Regularized Emission Image Reconstruction Using Imperfect Side Information

Jeffrey A. Fessler*, Neal H. Clinthorne, and W. Leslie Rogers
Division of Nuclear Medicine, University of Michigan

ABSTRACT

The inadequacy of the maximum-likelihood criterion for emission image reconstruction has spurred the development of several regularization methods. Despite the spatial variance of medical images, most of the proposed methods are spatially invariant. This paper reports a preliminary investigation of a spatially-variant penalized-likelihood method for tomographic image reconstruction based on a Gibbs penalty. The penalty weights are determined from structural side information, such as the locations of anatomical boundaries in high-resolution magnetic resonance images. Such side information will be imperfect in practice, and a simple simulation demonstrates the importance of accounting for the errors in boundary locations. We discuss methods for prescribing the penalty weights when the side information is noisy. Simulation results suggest that even imperfect side information is useful for guiding spatially-variant regularization.

I. INTRODUCTION

Many investigators have noted the inadequacy of the maximum-likelihood (ML) criterion for emission image reconstruction, and have proposed regularization techniques that stabilize the emission estimate. Most such methods are spatially invariant; however, the spatial variance typical of medical images argues for the use of spatially-variant reconstruction methods. This paper proposes using side information, such as the locations of anatomical boundaries obtained from magnetic resonance (MR) images, to control a spatially-variant penalized-likelihood method based on Gibbs functions. An important feature of this method is that it can accommodate imperfect side information.

The method described is a synthesis of three recent advances in emission image reconstruction. The measurement model includes the effects of attenuation and accidental coincidences, the importance of which is shown in [1]. The use of spatially-variant weights for a Gibbs penalty is analogous to the "weighted-splines" approach described in [2] for a Gaussian noise model. This leads to an optimality criterion with an analytically intractable M-step, so we apply a variant of the GEM [3] iterative method. These points are considered in detail in Section II.

The benefits of structural side information, such as might be derived from MR images, will clearly be task dependent. Therefore, in this paper we depart from the conventional global performance criteria, such as likelihood or global mean-square error, and focus on a specific local criterion: the accuracy of quantifying total uptake within a small region of interest (ROI) surrounded by regions of relatively higher activity. Since the results of such a study will be context dependent, we have chosen a context that has significance to clinical researchers at our institution: quantifying uptake within the globus pallidus and the putamen for patients with Huntington's Disease from position emission tomographic (PET) measurements of regional benzodiazepine receptor density obtained by injection of [^{11}C] flumazenil [4, 5]. These small neurological structures are poorly quantified by conventional filtered back-projection images, due to spill over from the surrounding cortex. However, the boundaries of these structures are obtainable from MR images, so an iterative reconstruction method that can exploit such side information could be beneficial.

Section II describes the basic method, which we have applied to both a representative one-dimensional profile and to a realistic two-dimensional computer phantom with regional activities corresponding to autoradiography results. Due to space limitations, only the one-dimensional simulation results are reported in Section III; preliminary two-dimensional results will be presented at the symposium. Section IV discusses the future directions of this research.

II. METHOD

A. Measurement model

Accurate quantification requires the use of an accurate measurement model. In particular, as shown in [1], the effects of attenuation and accidental coincidences in PET should be accounted for by including them in the measurement model, rather than by precorrecting the measurements. Precorrection destroys the Poisson nature of

*This work was supported in part by NCI Training Grant 5 T32 CA 09015 and by a DOE Alexander Hollaender Postdoctoral Fellowhip.

0-7803-0513-2/92$0.3.00 ©IEEE

the measurements.

For simplicity, we adopt the voxelized object model, and denote the rate of activity in the bth voxel by λ_b, $b = 1, \ldots, B$. The PET system consists of D detector pairs. Let p_{db} denote the point-spread function (PSF) of the PET system, i.e., p_{db} is the probability that, in the absence of attenuation, an event from the bth voxel is detected by the dth detector pair, conditioned on it being detected. Thus $\sum_{d=1}^{D} p_{db} = 1$. Let q_b denote the overall detection probability for an event originating in the bth voxel, in the absence of attenuation. Let μ_d denote the survival probability for the dth detector pair, i.e., the probability that both of an annihilation-produced pair of photons emitted along the d detector pair tube are detected (not attenuated). Let r_d denote the rate of accidental coincidences for the dth detector pair. Then if y_d denotes the number of events counted by the dth detector pair, we assume the y_d's have independent Poisson distributions:

$$y_d \sim \text{Poisson}(T \cdot (\sum_{b=1}^{B} a_{db}\lambda_b + r_d)), \qquad (1)$$

where T is the product of the scan time and correction factors such as that for radioactive decay, and $a_{db} = \mu_d p_{db} q_b$. For simplicity, we absorb T into λ_b and r_d.

B. Regularizing Penalty Function

Although one could use (1) to define an estimation method based on the ML criterion, the resulting estimates have excessive variance for the specific tasks we are considering, as we show in simulations below. By considering instead an optimality criterion that is the difference between the log-likelihood and a penalty function that discourages excessive variation between neighboring voxels, one can significantly reduce the variance with only a small increase in bias, thereby reducing the total RMS error. How much bias is tolerable is clearly task dependent, and is a subject needing further investigation.

One's choice for the penalty function again will be task dependent. Our hypothesis was that the task of quantifying uptake within a small cold spot would benefit from a spatially-variant penalty function, so a weighted Gibb's function is a natural choice. Specifically, we consider the following optimality criterion:

$$\Phi(\boldsymbol{\lambda}) = -\mathbf{1}'\mathbf{A}\boldsymbol{\lambda} + \mathbf{y}'\log(\mathbf{A}\boldsymbol{\lambda} + \mathbf{r}) - \alpha V(\boldsymbol{\lambda}) \qquad (2)$$

where $\mathbf{1}$ is the column vector of ones of length D, $\mathbf{A} = \{a_{db}\}$, $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_B]$, $\mathbf{r} = [r_1, \ldots, r_D]$, and $\mathbf{y} = [y_1, \ldots, y_D]$. For simplicity, we focus on the one-dimensional case hereafter, and define the Gibb's function V by:

$$V(\boldsymbol{\lambda}) = \frac{1}{2}\sum_{b=1}^{B-1} w_b(\lambda_{b+1} - \lambda_b)^2.$$

The weights w_b control the influence of the penalty. If the presence of an anatomical boundary in an MR image implies that the activities in two neighboring voxels are likely to be disparate, then the corresponding weight should be set relatively small, so as to avoid penalizing the discrepancy. Such a scheme reduces the "edge artifact" of spatially-invariant regularization.

In practice, side information will be imperfect due to noise in MR images and to registration errors. It is essential to account for these errors, and again the Gibbs function approach offers a convenient method. For example, one could first use the side information to generate weights that would be optimal if the boundaries were perfect, and then convolve the weights with a kernel whose width corresponds to the uncertainty in the side information.

C. Iterative Algorithm

Historically, the use of objective functions such as (2), with its Gibbs penalty, has been hampered by the slow convergence of the associated stochastic maximization procedures or by the uncertain behavior of locally convergent methods. Following the usual estimate-maximize (EM) algorithm derivation, one can show that the E step for (2) is:

$$\hat{y}_d(\boldsymbol{\lambda}^i) = \sum_{b=1}^{B} a_{db}\lambda_b^i + r_d,$$

$$n_{db}^i = y_d a_{db}\lambda_b^i / \hat{y}_d(\boldsymbol{\lambda}^i),$$

where $\boldsymbol{\lambda}^i$ denotes the emission estimate after the ith iteration, and n_{db}^i is the conditional expectation of the number of events in the dth detector due to the bth voxel. The M step requires maximizing:

$$\sum_{d=1}^{D}\sum_{b=1}^{B}\left[-a_{db}\lambda_b^{i+1} + n_{db}^i \log(a_{db}\lambda_b^i)\right] - \alpha V(\boldsymbol{\lambda}^{i+1}) \qquad (3)$$

over $\boldsymbol{\lambda}^{i+1}$. The resulting coupled set of equations appears to have no analytical solution. We experimented with the "one step late" (OSL) method of Green [6], but found the necessity of line-search operations [7] to be computationally prohibitive. We have adopted the generalized estimate-maximize (GEM) method of Hebert [3, 8], which, although originally applied to SPECT, is also applicable to the PET measurement model (1). Zeroing the derivative of (3) with respect to λ_b^{i+1} yields:

$$0 = -\left(\sum_{d=1}^{D} a_{db}\right) + \left(\sum_{d=1}^{D} n_{db}^i\right)/\lambda_b^{i+1}$$

$$- \alpha w_{b-1}(\lambda_b^{i+1} - \lambda_{b-1}^{i+1}) - \alpha w_b(\lambda_b^{i+1} - \lambda_{b+1}^{i+1}). \qquad (4)$$

The GEM approach is to first set $\lambda_b^{i+1} = \lambda_b^i$ for all b, then to loop through the b's in some order and to replace λ_b^{i+1} with the unique positive root of (4). After considering the discussion in [9], we chose the following ordering: on even iterations, the even voxels are updated first, and then the odd voxels are updated; the opposite order holds for odd

iterations. Unlike the method in [9], this is guaranteed to increase the penalized likelihood each iteration [3].

The convergence of these methods has been addressed by Lange [7], under the assumption that the penalized log-likelihood (2) is a strictly convex function of $\boldsymbol{\lambda}$. This was established in [7] by showing that the log-likelihood (for a SPECT measurement model) is convex, and the penalty term is strictly convex. In our case, it is possible that several weights could be set to zero, in which case the penalty term may not be strictly convex (although it would remain convex). Fortunately, the presence of accidental coincidences in PET ensures that the log-likelihood term is strictly convex, provided the matrix $\mathbf{A}$ has full column rank [10].

III. 1D SIMULATION

To explore the possible benefits of spatially-variant regularization in the presence of imperfect structural side information, we performed simulations based on the one-dimensional profile shown in Figure 1. This profile is representative of the flumazenil quantification task. The cold spot represents the putamen, which is adjacent to the globus pallidus and the cortex, both of which have significantly higher activity. For this simulation, the diameter of the cold spot is 7 pixels. The system matrix $\mathbf{A}$ corresponds to a triangular point spread function with a FWHM of 5 pixels. A typical measurement realization $\mathbf{y}$ is shown in Figure 2.

Our task is to quantify the total uptake within the cold spot. This task requires two components: 1) reconstructing the activity distribution from the noisy measurements, and 2) identifying the boundaries of the region of interest (ROI) and integrating the activity within that ROI. Here we focus only on the first task by integrating the activity within the *true* ROI (pixels 33 through 39) for all simulations.

We have examined the performance of the GEM reconstruction method as a function of the regularization parameter α in four scenarios. For each scenario and for each value of α, 50 measurement realizations were reconstructed via 1000 iterations of the GEM algorithm.

- Case 1: No side information was available, so all the weights w_b were set to 1. This corresponds to spatially-invariant regularization.
- Case 2: Perfect side information corresponding to the edge locations was used. Specifically, $w_{32} = w_{39} = 0$, and all other weights were set to 1. Thus, the weights for pixel pairs that straddled an edge of the cold spot were set to 0, thereby reducing spill over.
- Case 3: Imperfect side information was simulated. Let b_l and b_r denote the left and right edges of the cold spot as they might be determined from an MR image. (These two values are the side information). For each realization, b_l was randomly selected from the set $\{31, 32, 33\}$, and b_r was randomly selected from the set $\{38, 39, 40\}$. For Case 3, the side information was applied "blindly," i.e., we set $w_{b_l} = w_{b_r} = 0$ and all other weights to 1, despite the fact that b_l and b_r are usually incorrect.
- Case 4: Imperfect side information with the same error distribution as in Case 3. However, for Case 4 we made the following heuristic attempt to account for the errors in b_l and b_r: we set $w_{b_l-1} = w_{b_l} = w_{b_l+1} = w_{b_r-1} = w_{b_r} = w_{b_r+1} = 0.01$, and all other weights to 1. This small band of weights allows for a rapid activity transition in the neighborhood of edge location specified by the noisy side information.

Figure 3 displays the percent root mean-square (RMS) error in total uptake within the cold spot as a function of the regularization parameter α. The optimal performance of each method is summarized in Table I. It is useful to compare the mean of the 50 reconstructions for each case with the true activity distribution; these are shown in Figures 4—7. The dotted bands around each curve represent one standard deviation above and below the mean. The curve for each case corresponds to the value α_{opt} shown in Table I that minimized the RMS error. The four cases illustrate the tradeoff between bias and variance: Case 1 has high variance, but additional smoothing would cause more bias increase than variance decrease. Case 2 has low variance and low bias since the edges are known perfectly. Case 4 improves over Case 3 by significantly reducing the bias with some variance increase.

IV. DISCUSSION

The results summarized in Table I demonstrate that the use of structural side information, in conjunction with a spatially variant reconstruction method, can significantly reduce RMS error over spatially invariant regularization. In this case, the RMS error was reduced by almost a factor of 3 with the use of perfect side information. However, we must also make the sobering conclusion that if the side information is imperfect, then using it "blindly" is unlikely to be significantly beneficial for quantification tasks. By applying a simple heuristic scheme that attempts to account for the uncertainty in the side information, we were able to recover some of the benefits of the side information, despite its imperfections. More investigation into how to efficiently use noisy side information is clearly needed.

V. ACKNOWLEDGEMENT

The authors gratefully acknowledge the contributions of G. Hutchins and R. Koeppe.

	α_{opt}	% RMS error
Case 1	-4	30.8 ± 1.9
Case 2	-2	11.0 ± 1.1
Case 3	-3	28.8 ± 2.1
Case 4	-2.5	17.4 ± 1.7

Table I: Percent RMS error for the optimal α for the four cases (see text).

REFERENCES

[1] D. G. Politte and D. L. Snyder, "Corrections for accidental coincidences and attenuation in maximum-likelihood image reconstruction for positron-emission tomography," *IEEE Transactions on Medical Imaging*, vol. 10, pp. 82–89, Mar. 1991.

[2] X. Wang, B. G. Schunck, and W. L. Rogers, "A iterative reconstruction algorithm using spatially variant regularization for emission tomography," *Journal of Nuclear Medicine*, vol. 32, pp. 936–937, May 1991.

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

[4] R. A. Koeppe, V. A. Holtoff, K. A. Frey, M. R. Kilbourn, and D. E. Kuhl, "Compartmental analysis of [^{11}C] flumazenil kinetics for the estimation of ligand transport rate and receptor distribution using positrom emission tomography," *Journal of Cerebral Blood Flow and Metabolism*, vol. 11, no. 5, pp. 735–744, 1991.

[5] V. A. Holtoff, R. A. Koeppe, K. A. Frey, J. B. Penney, D. S. Markel, D. E. Kuhl, and A. B. Young, "PET measures of benzodiazepine receptors in Huntington's disease," 1991. In review.

[6] P. J. Green, "Bayesian reconstructions from emission tomography data using a modified EM algorithm," *IEEE Transactions on Medical Imaging*, vol. 9, pp. 84–93, Mar. 1990.

[7] K. Lange, "Convergence of EM image reconstruction algorithms with Gibbs smoothing," *IEEE Transactions on Medical Imaging*, vol. 9, pp. 439–446, Dec. 1990. Corrections, June 1991 TMI.

[8] T. Hebert and R. Leahy, "A Bayesian reconstruction algorithm for emission tomography using a Markov random field prior," in *Proc. SPIE 1092, Medical Imaging III: Image Processing*, pp. 458–4662, 1989.

[9] G. T. Herman and D. Odhner, "Performance evaluluation of an iterative image reconstruction algorithm for positron-emission tomography," *IEEE Transactions on Medical Imaging*, vol. 10, pp. 336–346, Sept. 1991.

[10] J. A. Fessler, N. H. Clinthorne, and W. L. Rogers, "On complete-data spaces for pet reconstruction algorithms," 1991. Submitted to IEEE Transactions on Medical Imaging.

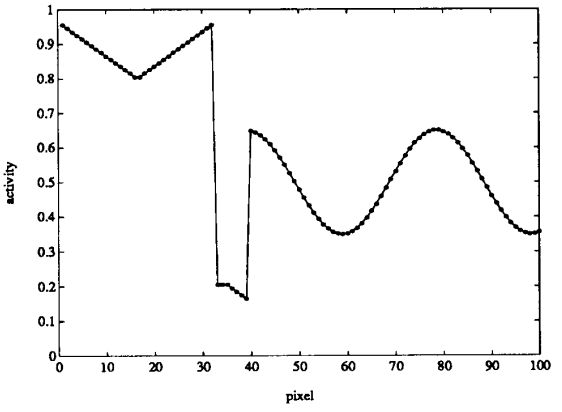

Figure 1: Simulated one-dimensional activity distribution.

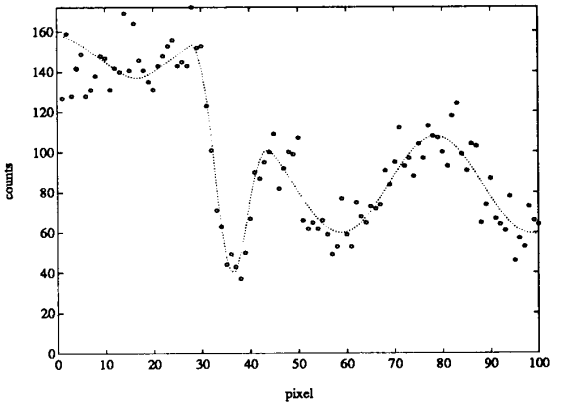

Figure 2: Measurement realization with 10^4 counts.

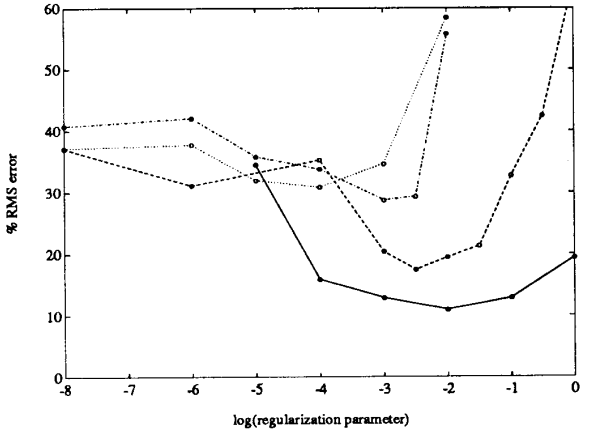

Figure 3: Percent RMS error for the four cases versus α. Case 1 (.), Case 2 (-), Case 3 (-.), Case 4 (- -).

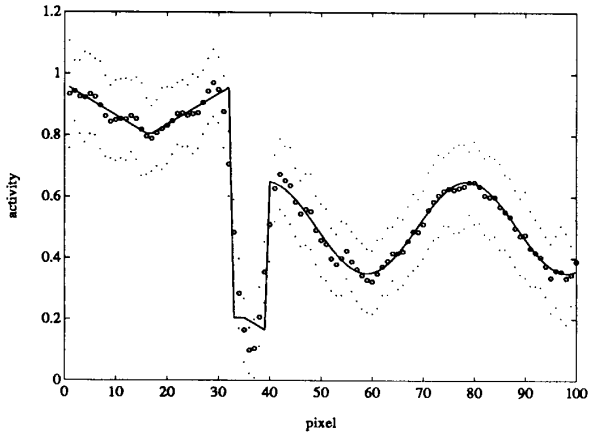

Figure 4: True activity (-), mean reconstruction (o), and pointwise standard deviation (.) for Case 1.

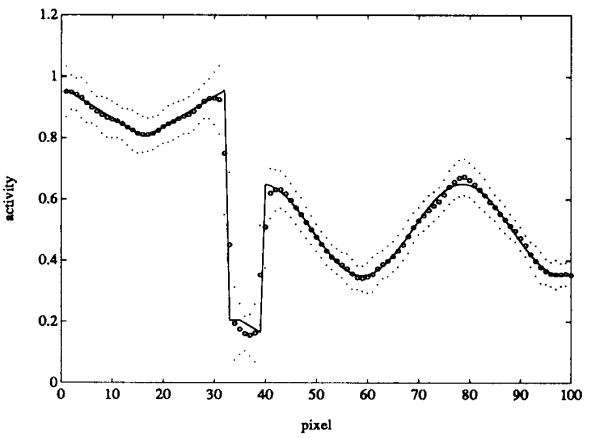

Figure 6: True activity (-), mean reconstruction (o), and pointwise standard deviation (.) for Case 3.

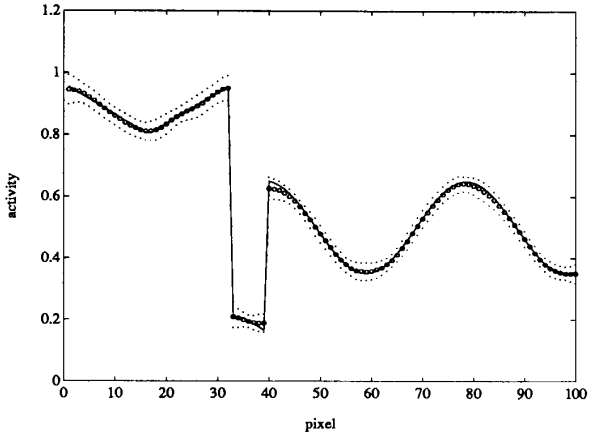

Figure 5: True activity (-), mean reconstruction (o), and pointwise standard deviation (.) for Case 2.

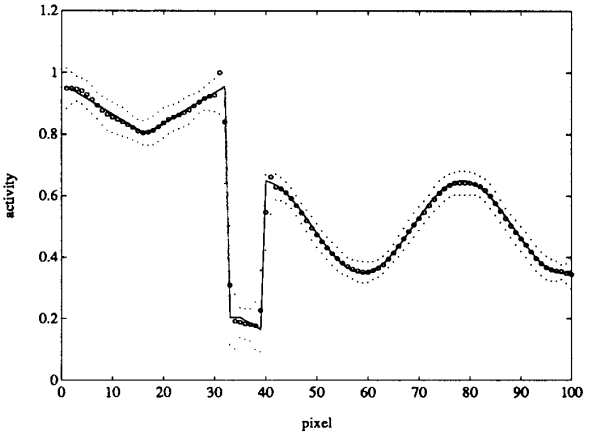

Figure 7: True activity (-), mean reconstruction (o), and pointwise standard deviation (.) for Case 4.