

GRAPH CUT: APPLICATION TO BAYESIAN EMISSION TOMOGRAPHY RECONSTRUCTION

Martin Bonneville[†], Jean Meunier[†], Sébastien Roy^{*†}

[†] DIRO, Département d'Informatique et de Recherche Opérationnelle, P.O. 6128, Montréal (Quebec), H3C 3J7.

^{*} NEC Research Institute, Princeton, New Jersey

e-mail: <bonnevil, meunier, roys>@iro.umontreal.ca

ABSTRACT

We present an application of graph cuts to Bayesian emission tomography (ET) reconstruction. The method is built on the expectation-maximization (EM) maximum a posteriori (MAP) reconstruction. In general, MAP estimates are hard to assess. For instance, methods such as simulated annealing cannot be employed, because of the computational complexity of bayesian ET reconstruction. We propose to perform a part of the M-step by a maximum-flow computation in a particular flow graph. Because the possible priors (in a maximum-flow approach) are limited to linear function, we have incorporated the estimation of a line process that will preserve discontinuities in the reconstructions. It is the iterative nature of EM that allows the introduction of the line process. The method is illustrated first over synthetic data and then over the Hoffman brain.

1. INTRODUCTION

Many image restoration algorithms are based on statistical models. The key to successfully use these algorithms resides in constructing a suitable model of the degradation (likelihood). However, the model must be regularized to alleviate overfitting and process outliers. One natural way to achieve this is to incorporate prior information about the desired restoration. This can be easily done in a Bayesian framework. For instance, Geman and Geman [3] proposed to incorporate prior information about the correlation of neighboring pixels in the image, using Gibbs distribution. A labelling of the image is then obtained by computing the maximum a posteriori (MAP) estimate.

Unfortunately, the MAP estimate of an image cannot usually be computed efficiently. Many methods have been proposed to find the MAP estimate or an approximation of it. Simulated annealing [3] can, in theory, compute the MAP estimate, but it is not always computationally practical. Iterated conditional

modes (ICM) developed by Besag [1] and gradient descent techniques (in the case of continuous labelling) are methods that do not guarantee convergence to a global optimum.

Recent developments were done in graph formulation of MAP estimate and related energy minimization problems [2, 7, 8, 10]. All these algorithms are based on the computation of maximum-flow and the related minimum-cut in a specific type of flow graph. They can be split into two categories. In the first, the graph formulation provides an efficient method for computing the exact global optimum [7, 8, 10]. However, these algorithms are limited to linear prior (linear clique potential) which can oversmooth boundaries, with the exception of [7] which is limited to binary images. The second category consists of fast approximation via approximated multi-way cut formulation [2]. One major limitation of this formulation is the restriction of the prior to Pott's model (unordered labelling).

In this paper, we propose an original application of a maximum-flow method to compute the maximization step in the Expectation-Maximization Bayesian emission tomography (ET) reconstruction. The approach is based on the computation of the maximum-flow and related minimum-cut in a graph similar to [2, 8, 10]. Because the maximum-flow approach is limited to the use of linear prior, we show how to incorporate a line process, that preserves discontinuities, in the iterative EM scheme. The line process is closely related to the introduction of contextual information that are obtained from a previous iteration of the EM algorithm.

After a general description of the Bayesian framework, We will present the description of our optimization method based on graph cut. We then provide an application to Bayesian ET image reconstruction. In this application, methods such as simulated annealing are precluded because of the size of the calculation. We provide a statistical model based on the seminal work of Shepp and Vardi [11] on maximum likelihood reconstruction for emission tomography and provide

Bayesian formulation similar to Green [6] and Geman and McClure [4], with a different prior.

2. BAYESIAN FRAMEWORK

For the purpose of the analysis, the image space is represented by a regular array of pixels indexed by i , where $i \in \mathcal{S} = \{0, \dots, N-1\}$. We also need to define a neighborhood system $\mathcal{N} = \{\mathcal{N}_i \mid i \in \mathcal{S}\}$. Typically, \mathcal{N} is the 8-neighbors system. The restoration is formulated as a labelling problem. The restored image corresponds to a configuration $f = \{f_i \mid i \in \mathcal{S}\}$ taking discrete values in the set of labels $\mathcal{L} = \{0, \dots, M-1\}$. To enforce piecewise smoothness, we introduce an unobservable line process l , taking discrete values in $\{0, 1\}^M$, into the image model to preserve discontinuities in the restoration [3]. For instance $l_{ii'} = 1$ if there is a discontinuity between i and i' . Let's formulate our restoration problem from the Bayes theorem:

$$Pr(f, l \mid x) \propto Pr(x \mid f, l) Pr(f, l),$$

where f is a configuration, l the line process and x is the observation. We are interested to estimate the configuration \hat{f} that maximizes the posterior probability $Pr(f, l \mid x)$,

$$\begin{aligned} \hat{f} &= \arg \max_f Pr(f, l \mid x) \\ &= \arg \max_f Pr(x \mid f) Pr(f, l). \end{aligned}$$

This last equation comes from the fact that x is conditionally (given f) independent of l . For the purpose of optimization, we reformulate everything into an energy minimization problem; the likelihood and prior energy are

$$\begin{aligned} E_{likelihood}(x \mid f) &= -\log Pr(x \mid f), \\ E_{prior}(f, l) &= -\log Pr(f, l). \end{aligned} \quad (1)$$

Thus, the MAP estimate is obtained by minimizing the following posterior energy,

$$E(f, l \mid x) = E_{likelihood}(x \mid f) + E_{prior}(f, l). \quad (2)$$

2.1. Selecting a prior

The prior is constructed as a Gibbs distribution and is specified by an energy function,

$$Pr(f, l) = \frac{e^{-E_{prior}(f, l)}}{Z},$$

where Z is a normalizing constant, called the partition function. The energy function $E_{prior}(f, l)$ is the prior

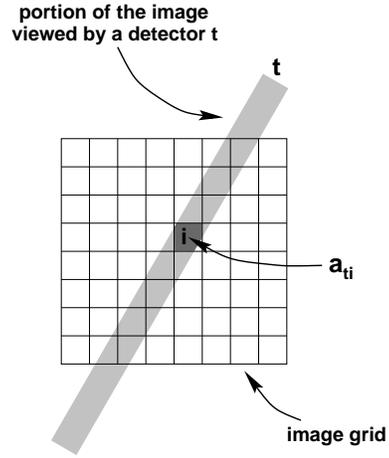


Figure 1: Geometrical model scheme employed for ET. The grid corresponds to the image space \mathcal{S} . The grey rectangle corresponds to the portion of the image viewed by some bin $t \in \mathcal{B}$. The portion of pixel i viewed by bin t is a_{ti} and is outlined in dark grey.

energy and it is designed so that the expected configurations are those for which typical neighboring pixels have similar labels. Moreover, we choose a prior that favors piecewise constant reconstructions and that preserves discontinuities. All these are local constraints that can be conveniently modeled by a local composed function of f and l ,

$$\begin{aligned} E_{prior}(f, l) &= \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{N}_i} \beta \phi(f_i - f_{i'}) (1 - l_{ii'}) \\ &\quad + \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{N}_i} \alpha l_{ii'}, \end{aligned} \quad (3)$$

where α and β are positive constants and the function $\phi(u)$ is nonnegative, even, monotonically increasing and minimized at $u = 0$. The choice of ϕ is made in order to achieved the desired properties of the configuration. It must reflect some qualitative features about the desired restoration. One feature is smoothness within homogeneous regions of the image. Some obvious choices are $\phi(u) = |u|$ and $\phi(u) = u^2$. The choice of finding the exact MAP estimate by a maximum-flow approach limits our choice to $\phi(u) = |u|$. At discontinuities (natural boundaries) the energy potential of pixel i , $\sum_{i' \in \mathcal{N}_i} \beta \phi(f_i - f_{i'})$ tends to be high and makes boundaries oversmoothed. To correct this problem, we have incorporated a discrete line process l . This way, the chosen prior will favor smoothness except when a discontinuity occurs.

2.2. Statistical model for ET

In ET, the image reconstruction task consists in recovering emitter densities from a sinogram (projection data¹). To accomplish this, we must construct a model of the image space \mathcal{S} and of the degradation that occurs in the observation (projection data). The projection data are indexed by $t \in \mathcal{B} = \{0, \dots, T\}$, where $T = \text{number of angles} \times \text{number of bins}$. In figure 1, we see a schematic representation of our geometrical model for ET. The degradation model assumes the emission from pixel i to be completely random. Therefore, the number of photons emitted from i and detected in bin t forms independent Poisson processes in i and t , because each photon is detected by at most one bin and the emissions are independent. Because the projection data are the superposition of independent Poisson processes, it follows that x_t , a projection data, is independent Poisson distributed random variables,

$$x_t \sim \text{Poisson} \left(\sum_{i \in \mathcal{S}} a_{ti} f_i \right). \quad (4)$$

As a result, the likelihood, which is the probability of the observation (projection data) knowing the emitter density, is defined as,

$$Pr(x | f) = \prod_{t \in \mathcal{B}} \frac{(\sum_{i \in \mathcal{S}} a_{ti} f_i)^{x_t} \exp(-\sum_{i \in \mathcal{S}} a_{ti} f_i)}{x_t!}.$$

In this model, the coefficients a_{ti} represent the probability that each emission from pixel i is detected in detector t (see figure 1). They are assumed known and they model the geometry of the detection system. Other major physical factors in ET, such as attenuation and scatter, can also be included in a_{ti} .

Instead of considering the model of equation (4), it is easier for computational reasons, to treat it as an incomplete data problem. In fact, it is more direct to estimate f if the unobserved data z_{ti} (the number of photons emitted from pixel i and recorded in bin t) is known. Since an estimate of f allows the distribution of the missing data z to be easily specified, this suggests the use of an EM iterative scheme. Iteratively, we perform successive estimations of z (E-step) and MAP estimates of f according to the posterior distribution $Pr(f | z)$ (M-step). Notice that the Poisson model still holds for the missing data z , so we have

$$z_{ti} \sim \text{Poisson}(a_{ti} f_i).$$

¹We prefer the term projection data to sinogram, because the action of detecting in bin t an emitted photon from pixel i is similar to a projection.

Procedure MAP EM Algorithm(x : projection data)
 $\hat{f} :=$ some initial configuration
Repeat
 $\left[\begin{array}{l} \text{update the complete data } z \\ \text{according to equation (7)} \end{array} \right]$ (E-step)
 $\left[\begin{array}{l} \hat{l} = \arg \min_l E(\hat{f}, l | z) \\ \hat{f} = \arg \min_f E(f, \hat{l} | z) \end{array} \right]$ (M-step)
Until convergence

Figure 2: Pseudo-code MAP EM algorithm. Each repeat loop is considered to be an iteration.

and the likelihood can be restated as

$$Pr(z | f) = \prod_{i \in \mathcal{S}, t \in \mathcal{B}} \frac{(a_{ti} f_i)^{z_t} \exp(-a_{ti} f_i)}{z_t!}. \quad (5)$$

The corresponding likelihood energy is defined by

$$E_{\text{likelihood}}(z | f) = \sum_i \sum_t (a_{ti} f_i - z_{ti} \log(a_{ti} f_i)) + \text{constant}. \quad (6)$$

As shown in [6, 9], we can estimate (E-step) the missing data z_{ti} by setting it to its conditional expectation given x and \hat{f} :

$$\begin{aligned} \hat{z}_{ti} &= E(z_{ti} | x, \hat{f}) \\ &= \frac{x_t a_{ti} \hat{f}_i}{\sum_{i' \in \mathcal{S}} a_{ti'} \hat{f}_{i'}}, \end{aligned} \quad (7)$$

where \hat{f} stands for the previous iteration estimate of f . Finally, the M-step consists in minimizing the posterior energy for f and l ,

$$\begin{aligned} E(f, l | \hat{z}) &= E_{\text{likelihood}}(\hat{z} | f) + E_{\text{prior}}(f, l) \\ &= \sum_{i \in \mathcal{S}} \sum_{t \in \mathcal{B}} a_{ti} f_i - \hat{z}_{ti} \log(a_{ti} f_i) + \\ &\quad \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{N}_i} \beta |f_i - f_{i'}| (1 - l_{ii'}) + \\ &\quad \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{N}_i} \alpha l_{ii'} + \text{constant}. \end{aligned} \quad (8)$$

3. MINIMIZING THE POSTERIOR ENERGY

We are interested to choose (\hat{f}, \hat{l}) that minimizes the posterior energy (8), since MAP estimates are difficult to compute in general, we propose a deterministic algorithm which is similar to ICM, but differs in the way

it updates the components f_i of a configuration f all at the same time. The algorithm (shown in figure 2) starts with an initial labelling $\hat{f}^{(0)}$ (typically a maximum likelihood estimation is used). Iteratively (like in ICM), we first update the line process while keeping f frozen to $\hat{f}^{(n-1)}$ (in equation (8)),

$$\hat{l}^{(n)} = \arg \min_l E(\hat{f}^{(n-1)}, l | \hat{z}).$$

Notice, that the energy is minimized by setting

$$\hat{l}_{ii'}^{(n)} = \begin{cases} 0, & \sum_{i' \in \mathcal{N}_i} |\hat{f}_i^{(n-1)} - \hat{f}_{i'}^{(n-1)}| \leq \alpha/\beta \\ 1, & \text{otherwise.} \end{cases}$$

It remains to minimize over f . We update the line process to $\hat{l}^{(n)}$ in equation (8) and then we determine $\hat{f}^{(n)}$ that minimizes

$$\begin{aligned} E(f, \hat{l}^{(n)} | \hat{z}) &= \sum_{i \in \mathcal{S}} \sum_{t \in \mathcal{B}} a_{ti} f_i - \hat{z}_{ti} \log(a_{ti} f_i) \\ &+ \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{N}_i} \Theta_{ii'} |f_i - f_{i'}| \\ &+ \text{constant,} \end{aligned} \quad (9)$$

where

$$\Theta_{ii'} = \begin{cases} 0, & \hat{l}_{ii'}^{(n)} = 1 \\ \beta, & \text{otherwise,} \end{cases} \quad (10)$$

is introduced to better fit our graph formulation of the next section. This minimization can be solve globally for f by computing the minimum-cut in a particular flow graph [2, 8, 10]. In the next section we describe the flow graph construction and explain how the minimum cut gives the desired configuration.

3.1. Flow graph

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a flow graph (see figure 3), where \mathcal{V} is the set of vertices and \mathcal{E} is the set of weighted edges. The set \mathcal{V} is composed of two types of vertices: a set of configuration nodes and two distinct terminals called the source s and the sink t ,

$$\mathcal{V} = \mathcal{V}' \cup \{s, t\},$$

where $\mathcal{V}' = \{(i, k) : i \in \mathcal{S}, k \in \mathcal{L}\}$ is the set of nodes corresponding to all possible assignments. The structure of the nodes and connections in the flow graph \mathcal{G} are illustrated in figure 3 for a 2-neighbor system \mathcal{N} . Two types of edges compose \mathcal{E} , likelihood edges $\mathcal{E}_{likelihood}$ and prior edges \mathcal{E}_{prior} . The set of likelihood edges generates paths from the source to the sink,

$$s \rightarrow (i, 0) \rightarrow \dots \rightarrow (i, k) \rightarrow \dots \rightarrow (i, M) \rightarrow t,$$

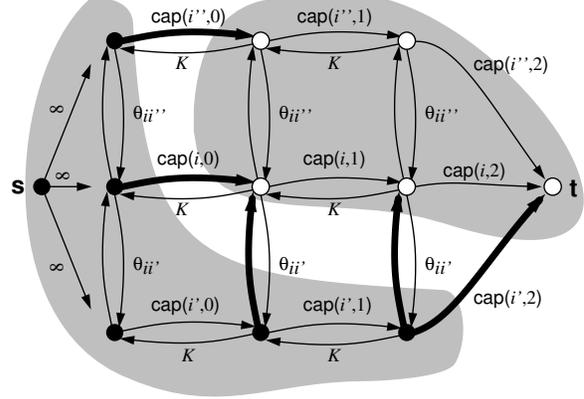


Figure 3: Flow graph representation of a 1D image with 3 pixels $\{i'', i, i'\}$ (from top to bottom) and 3 labels $\{0, 1, 2\}$. The minimum cut, composed of the thick edges, separates the graph into two parts (in grey), one linked to the source s (black nodes) and the other to the sink t (white nodes). The outgoing edges from the source have infinite capacities. A likelihood edge assigning label k to pixel i has capacity $cap(i, k)$ while its corresponding inverse edge has large constant capacity K . The prior edges have capacity $\Theta_{ii'}$. Here, label 0 is assigned to pixels i and i'' , and label 2 is assigned to pixel i' , yielding the configuration of emitter densities $f = \{0, 0, 2\}$. The discontinuity between pixels i and i' is revealed by the two prior edges in the cut, inducing a penalty cost $\Theta_{ii'}|0 - 2|$.

for all $i \in \mathcal{S}$. Outgoing edges from the source s have infinite capacity to insure they can never be saturated. All other likelihood edges $(i, k) - (i, k + 1)$ are assigned the capacity $cap(i, k)$ such that $cap(i, k) = \sum_{t \in \mathcal{B}} a_{ti}k - \hat{z}_{ti} \log(a_{ti}k)$ while the corresponding inverse edges $(i, k + 1) - (i, k)$ are given the capacity K , where K is chosen to be a large finite constant². This large value is required to insure that such edges are never saturated, therefore avoiding the possibility of multiple solutions to the labelling problem [2, 8]. The other type of edges present in the graph, prior edges \mathcal{E}_{prior} , is defined by the neighborhood system employed (\mathcal{N}) and generates connections between assignments (i, k) and (i', k) if $i' \in \mathcal{N}_i$ for all $i \in \mathcal{S}$. The capacity of prior edges is set to $\Theta_{ii'}$ of equation (10).

A cut $\mathcal{C} \subset \mathcal{E}$ is a set of edges such that the source s and the sink t are separated in the induced graph $\mathcal{G}_{\mathcal{C}} = (\mathcal{V}, \mathcal{E} - \mathcal{C})$. In short, the cut \mathcal{C} contains at least one edge of every path from the source s to the sink t . The capacity (cost) of a cut is simply the sum of the edge capacities in \mathcal{C} . We denote the cost of a cut as $|\mathcal{C}|$. The key part of our algorithm is based on the computation of the minimum cost cut \mathcal{C} . Its computation is achieved efficiently by computing the maximum flow between the source and the sink. We have chosen Golberg's preflow push relabel algorithm [5], which in our case features an almost linear average complexity.

As shown in [2, 8] for similar, but different flow graph formulations, the cost of the minimum cut $|\mathcal{C}|$ corresponds to the global optimum of equation (9). Because the cost of a cut is given by the summation of the edges it contains, the optimal configuration is provided by the likelihood edges contained in the minimum cut \mathcal{C} . In brief, cutting a likelihood edge $(i, k) - (i, k + 1)$ corresponds to assign label k to pixel i .

4. EXPERIMENTS AND RESULTS

In this section, we compare the performances of our method to maximum likelihood EM reconstruction and also to the standard ICM Bayesian ET reconstruction. Experiments are done on 2D synthetic data (figure 4) and over the Hoffman brain phantom (figure 5). The set of labels $\mathcal{L} = \{0, \dots, 255\}$ represents normalized photon counts.

4.1. Elliptic phantom

The image space \mathcal{S} is a regular grid of 64×64 pixels (figure 4). We performed 64 projection angles over 180° and used 64 detector bins for each projection. The bin

²For $k = M - 1$, the likelihood edge is $(i, k) - t$ and there is no inverse edge (see figure 3).

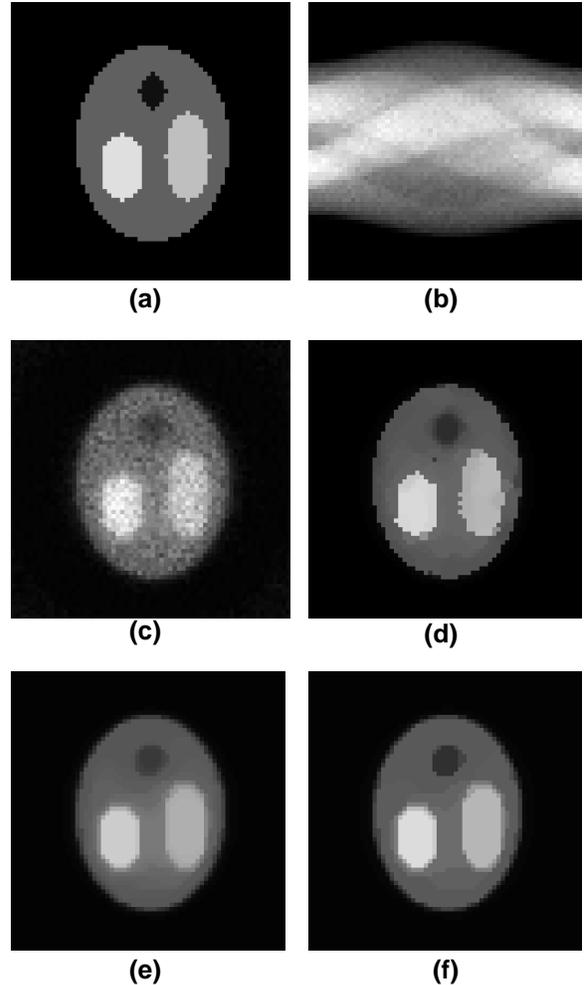


Figure 4: (a) Ground truth image. (b) A noisy projection data obtained from the noiseless projections, where the horizontal and vertical axes correspond to the angles and bins respectively. (c) Maximum likelihood EM reconstruction obtained after 20 iterations of the algorithm. The normalized posterior energy is 2924.0 and the RMS error is 26.4. (d) ICM MAP EM reconstruction with $\alpha = 10.0$ and $\beta = 2.0$. The normalized posterior energy is 66.4 and the RMS error is 12.4. (e) Maximum-flow MAP EM reconstruction with $\alpha = 0.0$ and $\beta = 2.0$. The normalized posterior energy is 77.6 and the RMS error is 16.0. (f) Maximum-flow MAP EM reconstruction with $\alpha = 10.0$ and $\beta = 2.0$. The normalized posterior energy is 34.2 and the RMS error is 11.4.

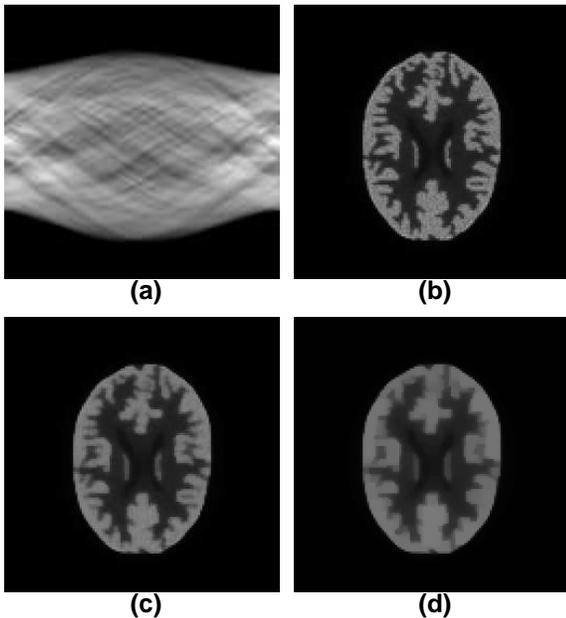


Figure 5: (a) Sinogram of the Hoffman brain with activity ratios of 4:1:0 for grey matter, white matter and CSF, respectively. (b) Maximum-flow MAP EM reconstruction with $\beta = 2.0$ and $\alpha = 10.0$. (c) Maximum-flow MAP EM reconstruction with $\beta = 2.5$ and $\alpha = 10.0$. (d) Maximum-flow MAP EM reconstruction with $\beta = 3.0$ and $\alpha = 10.0$.

and pixel width are the same. Finally, we simply add Poisson noise to each projection, using the noiseless projection values as the Poisson parameter. The phantom used and the noisy projection data are shown in figure 4 (a) and (b), respectively. In figure 4 (c), we see the noisy result obtained from the standard maximum likelihood EM reconstruction method. This image was generated by setting $\Theta_{ii'} = 0$ in equation (9) and then after running 20 iterations of the algorithm shown in figure 2 (starting with a flat labelling as initial value for \hat{f}). The RMS error between the original phantom and the reconstruction is 26.4. The image in figure 4 (d) corresponds to an ICM reconstruction with parameter $\beta = 2.0$ and $\alpha = 10.0$. The RMS error for the ICM reconstruction reduces to 12.4. Finally, in figure 4 (e) and (f), we present the reconstructions obtained with the maximum flow approach, one with parameter $\beta = 2.0$ and $\alpha = 0.0$ and the other for $\beta = 2.0$ and $\alpha = 10.0$. The RMS errors are 16.0 and 11.4, respectively. To compute the posterior energy of these reconstructions, we have set the complete data (z) of equation (9) to a flat image and the parameter were set

to $\alpha = 10.0$ and $\beta = 2.0$. The normalized ³posterior energy obtained are for the reconstruction in figure 4 (c) 2924.0, (d) 66.4, (e) 77.6 and (f) 34.2.

4.2. Hoffman brain

The Hoffman brain image space \mathcal{S} is a regular grid of 128×128 pixels (see figure 5). The ratios of activities in the Hoffman brain are 4:1:0 for grey matter, white matter and cerebrospinal fluid (CSF), respectively. 128 projection angles over 180° were performed, using 128 detector bins for each projection. The bin and pixel size are the same. Finally, Poisson noise was added to the sinogram (projection data) to simulate the degradation. In figure 5 we see the noisy sinogram (a) and three reconstructions obtained by the maximum-flow MAP EM algorithm: (b) with $\alpha = 10.0$ and $\beta = 2.0$, (c) with $\alpha = 10.0$ and $\beta = 2.5$ and finally (d) with $\alpha = 10.0$ and $\beta = 3.0$.

5. DISCUSSION AND CONCLUSION

We have presented a robust method based on graph cuts to compute Bayesian ET reconstruction using simple priors. The method showed great improvement over the maximum likelihood EM reconstruction by reducing the RMS error by 56.7%. It has also shown better results than ICM with reduction of 8.1% of the RMS error. With the maximum flow EM reconstruction method, flat regions within the images are almost perfectly recovered. However, if the line process is not used, as in figure 4 (e), the reconstruction does not behave as well as it looks (an increase in energy and RMS error). This problem happens when smoothing is done across adjacent regions (over boundaries) using a non-boundary preserving smoothing function such as $\phi(u) = |u|$. This problem was solved by introducing a line process in the iterative MAP-EM scheme. This process allows the introduction of contextual information and it could be improved by adding more information about boundaries, such as anatomical classification of brain tissue. The uses of a continuous line process estimated from a robust estimator, such as the truncated quadratic function could reduce some artefact created by the discreteness of the line process and the fact that they “freeze” while we optimize over f .

Even if the form of the prior is limited, our method is of great utility, because the complexity of ET precludes the use of techniques such as simulated annealing. Therefore, the energy function has to be minimized by some deterministic method. With other

³Translation and scaling with respect to the energy of the ground truth.

methods there are no guarantee to reach a global optimum, but just hope to escape a poor local optimum.

6. ACKNOWLEDGMENTS

We would like to thanks Neil Stewart for many interesting discussions. This work was supported by the FCAR (Québec).

7. REFERENCES

- [1] J. E. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society, Series B*, 48:259–302, 1986.
- [2] Y. Boykov, O. Veksle, and R. Zabih. Markov random fields with efficient approximations. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Santa Barbara, CA, 1998.
- [3] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- [4] S. Geman and D. E. McClure. Statistical methods for tomographic image reconstruction. *Bulletin of the International Statistical Institute*, LII-4:5–21, 1987.
- [5] A. V. Golberg and S. B. Rao. Length function for flow computation. Technical Report 97-055, NEC Research Institute, Princeton NJ, 1997.
- [6] P. J. Green. Bayesian reconstruction from emission tomography data using a modified em algorithm. *IEEE Transactions on Medical Imaging*, 9(1):84–93, 1990.
- [7] D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society, Series B*, 51(2):271–279, 1989.
- [8] H. Ishikawa and D. Geiger. Segmentation by grouping junctions. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Santa Barbara, CA, 1998.
- [9] V. E. Johnson, W. H. Wong, X. Hu, and C. T. Chen. Image restoration using gibbs priors: boundary modeling, treatment of blurring, and selection of hyperparameter. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(5):413–425, 1991.
- [10] S. Roy and I. Cox. A maximum-flow formulation of the n-camera stereo correspondence problem. In *6th International Conference on Computer Vision*, Bombay, India, 1998.
- [11] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122, 1982.