

NMR OBJECT BOUNDARIES: B-SPLINE MODELING AND ESTIMATOR PERFORMANCE

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

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

ABSTRACT

We give estimation error bounds and specify optimal estimators for continuous, closed boundary curves in an NMR image. The boundary is parameterized using periodic B-Splines. A Cramer-Rao lower bound on mean-square-estimate error in the presence of system smoothing and Gaussian noise is derived, and the performance of maximum likelihood and penalized maximum likelihood estimators is compared to this bound. Finally, we comment on the usefulness of estimates of the boundary for providing anatomical side information in the reconstruction of functional tomographic images like those of a PET or SPECT system.

1. INTRODUCTION

In the past few decades, the theory of splines has been applied to an increasing number of tasks including data interpolation, approximation of functions and functionals, solutions of differential equations, and numerous boundary/surface estimation problems [1],[2].

Recently it has been suggested that the use of anatomical boundaries taken from NMR images might improve the quality of functional tomographic image reconstructions [3],[4],[5]. Many organ and tumor boundaries are continuous, smooth curves that are well-suited to spline approximation. In this paper, we examine the performance of both the maximum likelihood (ML) and a penalized maximum likelihood (PML) estimator with respect to the Cramer-Rao lower bound on estimator mean squared error.

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

2. THEORY

2.1. The B-Spline Boundary Model

Let r be a continuous, closed curve which encloses a region $B \subset \mathbb{R}^2$. We define a continuous binary image, λ_r , over the plane as follows (see Figure 1):

$$\lambda_r = I_1 I_B(x, y) + I_0 (I(x, y) - I_B(x, y)) \quad (1)$$

for constants $I_1, I_0 \geq 0$. The functions $I_B(x, y)$ and $I(x, y)$ are indicator functions that are unity over B and $\mathbb{R}^2$, respectively.




Figure 1: Continuous Image Model

Let the boundary r be described in parameterized form $r(\phi)$ with respect to an origin (x_o, y_o) located in the interior of region B . We choose to model the curve $r(\phi)$ as a piecewise polynomial *spline* function. Specifically, let $\{k_i\}_{i=1}^n$ be a strictly increasing sequence of angles in $[0, 2\pi]$. A *periodic spline of degree d* defined with respect to $\{k_1, \dots, k_n\}$ is any function $f(\phi)$ which satisfies the following constraints:

- $f(\phi)$ is continuous on $[k_1, k_n]$

- $f(\phi)$ has at least $d - 1$ continuous derivatives on $[k_1, k_n]$
- $f(\phi)$ consists of a polynomial of degree d or less on the segments $[k_i, k_{i+1}]$, $i = 1 \dots n - 1$
- $f(k_1) = f(k_n)$

The set of all such splines is a linear space, which we will denote $S_d(k_1, \dots, k_n)$. In spline terminology the scalars $\{k_i\}_{i=1}^n$ are known as *knots*; therefore, we will refer to $S_d(k_1, \dots, k_n)$ as the space of all periodic splines of degree d defined on the set of knots $\{k_i\}_{i=1}^n$.

The usefulness of splines, especially when modeling curves in the physical world, is well-known. In addition, the polynomial representation of splines makes them computationally easy to handle, the integrals and derivatives of spline functions are themselves splines, and least-squares data fitting using splines preserves the first two moments of the data.

In our case, we assume that the boundary $r(\phi)$ can be expressed as a periodic spline curve of degree d with knots $\{k_1, \dots, k_n\}$ equiangularly spaced over $[0, 2\pi]$; i.e.

$$r(\phi) \in S_d(k_1, \dots, k_n), \quad k_i = \frac{2\pi}{n} i \quad (2)$$

The *B-spline* basis for $S_d(k_1, \dots, k_n)$ is a computationally tractable, nonorthogonal basis whose n elements are shifted versions of a single spline that is non-zero over only $d + 1$ knot intervals (these functions are depicted for the case $d = 2$ and $n = 6$ in Figure 2). Thus, if we let $B_i^{n,d}(\phi)$ denote the i th B-spline basis element of degree d defined over equispaced knots $\{k_1, \dots, k_n\}$, we can express the boundary r in terms of this basis as follows:

$$r_\theta(\phi) = \sum_{i=1}^n \theta_i B_i^{n,d}(\phi) \quad \forall \phi \in [0, 2\pi] \quad (3)$$

We will refer to the coefficients $\theta = \{\theta_1, \dots, \theta_n\}$ above as the *spline coefficients* corresponding to the boundary $r(\phi)$. Note that in Figure 2, as well as in the sequel, we have scaled the B-spline basis functions so that they sum to one, which makes for easy interpretation of coefficient magnitudes. For example, if we choose coefficients $\theta_i = 1, \forall i$, the resultant curve is $r(\phi) = 1, \forall \phi$ — a circle of radius one. In addition, we assume that $r_\theta(\phi)$ of equation (3) is defined with respect to origin (x_o, y_o) equal to image center. This restricts us to the large class of star-shaped regions which includes all convex objects.

2.2. The Imaging System Model

We will approximate the NMR imaging unit as a linear, spatially shift invariant system that convolves the input image λ_r of equation (1) with a 2D point response




Figure 2: Quadratic B-spline Basis Set

function $k(x, y)$, and adds Gaussian noise $N(x, y)$ of variance σ_n^2 to the result:

$$Y(x, y) = \lambda_s(x, y) + N(x, y), \quad (x, y) \in \mathbb{R}^2 \quad (4)$$

Function $\lambda_s(x, y) = \lambda_r(x, y) * k(x, y)$ is the smoothed image resulting from the convolution operation (which approximates the finite resolution of the physical system).

The point response $k(x, y)$ is approximated by a 2D symmetric Gaussian hill of variance σ_s^2 as in [6]; i.e.

$$k(x, y) = \frac{1}{2\pi\sigma_s^2} \exp \left[\frac{-(x^2 + y^2)}{2\sigma_s^2} \right] \quad (5)$$

2.3. Fisher Information and C-R Bound

The Cramer-Rao bound gives an expression for the minimum mean squared error or, equivalently, covariance of any estimator $\hat{\theta}$ of parameter vector θ . In the unbiased case, where $E[\hat{\theta}] = \theta$, this bound reduces to

$$\text{cov}(\hat{\theta}) \geq [F_\theta]^{-1}, \quad (6)$$

where F_θ is the Fisher information associated with estimating θ and is defined by

$$F_\theta = E [\nabla_\theta \ln f(Y; \theta; \lambda) \nabla_\theta^T \ln f(Y; \theta; \lambda)] \quad (7)$$

The vector λ above denotes the image pixel intensities. Our primary goal is to derive an analytical expression for F_θ , where θ is the n -tuple of equation (3) which describes the object boundary in terms of a prescribed B-spline basis set.

Since $N(x, y)$ is white noise, we have

$$\ln f(Y; \theta; \lambda) = \quad (8)$$

$$C + \left(\frac{-1}{2\sigma_n^2} \right) \iint_{\mathbb{R}^2} [Y(x, y) - \lambda_\theta^*(x, y)]^2 dx dy,$$

where $\lambda_\theta^*(x, y)$ above is the smoothed NMR image described earlier. Taking gradients yields the Fisher information matrix

$$F_\theta = \frac{1}{\sigma_n^2} \iint_{\mathbb{R}^2} \nabla_\theta \lambda_\theta^*(x, y) \nabla_\theta^T \lambda_\theta^*(x, y) dx dy \quad (9)$$

Note that in the above expression we have fixed the number of knots, knot positions, and spline degree.

Let T be a $p \times n$ matrix ($p \gg n$) whose columns are samples over $\phi \in [0, 2\pi]$ of the B-spline basis functions:

$$T_{ij} = B_j^{n,d}(\phi_i), \quad ; \quad \phi_i = \frac{2\pi}{p} i. \quad (10)$$

Given an estimate $\hat{\theta}$ of θ , we can obtain an estimate of $\mathbf{r} = \{r(\phi_j)\}_{j=1}^p$, the p samples of the curve corresponding to θ , through the linear matrix operation

$$\hat{\mathbf{r}} = (r_1, \dots, r_p) = T \hat{\theta}. \quad (11)$$

The bound on the covariance of the estimate $\hat{\mathbf{r}} = T \hat{\theta}$ is then

$$\text{cov}(T \hat{\theta}) \geq T [F_{\theta}]^{-1} T^T. \quad (12)$$

We will refer to the bound in (6) as the *coefficient bound* and that of equation (12) as the *radial bound*. It should be noted that our “radial bound” is *not* the CR bound on any unbiased estimator of $\mathbf{r}$; rather, it applies only to estimators of the form $T \hat{\theta}$.

2.4. Estimation of B-Spline Coefficients

The maximum likelihood estimator of the spline coefficients is identical to the nonlinear least squares estimator

$$\begin{aligned} \hat{\theta} &= \arg \max_{\theta} \ln f(Y; \theta) \\ &= \arg \min_{\theta} \|Y(x, y) - \lambda_{\theta}^s(x, y)\|^2 \end{aligned} \quad (13)$$

where $\|\bullet\|$ is the L_2 norm on the plane.

We generalize equation (1) to allow for slow variations of intensity within image regions by introducing spatially variant pixel intensities, λ , as unknown parameters. This presents us with an ill-conditioned deconvolution problem – one for which unregularized ML estimation would yield unacceptably high variance. For this reason, we introduce penalty function terms which promote smoothness and decrease variance of both the pixel intensity and boundary estimates. To this end, we consider the penalized maximum likelihood estimator

$$\begin{aligned} \hat{\theta}, \hat{\lambda} &= \arg \max_{\theta, \lambda} \ln f(Y; \lambda) - \beta V_{\theta}(\lambda) - \zeta U(\theta) \\ &= \arg \min_{\theta, \lambda} \sum_i (Y_i - \lambda_i)^2 + \beta V_{\theta}(\lambda) + \zeta U(\theta) \end{aligned}$$

where $V_{\theta}(\lambda)$ and $U(\theta)$ are the quadratic penalty terms below:

$$\begin{aligned} V_{\theta}(\lambda) &= \sum_j \sum_{i \in N_j} \omega_{ij}(\theta) (\lambda_j - \lambda_i)^2 \\ U(\theta) &= \sum_k \sum_{i \in N_k} \rho_{ik}(\theta_k - \theta_i)^2. \end{aligned} \quad (14)$$

Note that the sum over j in the upper equation traverses the entire image λ , while the lower sum over k covers the set of knots. The function $V_{\theta}(\lambda)$ is chosen to encourage neighboring pixels to be similar within and outside of, but not across, the boundary described by θ . This is implicit in the way in which we define $\{\omega_{ij}(\theta)\}_{i \in N_j}$, the set of weights that relate pixel j to the pixels in its neighborhood N_j . For a given pixel j , we choose $\omega_{ij}(\theta)$ to be the fraction of neighbor pixel i that is of the same type (‘interior’ or ‘exterior’, as defined by θ) as pixel j . Similarly, $\rho_{ik} = 1$ when knots i and k are adjacent and 0 otherwise, thereby penalizing variation among neighboring coefficients.

2.5. Simulations

For a given equiangular B-Spline model, 1000 radial samples of the continuous curve $r(\phi)$ were used to produce an 32×32 image $\lambda_{\theta}(x, y)$. This simulated NMR image was then smoothed by convolution with a discrete Gaussian surface of variance σ_s , and Gaussian noise was added. The minimization of the given objective function was performed using a coordinate descent algorithm on θ , and each variance measurement was computed as the sample variance of thirty independent parameter estimates.

Figure 3 shows the performance of the maximum likelihood estimator of θ vs. the coefficient CR bound as we vary σ_n , the noise standard deviation. In this case the true object boundary was a circle of radius 8 pixels, with $\lambda^I = 10$, $\lambda^E = 2$, and $\sigma_s = .01$ pixel. The spline model used was quadratic and employed six equiangularly spaced knots.




Figure 3: ML Estimator Performance vs. Noise Stdev.

At first glance, these results seem rather disappointing. Although the ML standard deviations nearly achieve the bound for very small values of σ_n , they more than double it as we move away from the high SNR region. The PML estimator standard deviations (shown in Figure 4 for $\beta = .005$ and $\zeta = .0002$) strayed even further from the bound, which we would expect due to the uncertainty added by the unknown pixel intensities. In the future, we expect to implement both a




Figure 4: PML Estimator Performance vs. Noise Stdev.

hybrid bound that will be tighter in the large error region, as well as modify the bound above for the case of unknown pixel intensities.

At least two facts, however, need to be kept in mind when considering these numbers. First, the bounds above have been calculated to high precision as the solutions to integrals — in deriving these integrals, the image and boundary were assumed continuous, rather than discrete. Also, note that the standard deviation values given by the bound are all much less than a pixel. This is relevant because under the current implementation of both algorithms, grey scale values for boundary pixels are produced by calculating the fraction of pixel area subtended by the boundary; as a result, boundaries that *differ* from each another over a given pixel — but subtend equal areas — will be indecipherable from one another. Thus for a given curve, variability of estimates will increase as the algorithms converge to many separate boundaries (i.e. coefficient estimates) that actually differ sub-pixelly.

3. APPLICATION TO TOMOGRAPHY

Many tomographic imaging systems such as PET and SPECT produce functional images (of a substance density rather than an anatomical structure) which, due to scatter, attenuation, and limited radiation dosage suffer from low SNR and large amounts of bias. The latter effect could be substantially reduced through the use concisely parameterized anatomical side information. Of the several works that have focused on using anatomical boundaries from NMR images for this purpose, many employ line-site or region labels to model the NMR boundary [4],[7]. Although the results of the previous section show clearly that more work needs to be done to determine the accuracy with which NMR boundaries modeled using splines can be estimated, the spline approach detailed above may eventually prove superior to line-site and label methods for a variety of reasons. First, the latter strategies typically have a pixel/parameter ratio of 1-1, whereas the spline strat-

egy can parsimoniously describe most common tumor and organ boundaries using less than 20 knots. Second, line-site and labeling methods assure only local continuity, which can result not only in missing edges and the resultant oversmoothing of interior regions, but also in problems estimating total region quantities such as radiation dose. Splines, on the other hand, offer an inherently continuous, closed representation of the region boundary; therefore, in situations where one is relatively convinced of the continuity of a boundary, splines would be a natural parameterization. Finally, spline parameters are continuous rather than discrete — this allows the use of stationary points to determine optimality if we choose to perform joint estimation of NMR and functional boundaries.

4. REFERENCES

- [1] T.N.E Greville, editor. *Theory and Applications of Spline Functions*. Academic Press, 1969.
- [2] J. H. Ahlberg. *The Theory of Splines and Their Applications*. Academic Press, 1967.
- [3] W. Ouyang. Incorporation of correlated structural images in pet image reconstruction, 1992. Submitted to Transactions on Medical Imaging.
- [4] Gene Gindi, Mindy Lee, Anand Rangarajan, and George Zubal. Bayesian reconstruction of functional images using anatomical information as priors. *IEEE Transactions on Medical Imaging*, 12(4), Dec 1993.
- [5] V E Johnson. A model for segmentation and analysis of noisy images. *Journal of the American Statistical Association*, 89(425):230-241, March 1994.
- [6] W.A. Edelstein. Intrinsic signal to noise ratio in nmr imaging. *Magnetic Resonance in Medicine*, 1986.
- [7] V E Johnson. A model for segmentation and analysis of noisy images. Technical Report 91-A15, Duke, 1992.