neighborhood, β_2 is often chosen as $\beta_2 = \beta_1/\sqrt{2}$. In the following, we will often write β for β_1 and assume $\beta_2 = \beta_1/\sqrt{2}$.

Substituting the prior (1.10) into the MAP and likelihood equations (1.8) and (1.4) for the emission case, we obtain the optimization criterion

$$\hat{x} = \arg\max_{x\in\mathcal{E}^N}\left\{\sum_{i=1}^{M}\left(-P_{i*}x \ + y_i \log(P_{i*}x)\right) - (\beta_1 t_1(x) + \beta_2 t_2(x))\right\}. \tag{1.11}$$

1.4 Optimization techniques

In order to compute the MAP reconstruction, we must perform the optimization of (1.11). Gradient methods such as steepest descent or conjugate gradient optimization are not directly applicable, since the discrete prior (1.10) is non-differentiable.

A method that is well suited for the MAP optimization is a discrete version of iterative coordinate descent (ICD) [16, 30]. The ICD method sequentially updates each pixel of the image. With each update, the current pixel is chosen to maximize the posterior probability (1.11). Therefore, the discrete ICD algorithm essentially implements the well known ICM optimization introduced by Besag [22]. However, while ICM was originally developed for image restoration tasks, the ICD implementation is specifically designed for the tomographic reconstruction problem. 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 instead of the ML initialization used by ICM. The ML estimate is not a good initialization for tomographic reconstruction problems and, since the pixel likelihoods are not independent, the ML estimate is computationally expensive to compute. The convolution backprojection, in comparison, is inexpensive to compute and captures most of the low-spatial frequency behavior of the reconstruction. This makes the CBP a suitable initialization, especially since coordinate-wise update methods have slow convergence for low spatial frequencies and fast convergence for high-spatial frequencies. In the following, we will show how the ICD can be used to efficiently compute the MAP estimate.

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

$$x_j^{n+1} = \arg\min_z\{-L(y|X_j = z, X_k = x_k^n, k \neq j) + (\beta_1 v_1(z, x_{\partial j}) + \beta_2 v_2(z, x_{\partial j}))\}. \tag{1.12}$$

In our notation, x^n is the image containing all previous pixel updates. Thus, the reconstruction x^{n+1} differs from x^n only at pixel j. A full update of the reconstruction requires N applications of (1.12).

Computation of the log-likelihood $L(y|z, x^n)$ using (1.4) for each pixel update would still lead to prohibitive computational complexity. This can be avoided by using only the change in the log-likelihood

$$\Delta L(z) = \sum_{i \in \mathcal{I}_j} \left(-P_{ij} z + y_i \log(P_{i*} x^n + P_{ij}(z - x_j^n)) - y_i \log(P_{i*} x^n) \right) \tag{1.13}$$

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

$$\begin{aligned} x_j^{n+1} = \arg\min_z \Bigg\{ & \sum_{i \in \mathcal{I}_j} (P_{ij} z - y_i \log(P_{i*} x^n + P_{ij}(z - x_j^n))) \\ & + (\beta_1 v_1(z, x_{\partial j}) + \beta_2 v_2(z, x_{\partial j})) \Bigg\}. \end{aligned} \tag{1.14}$$

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 (1.14). We store the M-dimensional state vector $S = Px$ between iterations. After a pixel x_j is updated, the components of S can be efficiently updated using

$$S_i^{n+1} = S_i^n + P_{ij}(x_j^{n+1} - x_j^n). \tag{1.15}$$

This update is necessary only for the components $i \in \mathcal{I}_j$ since for all other projections $P_{ij} = 0$.

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

$$M_0 = \frac{1}{N} \sum_{j=1}^{N} |\mathcal{I}_j| \ . \tag{1.16}$$

The computational cost for a pixel update is then on the order of KM_0 operations. The complexity of a full-update of the reconstruction is therefore NKM_0. This is quite reasonable, considering that due to the sparsity of P, M_0 is typically small compared to M, i.e. $M_0 \ll M$. In cases where K is large, it might be desirable to reduce computation by using a global second order approximation to the likelihood functions (1.2) and (1.4) as described in [16, 30]. Using these techniques, computation for a single pixel update can be reduced to order $K + M_0$, resulting in $N(K + M_0)$ complexity for a full reconstruction update.

1.5 Estimation of discrete levels

So far, we have assumed the set $\mathcal{E}$ of discrete emission rates or densities to be known. In practice, however, the exact emission rates corresponding to different regions in the cross-section may not be known. Even if a good initial guess is available, the accuracy of the emission rates is critical for the reconstruction. For illustration, assume that the emission rate of a particular region is over-estimated by some amount. For projections y_i which pass through this region, the forward projected reconstruction $P_{i*}x$ will be larger than the measured photon count, i.e. $P_{i*}x > y_i$. To compensate for this mismatch, the reconstruction algorithm may misclassify large numbers of pixels. Therefore, it is desirable to estimate the discrete emission rates as part of the reconstruction algorithm.

In this section, we show how the emission rates can be estimated concurrently with the reconstruction. We implement ML estimation of the emission rates by iteratively updating entire regions of pixels with equal emission rates [45]. We will show that this estimation is equivalent to a continuous-valued tomographic ML reconstruction problem with K pixels where K is the number of discrete levels in the reconstruction. The updates of the emission rates will be performed between each full ICD update of the reconstruction.

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. Given the region geometries, we can compute the entry Q_{ik} as the probability that an emission from the k^{th} region is registered by the i^{th} detector.

Assuming knowledge of Q, the likelihood for the emission rates can be computed analogously to the pixel likelihoods in (1.2) and (1.4). The resulting optimization problem is equivalent to a continuous-valued reconstruction with K pixels and projection matrix Q.

In practice, direct computation of Q from the geometry of the regions would be computationally involved and difficult to update. Instead, we can obtain an expression for the entries of Q by adding the contributions of all pixels in a region. We can rewrite the i'th forward projection $P_{i*}x$ as follows

$$P_{i*}x = \sum_{j=1}^{N} P_{ij} x_j$$

$$= \sum_{k=1}^{K} \left(\theta_k \sum_{\{j:x_j=\theta_k\}} P_{ij} \right)$$

$$= \sum_{k=1}^{K} \theta_k Q_{ik} \quad = \quad Q_{i*}\theta \tag{1.17}$$

where

$$Q_{ik} = \sum_{\{j:x_j=\theta_k\}} P_{ij}. \tag{1.18}$$

For the emission case, this yields the likelihood function

$$\log \mathcal{P}(Y=y|\theta) = \sum_{i=1}^{M} \left(-Q_{i*}\theta + y_i \log(Q_{i*}\theta) - \log(y_i!) \right). \tag{1.19}$$

This log-likelihood function is clearly of the same form as (1.4), except that the discrete-valued N-component vector x is replaced by the continuous-valued K-component vector θ. Also, the new projection matrix Q is of size $M \times K$ instead of $M \times N$. Thus, maximum likelihood estimation of θ is equivalent to a continuous-valued tomographic ML reconstruction with K pixels.

Since θ is continuous-valued, the optimization of (1.19) is different from the discrete case. In general, all methods proposed for continuous-valued tomographic reconstruction can be applied. Again, we will use 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 (1.14) given by

$$\theta_k^{n+1} = \arg\min_{v \geq 0} \left\{ \sum_{i \in \tilde{I}_k} \left(Q_{ik} v - y_i \log(Q_{i*}\theta^n + Q_{ik}(v - \theta_k^n)) \right) \right\} \tag{1.20}$$

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 (1.20) is well approximated by a quadratic, the optimization can be efficiently implemented using Newton's method.

Let ϕ_1 and ϕ_2 be the first and second derivatives of the log-likelihood function evaluated at the current emission rate θ_k. The Newton update for minimization of (1.20) is then given by

$$\theta_k' = \max \left\{ \theta_k - \frac{\phi_1}{\phi_2}\ ,\ 0 \right\} \tag{1.21}$$

where the derivatives ϕ_1 and ϕ_2 are computed as

$$\phi_1 \quad = \quad \sum_{i \in \tilde{I}_k} Q_{ik} \left(1 - \frac{y_i}{Q_{i*}\theta} \right) \tag{1.22}$$

$$\phi_2 \quad = \quad \sum_{i \in \tilde{I}_k} y_i \left(\frac{Q_{ik}}{Q_{i*}\theta} \right)^2 . \tag{1.23}$$

The Newton updates (1.21)-(1.23) are repeatedly applied until $|\phi_1| < \epsilon$. For our experimentation, we have found an accuracy of $\epsilon = 0.001$ to be sufficient. For efficient computation, we store and update the same state vector $S = Px = Q\theta$ as in the discrete MAP reconstruction (1.15). If θ_k^n is updated to θ_k^{n+1}, S can be updated as

$$S_i^{n+1} = S_i^n + Q_{ik}(\theta_k^{n+1} - \theta_k^n) \tag{1.24}$$

for all $i \in \tilde{I}_k$.

In order to assess computational complexity, we define H as the average number of Newton-iterations per class update (1.20). Furthermore, we use the number of projections M as a bound for the size of the sets $\tilde{I}_k$. Computation of a single emission rate θ_k then requires on the order of MH multiplies and divides. Typically, only a few Newton-iterations are necessary to obtain sufficient accuracy ($H < 2$). The computational complexity of updating the ML estimate of θ is then KMH. Note that KMH is typically small in comparison to a complete update of the discrete reconstruction x of order NKM_0.

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

$$\begin{aligned} Q'_{ik} &= Q_{ik} - P_{ij} \\ Q'_{il} &= Q_{il} + P_{ij} \end{aligned} \tag{1.25}$$

for all projections $i \in I_j$. This recursion results in a computationally efficient algorithm since it avoids recomputing Q using (1.18) 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 convolution backprojection reconstruction. One possibility is to extract the initial θ_k manually by taking the average value of approximately uniform regions in the CBP. This ensures that the estimated emission rates correspond to the regions of interest in the reconstruction and minimizes chances of the estimation getting trapped in local minima.

If, on the other hand, a fully unsupervised algorithm is desired, clustering techniques can be applied to the CBP reconstruction to estimate the initial emission rates θ_k. To do this, we used a clustering method based on Gaussian mixture models and the EM algorithm [46]. This method used the Rissanen criterion to estimate both the number of clusters K as well

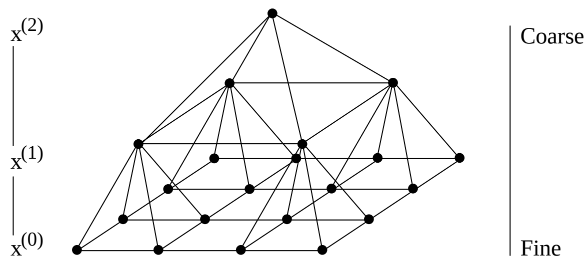

FIGURE 1.2. Illustration of multiresolution structure. Shown are pixels in three different resolution levels. We assume a quadtree structure in which each coarse scale pixel corresponds to four pixels at the next finer scale.

as the mean emission rate of each cluster. These estimates were then used to initialize the estimation of the emission rates which resulted in a fully unsupervised reconstruction algorithm.

1.6 Multiscale approaches

We now extend the previous results to a multiresolution framework. Multiresolution algorithms reconstruct the cross-section at different resolutions, typically starting at coarse resolution and progressing to the desired finest resolution. Fig. 1.2 illustrates the multiresolution structure used in our algorithm. We assume a quadtree model in which each coarse scale pixel corresponds to four pixels at the next finer scale. Each resolution level is half the size of the next finer level in each direction and therefore contains only $1/4^{th}$ the number of pixels. The small number of pixels at coarse scales implies lower computational complexity at these levels.

Compared to the fixed scale reconstruction, the multiscale approach has several significant advantages. While at first, reconstructing the cross-section at several resolutions might seem like additional overhead, the multiscale algorithm has substantially faster convergence behavior. The fixed scale ICD reconstruction algorithm updates one pixel at a time using a prior that only depends on a small pixel neighborhood. As a result, propagation of information per iteration is limited which results in slow convergence for low spatial frequencies. The multiscale version of the algorithm improves this by first reconstructing the image at coarse resolutions where local interactions are equivalent to large scale propagation at fine resolutions. The coarse scale reconstructions then serve as an initialization for the finer reconstructions. Since the coarse reconstructions already contain the large scale behavior of the solution, substantially fewer iterations are necessary at the finer scales. This, combined with the fact that the coarse scale reconstructions are of low complexity, makes the multiscale algorithm very efficient. In addition to increased efficiency, the multiresolution algorithm is more robust with respect to local minima in the optimization. This increased robustness holds for both the reconstruction and the estimation of the emission rates.

1.7 Multiscale MRF

The multiscale MRF model is a straightforward generalization of the fixed resolution model in section 1.3. For the multiresolution case, we essentially use the fixed scale algorithm for each resolution level and use the result to initialize the next finer level [33].

Let $x^{(n)}$ denote the reconstruction at resolution n, where $n = 0$ is the finest resolution 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 larger pixel size at level n. The matrix $P^{(n)}$ is of dimension $M \times 4^{-n}N$. The log-likelihood for the emission case is then given by

$$L^{(n)}(y|x^{(n)}) = \sum_{i=1}^{M} \left(-P_{i*}^{(n)} x^{(n)} \; + y_i \log(P_{i*}^{(n)} x^{(n)}) - \log(y_i!) \right) . \qquad (1.26)$$

This yields the MAP equation

$$\hat{x}^{(n)} = \arg\min_{x^{(n)}} \left\{ -L^{(n)}(y|x^{(n)}) + \beta_1^{(n)} t_1(x^{(n)}) + \beta_2^{(n)} t_2(x^{(n)}) \right\} . \qquad (1.27)$$

The remaining question is how to choose the coarse resolution parameters $\beta_1^{(n)}$ and $\beta_2^{(n)}$. An intuitive approach is to choose these parameters so that the cost functions for any two adjacent resolutions are equal when the finer reconstruction $x^{(n-1)}$ equals the coarser reconstruction $x^{(n)}$ [33]. This assumes that the finer reconstruction $x^{(n)}$ is constant on blocks of 2 by 2 pixels.

Let I denote the operator which interpolates by a factor of two using pixel-replication. The equality of adjacent levels can then be written as $x^{(n-1)} = Ix^{(n)}$. We now observe that a horizontal or vertical pixel difference in $x^{(n)}$ results in two horizontal or vertical plus two diagonal differences in the pixel-replicated $Ix^{(n)}$, and one diagonal pixel difference in $x^{(n)}$ yields one diagonal pixel difference in $Ix^{(n)}$. Therefore, $x^{(n-1)} = Ix^{(n)}$ implies

$$\begin{aligned} t_1^{(n-1)} &= 2\, t_1^{(n)} \\ t_2^{(n-1)} &= 2\, t_1^{(n)} \; + \; t_2^{(n)} \; . \end{aligned} \qquad (1.28)$$

Consequently, the fine and coarse resolution cost functions will be equal if

$$\begin{aligned} \beta_1^{(n)} &= 2(\beta_1^{(n-1)} \; + \; \beta_2^{(n-1)}) \\ \beta_2^{(n)} &= \beta_2^{(n-1)} \end{aligned} \qquad (1.29)$$

for all resolutions n. By using these parameters, minimization of (1.27) corresponds to the minimization of the original MAP equation (1.11) under the constraint that the solution be constant on the appropriately sized blocks.

Coarse resolution minimization using the parameters given by (1.29) will effectively minimize (1.11) if the correct segmentation is approximately block constant. However, this recursion for the parameters has an undesirable property. It implies that the MRF models for coarser resolution segmentations should have progressively higher spatial correlation, or alternatively, finer resolution segmentations should have lower correlation. This, of course, runs counter to normal assumptions of spatial coherence in images, and will tend to cause insufficient spatial correlation at finer resolutions or excessive correlation at coarse resolutions. A more reasonable approach is to assume that the spatial correlation is independent of the resolution since this avoids the problem of excessive correlation at coarse resolutions. Also, this assumption is appropriate when prior information is unavailable about the likely scale of regions in the image. Therefore, in all experimentation, we will fix the parameters of the MRF as a function of scale

$$\begin{aligned} \beta_1^{(n)} &= \beta_1 \\ \beta_2^{(n)} &= \beta_2 . \end{aligned} \tag{1.30}$$

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 the following method:
   (a) Update $x^{(n)}$ using fixed resolution MAP update
   (b) If no pixel change occurs, goto 4.
   (c) Else, perform six full updates of discrete levels θ, goto (a).
4. If $n = 0$ stop.
5. Initialize $x^{(n-1)}$ with pixel-replicated $x^{(n)}$.
6. Set $n = n-1$, goto 3.

The parameters β_1 and β_2 can be chosen manually to achieve the amount of regularization desired.

	Reconstruction Update	Emission Rate Update
Fixed Scale	NKM_0	KMH
Multiscale	$4^{-n}NKM_0^{(n)} \approx 2^{-n}NKM_0$	KMH

N	Number of pixels in the reconstruction
K	Number of discrete emission rates θ_k
M	Number of projections
$M_0^{(n)}$	Average number of projections intersecting a pixel at scale n, $M_0 = M_0^{(0)}$
H	Average number of Newton iterations for update of single θ_k

TABLE 1.1. Computational complexity for reconstruction and emission rate updates. While the complexity of the reconstruction is a function of the resolution n, the cost for the estimation of θ is constant. For the overall reconstruction, the time spent on estimation of θ is typically less than 10% of total execution time.

1.8 Computational complexity

Table 1.1 compares the computational complexity for one full update of the reconstruction to one update of the emission rates θ. The complexity of the reconstruction update NKM_0 depends on the number of pixels N and therefore on the resolution of the reconstruction. In the multiresolution framework, each level contains $1/4^{th}$ the number of pixels of the next finer level. The reconstruction complexity at scale n is therefore $4^{-n}NKM_0^{(n)}$, where $M_0^{(n)}$ replaces M_0. Due to the larger pixel size at coarse resolution, more projections intersect each coarse-scale pixel and $M_0^{(n)}$ increases with n. In order to compare the complexities for different scales, we can approximate $M_0^{(n)}$ as follows: As the size of a pixel doubles in each direction, we assume that the number of parallel projections intersecting the pixel at each angle doubles. The cost of a reconstruction update at scale n is therefore approximately $2^{-n}NKM_0$. Assuming that the multiscale reconstruction performs the same number of iterations at each scale as the fixed resolution algorithm, the multiscale overhead is bounded by a factor of 2. In most real applications, however, the multiscale algorithm performs considerably fewer computationally expensive iterations at fine scale than the fixed scale method.

The complexity of the emission rate update is not a function of resolution. To obtain a bound on the computational cost, we assume that all M projections intersect each region. Since the number of Newton iterations per class update H is usually small, the upper bound KMH for the emission rate update is small in comparison to the cost NKM_0 for a reconstruction update at finest or fixed scale. At coarse resolutions, however, the complexity for the emission rate update may become comparable to the cost for a reconstruction update. Again, the advantage of the multiscale method is that the emission rates often converge after the reconstruction of the coarser scales. Therefore, fewer iterations for reconstruction and emission rate estimation are necessary at finer scale.

In practice, the total cost of the emission rate estimation is usually small compared to that of the reconstruction updates. Performing six updates of θ between full reconstruction updates, we find that the cost of the emission rate updates is typically less than 10% of the total execution time. In addition, the multiscale algorithm is typically faster than the fixed resolution method.

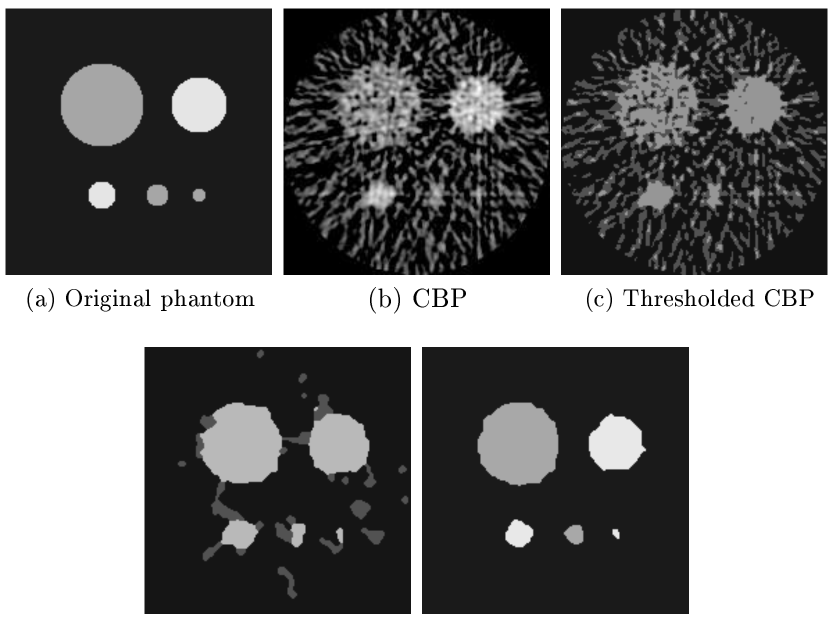

(a) Original phantom (b) CBP (c) Thresholded CBP

(d) Fix-Resolution MAP $\beta = 1.0$ (e) Multiresolution MAP $\beta = 1.0$

FIGURE 1.3. 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.

	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 1.2. 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}.

	Multi-Resolution	Fixed-Resolution
Phantom I	27	59
Phantom II	349	404

TABLE 1.3. CPU-time in seconds for fixed- and multiresolution algorithms. Both algorithms were run until convergence and terminated when no pixel change occurred in a discrete reconstruction update.

1.9 Results

Reconstructions using the fixed and multiscale algorithms on synthetic data are shown in Fig. 1.3. Fig. 1.3(a) shows the original cross-section of size 192 by 192 pixels where each pixel is of both width and height 3.13mm. The cross-section contains pixels with three different emission rates as shown in Table 1.2. The projection data was calculated at 16 evenly spaced angles each with 192 parallel projections. The projection beam was assumed to be infinitely thin. The data samples where formed by Poisson random variables with the appropriate means.

Fig. 1.3(b) shows the convolution backprojection (CBP) reconstruction using a generalized Hamming filter weighted by a Gaussian envelope. The CBP reconstruction was used to obtain initial values for the emission rates θ. The unsupervised clustering routine using a Gaussian mixture model 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. A first discrete-valued reconstruction can be obtained by thresholding the CBP reconstruction using the midpoints between the emission rates determined by the clustering routine. The resulting thresholded CBP is shown in Fig. 1.3(c). In addition to the errors in class estimates, the result contains noise and aliasing effects. Fig. 1.3(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 fixed resolution reconstruction was initialized

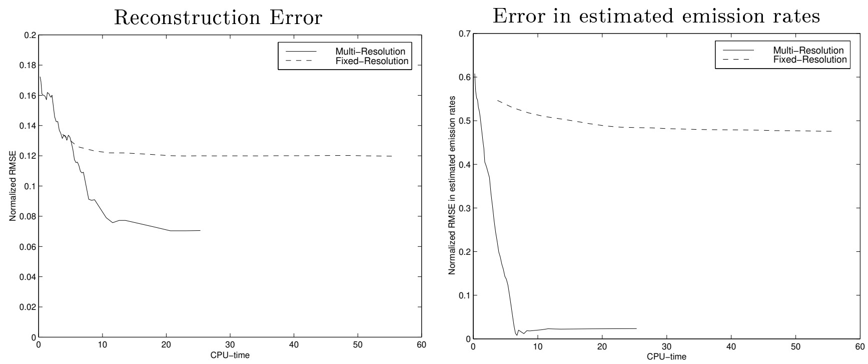

(a) (b)

FIGURE 1.4. Comparison of convergence for the fixed- and multiresolution MAP reconstructions of Fig. 1.3(d)-(e). Shown in (a) is the normalized reconstruction error as a function of CPU-time. The multiscale algorithm converges considerably faster than the fixed resolution method. This is partly due to the fast convergence of the emission rate estimates in the multiscale case (b). The multiresolution algorithm achieves lower final error in estimation and reconstruction.

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 from the clustering result. This results in the classification of all 5 discs into the same class of emission rates.

The reconstruction result using the multiscale algorithm with $L = 5$ resolution levels and $\beta = 1.0$ is shown in Fig. 1.3(e). The algorithm was initialized as in the fixed scale case. The estimated emission rates using the multiscale technique are very close to the true values as shown in Table 1.2. This results in correct classification of the 4 larger discs in the cross-section. Only the smallest disc is misclassified to a smaller area but higher emission rate than in the original phantom. This is not surprising, considering the high level of noise and the small size of the disc which practically eliminates it from coarser resolution levels. For $\beta = 1.0$, there is essentially no high frequency noise in the reconstruction. Overall, the quality of the multiscale reconstruction is superior to the fixed scale MAP reconstruction and the thresholded CBP. The multiscale method is particularly robust with respect to the estimation of the emission rates.

In addition to the superior reconstruction quality, the multiresolution method is faster than the fixed scale algorithm. Table 1.3 shows the execution times for both the fixed- and multiresolution reconstructions (Phantom I). Both algorithms were run until convergence and terminated when no pixel change occurred in a discrete reconstruction update. While the multiscale method terminated after 25 seconds, the fixed scale method needed

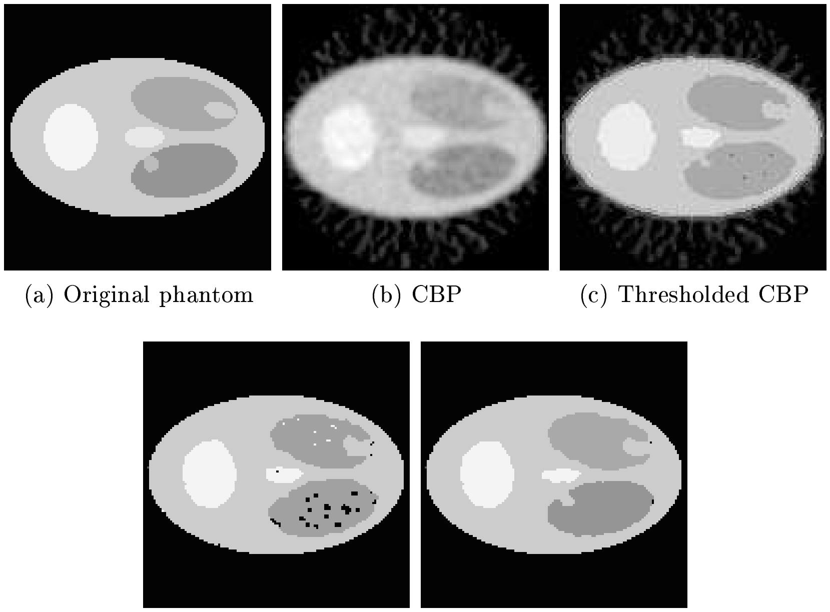

(a) Original phantom (b) CBP (c) Thresholded CBP

(d) Fix-Resolution MAP $\beta = 1.0$ (e) Multiresolution MAP $\beta = 1.0$

FIGURE 1.5. Results for second synthetic cross-section. The clustering only identifies 2 out of 7 classes correctly which is reflected in the thresholded CBP (c). The fixed resolution algorithm (d) estimates 3 of 7 classes correctly, and fails to divide the two gray ovals on right into two distinct classes, which results in compensation artifacts. The multiresolution algorithm (e) performs better, estimating 5 out of seven classes correctly. Neither algorithm divides the bright left and center spot into two distinct classes. Both also miss the emission rate of the small patch within the lower right oval.

56 seconds to converge. Fig. 1.4 compares the error convergence of the two algorithms. Shown in Fig. 1.4(a) is the reconstruction error as a function of CPU-time. The reconstruction error is calculated as normalized root mean square error, i.e.

$$E = \sqrt{\frac{\sum_{j=1}^{N} (\hat{x}_j - x_j)^2}{\sum_{j=1}^{N} x_j^2}} \tag{1.31}$$

where x denotes the original cross-section and $\hat{x}$ the current, if necessary interpolated, reconstruction. The multiscale algorithm converges considerably faster and achieves lower final reconstruction error than the fixed resolution method. Not shown is that the multiscale method achieves larger posterior likelihood, confirming that the fixed scale algorithm gets trapped in a local minimum.

	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7
Original Phantom	0.001	1.2	1.6	2.0	2.4	3.2	3.6
CBP Clustering	0.0005	0.028	0.094	0.307	1.606	2.359	3.335
Fixed-Res. MAP	0.001	0.028	0.094	0.338	1.445	2.403	3.574
Multi-Res. MAP	0.001	0.144	0.336	1.211	1.602	2.404	3.578

TABLE 1.4. Original and estimated emission rates for second synthetic cross-section. Both, the fixed and multiscale algorithm were initialized to the emission rates determined by the clustering algorithm. Units in mm^{-1}.

The difference in convergence speed and final reconstruction error between the fixed and multiscale algorithm is reflected in the convergence behavior of the emission rate estimates. Fig. 1.4(b) shows the convergence of the normalized root mean square error of the emission rate estimates. For the multiscale algorithm, the θ are essentially converged after only a few coarse scale iterations. This reduces the number of computationally expensive iterations at finer scales, thereby accelerating overall convergence substantially.

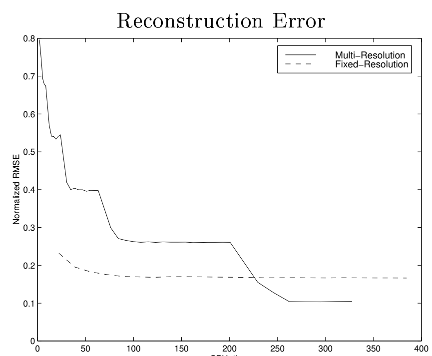

FIGURE 1.6. Error convergence for fixed- and multiresolution MAP reconstructions of the second phantom. While the multiscale error at coarser scales is comparably high, convergence at each resolution is rapid. Again, the multiresolution algorithm terminates earlier and achieves lower final reconstruction error than the fixed-resolution method.

Results for a second synthetic cross-section are shown in Fig. 1.5. The original cross-section in Fig. 1.5(a) contains 7 discrete levels of emission rates. The size of the phantom is 128 by 128 pixels where each pixel is of both width and height 1.56mm. The projection data was calculated at 128 evenly spaced angles, each with 128 parallel projections. Again, the data samples were obtained as Poisson random variables with the appropriate mean. Fig. 1.5(b) shows the convolution backprojection which is blurred and contains considerable amounts of noise. The clustering routine was used to obtain initial values for the emission rates. Due to the high noise in the CBP reconstruction the number of clusters was not estimated but manually set to seven. As shown in Table 1.4, the clustering only identified 2 out of the 7 classes within reasonable tolerance. The thresholded CBP shown in Fig. 1.5(c) shows that the two gray ovals on the right are erroneously classified to have the same emission rate. Similarly, the bright patch on the left is set to the same emission rate as the center spot. In addition, the small patch within the gray oval on the lower right is misclassified. Fig. 1.5(d) shows the reconstruction using the fixed resolution algorithm. As before, the emission rates were initialized to the cluster means and the reconstruction was initialized to the thresholded CBP. As indicated in Table 1.4, the fixed resolution algorithm terminates with class estimates close to the initial cluster values. While the background level of 0.001 is now correctly estimated, the emission rate for the two gray ovals with original rates 1.2 and 1.6 has moved between the two values to $\theta_5 = 1.4458$. Notice that the overestimation of the lower right oval yields a mismatch between observed projection counts y and the forward projected reconstruction Px. Since the pixel values x within the oval are too large, we obtain $P_{i*}x > y_i$ for many projections intersecting this region. To compensate for this, the reconstruction algorithm inserts a pattern of black spots into the region which lowers the average projection count. Analogously, the gray oval on the upper right contains white spots to compensate underestimation of the emission rate 1.6 by 1.4458. Effects such as these result from the algorithm's being trapped in local minima of the MAP cost function. This is particularly critical for fixed resolution reconstructions that include the estimation of emission rates.

The multiscale algorithm is less prone to being trapped in local minima. For this phantom, 5 out of 7 emission rates are estimated correctly. The reconstruction as shown in Fig. 1.5(e) contains little noise and classifies the two gray ovals correctly. However, the left and center bright regions with emission rates 3.6 and 3.2 are still both classified as having a single emission rate of 3.5783. This can be improved by initializing the emission rates closer to their true values. In general, by varying the initial estimate for the emission rates, it is often possible for the fixed scale algorithm to obtain reconstructions comparable to the multiscale reconstruction. However, the fixed resolution algorithm is less robust with respect to the emission rate estimation and close initialization does not guarantee a comparable reconstruction. In almost all cases, with and without estimation of emission rates, the multiresolution algorithm is faster than the fixed resolution method. As shown in Table 1.3 (Phantom II), the multiscale algorithm terminates after 349 seconds compared to 404 seconds for the fixed resolution method. Fig. 1.6 compares the error convergence for the fixed and multiscale reconstructions of the second phantom. The multiscale error at coarse scales is

	Multi-Resolution	Fixed-Resolution
Phantom I	6.48% (1.75/27s)	1.64% (0.97/59s)
Phantom II	8.70% (30.37/349s)	2.77% (11.21/404s)

TABLE 1.5. Percentage of CPU-time used for estimation of emission rates. In all cases, the complexity for estimating θ is smaller than 10% of the total execution time. Shown in brackets are the CPU-times for estimating θ over the total execution times in seconds.

comparably high since the coarse scale reconstructions cannot account for the phantom's fine structure. However, the multiscale algorithm converges rapidly at each scale, resulting in lower total execution time. Again, the multiresolution algorithm achieves lower final reconstruction error.

Finally, we examine the computational complexity of the estimation of emission rates in relation to the overall complexity of the algorithms. Table 1.5 shows the percentage of CPU-time used for the estimation of θ. For all reconstructions, six full updates of θ were performed between full reconstruction updates. In all cases, the emission rate updates make up less than 10% of the total CPU-time. The percentages are smaller for the fixed resolution algorithm than for the multiresolution algorithm. Since the fixed resolution algorithm performs more computationally expensive iterations at fine scale, the relative cost for reconstruction updates is higher than for the multiscale method. Added over all resolutions, however, the multiscale algorithm performs more iterations than the fixed scale method. Since an iteration at any scale includes a fixed-cost update of θ, the multiscale method spends more absolute time on estimating the emission rates. This may be reduced by introducing a convergence criteria for the emission rate updates instead of running a fixed number of iterations between reconstruction updates.

In conclusion, the results indicate that the multiresolution algorithm can achieve reconstruction results superior to the fixed scale method. It is particularly robust with respect to the initialization of emission rates. Furthermore, the multiscale method is typically faster than the fixed resolution algorithm. For both algorithms, the computational cost for estimating the emission rates is small in comparison to the reconstruction complexity.

1.10 Conclusion

In this work, we have described a fixed- and multiscale method for discrete-valued Bayesian reconstruction. The multiscale MRF reconstruction algorithm is a straightforward extension of the fixed scale model. Interaction between resolution levels is obtained by initialization of each reconstruction with the previous coarser 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.

Acknowledgments

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

References

[1] G. Herman, *Image Reconstruction from Projections: The Fundamentals of Computerized Tomography* (Academic Press, New York), 1980.

[2] G. C. McKinnon and R. H. T. Bates, "Towards imaging the beating heart usefully with a conventional CT scanner," *IEEE Trans. on Biomedical Engineering* **BME-28**, 123–127 (1981).

[3] J. G. Sanderson, "Reconstruction of fuel pin bundles by a maximum entropy method," *IEEE Trans. on Nuclear Science* **NS-26**, 2685–2688 (1979).

[4] T. Inouye, "Image reconstruction with limited angle projection data," *IEEE Trans. on Nuclear Science* **NS-26**, 2666–2684 (1979).

[5] G. H. Glover and N. J. Pelc, "An algorithm for the reduction of metal clip artifacts in CT reconstruction," *Med. Phys.* **8**, 799–807 (1981).

[6] S. K. Chang, "The reconstruction of binary patterns from their projections," *Communications of the ACM* **14**, 21–25 (1971).

[7] A. Shliferstein and Y. T. Chien, "Switching components and the ambiguity problem in the reconstruction of pictures from their projections," *Pattern Recognition* **10**, 327–340 (1978).

[8] A. Kuba, "The recognition of two-directionally connected binary patterns from their two orthogonal projections," *Comput. Vision Graphics and Image Process.* **27**, 249–265 (1984).

[9] S. K. Chang and C. K. Chow, "The reconstruction of three-dimensional objects from two orthogonal projections and its application to cardiac cineangiography," *IEEE Trans. on Computers* **C-22**, 18–28 (1973).

[10] M. Soumekh, "Binary image reconstruction from four projections," *Proc. of IEEE Int'l Conf. on Acoust., Speech and Sig. Proc.*, (IEEE, New York City NY), pp. 1280–1283, 1988.

[11] D. J. Rossi and A. S. Willsky, "Reconstruction from projections based on detection and estimation of objects – parts I and II: Performance analysis and robustness analysis," *IEEE Trans. on Acoustics, Speech, and Signal Processing* **ASSP-32**, 886–906 (1984).

[12] Y. Bresler, J. A. Fessler, and A. Macovski, "A Bayesian approach to reconstruction from incomplete projections of a multiple object 3D domain," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **11**, 840–858 (1989).

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

[14] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors," *IEEE Trans. on Medical Imaging* **8**, 194–202 (1989).

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

[16] K. Sauer and C. A. Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. on Signal Processing* **41**, 534–548 (1993).

[17] B. Hunt, "Bayesian methods in nonlinear digital image restoration," *IEEE Trans. on Computers* **C-26**, 219–229 (1977).

[18] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **PAMI-6**, 721–741 (1984).

[19] S. Geman and D. McClure, "Statistical methods for tomographic image reconstruction," *Bull. Int. Stat. Inst.* **LII-4**, 5–21 (1987).

[20] C. A. Bouman and K. Sauer, "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. on Image Processing* **2**, 296–310 (1993).

[21] H. Derin, H. Elliot, R. Cristi, and D. Geman, "Bayes smoothing algorithms for segmentation of binary images modeled by Markov random fields," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **PAMI-6**, 707–719 (1984).

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

[23] L. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. on Medical Imaging* **MI-1**, 113–122 (1982).

[24] H. Hart and Z. Liang, "Bayesian image processing in two dimensions," *IEEE Trans. on Medical Imaging* **MI-6**, 201–208 (1987).

[25] Z. Liang and H. Hart, "Bayesian image processing of data from constrained source distributions–I: Non-valued, uncorrelated and correlated constraints," *Bull. Math. Biol.* **49**, 51–74 (1987).

[26] G. T. Herman and D. Odhner, "Performance evaluation of an iterative image reconstruction algorithm for positron emission tomography," *IEEE Trans. on Medical Imaging* **10**, 336–346 (1991).

[27] G. T. Herman, A. R. De Pierro, and N. Gai, "On methods for maximum a posteriori image reconstruction with normal prior," *J. Visual Comm. Image Rep.* **3**, 316–324 (1992).

[28] T. Hebert and S. Gopal, "The GEM MAP algorithm with 3-D SPECT system response," *IEEE Trans. on Medical Imaging* **11**, 81–90 (1992).

[29] A. De Pierro, "A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography," *IEEE Trans. on Medical Imaging* **14**, 132–137 (1995).

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

[31] 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**, 41–63 (1994).

[32] F. Heitz, P. Perez, and P. Bouthemy, "Multiscale minimization of global energy functions in some visual recovery problems," *Comput. Vision Graphics and Image Process.* **59**, 125–134 (1994).

[33] C. A. Bouman and B. Liu, "Multiple resolution segmentation of textured images," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **13**, 99–113 (1991).

[34] C. A. Bouman and M. Shapiro, "A multiscale random field model for Bayesian image segmentation," *IEEE Trans. on Image Processing* **3**, 162–177 (1994).

[35] S. S. Saquib, C. A. Bouman, and K. Sauer, "A non-homogeneous MRF model for multiresolution Bayesian estimation," *Proc. of IEEE Int'l Conf. on Image Proc.*, (IEEE, Lausanne Switzerland), pp. 445–448, 1996.

[36] D. Snyder and M. Miller, "The use of sieves to stabilize images produced with the EM algorithm for emission tomography," *IEEE Trans. on Nuclear Science* **NS-32**, 3864–3871 (1985).

[37] E. Veklerov and J. Llacer, "Stopping rule for the MLE algorithm based on statistical hypothesis testing," *IEEE Trans. on Medical Imaging* **MI-6**, 313–319 (1987).

[38] T. Hebert, R. Leahy, and M. Singh, "Fast MLE for SPECT using an intermediate polar representation and a stopping criterion," *IEEE Trans. on Nuclear Science* **35**, 615–619 (1988).

[39] A. Tikhonov and V. Arsenin, *Solutions of Ill-Posed Problems* (Winston and Sons, New York), 1977.

[40] J. Besag, "Spatial interaction and the statistical analysis of lattice systems," *Journal of the Royal Statistical Society B* **36**, 192–236 (1974).

[41] R. Kindermann and J. Snell, *Markov Random Fields and their Applications* (American Mathematical Society, Providence), 1980.

[42] R. Kashyap and R. Chellappa, "Estimation and choice of neighbors in spatial-interaction models of images," *IEEE Trans. on Information Theory* **IT-29**, 60–72 (1983).

[43] D. Pickard, "Inference for discrete Markov fields: The simplest nontrivial case," *Journal of the American Statistical Association* **82**, 90–96 (1987).

[44] R. Dubes and A. Jain, "Random field models in image analysis," *Journal of Applied Statistics* **16**, 131–164 (1989).

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

[46] C. A. Bouman, "Cluster: an unsupervised algorithm for modeling Gaussian mixtures." Available from http://www.ece.purdue.edu/~bouman, 1997.

Index