Published on *Journal of the Royal Statistical Society Ser. C, Applied Statistics* (2009), 58 (3), 285-306.
Publisher's version available at http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9876.2008.00653.x/pdf

Efficient estimation of 3-dimensional curves and their derivatives by free knot regression splines, applied to the analysis of inner carotid artery centerlines *

L.M. Sangalli, P. Secchi, S. Vantini
Politecnico di Milano, Italy
A. Veneziani
Emory University, USA

Abstract

We deal with the problem of efficiently estimating a 3D curve and its derivatives, starting from a discrete and noisy observation of the curve. This problem is now arising in many applicative contexts, thanks to the advent of devices that provide 3D images and measures, such as 3D scanners in medical diagnostics. Our research, in particular, stems from the need for accurate estimation of the curvature of an artery, from image reconstructions of 3D angiographies. This need has emerged within AneuRisk Project, a scientific endeavor which aims at investigating the role of vessel morphology, blood fluid-dynamics, and biomechanical properties of the vascular wall, on the pathogenesis of cerebral aneurysms.

We develop a regression technique that exploits free knot splines in a novel setting, to estimate 3D curves and their derivatives. We thoroughly compare this technique to a classical regression method, local polynomial smoothing, showing that 3D free knot regression splines yield more accurate and efficient estimates.

Keywords and phrases. Functional data analysis, smooth curve fitting, free knot regression splines, local polynomial smoothing.

1 Introduction

With the advent of devices that provide (2D or) 3D images and measures, the problem of curve fitting in more than one dimension is progressively becoming of interest. As in the simpler 1D case, this problem requires special care when the curve estimate

* *Address for correspondance*: Piercesare Secchi, MOX, Dipartimento di Matematica, Politecnico di Milano, P.za Leonardo da Vinci 32, 20133 Milano, Italy.
E-mail: piercesare.secchi@polimi.it

is asked, not only to provide a good smoothing of the data, but also to reflect the features of the curve represented by its derivatives. In some cases, some of the curve derivatives, or their functions, are themselves objects of analysis; moreover, estimates of curve derivatives may be of help for further processing of the data, such as in curve alignment procedures. See e.g. the many illuminating discussions and examples given in Ramsay and Silverman (2005).

Our research, in particular, stems from the quest for accurate estimation of centerlines of blood vessels and their curvature functions. We show how to efficiently estimate these 3D curves and their derivatives (and thus their curvature functions), starting from discrete and noisy observations of the curves. In particular, we develop a technique that exploits in a novel setting *free knot regression splines*. These are regression splines where the number and position of knots are not fixed in advance, but chosen in a way to minimize a penalized sum of squared errors criterion; see e.g. Friedman (1991), Stone, Hansen, Kooperberg, and Truong (1997), Luo and Wahba (1997), Zhou and Shen (2001), and Mao and Zhao (2003). We show that the proposed technique has many comparative advantages over a classical regression method, local polynomial smoothing. First of all, 3D free knot regression splines are more efficient than local polynomials, in the sense that they are able to attain lower approximation errors fitting more parsimonious models. Moreover, they better capture the salient features of the curve, providing better estimates of curve derivatives. Finally, the proposed technique leads to data dimension reduction, a fundamental issue for highly dimensional datasets, such as those resulting from image reconstructions.

We illustrate our methodology directly on the applied problem that has motivated our research. This is described in detail in Section 2. The same regression technique can be easily applied to other curve fitting problems in 2D and 3D. In Section 3 we present our dataset and the main features of the statistical model. In Section 4 we recall the essentials of free knot regression splines for the case of functions with one dependent variable, and then describe the extension to the 3D case. In Section 5 we illustrate the results and discuss the properties of the obtained estimates. Section 6 is devoted to an extensive comparison between the proposed technique and local polynomial smoothing. Some conclusive considerations are drawn in Section 7.

2 Motivating applied problem: aneurysms and AneuRisk

Cerebral aneurysms are lesions of cerebral vessels characterized by a bulge of the vessel wall. The origin of this pathology is still unclear. Possible explanations discussed in the medical literature focus on the interactions between biomechanical properties of artery walls and hemodynamic factors, such as wall shear stress and pressure; the hemodynamics is in turn strictly dependent on vascular geometry. See e.g. Hoi et al. (2004), Hassan et al. (2005), Castro, Putman, and Cebral (2006). The study of these

interactions is the main goal of AneuRisk Project, a scientific endeavor which joins researchers of different scientific fields, ranging from neurosurgery and neuroradiology to statistics, numerical analysis and bio-engineering.

This paper, in particular, is devoted to the accurate estimation of vessel centerlines and their curvature functions. It is well known in fact that vessel curvature, i.e. the curvature of the vessel centerline, strongly affects the local hemodynamics. See for instance Caro, Doorly, and Tarnakasky (1996) and Chandran (1993). In general, the relevance of the curvature of a pipe, on the inner fluid motion, has been extensively investigated since the pioneering works by Dean. See e.g. Dean (1927a), Dean (1927b), Smith (1976), Berger, Talbot, and Yao (1983), and Chandran (1993). In particular, an adimensional index called *Dean number* has been introduced, as an extension of Reynolds number, in order to quantify fluid stability in a circularly curved pipe. Let R be the radius of the pipe section, and let R_{curv} be the radius of curvature of the pipe centerline; moreover, denote by μ the viscosity of the fluid at hand, by ρ its density, and by U_0 its mean velocity. The Dean number is

$$D = \frac{2\,\rho\,U_0\,R}{\mu}\left(\frac{R}{R_{curv}}\right)^{1/2}.$$

Blood dynamics can be really different for different values of D. This explains why accurate estimation of the profile of the curvature radius of the arteries (or equivalently, the profile of the curvature $curv = 1/R_{curv}$) is of fundamental importance for the analysis of hemodynamics, and its possible consequences on aneurysm pathology.

The dataset of AneuRisk Project is based on a set of 3D angiographic images taken from 65 patients at Niguarda Ca' Granda Hospital (Milan), suspected to be affected by cerebral aneurysms. Some of these patients have in fact an aneurysm on the Internal Carotid Artery (ICA), other patients have an aneurysm downstream the ICA, in the so-called Willis circle (see Ustun (2005)); finally, a few patients are healthy. The analyses conducted so far within the AneuRisk Project focus on the ICA, which is clearly recognizable in each of the 65 angiographies. This artery features a great curvature variability, not only within different parts of the vessel for the same patient, but also between different patients. Starting from the 3D array of grey-scaled pixels generated by the angiography (with lighter pixels showing presence of flowing blood), the artery lumen, i.e. the volume occupied by flowing blood, is identified by a reconstruction algorithm coded in the *Vascular Modeling ToolKit* (VMTK). See Antiga, Ene-Iordache, and Remuzzi (2003), and Piccinelli et al. (2007). The outcome of this process yields a set of points in a three dimensional space, that are computed as the centers of maximal spheres inscribed in the artery lumen. This set of points defines the vessel centerline. Moreover, the reconstruction algorithm provides the radius of lumen sections, computed as the radius of maximal inscribed spheres. Figure 1 shows the draw of the reconstruction of an ICA with its centerline (patient 1). Due to measurement and reconstruction errors, reconstructed centerlines may be quite wiggly

Figure 1: 3D image of an internal carotid artery with an aneurysm [patient 1].

and thus need to be smoothed, and so do the estimates of their derivatives, in order to obtain sensible estimates of their curvature functions. Sangalli, Secchi, Vantini, and Veneziani (2008b) obtained smooth pointwise estimates of vessel centerlines by means of local polynomials, and carried out exploratory analyses of vessel curvature and radius profiles which supported the existence of a strong relationship between vessel geometry and aneurysm location. Using the more accurate estimates described in the present paper, which are based on free knot spline regression, we can confirm this conclusion. See Sangalli, Secchi, and Vantini (2008a). This regression technique has also the advantage of providing to the Project highly accurate estimates of each individual ICA centerline and curvature function. These will be used in the study of the corresponding reconstructed ICA geometry, in order to efficiently evaluate the role of vessel curvature on the hemodynamics.

3 Data and model

For every patient i in our dataset ($i = 1, \ldots, 65$), VMTK reconstruction of ICA centerline is a set of points in $\mathbb{R}^3$, $\{(x_{ij}, y_{ij}, z_{ij}) : j = 1, 2, \ldots, n_i\}$, where x, y and z denote the three space coordinates of each point. Points are ordered moving downward along the ICA, from the point closest to its terminal bifurcation (detected by VMTK) towards the proximal districts, i.e. aorta and heart. The reason for this choice is that the terminal bifurcation of the ICA is present in each angiography, even if the portion of ICA captured by the angiography varies from patient to patient (depending on where the angiographic image has been centered). For each patient i, we associate the set of space coordinates with an index set $\{s_{ij} : j = 1, 2, \ldots, n_i\}$, which measures an approximate distance along the ICA. More precisely, $-s_{i1}$ is the distance of the point (x_{i1}, y_{i1}, z_{i1}) from the terminal bifurcation of the ICA (as determined by VMTK), and

$$s_{ij} - s_{ij-1} = -\sqrt{(x_{ij} - x_{ij-1})^2 + (y_{ij} - y_{ij-1})^2 + (z_{ij} - z_{ij-1})^2}\,, \quad \text{for } j = 2, \ldots, n_i.$$

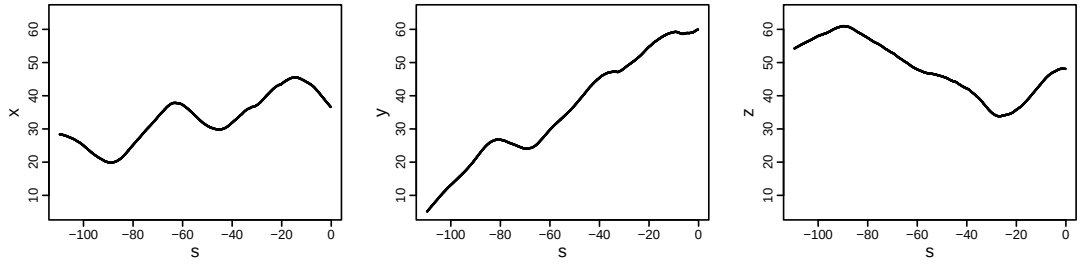

Figure 2: Reconstructed space coordinates of ICA centerline for patient 1, (x_{1j}, y_{1j}, z_{1j}), versus the abscissa parameter s_{1j}, for $j = 1, \ldots, n_1$ ($n_1 = 1350$).

The conventional negative sign highlights that we are moving upstream, i.e. in opposite direction with respect to blood flow. Figure 2, for instance, displays the reconstructed space coordinates of ICA centerline for patient 1, (x_{1j}, y_{1j}, z_{1j}), versus the abscissa parameter s_{1j}, for $j = 1, \ldots, n_1$ ($n_1 = 1350$). The number n_i of data points available for each patient ranges from 350 to 1380, and is almost perfectly correlated to the approximate length $|s_{i\,n_i} - s_{i\,1}|$ of the reconstructed centerlines (correlation coefficient=0.999), which in turn varies from 27.219mm to 110.136mm. In other words, the grid density of the 65 reconstructions is the same, even if the 65 grids are different. The average step of these grids is 0.079mm.

Our task is to estimate, for each patient i, the true ICA centerline $\mathbf{c}_i(s) = (x_i(s), y_i(s), z_i(s))$ and its curvature function, starting from the noisy reconstruction $\{(x_{ij}, y_{ij}, z_{ij}) : j = 1, 2, \ldots, n_i\}$. For each patient i we carry out an independent estimation process (from now on we drop the subscript i), considering the statistical model

$$\begin{aligned} (x_j, y_j, z_j) = \mathbf{c}(s_j) + \mathbf{e}_j\,, \qquad & j = 1, \ldots, n\,; \\ \mathbf{e}_j = (e_j^{[x]}, e_j^{[y]}, e_j^{[z]})\,, \qquad & \text{with } \mathbf{e}_j \text{ and } \mathbf{e}_{j'} \text{ independent for } j \neq j', \\ & E(\mathbf{e}_j) = \mathbf{0}\,,\ Var(\mathbf{e}_j) = \sigma^2 \mathbf{I}\,. \end{aligned} \tag{1}$$

This model generalises to the 3D case the models considered for 1D curve fitting problems; see e.g. Zhou and Shen (2001), who present the methodology that we extend here to the multidimensional case, or Zhang (2003), who considers various smoothing methods. The value of σ^2, in the variance matrix of the error term, is unknown, but assumed to be the same for each patient i, for $i = 1, \ldots, 65$. This assumption is justified by the fact that both the machine used to take the 3D-angiographies and the reconstruction algorithm are the same for each patient. Denoting by $\times$ the vector product, the curvature of $\mathbf{c}(s) = (x(s), y(s), z(s))$ is defined by

$$curv_{\mathbf{c}}(s) = \frac{|\mathbf{c}'(s) \times \mathbf{c}''(s)|}{|\mathbf{c}'(s)|^3} =$$

$$\frac{\sqrt{\big(x'(s)y''(s) - x''(s)y'(s)\big)^2 + \big(y'(s)z''(s) - y''(s)z'(s)\big)^2 + \big(z'(s)x''(s) - z''(s)x'(s)\big)^2}}{\big((x'(s))^2 + (y'(s))^2 + (z'(s))^2\big)^{3/2}}$$

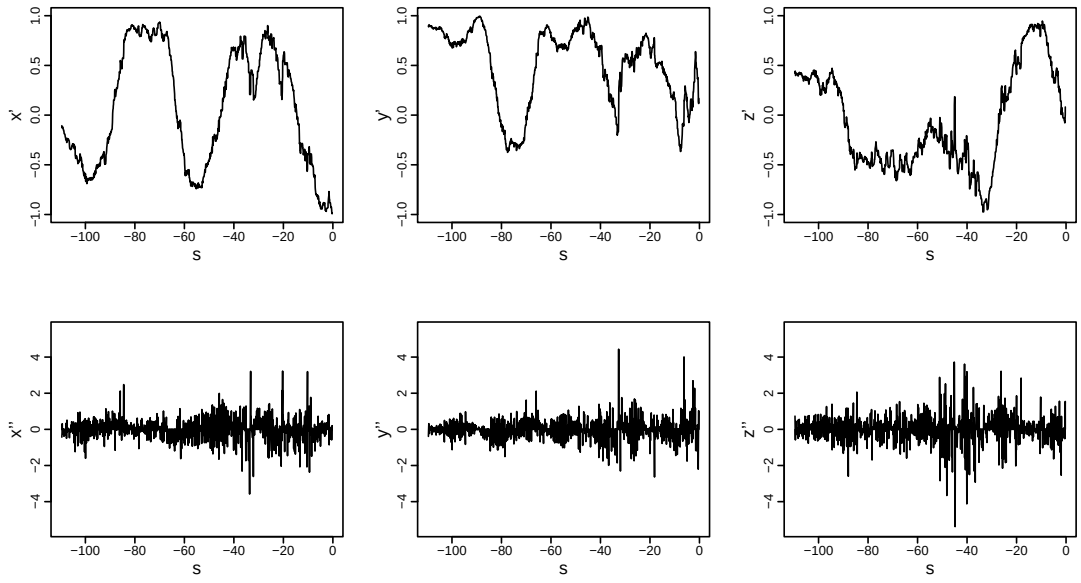

Figure 3: First and second differences (top and bottom, respectively) of reconstructed space coordinates of ICA centerline for patient 1.

where $\mathbf{c}'(s) = (x'(s), y'(s), z'(s))$ and $\mathbf{c}''(s) = (x''(s), y''(s), z''(s))$ are the first and second derivatives of $\mathbf{c}$. Rough pointwise estimates of first derivatives are given by the *first central differences*

$$Dx(s_j) = \frac{x_{j+1} - x_{j-1}}{s_{j+1} - s_{j-1}}, \qquad Dy(s_j) = \frac{y_{j+1} - y_{j-1}}{s_{j+1} - s_{j-1}}, \qquad Dz(s_j) = \frac{z_{j+1} - z_{j-1}}{s_{j+1} - s_{j-1}},$$

for $j = 2, \ldots, n-1$; the *second central differences*

$$D^2x(s_j) = 2\frac{\frac{x_{j+1}-x_j}{s_{j+1}-s_j} - \frac{x_j-x_{j-1}}{s_j-s_{j-1}}}{s_{j+1} - s_{j-1}}, \qquad D^2y(s_j) = 2\frac{\frac{y_{j+1}-y_j}{s_{j+1}-s_j} - \frac{y_j-y_{j-1}}{s_j-s_{j-1}}}{s_{j+1} - s_{j-1}},$$

$$D^2z(s_j) = 2\frac{\frac{z_{j+1}-z_j}{s_{j+1}-s_j} - \frac{z_j-z_{j-1}}{s_j-s_{j-1}}}{s_{j+1} - s_{j-1}}$$

are instead too noisy to be of any use. See for instance Figure 3, which displays first and second central differences of reconstructed space coordinates of ICA centerline for patient 1.

We propose a regression technique, for the estimation of the true centerline $\mathbf{c}$, based on free knot regression splines, i.e. regression splines where the number and position of knots are not fixed in advance, but chosen in a way to minimize a penalized sum of squared errors criterion. Since our data are 3D, the idea is to fit simultaneously the three space coordinates of the centerline $\big(x(s), y(s), z(s)\big)$, looking for the optimal spline knots along the abscissa parameter s. Estimates of $\mathbf{c}'$ and $\mathbf{c}''$, and hence of $curv_{\mathbf{c}}$, are thus obtained by differentiation of the fitted regression spline.

4 3D free knot regression splines

Consider an interval $[a, b]$, split into subintervals by knots $k_1, \ldots, k_{n_k}$ (with $a < k_1 < \ldots < k_{n_k} < b$). Assume, for the moment, that the number and position of knots are given. An order-m spline over $[a, b]$, with knot vector $\mathbf{k} = (k_1, \ldots, k_{n_k})$, is a piecewise-polynomial of degree $m-1$, with continuous derivatives of order $m-2$ at the knots. The set of all order-m splines over $[a, b]$, with knot vector $\mathbf{k}$, forms a vector space. The dimension of such space is $m + n_k$. This is given by the balance between the number of parameters needed to identify $(n_k + 1)$ polynomials of order m (one polynomial over each of the subintervals), and the continuity constraints, prescribing $m-1$ equations at each of the n_k knots: $m(n_k + 1) - n_k(m-1) = m + n_k$. A computationally convenient basis system for this space is the b-spline basis system $\{b^{[\mathbf{k}]}_{r,m}(s) : r = 1, \ldots, m + n_k\}$. Any order-$m$ spline over $[a, b]$ with knot vector $\mathbf{k}$ can thus be represented by the expansion $f(s) = \sum_{r=1}^{m+n_k} \lambda_r b^{[\mathbf{k}]}_{r,m}(s)$, for some coefficient vector $\lambda = \left(\lambda_1, \ldots, \lambda_{m+n_k}\right)^T$. See de Boor (1978) for a systematic introduction to splines. Note that derivatives of splines are still splines (of appropriate order), with the same knot vector and coefficients directly computed from the coefficients of the original spline. In particular, with the b-spline basis system, denoting by $\mathbf{v}$ the $(2m + n_k)$-vector with entries

$$v_1 = \ldots = v_m = a, \quad v_{m+j} = k_j \text{ for } j = 1, \ldots, n_k, \quad v_{m+1+n_k} = \ldots = v_{2m+n_k} = b,$$

the first and second derivatives of $f(s)$ are given by $f'(s) = \sum_{r=1}^{m-1+n_k} \lambda^{[1]}_r b^{[\mathbf{k}]}_{r,m-1}(s)$ and $f''(s) = \sum_{r=1}^{m-2+n_k} \lambda^{[2]}_r b^{[\mathbf{k}]}_{r,m-2}(s)$, where

$$\lambda^{[1]}_r = \frac{\lambda_{r+1} - \lambda_r}{v_{m+r} - v_{r+1}}, \qquad \text{for } r = 1, \ldots, m-1+n_k,$$

and

$$\lambda^{[2]}_r = \frac{\lambda^{[1]}_{r+1} - \lambda^{[1]}_r}{v_{m+r} - v_{r+2}}, \qquad \text{for } r = 1, \ldots, m-2+n_k,$$

for $m \geq 2$ and $m \geq 3$ respectively.

Now consider the problem of estimating the function $f : \mathbb{R} \to \mathbb{R}$ in the model

$$w_j = f(s_j) + e_j, \qquad j = 1, \ldots, n, \text{ with } e_j \text{ and } e_{j'} \text{ independent for } j \neq j',$$
$$E(e_j) = 0, \; Var(e_j) = \sigma^2,$$

by means of an order-m spline with knot vector $\mathbf{k}$. The regression spline is the spline

$$\hat{f}(s) = \sum_{r=1}^{m+n_k} \hat{\lambda}_r b^{[\mathbf{k}]}_{r,m}(s)$$

where $\hat{\lambda} = \left(\hat{\lambda}_1, \ldots, \hat{\lambda}_{m+n_k}\right)^T$ minimizes the sum of squared errors:

$$SSE(\lambda) = \sum_{j=1}^{n} \left(w_j - \sum_{r=1}^{m+n_k} \lambda_r b^{[\mathbf{k}]}_{r,m}(s_j)\right)^2. \tag{2}$$

In matrix notation, writing $\{B_m^{[\mathbf{k}]}\}_{j,r} = b_{r,m}^{[\mathbf{k}]}(s_j)$, $\mathbf{w} = (w_1, \ldots, w_n)^T$ and $\hat{\mathbf{w}} = (\hat{f}(s_1), \ldots, \hat{f}(s_n))^T$,

$$\begin{aligned} SSE(\lambda) &= (\mathbf{w} - B_m^{[\mathbf{k}]}\lambda)^T(\mathbf{w} - B_m^{[\mathbf{k}]}\lambda)\,, \\ \hat{\lambda} &= \left(B_m^{[\mathbf{k}]^T} B_m^{[\mathbf{k}]}\right)^{-1} B_m^{[\mathbf{k}]^T}\mathbf{w}\,, \\ \hat{\mathbf{w}} &= B_m^{[\mathbf{k}]}\left(B_m^{[\mathbf{k}]^T} B_m^{[\mathbf{k}]}\right)^{-1} B_m^{[\mathbf{k}]^T}\mathbf{w} = H\,\mathbf{w}\,, \end{aligned}$$

where $H = B_m^{[\mathbf{k}]}\left(B_m^{[\mathbf{k}]^T} B_m^{[\mathbf{k}]}\right)^{-1} B_m^{[\mathbf{k}]^T}$ is an orthogonal projection matrix, with $tr(H) = m + n_k$. Thus, $\hat{\mathbf{w}}$ is obtained as orthogonal projection of $\mathbf{w}$ on the space generated by the vectors of basis evaluations $\left(b_{r,m}^{[\mathbf{k}]}(s_1), \ldots, b_{r,m}^{[\mathbf{k}]}(s_n)\right)^T$, for $r = 1, \ldots, m + n_k$.

The order of the spline is usually chosen according to the problem at hand and the scope of the analysis. Commonly used splines are of order 4, corresponding to piecewise cubic polynomials. They provide a smooth estimate of the function f, for which knot-discontinuity is not visible to human eye, as it is often claimed. Moreover, order-4 splines also provide smooth estimates of the first derivative of f. When higher order derivatives are also of interest, it is usually convenient to use order $p + 3$ (or at least $p + 2$) splines, where p is the highest derivative order that has to be estimated.

To improve spline estimation, number and position of knots can be selected by minimizing a suitable error functional, which takes into account the dimension of the model being fitted. A possible model selector, in this case, is provided by the so-called *Stein's unbiased risk estimate* (see Stein (1981))

$$\begin{aligned} pSSE(\lambda, \mathbf{k}) &= \sum_{j=1}^{n}\left(w_j - \sum_{r=1}^{m+n_k} \lambda_r b_{r,m}^{[\mathbf{k}]}(s_j)\right)^2 + C\hat{\sigma}^2(m + n_k) \\ &= (\mathbf{w} - B_m^{[\mathbf{k}]}\lambda)^T(\mathbf{w} - B_m^{[\mathbf{k}]}\lambda) + C\hat{\sigma}^2(m + n_k) \end{aligned} \tag{3}$$

where C is a penalization constant, and $\hat{\sigma}^2$ is an estimate of error variance. Regression splines where the number and position of knots are not fixed in advance, but chosen by a data-driven criterion such as (3), are called free knot regression splines. Note that, once the optimal knot vector $\hat{\mathbf{k}}$ has been selected, and thus the model has been chosen, the vector of fitted values $\hat{\mathbf{w}}$ is again obtained as orthogonal projection of $\mathbf{w}$ on the space generated by the vectors of basis evaluations, according to $\hat{\mathbf{w}} = H\,\mathbf{w}$, where $H = B_m^{[\hat{\mathbf{k}}]}\left(B_m^{[\hat{\mathbf{k}}]^T} B_m^{[\hat{\mathbf{k}}]}\right)^{-1} B_m^{[\hat{\mathbf{k}}]^T}$. Algorithms for the search of the optimal knot vector, for the regression technique described above or other free knot spline procedures, have been proposed e.g. by Friedman (1991), Stone et al. (1997), Luo and Wahba (1997), and more recently by Zhou and Shen (2001), and Mao and Zhao (2003). In particular, the algorithm developed by Zhou and Shen (2001), by including knot relocation moves, strongly improved the previous stepwise forward/backward knot selection procedures, which suffered from knot confounding problems. The choice of the smoothness parameter C, in the model selector (3), is of paramount importance. High values of C yield more parsimonious models, i.e. splines with fewer knots, but with higher sum

of squared errors, and vice versa. Several researchers suggest particular values of the smoothness parameter in the model selector they are using. For example, Zhou and Shen (2001) note that $C = 2$ in (3) works well in simulations with some smooth and nonsmooth functions; Friedman (1991) finds that an effective value of the smoothness parameter in his MARS is 3; Luo and Wahba (1997) instead set it equal to 1.2 for their Hybrid Adaptive Splines. In practice, the choice of the smoothness parameter is made according to qualitative considerations, along a "Goldilocks approach" where the smoothness parameter is selected in order to neither under-smooth nor over-smooth the data, but lead to a "just right" degree of smoothness for the data under analysis.

4.1 Extension to the 3D case

Let us now come to the 3D problem described by model (1). The idea is to estimate the true centerline $\mathbf{c}(s)$ by simultaneously fitting the three coordinate functions, $x(s)$, $y(s)$ and $z(s)$, with three order-m splines having the same knot vector $\hat{\mathbf{k}}$, with knots chosen along the abscissa parameter s:

$$\hat{x}(s) = \sum_{r=1}^{m+n_{\hat{k}}} \hat{\lambda}_r^{[x]} b_{r,m}^{[\hat{\mathbf{k}}]}(s), \qquad \hat{y}(s) = \sum_{r=1}^{m+n_{\hat{k}}} \hat{\lambda}_r^{[y]} b_{r,m}^{[\hat{\mathbf{k}}]}(s), \qquad \hat{z}(s) = \sum_{r=1}^{m+n_{\hat{k}}} \hat{\lambda}_r^{[z]} b_{r,m}^{[\hat{\mathbf{k}}]}(s).$$

As before, denote by $B_m^{[\mathbf{k}]}$ the $(n \times (m + n_k))$-matrix of basis evaluations. Moreover, denote by W the $(n \times 3)$-matrix of observed values, $W = \big[\mathbf{x}\big|\mathbf{y}\big|\mathbf{z}\big]$, where $\mathbf{x} = (x_1, \ldots, x_n)^T$, $\mathbf{y} = (y_1, \ldots, y_n)^T$, $\mathbf{z} = (z_1, \ldots, z_n)^T$, and by $\hat{W}$ the analogous matrix of fitted values $\hat{W} = \big[\hat{\mathbf{x}}\big|\hat{\mathbf{y}}\big|\hat{\mathbf{z}}\big]$, where $\hat{\mathbf{x}} = \big(\hat{x}(s_1), \ldots, \hat{x}(s_n)\big)^T$, $\hat{\mathbf{y}} = \big(\hat{y}(s_1), \ldots, \hat{y}(s_n)\big)^T$, $\hat{\mathbf{z}} = \big(\hat{z}(s_1), \ldots, \hat{z}(s_n)\big)^T$. Finally, denote by Λ the $((m + n_k) \times 3)$-matrix of λ coefficients $\Lambda = \big[\lambda^{[x]}\big|\lambda^{[y]}\big|\lambda^{[z]}\big]$. For a fixed knot vector $\mathbf{k}$, the natural generalization of least squared errors criterion (2), for the 3D model (1), is given by

$$\begin{aligned} SSE3D(\Lambda) &= tr\big[(W - B_m^{[\mathbf{k}]}\Lambda)(W - B_m^{[\mathbf{k}]}\Lambda)^T\big] \\ &= \sum_{j=1}^{n} \big(x_j - \sum_{r=1}^{m+n_k} \lambda_r^{[x]} b_{r,m}^{[\mathbf{k}]}(s_j)\big)^2 + \sum_{j=1}^{n} \big(y_j - \sum_{r=1}^{m+n_k} \lambda_r^{[y]} b_{r,m}^{[\mathbf{k}]}(s_j)\big)^2 \\ &\quad + \sum_{j=1}^{n} \big(z_j - \sum_{r=1}^{m+n_k} \lambda_r^{[z]} b_{r,m}^{[\mathbf{k}]}(s_j)\big)^2, \end{aligned}$$

which leads to the least squared estimate $\hat{W} = B_m^{[\mathbf{k}]}\hat{\Lambda}$, where

$$\hat{\Lambda} = \big({B_m^{[\mathbf{k}]}}^T B_m^{[\mathbf{k}]}\big)^{-1} {B_m^{[\mathbf{k}]}}^T W$$

Thus, for the search of the optimal $(\hat{\Lambda}, \hat{\mathbf{k}})$, we will use the following model selector:

$$\begin{aligned} & pSSE3D(\Lambda, \mathbf{k}) \\ & = tr\big[(W - B_m^{[\mathbf{k}]}\Lambda)(W - B_m^{[\mathbf{k}]}\Lambda)^T\big] + \mathcal{C}(m + n_k) \\ & = \sum_{j=1}^{n} \big(x_j - \sum_{r=1}^{m+n_k} \lambda_r^{[x]} b_{r,m}^{[\mathbf{k}]}(s_j)\big)^2 + \sum_{j=1}^{n} \big(y_j - \sum_{r=1}^{m+n_k} \lambda_r^{[y]} b_{r,m}^{[\mathbf{k}]}(s_j)\big)^2 \\ & \quad + \sum_{j=1}^{n} \big(z_j - \sum_{r=1}^{m+n_k} \lambda_r^{[z]} b_{r,m}^{[\mathbf{k}]}(s_j)\big)^2 + \mathcal{C}(m + n_k)\,. \end{aligned} \tag{4}$$

Note that the penalization term in (4) does not explicitly take into account error variance, as was instead the case in (3). Hence, the penalization constant $\mathcal{C}$ in (4) cannot be directly compared with the penalization constant C in (3). This is related to the specific problem we deal with. Recall that the value of σ in the variance structure of the error term is assumed to be the same for each patient. Since we are here mainly interested in equally smoothing the ICA centerlines of the 65 patients, we will simply focus on choosing the best value of $\mathcal{C}$ for the overall dataset, without bothering with finding a sensible a priori estimate of σ. We discuss the choice of the penalization constant $\mathcal{C}$ in Section 5.1.

It should also be noted that, as in the 1D case, the matrix of fitted values $\hat{W}$ is obtained as orthogonal projection of the columns of the matrix of observed values W on the space generated by the vectors of basis evaluations $\big(b_{r,m}^{[\hat{\mathbf{k}}]}(s_1), \ldots, b_{r,m}^{[\hat{\mathbf{k}}]}(s_n)\big)^T$, for $r = 1 \ldots, m + n_{\hat{k}}$,

$$\hat{W} = B_m^{[\hat{\mathbf{k}}]}\big({B_m^{[\hat{\mathbf{k}}]}}^T B_m^{[\hat{\mathbf{k}}]}\big)^{-1} {B_m^{[\hat{\mathbf{k}}]}}^T W = H\,W.$$

In particular, the 3D spline estimator, like the 1D spline estimator, is a so-called linear estimator, i.e. there exists a linear operator S, independent of W, such that the fitted values $\hat{W}$ are obtained according to $\hat{W} = S\,W$. For these estimators it is common to take the trace of the linear operator S, or alternatively $tr(S^T S)$ or $tr(2S - S^T S)$, as Degrees of Freedom (DF) of the model. See e.g. Buja, Hastie, and Tibshirani (1989) and Hastie and Tibshirani (1990), who first introduced this notion. As commented in Ramsay and Silverman (2005), using the DF provides a uniform approach to compare different smoothing methods. We will thus consider, as DF of the spline estimator, the trace of the orthogonal projection matrix H, i.e. $m + n_{\hat{k}}$ (note that when the linear operator is an orthogonal projection matrix, as in this case, the three alternative definitions of DF coincide).

If the three coordinates in the error term were assumed to be correlated, and thus $Var(\mathbf{e}_j) = \sigma^2\mathbf{I}$ in (1) was changed in $Var(\mathbf{e}_j) = \Sigma$, for some symmetric positive-definite matrix Σ, then the natural generalization of least squared errors criterion (2) would become

$$SSE3D(\Lambda) = tr\big[(W - B_m^{[\mathbf{k}]}\Lambda)\Sigma^{-1}(W - B_m^{[\mathbf{k}]}\Lambda)^T\big]. \tag{5}$$

It is easy to show that

$$\begin{aligned} &argmin_{(\Lambda)}\{tr\big[(W - B_m^{[\mathbf{k}]}\Lambda)\Sigma^{-1}(W - B_m^{[\mathbf{k}]}\Lambda)^T\big]\} \\ &= argmin_{(\Lambda)}\{tr\big[(W - B_m^{[\mathbf{k}]}\Lambda)(W - B_m^{[\mathbf{k}]}\Lambda)^T\big]\} \\ &= \left(B_m^{[\mathbf{k}]^T} B_m^{[\mathbf{k}]}\right)^{-1} B_m^{[\mathbf{k}]^T} W \end{aligned}$$

so that the least square estimate $\hat{W} = B_m^{[\mathbf{k}]}\hat{\Lambda}$ does not depend on Σ. Thus, also in this more general case we could use the model selector (4). It turns out that the solution $\hat{\Lambda} = \left(B_m^{[\mathbf{k}]^T} B_m^{[\mathbf{k}]}\right)^{-1} B_m^{[\mathbf{k}]^T} W$, not only minimizes (5), for any symmetric positive-definite matrix Σ, but also minimizes the general variance $\left|(W - B_m^{[\mathbf{k}]}\Lambda)(W - B_m^{[\mathbf{k}]}\Lambda)^T\right|$, where $|\cdot|$ denotes the determinant.

4.2 Algorithm for the search of optimal knots

For the search of the optimal knot vector we modified Zhou and Shen (2001) algorithm, in order to deal with our 3D problem. For computational convenience, we look for optimal knots among the grid values $s_1, \ldots, s_n$. This is not a limitation, since, as noted in Section 3, the grids are very fine. We start with a set of initial knots which are evenly spaced along the grid. A more accurate selection of initial knots, along the lines of Zhou and Shen (2001), would speed up the algorithm, but is not necessary, as pointed out by the same authors. The algorithm then alternates between the following two steps.

Knot addition. Let $\mathbf{k} = (k_1, \ldots, k_{n_k})$ be the current knot vector, and Λ the corresponding estimate of the coefficient matrix. Set $k_0 = a$ and $k_{n_k+1} = b$. If $\mathbf{k}$ is the initial knot vector, then each of the subintervals $[k_r, k_{r+1}]$, for $r = 0, \ldots, n_k$, is checked for possible addition of one knot. Specifically, in each of the subintervals $[k_r, k_{r+1}]$, the knot s_j, with $k_r < s_j < k_{r+1}$, is added to the current knot vector, leading to the new knot vector $\mathbf{k}_j^*$ and to the new estimate of the coefficient matrix $\Lambda_j^* = \left(B_m^{[\mathbf{k}_j^*]^T} B_m^{[\mathbf{k}_j^*]}\right)^{-1} B_m^{[\mathbf{k}_j^*]^T} W$, if

$$\begin{aligned} &pSSE3D(\Lambda_j^*, \mathbf{k}_j^*) < pSSE3D(\Lambda, \mathbf{k}) \qquad \text{and} \\ &s_j = argmin_{\{s_l : k_r < s_l < k_{r+1}\}} pSSE3D(\Lambda_l^*, \mathbf{k}_l^*). \end{aligned}$$

If $\mathbf{k}$ is not the initial knot vector, then the subinterval $[k_r, k_{r+1}]$ is checked for possible addition of one knot only if at least one among the neighbor knots, $k_{r-2}, k_{r-1}, \ldots, k_{r+3}$, has been added in the previous iteration.

Knot relocation/deletion. Let $\mathbf{k} = (k_1, \ldots, k_{n_k})$ be the current knot vector. For $r = 1, \ldots, n_k$, if k_r is adjacent to a knot added in the preceding step, then it is checked for possible relocation/deletion. In particular, k_r is removed from the the current knot vector, or moved to the new position s_j, with $k_{r-1} < s_j < k_{r+1}$, if this improves the criterion (4), likewise in the knot addition move. The same procedure is

then applied to all knots that have been added in the preceding step.

When, for the first time, the knot addition move does not yield any new knot, all knots are searched for deletion/relocation. The resulting knot vector is set as new initial knot vector, and the algorithm is run again. When, for the second time, the knot addition move does not yield any new knot, all knots are searched for deletion/relocation, and the procedure ends.

The algorithm for the search of optimal knots, and all other data processing, have been coded in R (see R Development Core Team 2007).

5 Results

The 65 centerlines are estimated by means of 3D free knot regression splines of order $m = 5$, with penalization constant $\mathcal{C} = 4$. The choice of the penalization constant is discussed in Section 5.1. The order of the splines has been chosen to obtain smooth estimates of the first two derivatives. In the future, the torsion of the centerline may also be an object of study. In fact, the hemodynamics is influenced not only by vessel radius and curvature, but also by its torsion, although in a more complex and still unclear way. The torsion depends on the first three derivatives, and thus an order-6 spline would be recommended for its analysis. However, using splines of order $m = 5$, the third derivative is a piecewise linear continuous line, so that order-5 splines are enough for a continuous estimate of the torsion.

The top line of Figure 4 shows the estimates, $\hat{x}(s), \hat{y}(s), \hat{z}(s)$, of the three coordinate functions of ICA centerline for patient 1, obtained by order-5 free knot regression splines with $\mathcal{C} = 4$. The vertical lines show the position of the knots along the abscissa parameter s. The estimates are superimposed to the original $(s_j, x_j), (s_j, y_j), (s_j, z_j)$ (grey dots, almost completely hidden by the estimates). Center and bottom lines of the same figure show first and second derivatives of $\hat{x}(s)$, $\hat{y}(s)$ and $\hat{z}(s)$, superimposed to the rough estimates given respectively by first and second central differences (grey). Figure 5 is a 3D image of the fitted centerline, $\hat{\mathbf{c}}(s)$, and Figure 6 shows the corresponding curvature function.

It is of particular interest to identify the points of approximately zero curvature. In fact, these points can be taken as delimiters of artery bends or siphons, whose identification is important for the morphological analysis of the ICA. Preliminary investigations of the location of the aneurysms in the different bends, and of the position inside each specific bend, yielded interesting results relevant to the set up of prognostic indexes. See Piccinelli et al. (2007) and Sangalli et al. (2008b). We take as points of approximately zero curvature all points of local minimum for the curvature, with curvature smaller than the threshold 0.0145mm^{-1}, or equivalently with radius of curvature greater than $1/(0.0145\text{mm}^{-1})$=68.966mm, where 68.966mm is the mean length of reconstructed centerlines. Indeed, consider an arc of length l on

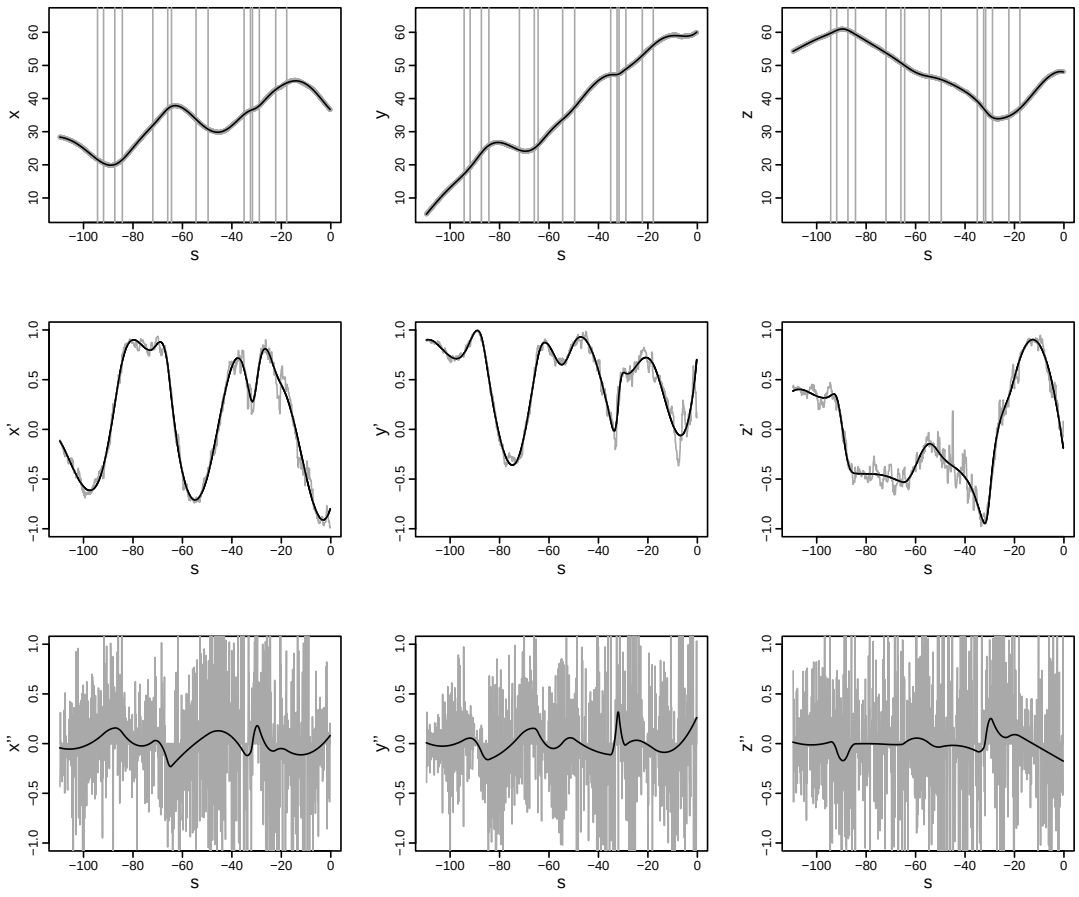

Figure 4: Top: fitted coordinates curves $\hat{x}(s), \hat{y}(s), \hat{z}(s)$, with vertical lines showing the position of the knots along the abscissa parameter s, superimposed to the original $(s_j, x_j), (s_j, y_j), (s_j, z_j)$ in grey [patient 1]. Center: first derivatives of $\hat{x}(s)$, $\hat{y}(s)$ and $\hat{z}(s)$, superimposed to first central differences in grey. Bottom: second derivatives of $\hat{x}(s)$, $\hat{y}(s)$ and $\hat{z}(s)$.

a circumference of radius R. The arc is a curve with constant curvature equal to $1/R$; the corresponding chord is an approximating curve with constant zero curvature, i.e. a straight line. The ratio between the length of the chord and the length of the arc is equal to $\frac{2R}{l}\sin(\frac{l}{2R})$: this ratio is greater than 0.95 when $R > l$. Hence for a curve belonging to a population of curves with mean length equal to $\bar{l}$, we propose to locally approximate the curve with a straight line in a neighborhood of a point of minimal curvature, if the radius of curvature R in the point is greater than $\bar{l}$. Figures 5 and 6 show the points of approximately zero curvature of the ICA centerline for patient 1.

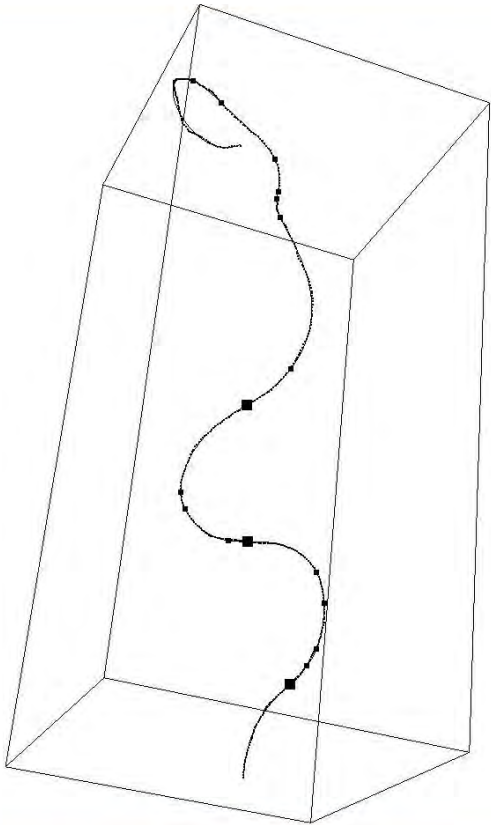

Figure 5: Fitted centerline (the little bullets show the positions of the spline knots), together with rough data [patient 1]. The big squares are points of approximately zero curvature.

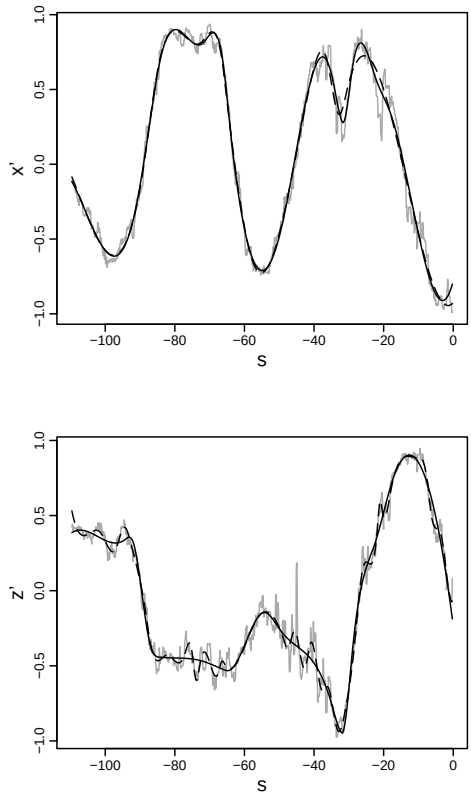

Figure 7: Top: estimates of x' obtained by free knot splines with $\mathcal{C} = 4$ (solid) and $\mathcal{C} = 9$ (dashed). Bottom: estimates of z' obtained by free knot splines with $\mathcal{C} = 4$ (solid) and $\mathcal{C} = 0.2$ (dashed) [patient 1].

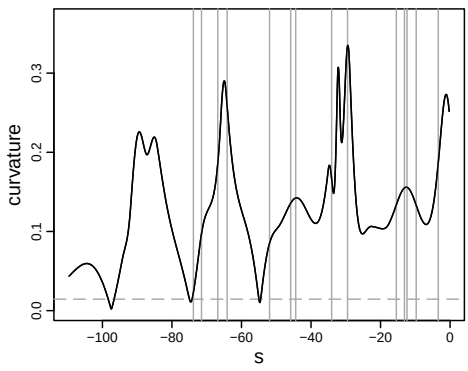

Figure 6: Curvature of $\hat{\mathbf{c}}(s)$ [patient 1]. The horizontal dashed line is the approximately zero curvature threshold.

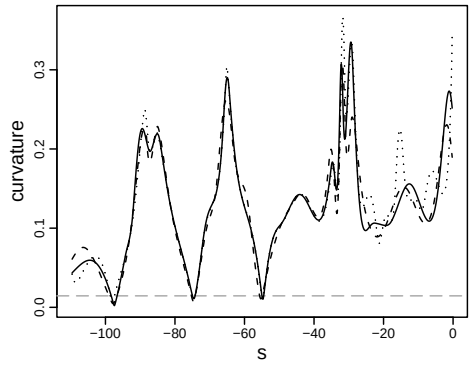

Figure 8: Estimated functions of centerline curvature [patient 1], obtained by free knot regression splines with $\mathcal{C} = 3, 4, 5$ (dotted, solid and dashed).

5.1 Choice of penalization constant

As goodness-of-fit measure for the curve estimate, we use the Average Squared Error (ASE):

$$\text{ASE} = \frac{1}{n}\left[\sum_{j=1}^{n}\big(x_j - \hat{x}(s_j)\big)^2 + \sum_{j=1}^{n}\big(y_j - \hat{y}(s_j)\big)^2 + \sum_{j=1}^{n}\big(z_j - \hat{z}(s_j)\big)^2\right].$$

Moreover, since we want to evaluate how well the curve estimate reflects the features of the curve represented by its derivatives, we also consider the Average Squared Error on first derivatives (ASEder) with respect to the rough estimates given by first central differences:

$$\begin{aligned}&\text{ASEder} = \\ &= \frac{1}{n}\left[\sum_{j=1}^{n}\big(Dx(s_j) - \hat{x}'(s_j)\big)^2 + \sum_{j=1}^{n}\big(Dy(s_j) - \hat{y}'(s_j)\big)^2 + \sum_{j=1}^{n}\big(Dz(s_j) - \hat{z}'(s_j)\big)^2\right].\end{aligned}$$

The average squared error on second derivatives is uninformative, since the second central differences are too noisy. Thus, we will use ASEder as a proxy measure of goodness-of-fit for derivatives in general.

The left panel of Figure 9 displays boxplots of the distribution of ASE, ASEder and also DF for the fits corresponding to the 65 patients, obtained by free knot regression splines with 5 different penalizations: $\mathcal{C} = 0.2, 3, 4, 5, 9$. When the penalization increases, the ASE and ASEder increase while the DF decrease. This tradeoff between average squared errors and DF translates the classical bias/variance tradeoff.

The value of $\mathcal{C}$ is chosen in order to neither under-smooth nor over-smooth the data, but provide a "just right" degree of smoothness. In particular, we set $\mathcal{C} = 4$: with this choice of the penalization constant, free knot regression splines can accurately estimate the salient features of the centerlines, without being too data-adapted. This can be better appreciated looking at spline estimates of first derivatives and comparing them with central difference estimates. Very high values of $\mathcal{C}$ can lead to centerlines estimates that cannot fully capture the peaks and troughs in the first (and thus subsequent) derivatives. See for example the top of Figure 7, which compares the estimates of x' obtained with $\mathcal{C} = 4$ and $\mathcal{C} = 9$, for patient 1 carotid centerline. Very low values of $\mathcal{C}$ may instead yield estimates where also the high frequency variation is fitted. See for example the bottom of Figure 7, which compares the estimates of z' obtained with $\mathcal{C} = 4$ and $\mathcal{C} = 0.2$, also for patient 1. These two opposite mistakes must be avoided to get sensible estimates of curvature profiles.

On the other hand, it is important to notice that the estimate of the curvature is stable with respect to variations of $\mathcal{C}$ over a reasonable span. Figure 8, for instance, compares the estimates of centerline curvature functions for the ICA of patient 1, obtained with $\mathcal{C} = 3, 4, 5$. In particular, the points of (approximately) zero curvature,

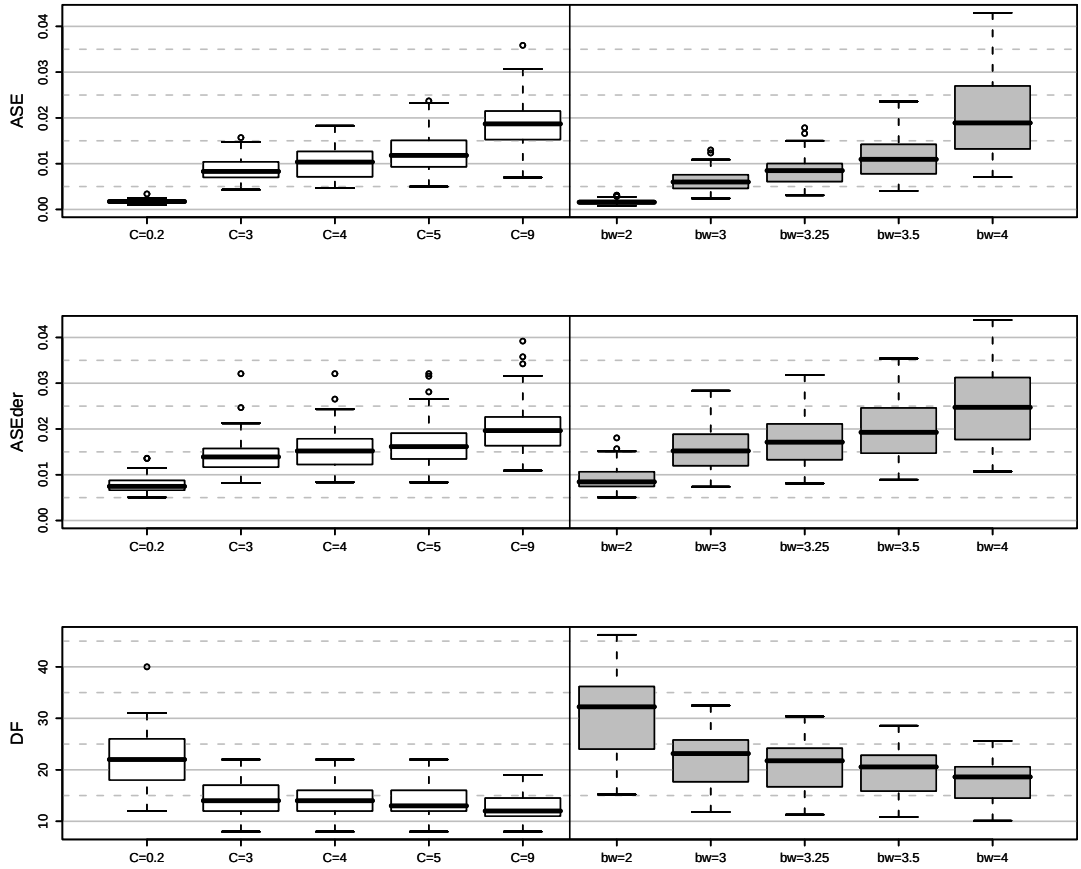

Figure 9: Left: boxplots of the distribution of ASE, ASEder, and DF for the fits corresponding to the 65 patients, obtained by free knot regression splines with penalizations $\mathcal{C} = 0.2, 3, 4, 5, 9$. Right: the same for the fits obtained by local polynomial smoothing with bandwidth $bw = 2, 3, 3.25, 3.5, 4$.

whose importance has been previously mentioned, do not change. To show the robustness of the estimation of zero-curvature points, with respect to the choice of the penalization constant in a reasonable span, we compare the results corresponding to the three different penalization $\mathcal{C} = 3, 4, 5$. For 28 out of the 65 patients, no zero-curvature point is found, in the three different estimates corresponding to $\mathcal{C} = 3, 4, 5$. For these patients just a portion of the first siphon is observed. For only 9 out of the 65 patients, the number of zero-curvature points is not the same in each of the three estimates: depending on the penalization, 0 points or 1 point are found, so that the three different estimates identify one full siphon, or just a portion of it. For the remaining 28 patients, each of the three estimates find the same number of zero-curvature points, either 1, or 2 or 3 depending on the patient, for a total of 34 points. Moreover we point out that not only the number of points found by the three different estimates are the same, but also their positions along the abscissa parameter are not statistically different. To show this fact we perform a principal component analysis of the matrix $U = \left[\boldsymbol{u}_1 | \boldsymbol{u}_2 | \boldsymbol{u}_3\right]$ that has as columns the vectors $\boldsymbol{u}_g = (u_{1g}, \ldots, u_{34g})^T$,

for $g = 1, 2, 3$, with the abscissae of the 34 zero-curvature points, respectively in the three different estimates corresponding to $\mathcal{C} = 3, 4, 5$ (points are ordered according to patient and, when more than 1 point is present per patient, according to the position of the point along the abscissa parameter s). The rows of U thus give the abscissae of corresponding zero-curvature points in the three different estimates. The first principal component of U is equal to $\frac{1}{\sqrt{3}}(1, 1, 1)$, up to the 3rd decimal digit, and the three eigenvalues are respectively 1487.539, 0.187, 0.021. We thus verify that, jointly, the second and third eigenvalue are not significantly different from $(0, 0)$. The p-value of the approximate chi-squared test is 0.20. This supports the assumption that the positions of the 34 zero-curvature points are essentially the same in the three different estimates.

6 Comparison with local polynomial smoothing

In this section we compare free knot spline regression with a more classical competitor, namely local polynomial smoothing. The latter technique was formerly used in Sangalli et al. (2008b), where pointwise estimates of true centerlines were obtained fitting order-5 polynomials with a Gaussian kernel and bandwidth equal to 3. Estimates of derivatives, and hence of curvatures, were thus obtained by pointwise differentiation of the fitted local polynomials.

Likewise free knot regression splines, also local polynomials are a linear estimator. Let m be the order of the polynomials, bw the bandwidth, and $\mathcal{K}(\cdot)$ the kernel. Moreover, denote by $K^{1/2}(s^*)$ the $(n \times n)$-diagonal matrix having j-th entry equal to $\big(\mathcal{K}\big((s^* - s_j)/bw\big)/bw\big)^{1/2}$, and set

$$S(s^*) = \begin{bmatrix} 1 & (s_1 - s^*) & \cdots & (s_1 - s^*)^{m-1} \\ \vdots & & \vdots & \\ 1 & (s_n - s^*) & \cdots & (s_n - s^*)^{m-1} \end{bmatrix}$$

Then the pointwise centerline estimate at s^*, given by local polynomials, is

$$(\hat{x}^*, \hat{y}^*, \hat{z}^*) = \mathbf{e}_{1,m}^T \hat{\Psi}(s^*) \tag{6}$$

where $\mathbf{e}_{r,m}$ represent the r-th column of an $(m \times m)$ identity matrix, and the $(m \times 3)$-matrix $\hat{\Psi}(s^*)$ is given by

$$\begin{aligned} \hat{\Psi}(s^*) &= argmin_{(\Psi)} \big\{ tr\big[K^{1/2}(s^*)^T (W - S(s^*)\Psi)(W - S(s^*)\Psi)^T K^{1/2}(s^*)\big]\big\} \\ &= \big[S(s^*)^T K^{1/2}(s^*) K^{1/2}(s^*) S(s^*)\big]^{-1} S(s^*)^T K^{1/2}(s^*) K^{1/2}(s^*)\, W \end{aligned}$$

where W is the matrix of observed values. Using this expression for $\hat{\Psi}(s^*)$, it is possible to rewrite (6) as

$$(\hat{x}^*, \hat{y}^*, \hat{z}^*) = \sum_{j=1}^{n} T\Big(s^*, \frac{s_j - s^*}{bw}\Big)(x_j, y_j, z_j)$$

where

$$T(u,v) = \mathbf{e}_{1,m}^T \big[S(u)^T K^{1/2}(u) K^{1/2}(u) S(u)\big]^{-1} diag\{1, v, \ldots, v^{m-1}\} \mathcal{K}(v)/bw$$

and $diag\{\cdot\}$ denotes a diagonal matrix. It follows that $\hat{W} = L\,W$, where the (i,j)-th entry of the linear operator L, mapping the observed values into the fitted values, is given by

$$L(i,j) = T\Big(s_i, \frac{s_j - s_i}{bw}\Big).$$

To compute the trace of the smoothing matrix L, i.e. the DF of the local polynomial estimator, we use the empirical formula provided by Zhang (2003):

$$tr(L) = m + \frac{n}{n-1}\,\mathcal{K}_0\,\frac{|s_n - s_1|}{bw}$$

where

$$\mathcal{K}_0 = \mathcal{K}(0)\,\mathbf{e}_{1,m}^T\,M^{-1}\,\mathbf{e}_{1,m}$$

and the $(m \times m)$-matrix M has (i,j)-th entry given by $M(i,j) = \mu_{i+j-2}$, with $\mu_d = \int t^d \mathcal{K}(t) dt$.

If the three coordinates in the error term in model (1) were assumed to be correlated, then it would be natural to look for the value of $\Psi(s^*)$ which minimises

$$tr\big[K^{1/2}(s^*)^T (W - S(s^*)\Psi)\Sigma^{-1}(W - S(s^*)\Psi)^T K^{1/2}(s^*)\big]$$

where Σ is the correlation matrix of the error term. It is easy to see that

$$\begin{aligned} &argmin_{(\Psi)}\big\{tr\big[K^{1/2}(s^*)^T (W - S(s^*)\Psi)\Sigma^{-1}(W - S(s^*)\Psi)^T K^{1/2}(s^*)\big]\big\} \\ &= argmin_{(\Psi)}\big\{tr\big[K^{1/2}(s^*)^T (W - S(s^*)\Psi)(W - S(s^*)\Psi)^T K^{1/2}(s^*)\big]\big\} \end{aligned}$$

so that the local least square estimate $\mathbf{e}_{1,m}^T \hat{\Psi}(s^*)$, like the free knot regression spline estimate, does not depend on Σ.

We now want to argue that, at the cost of an affordable increase in computational cost, estimation by free knot regression splines has many comparative advantages over local polynomial smoothing. To show some of these advantages, we first of all consider jointly the measures ASE, ASEder and DF. Figure 9 compares the boxplots of the distribution of ASE, ASEder, and DF, for the fits corresponding to the 65 patients, obtained with the two different techniques. Various values of $\mathcal{C}$ for spline estimates and bw for local polynomial estimates are considered, covering a wide range of ASE, ASEder and DF. Inspection of the boxplots shows immediately that is not possible to find couples of values (C, bw) that are equivalent in terms of the criteria ASE, ASEder and DF. For example, a value of C and a value of bw that are comparable in terms of ASE of the corresponding estimates, turn out to be not comparable in terms of ASEder nor in terms of DF. The non-existence of a map that converts the values of C for free knot regression splines in equivalent values of bw for local polynomials, or

in other words the non-equivalency of the two regression techniques with respect to the criteria ASE, ASEder and DF, gives us the opportunity to highlight some of the gains obtained by using free knot regression splines. In particular, it gives evidence of the higher efficiency and higher accuracy of free knot regression splines.

Efficiency. Smoothing by free knot regression splines is more efficient than smoothing by local polynomials, in the sense that the former technique attains lower ASE, and lower ASEder, using less DF. In fact, there is no value of the bandwidth for which local polynomials can beat free knot regression splines on both ASE and DF (or on both ASEder and DF); whereas for any value of bw it is always possible to find a full range of values of $\mathcal{C}$ for which splines estimates do better than local polynomials with respect to all three criteria: ASE, ASEder and DF. See Figure 9.

Accuracy. Free knot regression splines can better estimate the salient features of the curves, expressed by its first derivatives. Indeed, if spline estimates and local polynomial estimates are allowed similar data-adaptation, i.e. similar ASE, then the former provide better estimates of the curve derivatives, i.e. have lower ASEder; compare for instance splines with $C = 5$ and local polynomials with $bw = 3.5$, in Figure 9. Even with a higher ASE, splines can attain a lower ASEder; see e.g. splines with $C = 4$ and local polynomials with $bw = 3.25$. This means that spline estimates can simultaneously be more accurate and less data adapted. Note that the spline estimates with $C = 4$ are less data adapted than the local polynomial estimates with $bw = 3$, used in Sangalli et al. (2008b), having higher ASE (and of course less DF); but, nonetheless, they reach the same level of accuracy, having comparable ASEder. Spline estimates have higher accuracy thanks to their local adaptivity, which allows them to better detect sharp peaks and troughs in the first and second derivatives. Figure 10, for example, compares estimates of $x'(s)$ for patient 1 obtained by free knot regression splines with $\mathcal{C} = 4$ (ASE=0.0115, ASEder=0.0161, and DF=20), and by local polynomials with $bw = 3$ (ASE=0.0056, ASEder=0.0159, and DF=32.28) and $bw = 3.5$ (ASE=0.0087, ASEder=0.0183, and DF=28.38). Even if in this comparison the two local polynomial estimates have been advantaged by allowing them a higher data-adaptation, the spline estimate can better get the peaks and troughs in the first derivative, being otherwise smooth, whereas the two local polynomial estimates cannot fully capture these features. Also in subsequent derivatives, spline estimates exhibit more clear-cut peaks and troughs, being otherwise smooth, whereas local polynomial estimates show flatter local features, but are more wiggly over the whole range. Note that the ability of free knot regression splines of better detecting salient local features is particularly important in our problem, where a big interest lies in curvature landmarks such as peaks and zeros, and their possible influences on hemodynamics. It should be mentioned that this advantage of free knot regression splines over other smoothing methods has been also evidenced by Gervini (2006), who carried out comparative simulation studies of free knot regression splines versus smoothing splines.

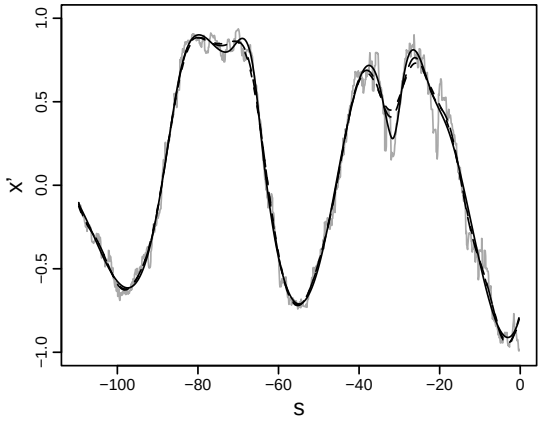

Figure 10: Estimates of x' obtained by free knot regression splines with $\mathcal{C}=4$ (solid), and by local polynomial smoothing with $bw=3, 3.5$ (dashed lines), superimposed to first central differences (grey) [patient 1].

Data dimension reduction and computational advantages. Other important gains obtained by using free knot regression splines are related to their functional nature. In particular, free knot regression splines naturally yield a dimension reduction of data, which is a fundamental issue for our highly dimensional dataset. Moreover, the closed functional form of their estimates can be exploited in subsequent analyses of the fitted curves, avoiding reiterated numerical approximations.

Sangalli et al. (2008b) obtained smooth pointwise estimates of vessel centerlines by means of local polynomials, and carried out exploratory analyses of vessel curvature and radius profiles which supported the existence of a strong relationship between vessel geometry and aneurysm location. In particular, a Functional Principal Component Analysis was used to identify the main uncorrelated modes of variability of the distribution of ICA curvature and radius profiles for the 65 patients, then a Quadratic Discriminant Analysis was carried out on the principal component scores, evidencing the existence of significant differences on the distribution of these two geometric features for patients having an aneurysm along the ICA and patients having an aneurysm downstream its terminal bifurcation. Note that local polynomial estimates were accurate enough to give good estimates of the first functional principal components of the curvature, and the conclusion reached in this exploratory study has been confirmed in Sangalli et al. (2008a) using centerline estimates obtained by free knot spline regression. The latter technique has now also made available highly accurate estimates of each individual ICA centerline and curvature function. It will be thus possible to use these estimates in the study of individual ICA geometries, in order to efficiently evaluate the role of vessel curvature on the hemodynamics, one of the main goals of AneuRisk Project.

7 Discussion

In this paper we have shown how to efficiently estimate a 3D curve and its derivatives by means of free knot regression splines. We have shown that the estimates obtained are more efficient than the estimates based on local polynomials, and that they can better estimate the local salient features of the curve. Very recently, Gluhovsky and Gluhovsky (2007) proposed a method to choose location-dependent bandwidths in local polynomials. A location-dependent bandwidth may also limit the drawbacks of local polynomials experienced in our 3D problem. An extension to the 3D case of the complex technique proposed by the two authors goes beyond the scope of the present paper. In the same paper, Gluhovsky and Gluhovsky express perplexities about the possibility of obtaining local adaptivity with regression splines methods. We believe our work dispels such doubts.

Acknowledgements. This research has been carried out within AneuRisk Project, a joint research program involving MOX Laboratory for Modeling and Scientific Computing (Dip. di Matematica, Politecnico di Milano), Laboratory of Biological Structures (Dip. di Ingegneria Strutturale, Politecnico di Milano), Istituto Mario Negri (Ranica), Ospedale Niguarda Ca' Granda (Milano), and Ospedale Maggiore Policlinico (Milano). The Project is supported by Fondazione Politecnico di Milano and Siemens Medical Solutions Italia, and partially supported by Ministero dell'Istruzione dell'Università e della Ricerca (research project "Metodi numerici avanzati per il calcolo scientifico" PRIN2006). We are especially grateful to Edoardo Boccardi (Ospedale Niguarda Ca' Granda), who provided the 3D-angiographies and motivated our reaserch by posing fascinating medical questions, and to Luca Antiga and Marina Piccinelli (Istituto Mario Negri), who performed the image reconstructions. Finally, we would also like to thank the associate editor and two anonymous referees for their constructive comments.

References

Antiga, L., Ene-Iordache, B., and Remuzzi, A. (2003), "Computational geometry for patient-specific reconstruction and meshing of blood vessels from MR and CT angiography," *IEEE Trans Med Imaging*, 22(5), 674–684.

Berger, S. A., Talbot, L., and Yao, L. S. (1983), "Flow in curved pipes," *Annual Review of Fluid Mechanics*, 15, 461–512.

Buja, A., Hastie, T., and Tibshirani, R. (1989), "Linear smoothers and additive models," *Ann. Statist.*, 17, 453–555.

Caro, C. G., Doorly, D. J., and Tarnakasky, M. (1996), "Non-planar curvature and branching of arteries and non-planar type-flow," *Proc Roy Soc*, 452, 185–197.

Castro, M. A., Putman, C. M., and Cebral, J. R. (2006), "Computational Fluid Dynamics Modeling of Intracranial Aneurysms: Effect of Parent Artery Segmentation on Intra-Aneurysmal Hemodynamics," *Am J Neuroradiol*, 27, 1703–1709.

Chandran, K. (1993), "Fluid Dynamics in the Human Aorta," *ASME J Biomech Engr*, 115, 611–616.

de Boor, C. (1978), *A practical guide to splines*, vol. 27 of *Applied Mathematical Sciences*, New York: Springer-Verlag.

Dean, W. R. (1927a), "Note on the motion of fluid in a curved pipe," *Phil. Mag.*, 4, 208–223.

— (1927b), "The streamline motion of fluid in a curved pipe," *Phil. Mag.*, 5, 673–695.

Friedman, J. H. (1991), "Multivariate adaptive regression splines," *Ann. Statist.*, 19, 1–141, with discussion and a rejoinder by the author.

Gervini, D. (2006), "Free-knot spline smoothing for functional data," *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68, 671–687.

Gluhovsky, I. and Gluhovsky, A. (2007), "Smooth location-dependent bandwith selection for local polynomial regression," *J. Amer. Statist. Assoc.*, 102, 718–725.

Hassan, T., Timofeev, E. V., Saito, T., Shimizu, H., Ezura, M., Matsumoto, Y., Takayama, K., Tominaga, T., and Takahashi, A. (2005), "A proposed parent vessel geometry-based categorization of saccular intracranial aneurysms: computational flow dynamics analysis of the risk factors for lesion rupture," *J. Neurosug.*, 103, 662–680.

Hastie, T. J. and Tibshirani, R. J. (1990), *Generalized additive models*, vol. 43 of *Monographs on Statistics and Applied Probability*, London: Chapman and Hall Ltd.

Hoi, Y., Meng, H., Woodward, S. H., Bendok, B. R., Hanel, R. A., Guterman, L. R., and Hopkins, L. N. (2004), "Effects of arterial geometry on aneurysm growth: three-dimensional computational fluid dynamics study," *J Nerusurg*, 101, 676–681.

Luo, Z. and Wahba, G. (1997), "Hybrid adaptive splines," *J. Amer. Statist. Assoc.*, 92, 107–116.

Mao, W. and Zhao, L. H. (2003), "Free-knot polynomial splines with confidence intervals," *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 65, 901–919.

Piccinelli, M., Bacigaluppi, S., Boccardi, E., Ene-Iordache, B., Remuzzi, E., Veneziani, A., and Antiga, L. (2007), "Influence of internal carotid artery geometry on aneurism location and orientation: a computational geometry study," Tech. rep., available at www.mathcs.emory.edu.

R Development Core Team (2007), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.

Ramsay, J. O. and Silverman, B. W. (2005), *Functional data analysis*, Springer Series in Statistics, New York: Springer, 2nd ed.

Sangalli, L. M., Secchi, P., and Vantini, S. (2008a), "Explorative Functional data analysis for 3D-geometries of the Inner Carotid Artery," in *Functional and Operatorial Statistics*, Physica-Verlag Springer, pp. 289–296, edited by S. Dabo-Niang and F. Ferraty.

Sangalli, L. M., Secchi, P., Vantini, S., and Veneziani, A. (2008b), "A Case Study in Exploratory Functional Data Analysis: Geometrical Features of the Internal Carotid Artery," *J. Amer. Statist. Assoc.*, to appear.

Smith, F. T. (1976), "Fluid Flow into a Curved Pipe," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 351, 71–87.

Stein, C. M. (1981), "Estimation of the mean of a multivariate normal distribution," *Ann. Statist.*, 9, 1135–1151.

Stone, C. J., Hansen, M. H., Kooperberg, C., and Truong, Y. K. (1997), "Polynomial splines and their tensor products in extended linear modeling," *Ann. Statist.*, 25, 1371–1470, with discussion and a rejoinder by the authors and Jianhua Z. Huang.

Ustun, C. (2005), "Dr. Thomas Willis' famous eponym: the circle of Willis," *J Hist Neurosci*, 14, 16–21.

Zhang, C. (2003), "Calibrating the degrees of freedom for automatic data smoothing and effective curve checking," *J. Amer. Statist. Assoc.*, 98, 609–628.

Zhou, S. and Shen, X. (2001), "Spatially adaptive regression splines and accurate knot selection schemes," *J. Amer. Statist. Assoc.*, 96, 247–259.