

IMPROVED PENALIZED LIKELIHOOD RECONSTRUCTION OF ANATOMICALLY CORRELATED EMISSION DATA

Stephen R. Titus, Alfred O. Hero III, Jeffrey A. Fessler

4401 EECS, University of Michigan
Ann Arbor, MI 48109
stitus@engin.umich.edu

ABSTRACT

This paper presents a method for incorporating anatomical NMR boundary side information into penalized maximum likelihood (PML) emission image reconstructions. The NMR boundary is parameterized as a periodic spline curve of fixed order and number of knots that is known *a priori*. Maximum Likelihood (ML) estimation of the spline coefficients yields an "extracted" boundary, which is used to define a set of Gibbs weights on the emission image space. These weights, when coupled with a quadratic penalty function, create an edge-preserving penalty that incorporates our prior knowledge effectively. Qualitative analysis demonstrates that our method results in smooth images that do not suffer loss of edge contrast, while quantitative estimates of bias and variance for various values of the smoothing parameter show an improvement over standard quadratically penalized maximum likelihood.

1. INTRODUCTION

Emission computed tomography (ECT) image reconstructions, which are generated from data that consists of counts of detected radionuclide decay events, have been traditionally poor in terms of the sensitivity/resolution (bias/variance) tradeoff inherent in any image reconstruction problem. Many factors including dose limitations, scatter, attenuation, and detector efficiency contribute to the low quality of these images. Recently, however, various researchers have suggested that when functional and anatomical boundaries are likely to be spatially correlated, anatomical boundary curves from NMR images, when used as side information, might improve the quality of emission image reconstructions [1],[2], [3]. We focus on the case where a single boundary of interest is present in both the emission and NMR images. In this paper, we describe a new

spline-based method of incorporating NMR-derived anatomical boundary information into the penalized maximum likelihood reconstruction (PML) algorithm for estimating radionuclide concentration in the ECT image.

2. THEORY

2.1. NMR System and Boundaries

We make three common assumptions about the magnetic resonance imaging system: it is linear and spatially invariant, its point spread function is Gaussian, and the source of noise is additive thermal noise introduced solely by the electronic instrumentation. This noise is well-modeled by white Gaussian noise [4]. Under this model, the magnetic resonance image I corresponding to a proton spin density I_{true} is

$$I(x, y) = I_{true}(x, y) ** G(x, y) + N(x, y), \quad (1)$$

where the point spread function G is given by

$$G(x, y) = \frac{1}{\sqrt{2\pi}\sigma_s} e^{-\frac{(x^2+y^2)}{2\sigma_s^2}},$$

N is a zero-mean spatial white noise process whose variance we will denote by σ_n^2 , and the $**$ operator denotes 2D convolution. Figure 1 is a simulated discretized




Figure 1: Simulated NMR image

This work was supported by National Cancer Institute, ROI-CA-54362-02 and DOE Grant DE-FG02-87ER65061.

NMR image of an elliptical phantom of interior intensity 6.0 and background intensity 0 with $\sigma_s = 1.2$ pixels and $\sigma_n^2 = .20$.

We parameterize the NMR phantom (I_{true}) boundary using a periodic polar spline function $f(\phi)$ of known degree and known number of knots. A *periodic polar spline* of degree d defined with respect to strictly increasing scalar “knots” $\{k_1, \dots, k_n\} \in [0, 2\pi]$ is any function $f(\phi)$ which has $k_1 = 0$ and $k_n = 2\pi$, and is

- (i) Continuous on $[k_1, k_n]$
- (ii) Has $d - 1$ continuous derivatives on $[k_1, k_n]$
- (iii) Is a polynomial of d th degree or less on intervals $[k_i, k_{i+1}] \quad i = 1, \dots, n - 1$
- (iv) Satisfies $f^{(j)}(k_1) = f^{(j)}(k_n) \quad j = 0, \dots, d - 1$

The simplicity, continuity, and smoothness of splines has made them an obvious choice as boundary modeling functions in many applications [5], [6]. Since a polar spline is a single valued function it can be used to represent any closed curve that can be defined as a polar function with respect to some interior point (i.e. a star shaped region). This requires that the interior point be known (or estimated) *a priori*. We will assume in this paper that this polar origin is known, and we will use a quadratic spline model consisting of $n = 16$ equi-angularly spaced knots on $[0, 2\pi]$.

Any n -knot quadratic, periodic spline function on $[0, 2\pi]$ can be expressed as a linear combination of a set of n *b-spline basis functions* [6]:

$$f_{\theta}(\phi) = \sum_{j=1}^n \theta_j B_j(\phi) \quad \phi \in [0, 2\pi]$$

where each basis function $B_j(\phi)$ depends on the degree d as well as the number of knots n , and is defined in terms of divided differences of truncated polynomial functions (see [5]). With the b-spline parameterization the phantom I_{true} in (1) becomes a parametric function of the vector of coefficients $\theta = (\theta_1, \dots, \theta_n)$. We will extract the boundary side information from the discretized NMR image using ML estimation of θ . Since the noise model is Gaussian this is equivalent to solving the non-linear least squares problem:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \|I_{true}^{\theta} ** G - I\|^2. \quad (2)$$

Here I 's dependence on θ has been explicitly written, and the operator $**$ now represents discrete rather than continuous convolution. In order to minimize this objective function, we must generate a smoothed version $I_{true}^{\theta} ** G$ of the “true” image corresponding to boundary θ . For simplicity, we assume that the Gaussian

noise variance, point spread parameter σ_s , and the NMR image intensity inside the spline boundary are known (our method could easily be extended to include these quantities as parameters to be estimated).

2.2. Modified Penalty for Emission Data

In emission tomography, a patient is injected with radionuclide, and emissions are counted by pairs of detectors oriented around a particular anatomical slice. The reconstruction problem consists of estimating parameters $\lambda = (\lambda_1, \dots, \lambda_P)$, the radionuclide concentrations in the slice pixels, from a dataset \mathbf{Y} of detected counts. The emission measurements have independent Poisson distributions, and we assume that Y_n has mean

$$\bar{Y}_n(\lambda) = \sum_j a_{nj} \lambda_j + r_n$$

where the a_{nj} are proportional to the probability that an emission in pixel j is detected by the n th detector pair, and $\{r_n\}$ represents additive background events (e.g. random coincidences). Given a measurement realization (sinogram) $\mathbf{Y} = \mathbf{y}$, the goal in PML emission reconstruction is to compute the penalized maximum likelihood estimate $\hat{\lambda}$ of the emission intensities, defined by:

$$\hat{\lambda} = \operatorname{argmax}_{\lambda} \{\ln f(\mathbf{y}; \lambda) - P(\lambda)\} \quad (3)$$

where $f(\mathbf{y}; \lambda)$ is the Poisson probability distribution of random vector \mathbf{Y} , and $P(\lambda)$ is a roughness penalty. A common choice for the penalty $P(\lambda)$ is the function

$$P(\lambda) = \beta \frac{1}{2} \sum_k \frac{1}{2} \sum_j \omega_{kj} (\lambda_k - \lambda_j)^2, \quad (4)$$

where the summations above are over pixels, and $\{\omega_{kj}\}$ is a set of penalty weights that for fixed k promote smoothing within the neighborhood of the k th pixel. We use a Gibbs weighting scheme, where a first-order 2D pixel neighborhood consists of four (up, down, left, right) adjacent neighbors, a second-order neighborhood includes all eight surrounding pixels, and the weights are chosen to be symmetric ($\omega_{kj} = \omega_{jk}$) with ω_{kj} non-zero only in the neighborhood of the k th pixel.

Since large P decreases our objective function, it is clear that the maximization of (3) will attempt to some extent to minimize P by encouraging pixels with non-zero weights to take on similar values in the reconstruction $\hat{\lambda}$. The smoothing parameter β controls the tradeoff between the conflicting goals of maximizing $\ln f(\mathbf{y}; \lambda)$ and minimizing $P(\lambda)$. The objective function maximum can not be found analytically; therefore, an iterative method must be used. In the sequel we use the SAGE3 algorithm [7] to maximize (3). This

EM-type algorithm is characterized by monotonic, fast convergence.

Choosing a uniform weight scheme (e.g. all $\omega_{kj} = 1$) results in a global smoothing of the reconstructed image that in turn causes severe bias as image detail is blurred indiscriminantly across region boundaries. Our side information θ can be used to encourage more intelligent smoothing that occurs within, but not across, the region boundary present in the NMR image (which is assumed correlate highly with the functional region boundary). We encode this information into the Gibbs weights as follows. First, a "soft limited" image I^θ is generated whose pixel intensities lie in $[0, 1]$ — each intensity is equal to the normalized area of that pixel which belongs to the interior of closed curve θ . Then, the set of weights $\{\omega_{kj}^\theta\}$ is assigned using I^θ as follows:

$$\omega_{jk}(I^\theta) = \begin{cases} \min(I_j^\theta, I_k^\theta) & : I_j^\theta, I_k^\theta \geq .50 \\ 1 - \max(I_j^\theta, I_k^\theta) & : I_j^\theta, I_k^\theta < .50 \\ 0 & : \text{else} \end{cases}$$

This scheme divides image pixels into two groups, interior and exterior, and associates all pixels whose intensities are $\geq .50$ to the interior and the remainder to the exterior. The strength of the bond between neighbors of the same type varies between .50 and 1, and is equal to the value of the smaller neighbor (interior) or one minus the larger neighbor (exterior).

2.3. Mean Uptake Measurement

To compare the performance of PML emission reconstruction with and without side information, it is crucial to examine the tradeoff between the bias and variance components of mean squared estimate error. In this paper, we focus on mean uptake in a region of interest (ROI) defined as the pixels where the elliptical hot spot of the phantom pictured in Figure 2(a) is non-zero. If we let \mathbf{v} be a vector whose length is the same as λ and is 1 in the specified ROI and 0 elsewhere, we can estimate mean uptake as

$$\hat{\alpha} = \frac{\hat{\lambda}' \mathbf{v}}{M}$$

where M is the number of pixels in the ROI and operator $'$ denotes vector transpose. The bias and standard deviation of our mean uptake estimate are defined as

$$\begin{aligned} b(\hat{\alpha}) &= E[\hat{\alpha}] - \alpha \\ s(\hat{\alpha}) &= \sqrt{E[(\hat{\alpha} - E[\hat{\alpha}])^2]} \end{aligned}$$

In the simulations that follow, we quantify estimator performance by examining estimates of the bias and standard deviation of α for a wide range of values of β , the smoothing parameter.

3. RESULTS

3.1. Qualitative

To qualitatively assess the differences between reconstructions performed with no side information ($\omega_{jk} = 1$ for all neighbors), noiseless (ideal) side information, and noisy (extracted) side information, we generated simulated PET emission data for the 64 x 64 phantom shown in Figure 2(a) using 64 bins and 60 angles covering an orbit of 180 degrees. Total counts were normalized to 10^6 , random coincidences made up 15% of total counts, and effects of both attenuation and limited detector efficiency were ignored.

Figure 2 (b)-(d) shows reconstructions that were obtained by maximizing the PML objective via 40 iterations of the SAGE3 algorithm. A 2nd order Gibbs neighborhood was used, and the algorithm was initialized with a filtered backprojection reconstruction (Hanning filter with a discrete frequency cutoff of .50). Figure 2(b) is the reconstructed image when $\beta = 0$ and no side information is used: as one would expect, the boundary of the ROI ellipse is severely blurred. In figure 2(c), we kept β constant, but encoded ideal boundary side information into the Gibbs weights. Note that despite the smoothing that has been achieved in both background and interior regions, the ellipse edges are perfectly clean. Figure 2(d) shows the result using side information that has been extracted from the simulated NMR image of Figure 1. Twenty iterations of a standard conjugate gradient algorithm were used to perform the minimization of (2). Some boundary estimation error is evident; however, considering the level of smoothing and blurring present in the NMR image of Figure 1, the reconstruction quality seems quite good.

3.2. Quantitative

Bias and standard deviation estimates for our measure $\hat{\alpha}$ of the mean uptake in our ROI were obtained by reconstructing 100 independent realizations of emission data via 40 SAGE3 iterations for a range of β values. Sample mean and standard deviation values were then used to calculate bias and standard deviation estimates, which are expressed in our plots as a percentage of the true mean uptake value for our region of interest. Figure 3 shows the performance of uniform versus ideal side information. The uniform and ideal curves shown were swept out from upper right to lower left by increasing β . We can immediately see that the use of side information has improved reconstruction performance with respect to both bias and % standard deviation for nearly all values of β (note that the perfect performance would be achieved at the origin, where both bias




Figure 2: (a) emission phantom with background ellipse of intensity 3 and ROI ellipse of intensity 9. (b) emission reconstruction using 40 sage3 iterations with no side information ($\beta = 0$). (c) using ideal spline side information. (d) using a spline extracted from a single noisy NMR image via MLCG estimation.

and variance are 0). For nearly all levels of standard deviation, the selective smoothing afforded by the ideal side information resulted in significantly less bias than was present without side information. The difference is most striking for smaller values of standard deviation. For example, at a standard deviation level of .25%, increased smoothing has pushed bias up to almost -10% for the uniform case, while the reconstructions that employed ideal side information had bias of less than 2%. Results for the case of extracted side information were similar to those for ideal side information, with a curve that looked nearly identical to that of the ideal case, but with slightly higher bias and variance due to the estimation error of the extraction process.




Figure 3: Bias variance estimates for reconstructions of Figure 2 (b) and (c) as a function of β

4. REFERENCES

[1] W. Ouyang. Incorporation of correlated structural im-

ages in PET image reconstruction, 1992. Submitted to Transactions on Medical Imaging.

[2] G. Gindi, M. Lee, A. Rangarajan, and G. Zubal. Bayesian reconstruction of functional images using anatomical information as priors. *IEEE Transactions on Medical Imaging*, 12(4), December 1993.

[3] V. E. Johnson. A model for segmentation and analysis of noisy images. *Journal of the American Statistical Association*, 89(425):230-241, March 1994.

[4] W.A. Edelstein. Intrinsic signal to noise ratio in NMR imaging. *Magnetic Resonance in Medicine*, 1986.

[5] T.N.E. Greville, editor. *Theory and Applications of Spline Functions*. Academic Press, 1969.

[6] J. H. Ahlberg. *The Theory of Splines and Their Applications*. Academic Press, 1967.

[7] J. A. Fessler and A.O. Hero. Penalized maximum likelihood image reconstruction using space alternating generalized em algorithms. *IEEE Transactions on Image Processing*, 4(10), October 1995.