Published on *Journal of the American Statistical Association* (2009), 104, 485, 37-48.
Publisher's version available at http://pubs.amstat.org/doi/pdf/10.1198/jasa.2009.0002

A Case Study in Exploratory Functional Data Analysis: Geometrical Features of the Internal Carotid Artery

Laura M. Sangalli, Piercesare Secchi,
Simone Vantini, Alessandro Veneziani *

*Laura M. Sangalli is Post-Doctoral Fellow (email: laura.sangalli@polimi.it), Piercesare Secchi is Professor (email: piercesare.secchi@polimi.it) and Simone Vantini is PhD student (email: simone.vantini@polimi.it), MOX - Dipartimento di Matematica, Politecnico di Milano, P.za Leonardo da Vinci 32, 20133 Milano, Italy; Alessandro Veneziani is Associate Professor, Department of Mathematics and Computer Sciences, Emory University, 1131-002-1AC, Atlanta GA 30322, USA (ale@mathcs.emory.edu). This research has been carried out within AneuRisk Project, a joint research program involving MOX Laboratory for Modeling and Scientific Computing (Dipartimento di Matematica, Politecnico di Milano), Laboratory of Biological Structures (Dipartimento 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 research by posing fascinating medical questions, and to Luca Antiga and Marina Piccinelli (Istituto Mario Negri), who performed

Abstract

This pilot study is a product of the AneuRisk Project, a scientific program that aims at evaluating the role of vascular geometry and hemodynamics in the pathogenesis of cerebral aneurysms. By means of functional data analyses, we explore the AneuRisk dataset to highlight the relations between the geometric features of the internal carotid artery, expressed by its radius profile and centerline curvature, and the aneurysm location. After introducing a new similarity index for functional data, we eliminate ancillary variability of vessel radius and curvature profiles, through an iterative registration procedure. We then reduce data dimension by means of functional principal components analysis. Finally a quadratic discriminant analysis of functional principal components scores allows us to discriminate patients with aneurysms in different districts.

Keywords: Curve Registration, Functional Principal Components Analysis, Aneurysm Classification, Hemodynamics.

Cerebral aneurysms are deformations of cerebral vessels characterized by a bulge of the vessel wall. This is a common pathology in adult population, usually asymptomatic and not disrupting: epidemiological statistics (Rinkel et al., 1998) suggest that between 1% and 6% of adults develop a cerebral aneurysm during their lives. On the other hand, the rupture of a cerebral aneurysm, even if quite uncommon - about 1 event every 10000 adults per year - is usually a tragic event. Unfortunately, rupture preventing therapies, both endovascular and surgical treatment, are not without risks; this adds to the fact that in clinical practice general indications about rupture risk are still missing.

Even the origin of the aneurysmal pathology is still unclear. Possible explanations, discussed in the medical literature, focus on the interactions between biomechanical

the image reconstructions. Finally, we would also like to thank the associated editor and three anonymous referees for their constructive comments.

properties of artery walls and hemodynamic factors, such as wall shear stress and pressure; the hemodynamics are in turn strictly dependent on vascular geometry. See e.g. Hoi et al. (2004), Hassan et al. (2005), Castro et al. (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.

Arteries are basically hollow cylindrical pipes, featuring three-dimensional bends, branchings, bifurcations and progressive narrowing ("tapering") from proximal district (heart) to distal districts (peripheral circulation). Impact of morphology on fluid dynamics has been extensively investigated (e.g. Berger et al., 1983). An adimensional index, called Dean Number $\mathcal{D}$ (e.g. Jitchote and Robertson, 2000), has been proposed in order to describe different possible flow situations. The $\mathcal{D}$ depends on blood viscosity and density (quite easy to measure), mean velocity (to be computed by numerical simulations) and two geometric quantities: vessel radius and curvature. Hemodynamics induced by these features are supposed to play a relevant role in aneurysmal pathogenesis.

The present work stems from a conjecture grounded on practical experience of neuroradiologists at *Niguarda Ca' Granda Hospital* (E. Boccardi, personal communication): cerebral arteries of patients with an aneurysm at the terminal bifurcation of the Internal Carotid Artery (ICA), or after it, show peculiar geometrical features. We support this conjecture through the exploration of the relations between aneurysm location and the radius and curvature of the ICA for the 65 patients included in AneuRisk dataset. In brief: we highlight significant differences in the geometry of the last 3 cm of ICA of patients with an aneurysm located at or after the terminal bifurcation of the ICA compared to patients having an aneurysm before the terminal bifurcation or who are healthy. The former patients have significantly wider, more tapered and less curved ICA's. Moreover, within this group, there is a lower variability

of radius and curvature of the ICA.

Our data analysis follows a two-stage approach, along a paradigm advocated by J. O. Ramsay (see Ramsay and Silverman, 2005, for an exhaustive account on this approach to functional data analysis). In fact, our data are functional in nature, characterized by very high within-patient signal-to-noise ratio and large within-patient sample size. Analogous two-stage methods have been favored, for instance, by Sheehy et al. (2000a) and Sheehy et al. (2000b) for the analysis of human growth curves, and by Ramsay (2000) for the analysis of handwriting data. Indeed, given the complexity of the objects that we study, our work could also be seen as an instance of Object Oriented Data Analysis, as defined in Wang and Marron (2007).

The first stage of the analysis is covered in Sections 1 and 2. In Section 1, after a brief description of the dataset and its elicitation, we smooth - separately for each patient - the rough data, by means of local polynomial regression. We thus carry out, in Section 2, a novel functional data registration procedure, that enables meaningful comparisons across patients. The second stage is covered in Sections 3 and 4, where we analyze the registered data. In particular, in Section 3, we reduce the dimensionality of data, finding the main uncorrelated modes of variability of registered radius and curvature profiles, by means of functional principal component analysis. In Section 4, a quadratic discriminant analysis of principal component scores identifies the optimal number of principal components that best discriminate the patients with an aneurysm located at or after the terminal bifurcation of the ICA from the remaining patients. This also allows us to select representative geometries for numerical simulations. Conclusions are drawn in Section 5.

1 Data Capture and Elicitation

The dataset of the AneuRisk project is based on three-dimensional angiographies of 65 patients hospitalized at the *Neuroradiology Department* of *Niguarda Ca' Granda Hospital*, Milano, from September 2002 to October 2005. Among these, 33 patients have an aneurysm at the terminal bifurcation of the ICA or after it, 25 patients have an aneurysm along the left or right ICA, and finally 7 patients are healthy. None of the patients has other severe diseases affecting the cerebral vascular system, apart from the possible aneurysm. Percentages of females and males and of right and left ICA's do not differ significantly from 1/2 (the p-values of the test for equal proportions are .14 and .78 respectively). Ages - apart from a superior outlier - appear normally distributed (the p-value of Shapiro-Wilk test is .29) with sample mean equal to 55.85 years and sample standard deviation equal to 13.45 years. Gender, ICA side and age will not be included in the statistical analysis because they are supposed to be related to the aneurysmal pathology only through their effect on geometry.

The Integris Philips Allura Biplanar Unit (year 2001) working at the *Neuroradiology Department* of *Niguarda Ca' Granda Hospital* is an advanced rotational angiography system based on X-Ray Dual Fluoroscopy technology, able to provide 3D visualizations of the scanned volume. The Unit produces, for each patient, a three-dimensional array of gray-scaled pixels: lighter pixels show presence of flowing blood in the related volume while darker pixels show absence of flowing blood. This array is automatically generated by back-projection of 100 bi-dimensional angiographies (of 512×512 pixels) taken spanning a total angle of 240^o, facing the patient, in a period of less than 5 seconds. During image acquisition, 18 ml of nonionic hydrosoluble contrast agent is injected in the ICA at a rate of 4 ml/s.

Our data come from image reconstruction of these 3D angiographies, by means of an algorithm devised and implemented at the *Mario Negri Institute* (coded in the

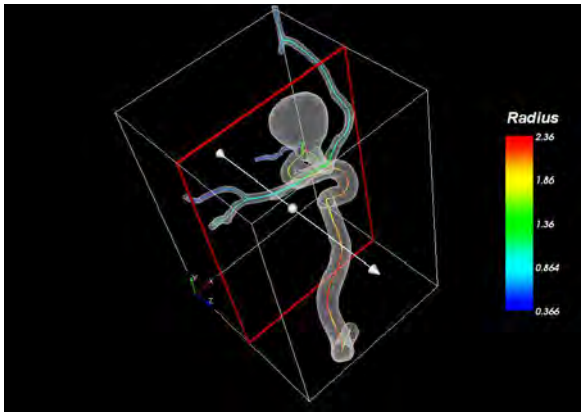

Figure 1: Example of a reconstructed vessel with an aneurysm. The transparent grid represents the reconstructed surface, the colored lines represent the reconstructed centerlines (with color referring to the maximal inscribed sphere radius) and the main cube represents the scanned region.

Vascular Modelling Toolkit, available at `http://vmtk.sourceforge.net`). The algorithm relies on the Level Set method for image segmentation and on differential geometry methods for surface parameterization, with the aim of providing quantitative descriptions of the geometries at hand. In particular, it identifies the lumen of the ICA (the volume occupied by flowing blood) and provides the three spatial coordinates of its centerline (computed as the set of centers of maximal spheres that can be inscribed in the vessel lumen) and the radius of lumen sections (computed as the radius of the maximal inscribed spheres) (Figure 1). Details about the elicitation of these features are in Antiga et al. (2003), Antiga and Steinman (2004) and Piccinelli et al. (2007). Alternative methods, that generate image reconstructions suitable for object oriented data analysis, are also available; for instance, the blood vessel tree data studied in Wang and Marron (2007) have been elicited with a method proposed in Bullitt and Aylward (2002).

The ith patient is represented by the function:

$$\mathbf{f}_i : S_i \subset \mathbb{R} \longrightarrow \mathbb{R}^4$$

$$s \longmapsto \mathbf{f}_i(s) = (x_i(s), y_i(s), z_i(s), R_i(s)) .$$

The abscissa parameter s measures an approximate distance along the ICA, from its terminal bifurcation towards the heart. For conventional reasons, this abscissa parameter takes negative values, to highlight that the direction is opposite with respect to blood flow. Functions $x_i(s)$, $y_i(s)$ and $z_i(s)$ map s into the left-right, up-down and front-back coordinates of the corresponding point of the centerline. Note that these coordinates are not absolute but relative to the cubic volume analyzed during the angiography. Moreover left carotids are left-right reflected to make all ICA's comparable. Finally, $R_i(s)$ is the radius of the maximal inscribed sphere centered in $(x_i(s)$, $y_i(s)$, $z_i(s))$.

As a matter of fact, the reconstruction algorithm provides centerlines and radius profiles only on a fine grid of points; the number of points available for each patient ranges from 350 to 1380, and is almost perfectly correlated to the approximate length of the reconstructed centerlines, which in turn varies from 27.219 mm to 110.136 mm. Moreover, data are affected by acquisition and reconstruction errors, even though within-patient signal-to-noise ratio is always high. Regression techniques are necessary to obtain continuous and differentiable estimates of the centerline functions and thus to estimate their curvature profiles, which are functions of first and second derivatives. In particular, we use gaussian kernel local polynomial regression. We set to 4 the degree of the local polynomials, according to a well accepted rule of thumb that fixes it to $p + 2$, where p is the highest derivative of interest. The bandwidth is set equal to 3 following a "Goldilocks approach". Figure 2 shows that this value for the bandwidth neither under-smooths nor over-smooths the data, but leads to a "just right" degree of smoothing: on the one hand, the average squared error is low, and on the other hand, the measure of roughness, given for each curve by $\frac{1}{|S_i|}\int_{S_i}(x_i''(s)^2 + y_i''(s)^2 + z_i''(s)^2)\, ds$, is also low. This fact can be even better appreciated by comparing the estimates of first and second derivatives to first and second

central differences of rough data, as it is shown for instance in Figure 2 for the x coordinate of patient 1: estimates obtained with a low bandwidth (e.g. 0.1) are too data adapted and fit high frequency variations, whereas estimates obtained with a high bandwidth (e.g. 6) cannot fully track the peaks and troughs in the first and second derivatives. However, variations of the bandwidth in a reasonable neighborhood of 3 do not appreciably affect the subsequent analyses.

A different method for data smoothing, based on free knot regression splines, is explored in Sangalli et al. (2007).

2 Data Registration

Centerline coordinates $x_i(s)$, $y_i(s)$ and $z_i(s)$ depend on the location of the scanned volume. This nuisance could be simply removed by considering the first derivatives $x_i'(s)$, $y_i'(s)$ and $z_i'(s)$, instead of $x_i(s)$, $y_i(s)$ and $z_i(s)$, the only lost information being, in fact, the location of the scanned volume. Looking at first derivatives (Figure 3) it becomes apparent that data display two types of variability: a *phase variability* and an *amplitude variability*. The former is strongly dependent on the dimensions and proportions of patients' skulls. In order to make correct comparisons among the features $\mathbf{f}_i$, observed in different patients, we need to separate these two types of variability (Ramsay and Silverman, 2005) and look for a new parameterization of each of the $n = 65$ centerlines. This can be achieved by means of a registration procedure that, optimizing a similarity criterion, finds 65 warping functions h_i of the abscissa, leading to the new registered feature functions $\tilde{\mathbf{f}}_i$:

$$\tilde{\mathbf{f}}_i = \mathbf{f}_i \circ h_i^{-1} \qquad \forall i = 1, \ldots, n = 65$$

or equivalently:

$$\tilde{\mathbf{f}}_i \circ h_i = \mathbf{f}_i \qquad \forall i = 1, \ldots, n = 65.$$

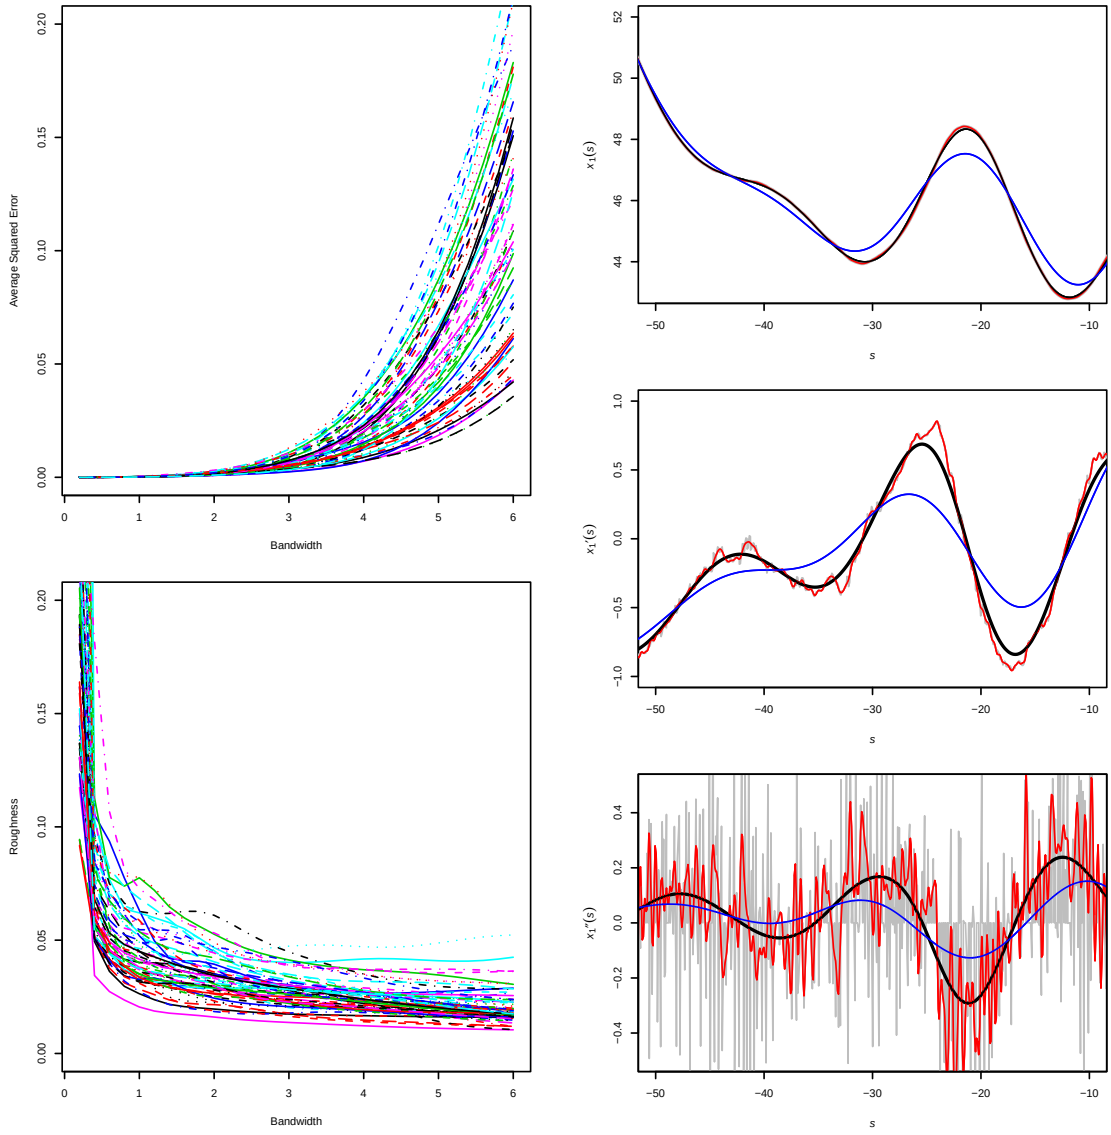

Figure 2: Top-left: average squared error for the 65 estimated ICA centerlines as a function of local polynomial bandwidth. Bottom-left: measure of roughness for the 65 estimated ICA centerlines as a function of local polynomial bandwidth. Right: estimates x_1, x_1' and x_1'' obtained with bandwidth equal to 0.1 (red), 3 (black) and 6 (blue), superimposed to rough data, first central differences of rough data and second central differences of rough data, respectively (gray).

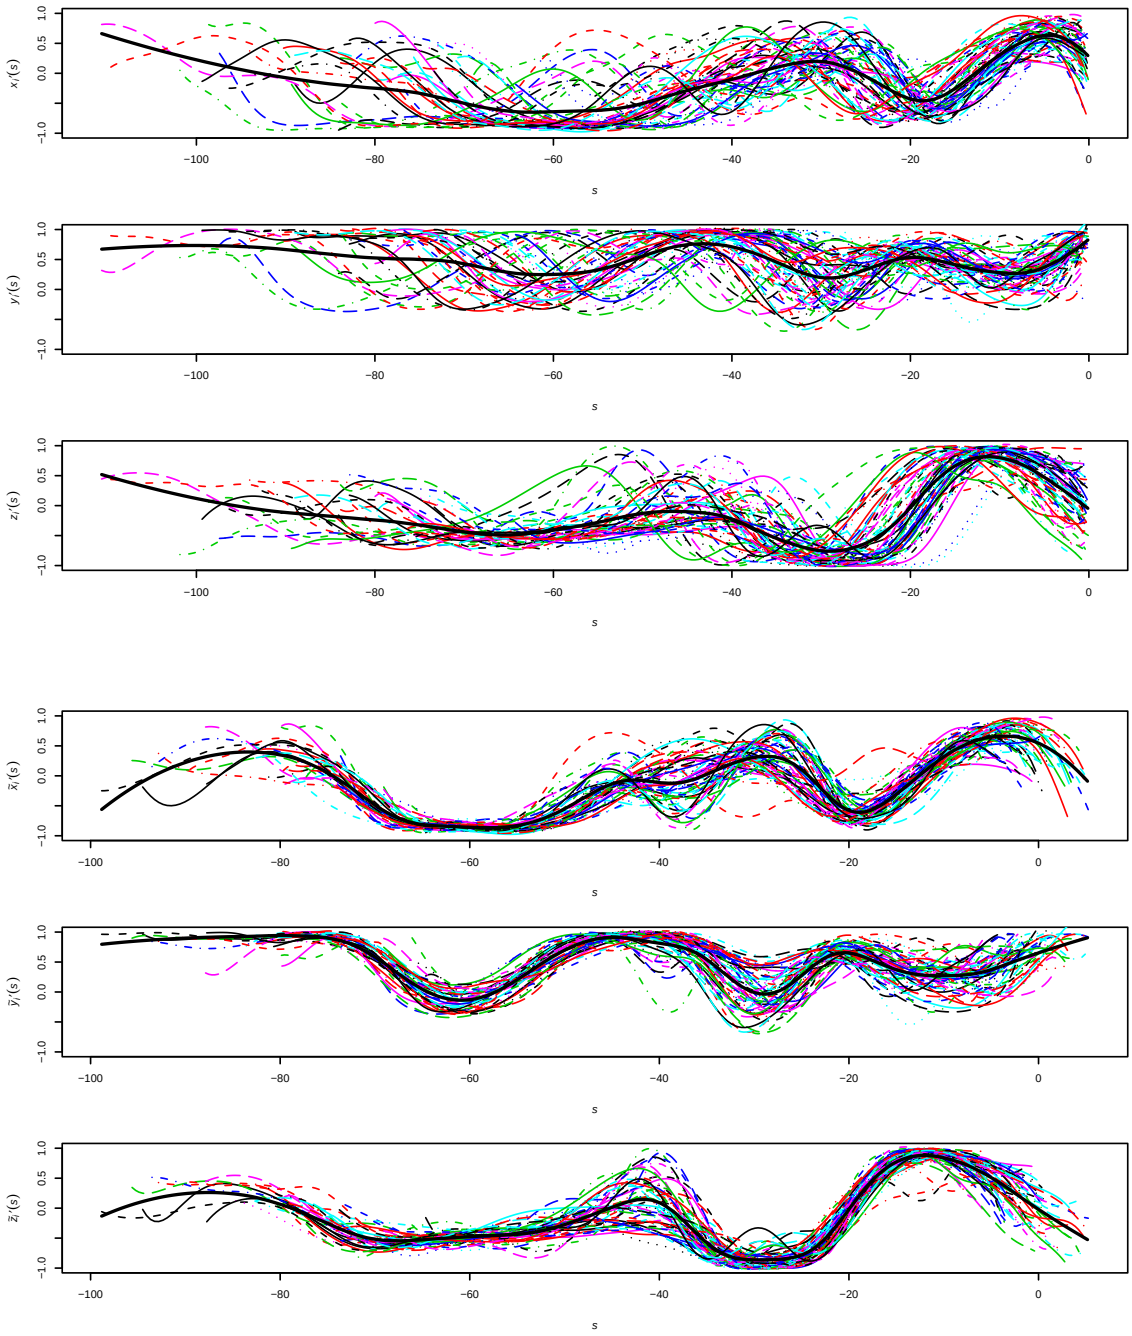

Figure 3: The 65 first derivatives $x_i'(s)$, $y_i'(s)$ and $z_i'(s)$ before registration (top) and the 65 first derivatives $\tilde{x}_i'(s)$, $\tilde{y}_i'(s)$ and $\tilde{z}_i'(s)$ after registration (bottom). The superimposed solid black lines are first derivatives of the reference centerline - as estimated by Loess - before and after registration (top and bottom respectively). The 65 radius and curvature profiles, before and after registration, are shown in Figure 5.

Note that the registered features $\tilde{\mathbf{f}}_i$ are obtained by moving the observed features $\mathbf{f}_i(s)$ to their "correct" location $h_i(s)$. The registration procedure thus separates the amplitude variability, captured by the 65 registered functions $\tilde{\mathbf{f}}_i$, from the phase variability, captured by the 65 warping functions h_i, without loss of information. The function $\tilde{\mathbf{f}}_i$ will be the main object of our study; we shall show that the information captured by the warping functions $h_i(s)$, i.e., the phase variability, is ancillary with respect to the scope of our analyses.

Analogously to Ramsay and Silverman (2005), the 65 warping functions h_i are elicited by maximizing, with respect to h_i, a similarity index between each centerline and a reference centerline. See also the recent paper by James (2007). In the following section, we introduce a novel similarity index between curves, inspired by the one used in Ramsay and Silverman (2005), but with more interesting mathematical properties and suitable for managing curves defined on different supports, as are the ones we deal with. It will be clear from the following discussion that the choice of the similarity index ρ and the class W of warping functions h_i are intrinsically connected, and that the couple ρ and W jointly defines what is meant by phase variability.

The problem of data registration is also encountered in the analysis of longitudinal data. For example, Lawton et al. (1972) use the framework of self-modelling non-linear regression to face this problem; more recently, Altman and Villarreal (2004) extend the former technique by including time-invariant covariates in the regression model; Lindstrom and Bates (1990) instead approach the problem in the light of non-linear mixed-effects models; Ke and Wang (2001) merge the above approaches in a unique framework, proposing semiparametric non-linear mixed-effects models. However, we note that the methods proposed in this literature are not ideal for data that have high within-subject signal-to-noise ratio and large within-subject sample size. See the discussion following Ke and Wang (2001) for illuminating comments about this issue.

2.1 Similarity Index

Two centerlines are said to have maximal similarity if they are identical except for a shifting and/or a dilation along the main axes x, y and z. Since location of the scanned volume and proportions of the skull change across patients, different shifting and/or dilation for each axis must be permitted for centerlines to have maximal similarity.

A similarity index for two parametric curves in $\mathbb{R}$, namely two functions from $\mathbb{R}$ into $\mathbb{R}$, will be now introduced. Later this index will be generalized for the evaluation of the similarity between two parametric curves in $\mathbb{R}^3$, i.e., functions from $\mathbb{R}$ into $\mathbb{R}^3$.

Let $g_i \in L^2(S_i \subset \mathbb{R}; \mathbb{R})$ and $g_j \in L^2(S_j \subset \mathbb{R}; \mathbb{R})$ be differentiable with $g_i' \in L^2(S_i \subset \mathbb{R}; \mathbb{R})$ and $g_j' \in L^2(S_j \subset \mathbb{R}; \mathbb{R})$, and let the domains $S_i \subset T$ and $S_j \subset T$ be closed intervals included in $\mathbb{R}$ such that $S_{ij} = S_i \cap S_j$ has positive Lebesgue measure. Note that Sobolev embedding theorem (Adams, 1975) guarantees that $g_i \in C^0(S_i \subset \mathbb{R}; \mathbb{R})$ and $g_j \in C^0(S_j \subset \mathbb{R}; \mathbb{R})$. Assuming that $||g_i'||_{L^2(S_{ij})} \neq 0$ and $||g_j'||_{L^2(S_{ij})} \neq 0$, the similarity index between g_i and g_j is defined as:

$$\rho(g_i, g_j) = \frac{\int_{S_{ij}} g_i'(s) g_j'(s) ds}{\sqrt{\int_{S_{ij}} g_i'(s)^2 ds} \sqrt{\int_{S_{ij}} g_j'(s)^2 ds}}. \tag{1}$$

This is the cosine of the angle θ_{ij} between first derivatives of the functions g_i and g_j, when the inner product $\int_{S_{ij}} g_i'(s) g_j'(s) ds$ is introduced. Index (1) can also be interpreted as a continuous version of Pearson's uncentered correlation coefficient for first derivatives.

The following useful properties of the symmetric similarity index ρ hold for any g_i, g_j, S_i and S_j for which $\rho(g_i, g_j)$ is defined:

(i) From Cauchy-Schwartz inequality it follows that: $|\rho(g_i, g_j)| \leq 1$.

(ii) Moreover: $\rho(g_i, g_j) = 1 \Leftrightarrow \exists A \in \mathbb{R}^+, B \in \mathbb{R} : \; g_i = A g_j + B$.

(iii) For all invertible affine transformations of g_i and g_j, say $r_1 \circ g_i = A_1 g_i + B_1$ and $r_2 \circ g_j = A_2 g_j + B_2$ with $A_1, A_2 \neq 0$,

$$\rho\left(g_i, g_j\right) = \operatorname{sign}(A_1 A_2)\, \rho\left(r_1 \circ g_i, r_2 \circ g_j\right).$$

(iv) For all invertible affine transformations of the abscissa s, say $r_1(s) = m_1 s + p_1$ and $r_2(s) = m_2 s + p_2$ with $m_1, m_2 > 0$,

$$\rho\left(g_i \circ r_1, g_j \circ r_2\right) = \rho\left(g_i \circ r_1 \circ r_2^{-1}, g_j\right) = \rho\left(g_i, g_j \circ r_2 \circ r_1^{-1}\right).$$

The similarity index $\rho(g_i, g_j)$ can also be interpreted as a modified version of the eigenvalue criterion used in Ramsay and Silverman (2005). Note that here, differently from Ramsay and Silverman (2005), the maximal value of the similarity index between two functions g_i and g_j is always 1, regardless of the measure of S_{ij} or the magnitude of the observed features (properties (i) and (ii)). This is crucial in our analysis because the functions $\mathbf{f}_i$ have different domains, and the measure of these domains is modified by the registration procedure.

For our purposes, a suitable generalization of the similarity index (1), for two vectorial functions $\mathbf{g}_i$ and $\mathbf{g}_j$ from $\mathbb{R}$ into $\mathbb{R}^3$, is:

$$\rho(\mathbf{g}_i, \mathbf{g}_j) = \frac{1}{3} \cdot \left[\rho(g_{xi}, g_{xj}) + \rho(g_{yi}, g_{yj}) + \rho(g_{zi}, g_{zj})\right]. \tag{2}$$

Properties (i) and (iv) still hold. Properties (ii) and (iii) hold with respect to a different affine transformation on each component. In particular (ii) becomes:

(ii)'

$$\rho(\mathbf{g}_i, \mathbf{g}_j) = 1 \Leftrightarrow \exists \mathbf{A} \in (\mathbb{R}^+)^3, \mathbf{B} \in \mathbb{R}^3 : \begin{cases} g_{xi} = A_x g_{xj} + B_x \\ g_{yi} = A_y g_{yj} + B_y \\ g_{zi} = A_z g_{zj} + B_z \end{cases}$$

Note that (ii)' holds for any vectorial generalization of the index (1) which depends on $\mathbf{g}_i$ and $\mathbf{g}_j$ only through $\rho(g_{xi}, g_{xj})$, $\rho(g_{yi}, g_{yj})$ and $\rho(g_{zi}, g_{zj})$, and is equal to 1 if and only if $\rho(g_{xi}, g_{xj})$, $\rho(g_{yi}, g_{yj})$ and $\rho(g_{zi}, g_{zj})$ are all equal to 1. This property, instead, does not hold for the natural generalization:

$$\frac{\int_{S_{ij}} \langle \mathbf{g}_i'(s); \mathbf{g}_j'(s) \rangle \, ds}{\sqrt{\int_{S_{ij}} \langle \mathbf{g}_i'(s); \mathbf{g}_i'(s) \rangle \, ds} \sqrt{\int_{S_{ij}} \langle \mathbf{g}_j'(s); \mathbf{g}_j'(s) \rangle \, ds}}$$

where brackets $\langle \, \rangle$ refer to the euclidean inner product in $\mathbb{R}^3$. In this case, property (ii)' would hold if $A_x = A_y = A_z$, i.e., if the dilation factor is the same along all three axes. As explained at the beginning of this section, this is not appropriate for our problem.

2.2 Registration Criterion

The ICA centerline of the ith patient is a curve in $\mathbb{R}^3$ that is described by the function $\mathbf{c}_i(s) = (x_i(s), y_i(s), z_i(s))$. Since centerlines are regular curves in S_i, namely $\mathbf{c}_i \in C^1(S_i \subset \mathbb{R}; \mathbb{R}^3)$, and $||x_i'||_{L^2(S_j)}$, $||y_i'||_{L^2(S_j)}$ and $||z_i'||_{L^2(S_j)}$ are different from zero, the similarity index between two ICA centerlines is always computable.

Given the similarity index ρ between two curves, registering a curve $\mathbf{c}_i$ with respect to another curve $\mathbf{c}_j$ means finding the function h in a class of warping functions W, that maximizes:

$$\rho(\mathbf{c}_i \circ h^{-1}, \mathbf{c}_j).$$

It is natural to ask that the registration of $\mathbf{c}_i$ with respect to $\mathbf{c}_j$ is equivalent to the registration of $\mathbf{c}_j$ with respect to $\mathbf{c}_i$, i.e.:

$$\sup_{h \in W} \rho(\mathbf{c}_i \circ h^{-1}, \mathbf{c}_j) = \sup_{h \in W} \rho(\mathbf{c}_i, \mathbf{c}_j \circ h^{-1}). \tag{3}$$

It is also natural to require that a warping function $h \in W$ applied simultaneously to the curves $\mathbf{c}_i$ and $\mathbf{c}_j$, does not change the similarity index, i.e.:

$$\rho\left(\mathbf{c}_i, \mathbf{c}_j\right) = \rho\left(\mathbf{c}_i \circ h^{-1}, \mathbf{c}_j \circ h^{-1}\right) \qquad \forall h \in W. \tag{4}$$

In words: using a function $h \in W$ to simultaneously warp two curves does not lead to a fictitious increment of the similarity between the two. Property (iv) assures that (3) and (4) hold if W is the group of strictly increasing affine functions. We will thus take:

$$W = \left\{h : h(s) = ms + p \;\text{ with }\; m \in \mathbb{R}^+, p \in \mathbb{R}\right\}.$$

This joint choice for ρ and W ensures that the similarity between two curves is invariant to changes in scale and/or location of the abscissa parameter s. Hence the similarity index between two curves does not depend on the velocity used to track them. The group structure of W, in particular the fact that W is closed with respect to composition, supports the iterative procedure presented in the next section.

2.3 Iterative Procedure

If a template ICA centerline $\mathbf{c}_0$, defined on the interval $S_0 = \bigcup_{i=1}^{n} S_i$, were given, the registration procedure would consist in finding, for each patient i, the function $h_i \in W$ that maximizes:

$$\rho(\mathbf{c}_i \circ h_i^{-1}, \mathbf{c}_0).$$

Unfortunately there is no a template ICA centerline. Therefore, as suggested in Ramsay and Li (1998) and Kneip et al. (2000) we will find both a reference centerline, acting as the template $\mathbf{c}_0$, and the 65 warping functions h_i by means of a *Procrustes fitting criterion*, implemented by alternating expectation and maximization steps:

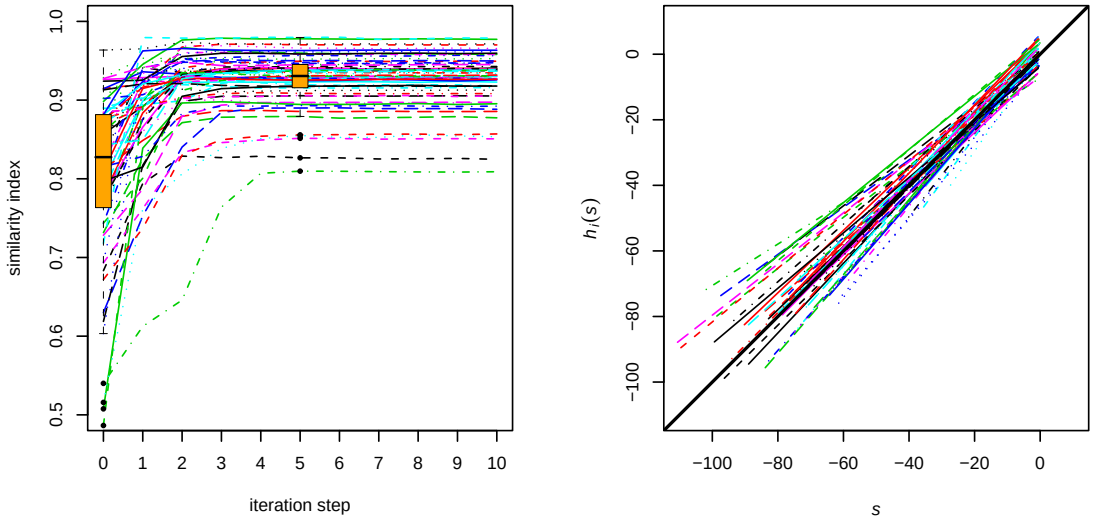

Figure 4: Left: similarity index between the 65 curves and the reference curve, after each iteration step. Boxplots of the distribution of the similarity index before registration and after 5 iterations are displayed. Right: optimal warping functions $h_i(s)$, each represented only on its domain S_i. The identity function is plotted in black.

1. *Expectation step*:

   The reference curve is estimated using all the curves obtained at the previous iteration. A new reference curve is obtained.

2. *Maximization step*:

   Each curve is shifted and dilated in order to maximize its similarity with the estimated reference curve. New curves are obtained.

The warping functions h_i (Figure 4) are simply given by the composition of the optimal warping functions found at each iteration:

$$h_i = h_{i_{iterK}} \circ \ldots \circ h_{i_{iter2}} \circ h_{i_{iter1}}.$$

The registered centerline is then defined as $\tilde{\mathbf{c}}_i = \mathbf{c}_i \circ h_i^{-1}$. Note that the group

structure of W ensures that $h_i \in W$. Technical details about the iterative procedure are reported in the last part of this section.

The registration allows to obtain a high value of the similarity index (2) for each patient (Figure 4). The sample mean of the similarity index, between patients' centerlines and the estimated reference curve, increases from 0.80 to 0.93; moreover, its standard deviation decreases from 0.11 to 0.03. The fact that, for each centerline, by means of a unique warping function of the abscissa it is possible to simultaneously and effectively align its three spatial coordinates, is strong evidence that the registration procedure is sound.

Visual inspection of first derivatives before and after registration (Figure 3) confirms that registered curves are much more similar than unregistered ones. A close look at Figure 3 shows that the variability in $\tilde{\mathbf{c}}_i'$ is mostly concentrated in the interval between values of abscissa -50 and -20. The presence of very different behaviors in this region agrees with the fact (Krayenbuehl et al., 1982) that some patients have ICA's with two siphons (S-shape ICA), others with only one (Ω-shape ICA), and others with no siphon at all (Γ-shape ICA). Here, a siphon is defined as a segment of the ICA included between two points of approximately zero curvature of the centerline.

2.4 Technical details

During each expectation step, we estimate the first derivatives of the reference centerline, from the first derivatives of the 65 patients' centerlines, by means of Loess with gaussian kernel and smoothness parameter α equal to 20% (see for example Cleveland et al., 1992). This adaptive fitting method has been preferred here in order to keep the variance of the estimator of the reference centerline as constant as possible along the ICA (Hastie and Tibshirani, 1990); indeed, the domains of the 65 centerlines are different, and very few curves are defined at high distance from the terminal bifurcation. Even if the estimated reference centerline is slightly influenced

by the value of α, the 65 final warping functions are robust with respect to this choice within a reasonable neighborhood of $\alpha = 20\%$.

Note that the computation of the index $\rho(\tilde{\mathbf{c}}_i, \mathbf{c}_0)$ does not require a new differentiation of the warped function $\tilde{\mathbf{c}}_i$, since this can be simply obtained by the identity:

$$\tilde{\mathbf{c}}_i'(s) = \mathbf{c}_i'(h_i^{-1}(s))\frac{1}{m_i}.$$

For ease of computation, during each maximization step curves are constrained to be shifted forward/backward no more than ± 5 mm and to be inflated/deflated no more than $\pm 10\%$. In any case, these constraints do not affect the final optimum.

After each maximization step, a global affine transformation is applied to all warping functions in order to have:

$$\frac{\sum_{i=1}^{n} h_i(s)}{n} = s \tag{5}$$

or equivalently $\sum_{i=1}^{n} m_i/n = 1$ and $\sum_{i=1}^{n} p_i/n = 0$. The reason for this rescaling is that, since no template curve exists, then it is desirable to have no global drift, in terms of shifting or dilating. Note that property (iv) guarantees that the similarity between pairs of curves does not change as long as the 65 curves are shifted and dilated all together with the constraint (5).

The iterative algorithm is stopped when the increments of the 65 similarity indexes are all lower than 0.01 in the maximization step; from (i) this corresponds to 1% of the achievable maximum for each index. This occurred after 5 iterations of the algorithm (Figure 4).

3 Data Analysis

The following analyses will involve maximal inscribed sphere radius functions $\widetilde{R}_i$ and centerline curvature function $\widetilde{C}_i$ obtained after registration of the original functions

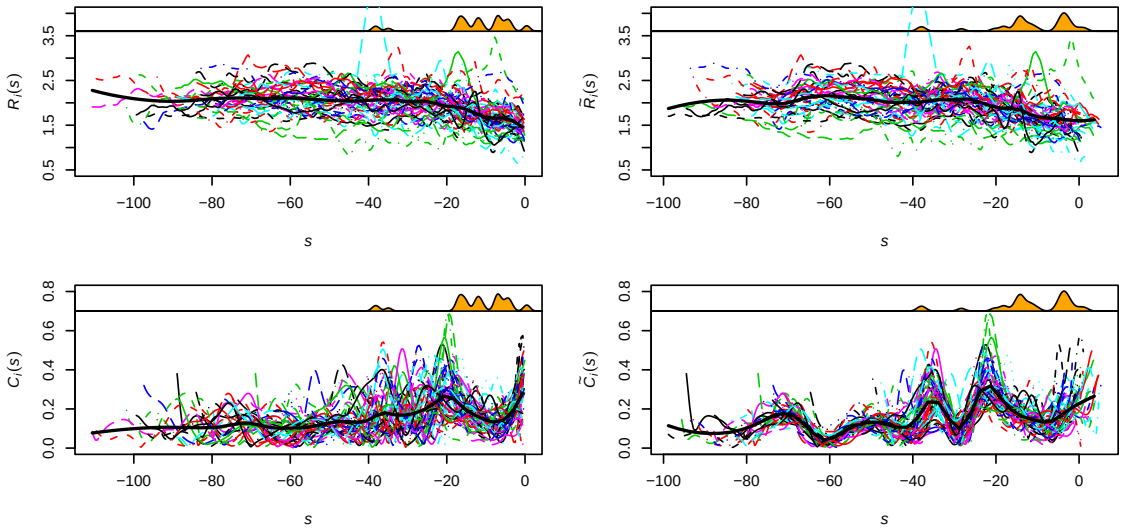

Figure 5: Radius (top) and curvature (bottom) profiles of the 65 patients respectively before (left) and after (right) registration. Solid black lines show mean curves, as estimated by Loess. On top of each picture is also displayed the estimate of the probability density function of the location of aneurysms along the ICA.

R_i and C_i along the optimal warping functions h_i shown in Figure 4:

$$\begin{aligned} \widetilde{R}_i(s) &= R_i(h_i^{-1}(s)), \\ \widetilde{C}_i(s) &= C_i(h_i^{-1}(s)). \end{aligned}$$

The curvature $C_i(s)$ is computed as follows:

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

where the symbol $\times$ refers to the vector product in $\mathbb{R}^3$ and $||\ ||$ is the euclidean norm in $\mathbb{R}^3$. Note that the registered curvature $\widetilde{C}_i$ can be obtained either by warping, along h_i, the curvature of the centerline $\mathbf{c}_i$, or by computing the curvature of the registered centerline $\tilde{\mathbf{c}}_i$.

Figure 5 shows the tapering of the ICA, i.e., the progressive reduction of the average radius of the carotid toward the end (values of the abscissa roughly greater

than -30). Moreover, it shows that two peaks of curvature (the siphon centers), are usually present at values of the abscissa of about -35 and -20. The same figure also displays a gaussian kernel estimate of the probability density function of aneurysm's location along the ICA: most aneurysms are clustered in two groups, both located in the terminal part of the ICA, where tapering is evident, and one located just after the last peak of curvature. These results provide evidence of a link between morphology and aneurysms onset, likely induced by hemodynamics.

We now analyze the autocovariance of radius profiles and of curvature profiles, in order to more deeply investigate their variability structure. The autocovariance function Σ_G of a generic process G is defined as:

$$\Sigma_G(t,s) = E\left[\ (G(t) - E\left[G(t)\right])\ (G(s) - E\left[G(s)\right])\ \right]$$

when the expected value exists. Since the 65 curves are known on different intervals of the abscissa, the following analyses will focus on the interval where all curves are available, i.e., for values of the abscissa between -33.7 and -8.0. Figure 6 shows the sample autocovariance function (separated in sample autocorrelation and sample standard deviation) of registered radius profiles $\hat{\Sigma}_{\widetilde{R}}$ and registered curvature profiles $\hat{\Sigma}_{\widetilde{C}}$:

$$\hat{\Sigma}_{\widetilde{R}}(t,s) = \frac{1}{n-1}\sum_{i=1}^{n}\left[(\widetilde{R}_i(t) - \overline{\widetilde{R}}(t))(\widetilde{R}_i(s) - \overline{\widetilde{R}}(s))\right]$$

$$\hat{\Sigma}_{\widetilde{C}}(t,s) = \frac{1}{n-1}\sum_{i=1}^{n}\left[(\widetilde{C}_i(t) - \overline{\widetilde{C}}(t))(\widetilde{C}_i(s) - \overline{\widetilde{C}}(s))\right].$$

Some details of the structure of the radius sample autocovariance function are amenable of an anatomical interpretation. First of all, the local minimum of the variance of the radius near the value of the abscissa -13, and the weak correlation of the radius measurements in close opposite neighborhoods of this point (block structures), suggest that this is the average position of the dural ring the ICA goes through before

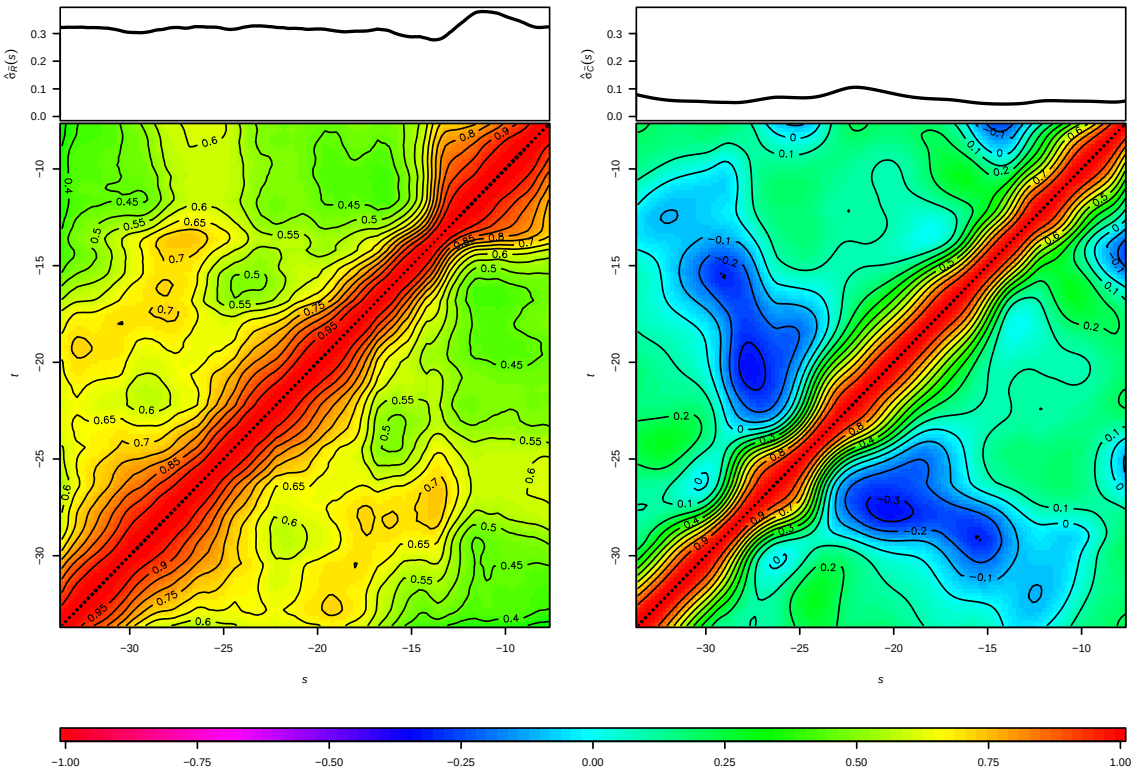

Figure 6: Isosurfaces of the sample autocorrelation function of registered radius profiles (left) and of registered curvature profiles (right). On top, sample standard deviations.

its terminal bifurcation. This is consistent with the presence of two clusters of aneurysms locations along the ICA, before and after this point, as evidenced in Figure 5. Note that this ring cannot be directly detected through angiographies. Moreover, sample autocorrelation functions of radius and curvature show that close points of the ICA have a weaker correlation of the curvature than of the radius. Finally, there is negative correlation of the curvature between points in proximity of the last peak of curvature and points in the region of lower curvature between the two peaks. This means that, if there is a segment of the centerline with very low curvature, a marked elbow is likely to occur just afterward, in order to enable the correct positioning of the

final bifurcation of the ICA. The registration procedure thus highlights some physical features common throughout the patients.

The autocovariance structures of radius and curvature profiles have been separately explored by means of Functional Principal Component Analysis (FPCA) (Ramsay and Silverman, 2005) in order to estimate the main uncorrelated modes of variability of these two geometric quantities, and to find their effective dimensionality. In this work, the main purpose of FPCA is dimension reduction, hence an analysis based on the autocovariance function is preferred to the alternative analysis based on the autocorrelation function.

The notation $\hat{\beta}_{Gk}$ and $\hat{\lambda}_{Gk}$ will respectively indicate the estimate of the kth eigenfunction and of the kth eigenvalue of the autocovariance function Σ_G. The score corresponding to the ith observed curve g_i and the kth estimated eigenfunction $\hat{\beta}_{Gk}$ is defined as the component along $\hat{\beta}_{Gk}$ of the ith observed curve g_i centered around the sample mean $\overline{g}$:

$$\int_S (g_i(s) - \overline{g}(s))\, \hat{\beta}_{Gk}(s) ds.$$

From now on, the analysis will focus only on the first and the second eigenfunctions of radius and curvature autocovariances. The reason for this choice is related to the Quadratic Discriminant Analysis (QDA) that will be presented in the next section. Figure 7 shows the estimates of the first and second eigenfunctions of radius and curvature. As suggested in Ramsay and Silverman (2005), the eigenfunctions are not directly plotted; instead, sample means of radius and curvature are plotted (solid lines), together with two curves obtained by adding/subtracting, to the sample means, the estimated normalized eigenfunctions multiplied by the estimated standard deviation of the corresponding scores. The first and second eigenfunctions for radius profiles $\hat{\beta}_{R1}$ and $\hat{\beta}_{R2}$, explain, respectively, 65.6% and 13.0% of the total variance (cumulative 78.6%). The first and second eigenfunctions for curvature profiles $\hat{\beta}_{C1}$ and

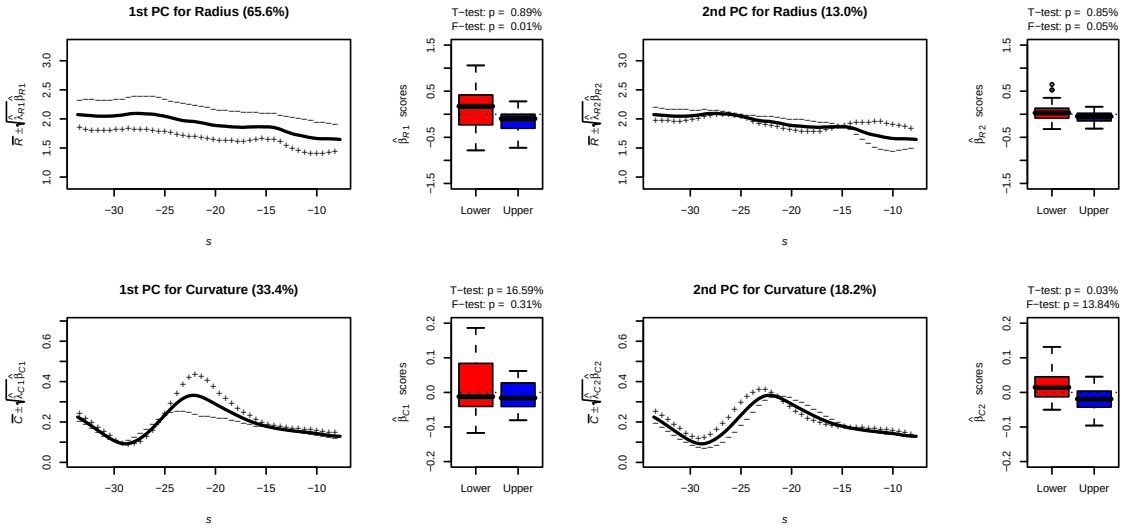

Figure 7: Estimates of the first (left) and the second (right) eigenfunctions for radius (top) and curvature (bottom), and boxplots of the corresponding scores for the two groups (red for Lower group and blue for Upper group).

$\hat{\beta}_{C2}$, explain, respectively, 33.4% and 18.2% of the total variance (cumulative 51.6%).

Eigenfunctions and corresponding scores can be interpreted as follows. Scores corresponding to $\hat{\beta}_{R1}$ quantify the overall width of the ICA: lower values are associated to wider ICA's, and higher values to narrower ICA's. Scores corresponding to $\hat{\beta}_{R2}$ quantify the tapering effect: lower values are associated to more tapered ICA's, and higher values to less tapered ICA's. Scores corresponding to $\hat{\beta}_{C1}$ quantify the curvature of the ICA in proximity of the last peak of curvature: lower values are associated to less curved siphons, and higher values to more curved siphons. Finally, scores corresponding to $\hat{\beta}_{C2}$ quantify the curvature along the segment of the ICA between the two peaks of curvature: lower values are associated to less curved segments, and higher values to more curved segments. As a referee pointed out, $\hat{\beta}_{C2}$ could also be interpreted as controlling shift in the location of the peak of curvature.

4 Data Classification and Selection

In order to support the conjecture outlined in the introduction, we now focus on discriminating the 33 patients with an aneurysm downstream of the ICA (Upper group) from the remaining 32 patients which represent the controls. We will refer to the latter patients as the Lower group; this includes 25 patients with an aneurysm located before the terminal bifurcation of the ICA and 7 healthy patients. A further split into two subgroups was not deemed of primary interest by physicians.

It is evident from inspection of the boxplots in Figure 7 that, for all four eigenfunctions, the distributions of scores have different means and/or variances in the two groups. This is confirmed by the F-tests for equal variances and Student's t-tests for equal means (degrees of freedom are computed according to Welch approximation and normality assumptions are verified by means of Shapiro-Wilk tests). Figure 7 reports the p-values of these tests. In particular, except for $\hat{\beta}_{C2}$, variances of scores of the Upper group are significantly lower than the ones of the Lower group. Moreover, mean values of scores corresponding to $\hat{\beta}_{R1}$, $\hat{\beta}_{R2}$ and $\hat{\beta}_{C2}$ are significantly lower in the Upper group than in the Lower group.

According to the proposed interpretations for the eigenfunctions $\hat{\beta}_{R1}$, $\hat{\beta}_{R2}$, $\hat{\beta}_{C1}$ and $\hat{\beta}_{C2}$, these differences can be interpreted as follows:

1. The geometrical features described by $\hat{\beta}_{R1}$, $\hat{\beta}_{R2}$, $\hat{\beta}_{C1}$ and $\hat{\beta}_{C2}$ have smaller variances in the Upper group than in the Lower group.

2. Patients in the Upper group tend to have a wider and more tapered ICA's than those in the Lower group. Moreover they present a less curved ICA between the two peaks of curvature.

Some of the differences detected through the analysis of FPCA scores can be roughly retrieved by visual inspection of the different distributions of radius and curvature profiles illustrated in Figure 8.

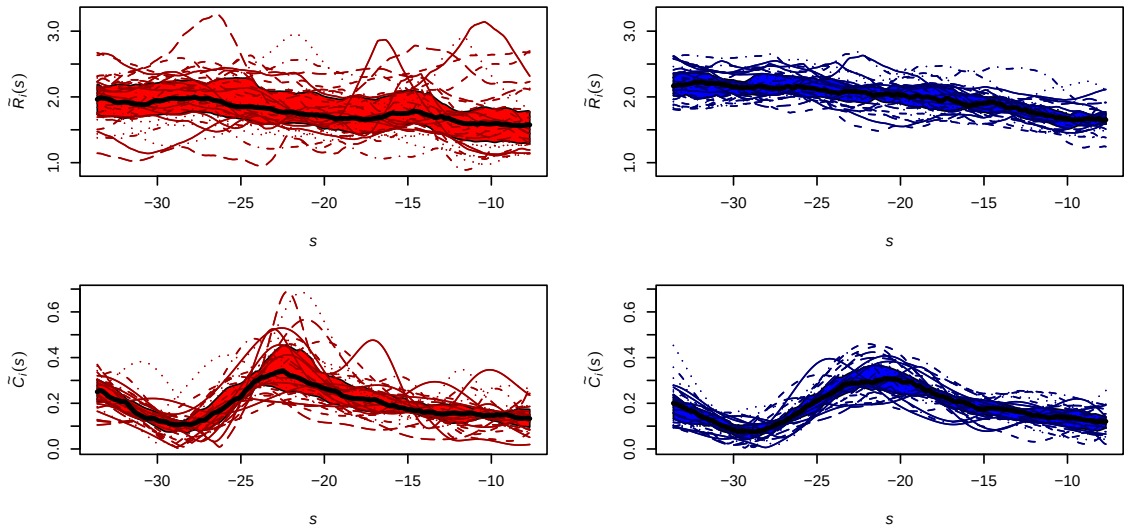

Figure 8: Radius (top) and curvature (bottom) profiles for the 32 patients of the Lower group (left) and for the 33 patients of the Upper group (right). Pointwise sample medians (black lines) and pointwise sample interquartile regions (colored regions) are displayed.

We now perform a Quadratic Discriminant Analysis (QDA) (Hand, 1981) of FPCA scores, in order to investigate the relationship between geometrical features and membership to the two groups.

Three issues have to be taken into account when selecting the eigenfunctions whose scores will be considered in the QDA:

1. A small value of the Actual Error Rate (AER) - the probability for a new case to be misclassified - is required.

2. Variances of scores are monotonically decreasing with respect to the index of eigenfunctions.

3. The efficiency of the estimates of the eigenfunctions and of the corresponding scores is decreasing with respect to the index of eigenfunctions (Monte Carlo simulations, not reported here, suggest this trend). This means that estimates of

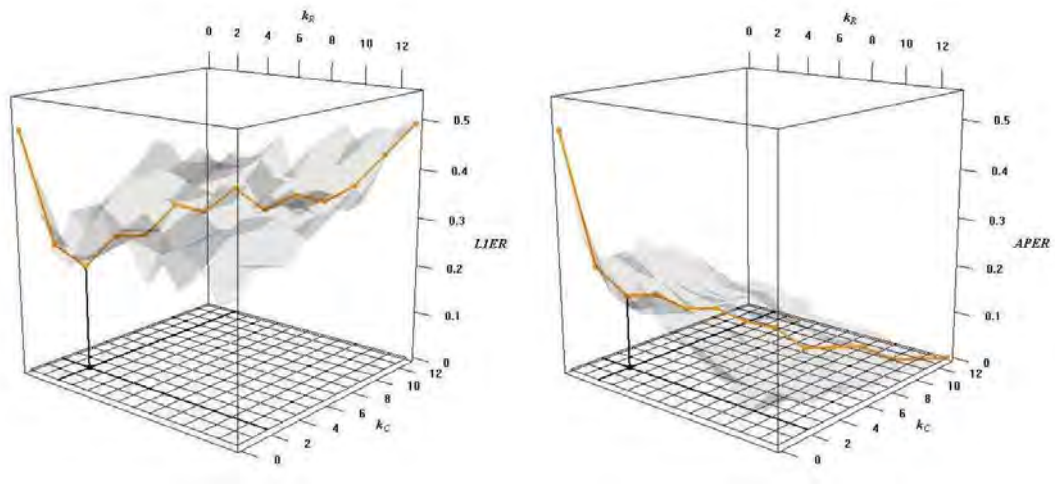

Figure 9: L1ER (left) and APER (right), as function of k_R and k_C. The broken lines correspond to points such that $k_R = k_C$.

the eigenfunctions and of the corresponding scores become progressively worse as the index of eigenfunctions increases.

Hence, it is natural to select the optimal set of eigenfunctions to be used for QDA among those of the form:

$$(\hat{\beta}_{R1}, \hat{\beta}_{R2}, \ldots, \hat{\beta}_{Rk_R}, \hat{\beta}_{C1}, \hat{\beta}_{C2}, \ldots, \hat{\beta}_{Ck_C})$$

with k_R and k_C small enough.

The performance of the prediction rule induced by QDA is measured by an estimate of its AER. We use two different estimators for error rate evaluation:

1. Apparent Error Rate (APER), i.e., the number of sample cases that are misclassified according to the prediction rule obtained by using the entire sample; this estimator is known to be overoptimistic, especially when sample size is not large (Efron and Gong, 1983).

2. Leave-One-Out Error Rate (L1ER), i.e., the number of sample cases that are misclassified according to the prediction rule obtained by iteratively leaving out the observation to be predicted; for small sample sizes, this unbiased estimator should give the most accurate assessment among cross-validation estimators of AER (Ripley, 1996).

In Figure 9, the performance of the prediction rule induced by QDA is shown as function of k_R and k_C. For $k_R = 2$ and $k_C = 2$, the minimum of the L1ER is reached and a marked elbow is also present in the APER. These make $k_R = 2$ and $k_C = 2$ the joint optimal choice for k_R and k_C. As suggested by a referee, we validated this finding by estimating AER also with repeated 10-fold cross-validation (Burman, 1989). Hence, for our final QDA we use the scores relative to the first and second eigenfunctions for radius and curvature. According to L1ER, if $k_R = k_C = 2$ are used, 21.54% of the new patients would be misclassified; hence, the number of misclassified patients, using the prediction rule induced by QDA, is estimated to be less than half the number of patients that would be misclassified by randomly assessing patients to the Lower or Upper group without taking into account FPCA scores.

Inspection of the estimated membership probabilities for the 65 patients using $k_R = k_C = 2$ generates the following remarks:

1. estimated Lower group membership probabilities, for the 32 Lower group patients, range from 1 to almost 0. This means that a patient in the Lower group may have geometrical features similar to those characterizing the Upper group;

2. estimated Upper group membership probabilities for the 33 Upper group patients are all, but one, greater than 0.5. This means that (nearly) no Upper group patient has geometrical features similar to those characterizing the Lower group;

3. many Lower group patients have an estimated Lower group membership probability approximately equal to 1, while no Upper group patient has such a high Upper group membership probability.

In terms of geometrical characterization of the two groups, these facts show that the Lower group patients are more spread in the scores space, whereas the patients in the Upper group are concentrated in a smaller region, nested within the region covered by the Lower group. Upper group patients can thus be interpreted as a subpopulation characterized by better defined geometrical features.

These conclusions are confirmed by inspection of the conditional error rates (Table 1): according to L1ER, predicting correctly a Lower group patient is more difficult than predicting correctly an Upper group patient; in fact the probability of misclassifying a patient belonging to the Lower group (28.12%) is nearly twice as big as the probability of misclassifying a patient belonging to the Upper group (15.15%).

Estimated membership probabilities to the two groups are also used to select patients whose ICA has features that better distinguish them from the other group. The ICA geometries of these patients will be used for fluid-dynamic numerical simulations in order to investigate the impact of morphological features on hemodynamics.

Finally, Student's t-tests, F-tests and a Linear Discriminant Analysis (Hand, 1981), are used to analyze stretching factors m_i and shifting factors p_i determined in the registration phase of the analysis, as described in Section 2. The results of these analyses show that no significant difference exists between the two groups, relative to means and variances of stretching factors m_i and shifting factors p_i. Hence, the information captured by warping functions is ancillary to the problem of classification, confirming the effectiveness of the registration procedure.

L1ER = 21.54%

Abs.	Lower	Upper
Lower	23	9
Upper	5	28

Rel.	Lower	Upper
Lower	35.38%	13.85%
Upper	7.69%	43.08%

Cond.	Lower	Upper
Lower	71.88%	28.12%
Upper	15.15%	84.85%

APER = 15.38%

Abs.	Lower	Upper
Lower	23	9
Upper	1	32

Rel.	Lower	Upper
Lower	35.38%	13.85%
Upper	1.54%	49.23%

Cond.	Lower	Upper
Lower	71.88%	28.12%
Upper	3.03%	96.97%

Table 1: Absolute, relative and conditional confusion matrices estimated according to L1ER (left) and to APER (right). Rows labels refer to true classes, column labels refer to predicted classes. L1ER and APER are also reported on top.

5 Conclusions

The statistical analysis highlights significant differences in the geometry of the last 3 cm of Internal Carotid Arteries of patients with an aneurysm located at or after the terminal bifurcation of the ICA (Upper group) and patients with an aneurysm located before the terminal bifurcation of the ICA or who are healthy (Lower group). These differences refer to both radius and curvature of the ICA. The Upper group patients present significantly wider and more tapered ICA's than the patients of the Lower group; the segment of ICA between the two peaks of curvature is less curved in patients of the Upper group than in patients of the Lower group. Moreover, variability

related to geometrical features is much lower within the Upper group than within the Lower group. Geometrical features are well characterized through projections on the first two functional principal components of radius and curvature. A Quadratic Discriminant Analysis of principal components scores allows us to select cases for numerical simulations.

The results presented here, if confirmed on a larger dataset, can support clinical practice. Indeed, they could be appropriately included in the output generated by the image acquisition device, providing immediate decision support to medical doctors.

References

Adams, R. A. (1975), *Sobolev spaces*, Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London, pure and Applied Mathematics, Vol. 65.

Altman, N. S. and Villarreal, J. C. (2004), "Self-modelling regression for longitudinal data with time-invariant covariates," *Canad. J. Statist.*, 32, 251–268.

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.

Antiga, L. and Steinman, D. A. (2004), "Robust and Objective Decomposition and Maping of Bifurcating Vessels," *IEEE Transactions on Medical Imaging*, 23, 704–713.

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

Bullitt, E. and Aylward, S. (2002), "Volume rendering of segmented image objects," *IEEE Transactions on Medical Imaging*, 21, 998–1002.

Burman, P. (1989), "A comparative study of ordinary cross-validation, v-fold cross-validation and the repeated learning-testing methods," *Biometrika*, 76, 503–514.

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," *American Journal of Neuroradiology*, 27, 1703–1709.

Cleveland, W. S., Grosse, E., and Shyu, W. M. (1992), *Local regression models. Chapter 8 of Statistical Models in S*, eds Chambers, J. M. and Hastie, T. J., Wadsworth & Brooks, Pacific Grove, California.

Efron, B. and Gong, G. (1983), "A leisurely look at the bootstrap, the jackknife, and cross-validation," *Amer. Statist.*, 37, 36–48.

Hand, D. J. (1981), *Discrimination and Classification*, John Wiley & Sons, London.

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," *Journal of Neurosurgery*, 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," *Journal of Nerusurgery*, 101, 676–681.

James, G. M. (2007), "Curve alignment by moments," *The Annals of Applied Statistics*, 1, 480–501.

Jitchote, W. and Robertson, A. M. (2000), "Flow of second order fluids in curved pipes," *J Non-Newton Fluid Mech*, 90, 91–116.

Ke, C. and Wang, Y. (2001), "Semiparametric nonlinear mixed-effects models and their applications," *J. Amer. Statist. Assoc.*, 96, 1272–1298, with comments and a rejoinder by the authors.

Kneip, A., Li, X., MacGibbon, K. B., and Ramsay, J. O. (2000), "Curve registration by local regression," *Canad. J. Statist.*, 28, 19–29.

Krayenbuehl, H., Huber, P., and Yasargil, M. G. (1982), *Krayenbuhl/Yasargil Cerebral Angiography*, Thieme Medical Publishers, 2nd ed.

Lawton, W. H., Sylvestre, E. A., and Maggio, M. S. (1972), "Self Modeling Nonlinear Regression," *Technometrics*, 14, 513–532.

Lindstrom, M. J. and Bates, D. M. (1990), "Nonlinear mixed effects models for repeated measures data," *Biometrics*, 46, 673–687.

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," Available at www.mathcs.emory.edu.

Ramsay, J. O. (2000), "Functional Components of Variation in Handwriting," *Journal of the American Statistical Association*, 95, 9–15.

Ramsay, J. O. and Li, X. (1998), "Curve registration," *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 60, 351–363.

Ramsay, J. O. and Silverman, B. W. (2005), *Functional Data Analysis*, Springer New York NY, 2nd ed.

Rinkel, G. J., Djibuti, M., Algra, A., and Van Gijn, J. (1998), "Prevalence and risk of rupture of intracranial aneurysms: A systematic review," *Stroke*, 29, 251–256.

Ripley, B. D. (1996), *Pattern Recognition and Neural Networks*, Cambridge University Press, Cambridge.

Sangalli, L. M., Secchi, P., Vantini, S., and Veneziani, A. (2007), "Efficient estimation of 3-dimensional centerlines of inner carotid arteries and their curvature functions by free knot regression splines," Tech. Rep. 23/2007, MOX, Dipartimento di Matematica, Politecnico di Milano.

Sheehy, A., Gasser, T., Molinari, L., and Largo, R. H. (2000a), "An analysis of variance of the pubertal and midgrowth spurts for length and width," *Annals of Human Biology*, 26, 309–331.

— (2000b), "Contribution of growth phases to adult size," *Annals of Human Biology*, 27, 281–298.

Wang, H. and Marron, J. S. (2007), "Object oriented data analysis: sets of trees," *Ann. Statist.*, 35, 1849–1873.