A MIXTURE-SITE MODEL FOR EDGE-PRESERVING IMAGE RESTORATION

Jeffrey A. Fessler

3480 Kresge III, Box 0552
University of Michigan
Ann Arbor, MI 48109-0552
fessler@umich.edu

ABSTRACT

This paper summarizes a new Bayesian method for edge-preserving image restoration from noisy measurements. The line-site method of Geman and Geman forces region boundaries to lie along pixel boundaries, which is unnatural, particularly for 3D data. Here, we augment the intensity process with a binary "mixture site" process, which has one parameter for each pixel indicating the presence of a boundary a some unknown location within that pixel. The method was motivated by PET and SPECT transmission images with partial volume effects, and is easily extended to 3D data sets.

1. INTRODUCTION

Natural objects often consist of several distinct regions, where the object intensity is a relatively smooth function within each region, but may have large gradients along the boundaries between regions. Given noisy measurements of such a *piecewise smooth* intensity function, one would like to obtain an estimate of the underlying object. Ideally the estimation method should be *edge-preserving*, i.e., one smooths within regions to reduce variance, while preserving edges by not smoothing between regions. This paper summarizes a new Bayesian method for edge-preserving image restoration from noisy measurements.

Our method was motivated by the medical imaging application of estimating attenuation maps from noisy positron emission tomography (PET) or single photon emission computed tomography (SPECT) transmission scans. At the photon energies of PET and SPECT imaging, the attenuation coefficient distributions in human subjects consist of distinct regions of air, lung, soft tissue, bone, etc. Within each tissue type the attenuation coefficient is fairly uniform (though not perfectly uniform). Therefore attenuation maps are an example of piecewise smooth functions. A limiting characteristic of PET and SPECT instruments (and virtually all imaging devices) is the *partial volume effect*: if the support of an image pixel overlays the boundary between two regions, then the measured intensity value for that pixel will consist of a mixture of partial contributions from both tissue types. The method proposed in this paper attempts to address these mixture characteristics in the PET and SPECT context, although it is generally applicable to images of piecewise smooth objects.

There are at least three categories of methods for edge-preserving image restoration. Perhaps the simplest approach is to make the stronger assumption that the object is *piecewise constant*, and then assign each image pixel to one of several classes (segmentation), using an algorithm such as ICM [1]. However, the piecewise-constant assumption is inappropriate for biological images that exhibit intensity variations within regions. A second approach is to use a penalized likelihood method with a *non-quadratic penalty*. The "broken parabola" function [2], the log cosh function [3], and the generalized Gaussian random field [4] are examples of non-quadratic penalty functions that encourage smoothness while allowing discontinuities. Most such penalty functions are non-convex, and as shown by Bouman and Sauer [4], non-convex objectives can lead to object estimates that are discontinuous functions of the measurements. Another disadvantage of non-quadratic penalty methods is that they do not explicitly enforce *continuity* of object boundaries, since they essentially act only pairwise on the object's pixels.

A third approach, more closely related to our proposed method, is to augment the intensity function (or *intensity process*) with an ancillary random process that controls the formation of object boundaries. The canonical example of such an approach is the *line-site* method of Geman and Geman [5], in which one augments the intensity process with a binary *line-site*

This work was supported in part by NIH Grant CA60711 and DOE Grant DE-FG02-87ER65061.

process. Each pair of pixels is connected by a line-site that takes values 1 or 0 to indicate the presence or absence respectively of an edge between those two pixels. One jointly estimates the line-sites and the intensity parameters from the noisy data.

The line-site method has the following drawbacks.

- Region boundaries are effectively constrained to lie along pixel edges. This constraint leads to unnaturally "blocky" restorations, and is not naturally suited to multi-resolution processing. Although in principle one could minimize this problem by using very small pixels, (i.e., smaller pixels than the discrete measurements) this increases the computation, and increases the likelihood of local extrema in the objective function when estimating more parameters than there are measurements.

- For 3D volume data, the extension of the line-site method requires 3 additional "edge-site" parameters for every intensity parameter. The computational burden and problems of over parameterization scale accordingly. PET and SPECT data is virtually always acquired over several contiguous slices, so a 3D approach is very natural, since the attenuation map is piecewise smooth both axially and transaxially. In fact, human attenuation maps are probably smoother axially than transaxially since many important structures (spine, skull, lungs, etc.) have a vertical structure.

Rather than augmenting the intensity process with a line-site process, we propose a new Bayesian model based on *mixture sites*, as illustrated by Figure 2. We augment each pixel or voxel with one additional binary mixture parameter that takes values 1 or 0 indicating the presence or absence respectively of a "mixture" of two or more regions within that pixel. The advantages of this approach are the following:

- The mixture-site model does not constrain the region boundaries to lie along pixel edges, thereby allowing the smooth region boundaries of natural objects. Our prior assumes that smooth boundaries of natural objects almost surely will not lie exactly between any pair of pixels. This assumption is completely opposite to that of the line-site method.

- The mixture-site model adds only one extra parameter per pixel, even for 3D problems.

- The mixture-site model is naturally suited to hierarchical or multi-resolution data processing. If a fine-grid pixel is a mixture of two or more region types, then any coarse-grid pixel that overlaps that fine-grid pixel must also be a mixture.

The mixture-site method does share with the line-site method the problems of non-convexity associated with models that have both discrete valued and continuous valued parameters.

Whether the philosophical advantages of the mixture-site method translate into more accurate restored images than the line-site method will not be addressed in this short paper. Here we simply introduce the model and illustrate its behavior with anecdotal results.

2. MODEL AND METHOD

2.1. Measurement Model

Let $\mu(\underline{x})$ denote the underlying (continuous) object intensity function, where $\underline{x}$ denotes either 2D or 3D spatial coordinates. For this paper we assume the following linear additive noise model for the (discrete) measurements:

$$y_n = \int h_n(\underline{x})\mu(\underline{x})\ d\underline{x} + \varepsilon_n,\ n = 1, \ldots, N, \qquad (1)$$

where $\varepsilon_n \sim \mathcal{N}(0, \sigma^2)$, and $h_n(\underline{x})$ denotes the point-spread function (PSF) for the nth detector element. In the most ideal case, the PSF would be a Dirac delta function located at the center $\underline{x}_n$ of the nth pixel, in which case we would have the oft-assumed model: $y_n = \mu(\underline{x}_n) + \varepsilon_n$. This is never achieved in real bandlimited imaging systems; at best the PSF $h_n(\underline{x})$ will be an indicator function over the (typically square) support of the nth pixel. Any such imaging system exhibits partial volume effects, i.e. if the nth measurement pixel overlaps a boundary between two regions, then the mean of y_n will be a mixture of more than one object intensity. Without multi-channel information [6], it would be very difficult to identify the relative contributions of each region in such mixture pixels. Nevertheless, we would like to avoid arbitrarily assigning the intensity of such pixels to that of one or the other of the two regions, which is essentially what the line-site method does.

To estimate μ, we must adopt a discrete parameterization. We use the conventional voxel basis:

$$\mu(\underline{x}) \approx \sum_{j=1}^{p} \mu_j I_j(\underline{x}), \qquad (2)$$

where $I_j(\underline{x})$ is the indicator function over the jth voxel. Vectorizing (1) and (2) yields:

$$\boldsymbol{y} = \mathbf{H}\boldsymbol{\mu} + \boldsymbol{\varepsilon},$$

where the (n, j)th element of $\mathbf{H}$ is $\int h_n(\underline{x}) I_j(\underline{x})\, d\underline{x}$. Our goal is to estimate the unknown object intensities $\boldsymbol{\mu} = [\mu_1, \ldots, \mu_p]'$ from the noisy measurements $\boldsymbol{y} = [y_1, \ldots, y_n]'$.

2.2. Mixture-site Process

We augment the intensity parameter μ_j for the jth pixel with a binary mixture-site parameter denoted m_j. We would like to have $m_j = 1$ if the jth pixel overlays a region boundary, and to have $m_j = 0$ otherwise. Of course the m_j's are unknown so we must jointly estimate the mixture-site process $\boldsymbol{m} = [m_1, \ldots, m_p]'$ and the intensity process $\boldsymbol{\mu}$. We model the prior distribution of $\boldsymbol{m}$ using a Gibb's distribution, with the six (rotationally invariant) clique types shown in Figure 1. Our log-prior for the mixture-site process $\boldsymbol{m}$ is

$$\Omega_2(\boldsymbol{m}) = -\log p(\boldsymbol{m}) = \sum_{c=0}^{5} \beta_c \#_c(\boldsymbol{m}),$$

where $\#_c(\boldsymbol{m})$ is the number of cliques in $\boldsymbol{m}$ of the cth type. Since the cliques are 2×2, the neighborhood structure is 3×3. We believe these are the simplest cliques that will still allow us to encourage some boundary continuity.

What should the "cost" be for each of the clique types shown in Fig. 1, i.e. how should one specify the parameters $\beta_0, \ldots, \beta_5$? A similar question applies to the line-site method, and Silverman *et al.*[7] have proposed an answer using an isotropy constraint. They minimize the ratio of the maximum to minimum cost for a boundary at angles between 0 and $\pi/4$. Applying this same approach to the mixture site penalties yields the following constraint:

$$\beta_1 + \beta_4 = \sqrt{2}\beta_3.$$

Clearly additional constraints are needed, and this is a subject for future research.

2.3. Intensity Process

We model the prior distribution of $\boldsymbol{\mu}$ conditioned on $\boldsymbol{m}$ to be a modified Gaussian Gibbs distribution, with log-prior given by:

$$\Omega_1(\boldsymbol{\mu}|\boldsymbol{m}) = -\log p(\boldsymbol{\mu}|\boldsymbol{m}) =$$
$$\sum_{j=1}^{p} \frac{1}{2} \sum_k w_{jk} \gamma(m_j, m_k) \frac{1}{2}(\mu_j - \mu_k)^2, \qquad (3)$$

where w_{jk} is 1 for horizontal and vertical neighbors, $1/\sqrt{2}$ for diagonal neighbors, and 0 otherwise, and

$$\gamma(m_k, m_j) = \begin{cases} \gamma_0, & m_k = m_j = 0 \\ \gamma_2, & m_k = m_j = 1 \\ \gamma_1, & \text{otherwise} \end{cases}.$$

Using $\gamma_0 > \gamma_2 > \gamma_1$ encourages smoothness within regions, but allows discontinuities between regions. Note that if $\gamma_0 = \gamma_1 = \gamma_2$, then Ω_1 degenerates to the classical auto-normal prior [1].

2.4. Posterior

We adopt a Bayesian (or penalized likelihood) approach, obtaining a point estimate of $\boldsymbol{\mu}$ by the maximum *a posteriori* (MAP) criterion:

$$(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{m}}) = \arg\max_{\boldsymbol{\mu}, \boldsymbol{m}} \log p(\boldsymbol{\mu}, \boldsymbol{m}|\boldsymbol{y}).$$

Using the fact $p(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{m}) = p(\boldsymbol{y}|\boldsymbol{\mu})$ and Bayes rule:

$$p(\boldsymbol{\mu}, \boldsymbol{m}|\boldsymbol{y}) = p(\boldsymbol{y}|\boldsymbol{\mu}) p(\boldsymbol{\mu}|\boldsymbol{m}) p(\boldsymbol{m}) / p(\boldsymbol{y}),$$

so the MAP estimate minimizes the objective:

$$\Phi(\boldsymbol{\mu}, \boldsymbol{m}) = \frac{\|\boldsymbol{y} - \mathbf{H}\boldsymbol{\mu}\|^2}{2\sigma^2} + \Omega_1(\boldsymbol{\mu}|\boldsymbol{m}) + \Omega_2(\boldsymbol{m}).$$

One could minimize this objective function using simulated annealing [5], deterministic or mean field annealing [8], or by a greedy descent algorithm like iterated conditional modes (ICM) [1]. In the examples below, we simply used ICM, where for each iteration each pair (μ_j, m_j) is *simultaneously* updated to monotonically decrease Φ.

As for all greedy algorithms for non-convex objectives, a good initial estimate is desirable. We first performed a classical ML estimate of $\boldsymbol{\mu}$, (under the piecewise uniform assumption). We used that ML estimate to initialize the ICM algorithm of Besag [1] with the pairwise interaction model (also piecewise uniform). The result of that estimation was then taken as the initial estimate $\boldsymbol{\mu}^0$ for minimizing $\Phi(\boldsymbol{\mu}, \boldsymbol{m})$, where $\boldsymbol{m}^0$ was taken simply to be $\mathbf{0}$.

3. RESULTS

To illustrate the mixture-site method, we performed simulations summarized by Figures 3 and 4. For these examples, we took $\mathbf{H} = \mathbf{I}$, but the true objects were first generated on a very fine grid and then reduced down to a coarse grid, so that pixels on object boundaries would have partial volume effects, as is usually the case with real measurements.

For the geometric phantom in Fig. 3, the object intensities in the rectangle and circle were 1 and 2 respectively, and $\sigma = 0.25$. From top to bottom, Fig. 3 shows:

- the noisy measurements $\boldsymbol{y}$,
- the ordinary ML estimate of $\boldsymbol{\mu}$,

- the pairwise-interaction estimate [1] of $\boldsymbol{\mu}$,
- the mixture-site model estimate of $\boldsymbol{\mu}$, and
- the mixture-site model estimate of $\boldsymbol{m}$.

Only 6 ICM iterations were required for the mixture-site model. There are 3 noise replications shown in Fig. 3, so that one can visually assess the variability of the methods.

The phantom in Fig. 4 represents an attenuation map of a human thorax, with attenuation coefficients $\mu = \{0, 0.025, 0.096, 0.165\}\mathrm{cm}^{-1}$ for air, lung, water, and bone respectively. In this case $\sigma = 0.03$. The order of the results in Fig. 4 is the same as that for Fig. 3.

The following table compares the relative improvement in RMS error of the methods.

Method	Geometric Phantom	Thorax Phantom
Data $\boldsymbol{y}$	1	1
ML Estimate	1.4	1.1
Pairwise Interaction ICM	2.9	3.5
Mixture-Site Model	4.7	4.1

Note that even though the underlying object is piecewise constant, the mixture-site method (which is based only on a piecewise-smooth assumption) outperforms the classical pairwise interaction ICM method. It is clear from examining the error images that the largest errors in the pairwise interaction ICM method are near region boundaries. That method, like the line-site method, makes the inaccurate implicit assumption that region boundaries lie along pixel edges, whereas the mixture-site method assumes the region boundaries lie somewhere *within* the pixels. We do not attempt to locate those boundaries within the pixels, but Figs. 3 and 4 demonstrate that the mixture-site method does a reasonable job of identifying which pixels contain region boundaries.

These preliminary results indicate the potential utility of the method. In the future we will perform more quantitative comparisons between our mixture-site method and the line-site method, and will evaluate the method on 3D data sets acquired using PET transmission scanning.

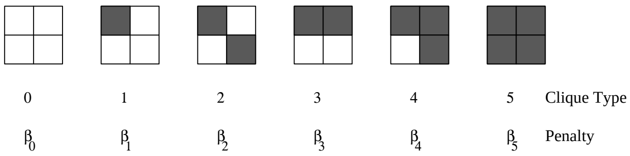

Figure 1: The six basic 2×2 cliques for mixture sites. There are 4 rotationally invariant versions of types 1, 3, and 4, and 2 rotations of type 3, for a total of 2^4 possibilities.

4. REFERENCES

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

[2] G Gindi, M Lee, A Rangarajan, and I G Zubal. Bayesian reconstruction of functional images using anatomical information as priors. *IEEE Trans. Med. Im.*, 12(4):670–680, December 1993.

[3] P J Green. On use of the EM algorithm for penalized likelihood estimation. *J. Royal Stat. Soc. Ser. B*, 52(3):443–452, 1990.

[4] C Bouman and K Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Trans. Im. Proc.*, 2(3):296–310, July 1993.

[5] S Geman and D Geman. Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images. *IEEE Trans. Patt. Anal. Mach. Int.*, 6(6):721–741, November 1984.

[6] H S Choi, D R Haynor, and Y Kim. Partial volume tissue classification of multichannel magnetic resonance images—a mixel model. *IEEE Trans. Med. Im.*, 10(3):395–407, September 1991.

[7] B W Silverman, C Jennison, J Stander, and T C Brown. The specification of edge penalties for regular and irregular pixel images. *IEEE Trans. Patt. Anal. Mach. Int.*, 12(10):1017–1024, October 1990.

[8] D Geiger and F Girosi. Parallel and deterministic algorithms from MRF's: Surface reconstruction. *IEEE Trans. Patt. Anal. Mach. Int.*, 13(5):401–412, May 1991.

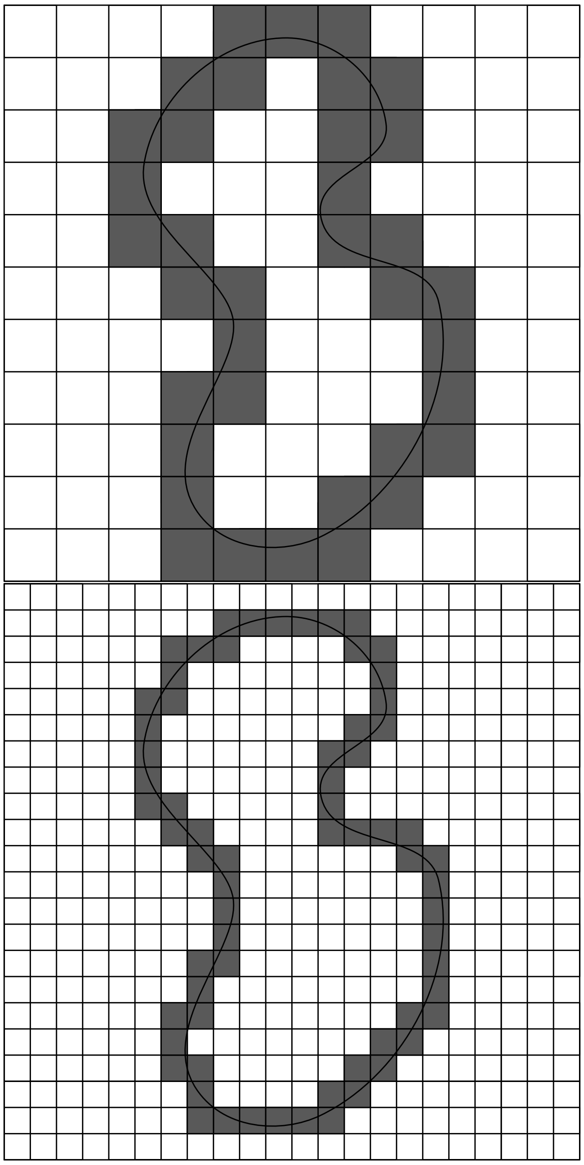

Figure 2: Coarse and fine grid representations of mixture sites for a scene containing two regions. There is a natural hierarchical relationship between coarse and fine grid mixture sites.

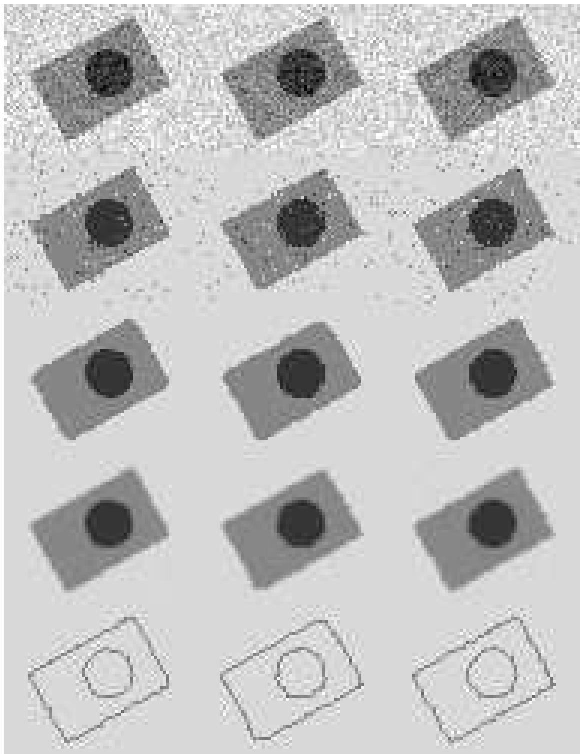

Figure 3: Geometrical phantom. See text.

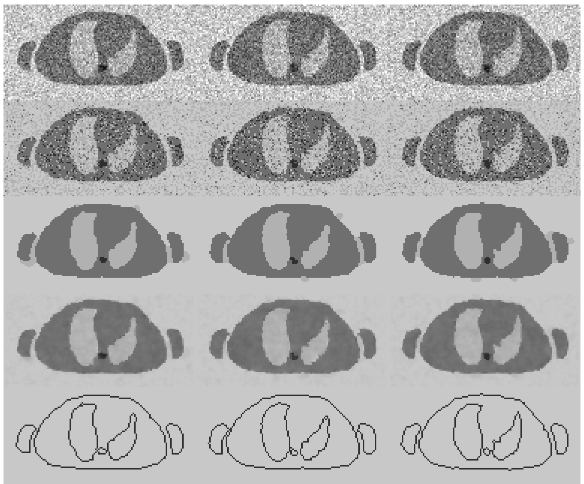

Figure 4: Human thorax phantom. See text.