Published on *Biometrika* (2016), 103 (1), 71-88.
Publisher's version available at http://biomet.oxfordjournals.org/content/103/1/71.abstract

Spatial regression models over two-dimensional manifolds

BY B. ETTINGER

Mathematics and Computer Science Department, Emory University, Atlanta, GA 30322, U.S.A.
betting@emory.edu

S. PEROTTO AND L. M. SANGALLI

Laboratory for Modeling and Scientific Computing MOX, Dipartimento di Matematica, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy
simona.perotto@polimi.it laura.sangalli@polimi.it

SUMMARY

We propose a regression model for data spatially distributed over general two-dimensional Riemannian manifolds. This is a generalized additive model with a roughness penalty term involving a differential operator computed over the non-planar domain. Owing to a semi-parametric framework, the model allows the inclusion of space-varying covariate information. Estimation can be performed by conformally parameterizing the non-planar domain and then by generalizing existing models for penalized spatial regression over planar domains. The conformal coordinates and the estimation problem are both computed by a finite element approach.

Some key words: generalized additive model; partial differential regularization; penalized regression; smoothing on manifolds.

1. INTRODUCTION

1·1. *Motivation*

We consider smoothing data distributed over general non-planar two-dimensional domains, and more generally regression for data distributed over non-traditional surface domains. The applied problem driving this research is the study of hemodynamic forces, such as shear stress and pressure, exerted by blood-flow on the wall of an internal carotid artery. The data used in this study are part of the AneuRisk project, a scientific endeavor that investigated the role of vessel morphology, blood fluid dynamics and biomechanical properties of the vascular wall in the pathogenesis of cerebral aneurysms; see Passerini et al. (2012) and Sangalli et al. (2009a). Figure 1(a) shows an internal carotid artery affected by an aneurysm, a deformation of the vessel characterized by a bulge of the vessel wall. The shear stress exerted by the blood flow on the wall of the artery, at the systolic peak, is represented by a colormap. Each value refers to a point $(x_{[1]}, x_{[2]}, x_{[3]})$ on the bi-dimensional and non-planar artery wall. Within the AneuRisk project, these data have been analyzed by simplifying the three-dimensional artery to a cylinder. A bijective cylindrical map, which implicitly fixes the radius to a constant and does not account for the curvature of the vessel, is used to flatten the artery wall; standard spatial methods are then applied in the resulting planar domain. Though very convenient from an analytical point of view, this approach is inaccurate, since the radius and curvature of the vessel are geometrical quantities that greatly influence the hemodynamics in an artery, and are shown to be discriminating factors of aneurysm presence and location (Sangalli et al., 2009a). Moreover, this approach is not able

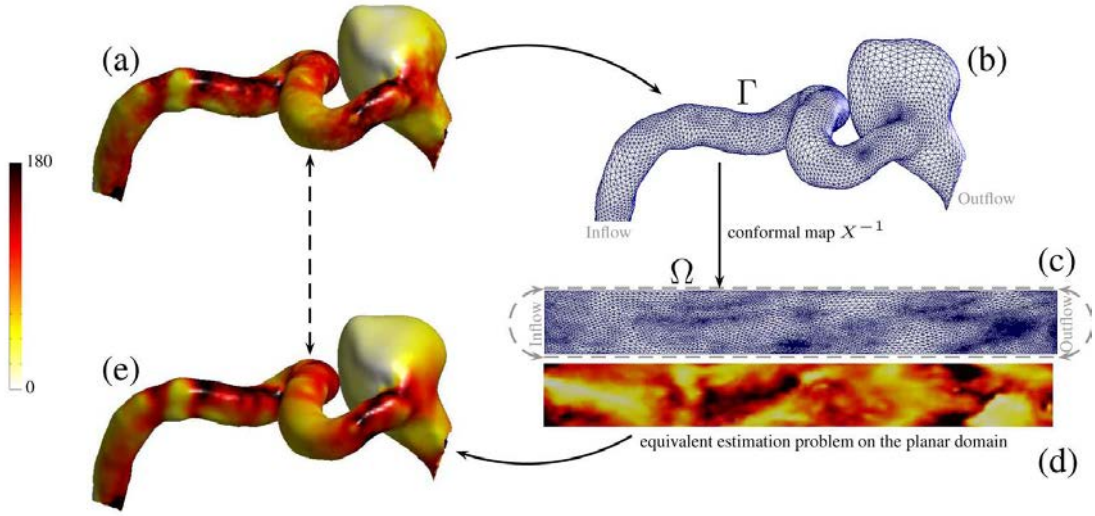

Fig. 1. A visual illustration of the method. (a): Shear stress (dyn/cm^2) exerted by blood flow on the wall of an internal carotid artery affected by an aneurysm. The colormap shows the modulus of the wall shear stress at the systolic peak. The colormap ranges from 0 dyn/cm^2 (white) to 180 dyn/cm^2 (red). (b): A triangular mesh reconstruction of the wall of the internal carotid artery in (a). (c): Planar triangular mesh generated by a conformal flattening of the mesh in (b). (d): Solution to equivalent estimation problem solved on the planar domain in (c). (e): Solution to estimation problem mapped back to the original manifold domain.

to consider vessels affected by a large aneurysm, as in Figure 1. For these types of domain geometries it is impossible to define a bijective cylindrical map unless the aneurysmal sac is removed. This suggests the development of an approach able to take into account the geometry of the domain. In addition, we also want to include the local vessel geometry as space-varying covariate information to better explore the relationship between the morphological features of the vessel and the hemodynamics.

Non-planar domains are usually approximated by three-dimensional triangular meshes characterized by varying distances and angles between neighbouring vertices. Figure 1(b) displays an example of a triangular mesh that approximates the artery wall in Figure 1(a). Few methods are available to deal with data on non-planar domains. Iterative schemes for nearest neighbor averaging have been developed to work on surface meshes (Hagler et al., 2006). In this approach, the value of the variable of interest at each vertex of the mesh is obtained by averaging the values at the neighboring vertices. This process is repeated several times to create a smoothing effect. A more sophisticated method, heat kernel smoothing, is presented, e.g., in Chung et al. (2005). This is a geodesic distance-based kernel smoothing approach for non-planar domains. However, neither nearest neighbor averaging nor heat kernel smoothing is well-suited for our applied problem, as they are not designed to include space-varying covariates. Among the other models devised to handle data over non-planar domains we cite the approaches devoted to specific types of Riemannian manifolds such as spheres, hyperspheres, e.g., the spherical splines introduced in Wahba (1981) and Baramidze et al. (2006), the models described by Jun & Stein (2007), Jun (2011) and Gneiting (2013), and methods that work on more general sphere-like domains, e.g., those of Alfeld et al. (1996) and Lindgren et al. (2011).

In this paper we propose a regression method that efficiently handles data distributed over general two-dimensional Riemannian manifolds and also includes space-varying covariate information, generalizing the work of Ramsay (2002), Sangalli et al. (2013) and Azzimonti et al. (2014b) to the case of non-planar domains. The resulting estimators are linear in the observed data values and have a penalized regression form. Moreover, the proposed model has a high com-

putational efficiency, thanks to the use of scientific computing techniques and in particular of the finite element approach. The excellent performance of the proposed methodology is shown by simulation studies on both test domains and domains from real artery geometries. The proposed models have important fields of application not only in the medical sciences, but also in the geosciences, environmental sciences and engineering.

1·2. *The model*

Consider n fixed data locations $\{x_i = (x_{i[1]}, x_{i[2]}, x_{i[3]}) : i = 1, \ldots, n\}$ lying on a non-planar domain Γ that is a surface embedded in $\mathbb{R}^3$. For each location x_i, a real-valued random variable of interest, z_i, is observed. We assume the model

$$z_i = f(x_i) + \epsilon_i \quad (i = 1, \ldots, n) \tag{1}$$

where ϵ_i are independent errors with zero mean and constant variance σ^2, and f is a twice continuously differentiable real-valued function defined on the surface domain Γ.

To estimate the function f in (1), we propose to minimize a penalized sum of squared error functional, with a roughness penalty involving a differential operator computed over Γ. In particular, we generalize to manifold domains the spatial regression models over planar domains in Ramsay (2002) and Sangalli et al. (2013), where the roughness penalty involves the standard Laplacian of the function to be estimated. In detail, let us consider a surface $s = s(u_{[1]}, u_{[2]})$ defined on a planar domain Ω, with $u = (u_{[1]}, u_{[2]})$ a generic point in Ω. Denote the gradient of s by $\nabla s(u) = \big(\partial s/\partial u_{[1]}(u), \partial s/\partial u_{[2]}(u)\big)^T$, where ^{T} is the transpose operator. Moreover, given a vector field $v(u) = \big\{v_{[1]}(u), v_{[2]}(u)\big\}^T$ on Ω, define the divergence of the vector field as $\mathrm{div}\, v(u) = \partial v_{[1]}/\partial u_{[1]}(u) + \partial v_{[2]}/\partial u_{[2]}(u)$. Then, the Laplacian of the surface s is defined as $\Delta s(u) = \mathrm{div}\{\nabla s(u)\} = \partial^2 s/\partial u^2_{[1]}(u) + \partial^2 s/\partial u^2_{[2]}(u)$. The Laplacian Δs provides a simple measure of the local curvature of the surface s and is invariant with respect to rigid transformations of the spatial coordinates of Ω. The generalization of the Laplacian to functions $f = f(x_{[1]}, x_{[2]}, x_{[3]})$ defined on a non-planar domain Γ, with $x = (x_{[1]}, x_{[2]}, x_{[3]}) \in \Gamma$, requires the computation of the gradient operator ∇_Γ and of the divergence operator div_Γ associated with Γ (see, e.g., Dierkes et al., 2010, and Section 2). It is thus possible to define the Laplace–Beltrami operator of f, represented by $\Delta_\Gamma f(x) = \mathrm{div}_\Gamma \nabla_\Gamma f(x)$. Similarly to the standard Laplacian, the Laplace–Beltrami operator provides a simple measure of the local curvature of the function f, and is invariant with respect to rotations, translations and reflections of the spatial coordinates of Γ. Thus, the use of the Laplace–Beltrami operator in a roughness penalty ensures that the degree of smoothness does not depend on the orientations of the coordinate system or of the domain Γ. As a consequence, we propose to estimate f by minimizing the functional

$$J_{\Gamma,\lambda}(f) = \sum_{i=1}^{n} \{z_i - f(x_i)\}^2 + \lambda \int_\Gamma \{\Delta_\Gamma f(x)\}^2 \, d\Gamma \tag{2}$$

where λ is a positive smoothing parameter; the higher λ is, the more we control the roughness of f, the smaller λ is, the more we allow flexibility to fit the local curvature of f. The Laplace–Beltrami operator is also used by heat kernel smoothing, although in a different framework, as detailed in Section 5. The spherical splines introduced in Wahba (1981), Alfeld et al. (1996) and Baramidze et al. (2006) consider the same penalized sum of squared errors functional in (2), or similar functionals where the Laplace–Beltrami operator is replaced by other differential quantities.

In this paper, we present a solution to the estimation problem (2) that combines tools from differential geometry, analysis and scientific computing. The key idea is summarized in Figure 1. We first conformally parameterize the non-planar domain Γ, so that we can conformally flatten Γ into a planar domain; see Figure 1, (b) to (c). In particular, the conformal parametrization allows us to reformulate the estimation problem (2) over a planar domain Ω, by providing an appropriate change of variable from the surface domain Γ to Ω. This leads to an estimation problem over Ω where the geometry of the original domain Γ is appropriately accounted for. In this new setting, the problem is solved by generalizing spatial regression models over planar domains; see Figure 1(d). From a computational viewpoint, we use finite elements to compute the conformal flattening map and to solve the equivalent estimation problem on the planar domain; see Figure 1, (b) to (d). Finally, by properly exploiting the conformal map between Γ and Ω, we are able to provide the estimator over the original non-planar domain; see Figure 1, (d) to (e). The proposed model has been implemented in R and Matlab.

2. FORMULATING EQUIVALENT ESTIMATION PROBLEM ON A PLANAR DOMAIN

2·1. *Conformal parametrization*

We consider a non planar domain Γ that is a two dimensional Riemannian manifold. We assume in the sequel that Γ admits a global parametrization

$$\begin{aligned} X : \Omega &\to \Gamma \\ u = (u_{[1]}, u_{[2]}) &\mapsto x = (x_{[1]}, x_{[2]}, x_{[3]}), \end{aligned} \tag{3}$$

where Ω is an open, convex and bounded set in $\mathbb{R}^2$ and the boundary $\partial\Omega$ of Ω is sufficiently regular. The map X provides a change of variable between the planar coordinates $u = (u_{[1]}, u_{[2]})$ and the non-planar coordinates $x = (x_{[1]}, x_{[2]}, x_{[3]})$. Such global parametrizations are available for a large class of bidimensional manifolds, including surfaces with self-intersections and holes and other non-trivial geometries; see, e.g., Gu & Yau (2003). In particular, global parametrizations are available for all the geometries considered throughout the paper. Our goal is to use X to rewrite the estimation problem (2) over a planar domain Ω; see Figure 1, (a) to (d). Consider the first order partial derivatives of X with respect to the planar coordinates $u_{[1]}$ and $u_{[2]}$, $\partial X/\partial u_{[1]}(u)$ and $\partial X/\partial u_{[2]}(u)$, which are column vectors in $\mathbb{R}^3$. Let $\langle\cdot,\cdot\rangle$ denote the Euclidean scalar product of two vectors and $\|\cdot\|$ the corresponding norm. Then, the map X is said to be conformal if $\left\|\partial X/\partial u_{[1]}(u)\right\|^2 = \left\|\partial X/\partial u_{[2]}(u)\right\|^2$ and $\langle \partial X/\partial u_{[1]}(u), \partial X/\partial u_{[2]}(u)\rangle = 0$, for any $u \in \Omega$. Such maps are unique up to dilations, rotations and translations; see, e.g., Hurdal & Stephenson (2009). They preserve angles and the most important geometrical features of the domain. The use of a map X with these properties allows us to obtain a simplified version of the estimation problem.

To rewrite problem (2) on Ω, we consider the Jacobian matrix of X, defined as $\nabla X(u) = \left\{\partial X/\partial u_{[1]}(u),\ \partial X/\partial u_{[2]}(u)\right\}$, for $u \in \Omega$; this matrix has maximal rank equal to two. Using the Jacobian, we define the symmetric positive definite matrix

$$G(u) = \nabla X(u)^T \nabla X(u) = \begin{pmatrix} \left\|\frac{\partial X}{\partial u_{[1]}}(u)\right\|^2 & \langle \frac{\partial X}{\partial u_{[1]}}(u), \frac{\partial X}{\partial u_{[2]}}(u)\rangle \\ \langle \frac{\partial X}{\partial u_{[1]}}(u), \frac{\partial X}{\partial u_{[2]}}(u)\rangle & \left\|\frac{\partial X}{\partial u_{[2]}}(u)\right\|^2 \end{pmatrix} \qquad u \in \Omega.$$

The matrix $G(u)$ is the space-varying metric tensor associated with the map X: for any point u in the planar domain Ω, $G(u)$ describes the local geometry of the non-planar domain Γ, at the corresponding point $x = X(u)$. The determinant of $G(u)$ is strictly positive; in particular, there

exists a positive constant η such that, for any $u \in \Omega$, $\mathcal{D}(u) = \det\{G(u)\}^{1/2} \geq \eta$. The quantity $\mathcal{D}(u)$ identifies the elemental area involved in the change of variable from the original non-planar coordinates $x \in \Gamma$ to the planar coordinates $u \in \Omega$, i.e, $d\Gamma = \mathcal{D}(u)d\Omega$. Finally, we introduce the matrix $K(u) = \mathcal{D}(u)\, G^{-1}(u)$, which is also symmetric positive definite for any $u \in \Omega$, since $\mathcal{D}(u)$ is positive and the inverse metric tensor $G^{-1}(u)$ is symmetric and inherits positive definiteness from $G(u)$.

After setting the geometrical framework, we now focus on the function f to be estimated in model (1). By exploiting the map X, instead of the function $f(x)$ defined on the manifold Γ, we can consider the function $f \circ X(u) = f\{X(u)\}$ defined over the planar domain Ω. The function $f \circ X$ provides a planar parametrization of f. Owing to the regularity assumptions on f and X, $f \circ X \in \mathcal{C}^2(\bar{\Omega})$. Using the map X, the Γ-gradient of f can be parametrized in terms of the planar coordinates u as $\nabla_\Gamma f(x) = \nabla X(u)\ G^{-1}(u) \nabla f\{X(u)\} \in \mathbb{R}^3$, for any $x \in \Gamma$ and with $u = X^{-1}(x)$, where $\nabla f\{X(u)\}$ denotes the standard gradient on Ω defined in Section 1·2; see, e.g., Dierkes et al. (2010). Similarly, the Γ-divergence of a vector field $v = \big\{v_{[1]}(x), v_{[2]}(x)\big\}^T$, defined on the non-planar domain Γ, can be parametrized in terms of the planar coordinates u as $\mathrm{div}_\Gamma v(x) = \mathcal{D}(u)^{-1}\big[\partial/\partial_{u_{[1]}} \mathcal{D}(u) v_{[1]}\{X(u)\} + \partial/\partial_{u_{[2]}} \mathcal{D}(u) v_{[2]}\{X(u)\}\big]$. Hence, the Laplace–Beltrami operator can be re-expressed in terms of the planar coordinates u as

$$\Delta_\Gamma f(x) = \mathrm{div}_\Gamma\big[\nabla_\Gamma f\{X(u)\}\big] = \frac{1}{\mathcal{D}(u)} \mathrm{div}\big[K(u) \nabla f\{X(u)\}\big], \tag{4}$$

where div denotes the standard divergence operator over Ω defined in Section 1·2. Thus, by exploiting (4), after setting $u_i = X^{-1}(x_i)$, we can reformulate the estimation problem (2) over the manifold Γ as an equivalent problem over the planar domain Ω as follows: find the function $f \circ X$, defined on Ω, that minimizes the functional

$$J_{\Omega,\lambda}(f \circ X) = \sum_{i=1}^{n} \big[z_i - f\{X(u_i)\}\big]^2 + \lambda \int_\Omega \frac{1}{\mathcal{D}(u)} \Big(\mathrm{div}\big[K(u) \nabla f\{X(u)\}\big]\Big)^2 d\Omega. \tag{5}$$

Problem (5) can be considered as a generalization of the estimation problem for spatial regression models over planar domains in Ramsay (2002) and Sangalli et al. (2013) to a more complex setting, characterized by the presence of the quantities $K(u)$ and $\mathcal{D}(u)$ in the roughness penalty term. These quantities appropriately account for the geometry of the original manifold Γ. As previously remarked, the map X and its corresponding planar domain Ω are not unique; see Section 2·3. However, the terms $K(u)$ and $\mathcal{D}(u)$ adjust for each map X considered and for the corresponding planar domain Ω. This leads to different planar parameterizations of the function f, as well as of the estimation problem (5), but all equivalent to problem (2) on the original manifold Γ.

In the special case of a conformal map X, we obtain a simplified estimation problem over Ω. In fact, for a conformal map we have $\mathcal{D}(u) = \big\|\partial X/\partial u_{[1]}(u)\big\|^2$, $G(u) = \mathcal{D}(u) I_2$, $G^{-1}(u) = \mathcal{D}(u)^{-1} I_2$ and $K(u) = I_2$, where I_m is the identity matrix of order m. The Laplace–Beltrami operator (4) also simplifies to $\Delta_\Gamma f\{X(u)\} = \mathcal{D}(u)^{-1} \Delta f\{X(u)\}$, where Δ is the standard Laplace operator defined on Ω. Thus, the estimation problem (5) simplifies to finding the function $f \circ X$, defined on Ω, minimizing the functional

$$J_{\Omega,\lambda}(f \circ X) = \sum_{i=1}^{n} \big[z_i - f\{X(u_i)\}\big]^2 + \lambda \int_\Omega \Big[\mathcal{D}(u)^{-1/2} \Delta f\{X(u)\}\Big]^2 d\Omega. \tag{6}$$

2·2. *Characterization of the estimation problem on the planar domain*

We hence characterize the solution to the estimation problem (5) on the planar domain; see Figure 1(d). We have to introduce a suitable function setting. Consider the Sobolev space $H^m(\Omega)$, i.e., the space of the functions which belong to $L^2(\Omega)$ along with all their distributional derivatives up to the order m (Lions & Magenes, 1973). Here we propose a variant of the function space $H^m(\Omega)$ given by the space $H^m_{n0,K}(\Omega) = \{h \in H^m(\Omega) \, : \nu^T K \nabla h = 0 \text{ on } \partial\Omega\} \subset H^m(\Omega)$, consisting of the $H^m(\Omega)$-functions with normal derivative $\nu^T K \nabla h$ identically equal to zero on the boundary $\partial\Omega$, where ν is the unitary outward vector normal to $\partial\Omega$. The requested condition on the derivative is equivalent to the condition that the normal derivative on the boundary of Γ vanishes. The functional $J_{\Omega,\lambda}(f \circ X)$ in (5) is well-defined since, for $(f \circ X) \in H^2(\Omega)$, the roughness and sum of squared errors terms are well-posed, owing to the inclusion $H^2(\Omega) \subset C^0(\bar{\Omega})$ that allows for pointwise evaluation of $f \circ X$. The boundary conditions in $H^2_{n0,K}(\Omega)$ ensure uniqueness of the solution to the estimation problem; see the Supplementary Material.

Let $z = (z_1, \ldots, z_n)^T$ be the vector collecting the observed values in (1) for the quantity of interest. For any function h defined on Γ, such that $h \circ X$ is defined on Ω, we denote the column vector of evaluations of the function h at the n data locations x_i by

$$\mathrm{h}_n = \big[h(x_1), \ldots, h(x_n)\big]^T = \big[h\{X(u_1)\}, \ldots, h\{X(u_n)\}\big]^T. \tag{7}$$

To simplify the notation, in the following we suppress the dependence on u. The following proposition states that the estimation problem is well posed and characterizes the estimator.

PROPOSITION 1. *The estimator $\hat{f} \circ X$ that minimizes (5) over $H^2_{n0,K}(\Omega)$ satisfies the relation*

$$\mathrm{q}_n^T \hat{\mathrm{f}}_n + \lambda \int_\Omega \frac{1}{\mathcal{D}} \operatorname{div}\big\{K \nabla (q \circ X)\big\} \operatorname{div}\big\{K \nabla (\hat{f} \circ X)\big\} d\Omega = \mathrm{q}_n^T z \tag{8}$$

for any function q defined on Γ such that $q \circ X \in H^2_{n0,K}(\Omega)$, with q_n and $\hat{\mathrm{f}}_n$ defined according to (7). Moreover, the estimator $\hat{f} \circ X$ exists and is unique.

For any fixed map X, Proposition 1 establishes the existence and uniqueness of the solution to the estimation problem (5) on Ω. Since, as shown previously, this problem is equivalent to the estimation problem (2) on Γ, Proposition 1 also guarantees existence and uniqueness of the solution to (2). Its proof is given in the Supplementary Material.

2·3. *Conformal flattening of non-planar domains*

In this section we describe how the map X and the corresponding planar domain Ω are determined; see Figure 1, (b) to (c). Conformal parameterizations, a very active area of research, are extensively used in graphics in a variety of applications from re-meshing to texture mapping. For surfaces without self-intersections and holes, the approaches for approximating a conformal parameterization fall into two categories: harmonic energy minimization (Angenent et al., 1999) or circle packing (Hurdal & Stephenson, 2009). Here, we employ the first approach. Since the application considered in this paper features domains topologically equivalent to a cylinder, we adopt the method for flattening tubular surfaces developed in Haker et al. (2000). We describe the construction of the conformal map X for a tubular surface with the help of the example in Figure 2. A tubular surface Γ has the same topology as an open ended cylinder and is embedded in $\mathbb{R}^3$. The open ends of the cylinder, denoted by b_0 and b_1, represent the boundary of Γ and are homeomorphic to a circle. With these assumptions, Γ is conformally equivalent to a rectangle in $\mathbb{R}^2$.

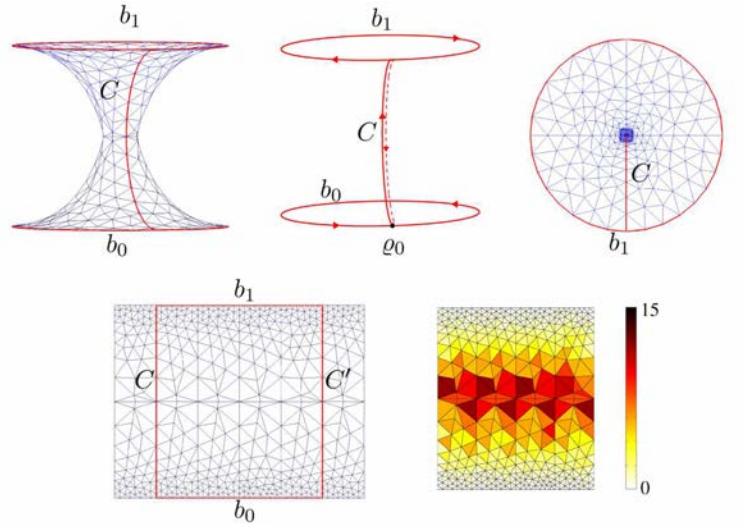

Fig. 2. Example of conformal flattening. Top left: three-dimensional triangular mesh approximating a non-planar test domain. Top center: the oriented boundary B. Top right: the annulus generated by the flattening procedure. Bottom left: the planar triangulated domain obtained as a final product of the conformal flattening of the test domain. Bottom right: colormap of the quantity $\mathcal{D}(u)$ associated with the conformal flattening map.

The map X in (3), having the conformal properties, is identified by the solution $(u_{[1]}, u_{[2]})$ of two coupled differential problems defined on the original manifold Γ. In particular, the map X is calculated in two steps, the first providing the conformal coordinate $u_{[1]}$ and the second providing the conformal coordinate $u_{[2]}$. In particular, the coordinate $u_{[1]}$ is determined by solving

$$\begin{cases} -\Delta_\Gamma u_{[1]} = 0 & \text{on } \Gamma, \\ u_{[1]} = 0 & \text{on } b_0, \\ u_{[1]} = 1 & \text{on } b_1. \end{cases} \tag{9}$$

To define the differential problem identifying the coordinate $u_{[2]}$, let us cut the surface Γ along a curve C that runs from b_0 to b_1 such that $u_{[1]}$ is strictly increasing along C; see Figure 2, top left. The maximum principle ensures the existence of such a cut (Lions & Magenes, 1973). The cut is not uniquely determined; Section 3·2 explains how to avoid any influence of the cut on the estimation problem. The boundaries b_0 and b_1, together with the cut C, form an oriented boundary, denoted by B, which circles around b_0, then along C, around b_1, and finally back down C in the opposite direction. Thus, the boundary B is a closed curve whose direction is determined by the orientation of the surface; see Figure 2, top center. At this point, we define $u_{[2]}$ as the solution to the differential problem

$$\begin{cases} -\Delta_\Gamma u_{[2]} = 0 & \text{on } \Gamma \\ u_{[2]}(\varrho) = \int_{\varrho_0}^{\varrho} \frac{\partial u_{[1]}}{\partial \nu}\, ds & \text{on } B, \end{cases} \tag{10}$$

where $\varrho_0 \in b_0$ is a designated starting point of the boundary B, $\varrho \in B$, $\partial u_{[1]}/\partial \nu$ is the normal derivative of $u_{[1]}$ and ds denotes the arc-length element along B.

The combined effect of problems (9) and (10) is visualized in Figure 2. The top left panel shows a three-dimensional tubular test surface approximated by a triangular mesh. The boundary B is highlighted in red and its orientation is shown in Figure 2, top center. The annulus generated during the flattening procedure is shown in Figure 2, top right. This annulus has b_0 as inner boundary, not clearly visible in the figure, and b_1 as outer boundary. The rectangle generated at

the end of the flattening is shown in Figure 2, bottom left, and it is enclosed by the red lines that correspond to the boundary B in Figure 2, top center. Moreover, the rectangular planar domain is properly scaled to maintain the same proportion as the original surface, i.e., the height of the rectangle is set to $2\pi/|C|$, where $|C|$ is the length of the cut.

3. COMPUTATION OF THE CONFORMAL MAP AND OF THE ESTIMATOR

3·1. *Finite Elements*

Problems (9)–(10) describing the conformal map X and problem (8) characterizing the estimator $\hat{f} \circ X$ are infinite-dimensional and cannot be solved analytically. A standard approach consists of recasting them into a finite-dimensional function space that approximates the original infinite-dimensional space. Then, the problem is solved in the finite-dimensional space, normally by exploiting a finite-dimensional basis spanning it. In particular, since (9)–(10) and (8) involve partial differential operators, we resort to finite element spaces and the associated basis. The finite element method is widely used in engineering applications to deal with problems involving partial differential equations (Gockenbach, 2006; Quarteroni, 2014). Finite elements are a generalization of the widely-employed univariate spline functions, even though finite elements are usually less regular than splines. Similarly to univariate splines, finite elements partition the domain of interest, either planar or non-planar, into subdomains. Then, the solution is approximated via a globally continuous function which coincides with a polynomial of a certain degree on each subdomain, i.e., a piecewise polynomial. In particular, analogously to the basis for univariate splines used to describe piecewise polynomial curves, the finite element basis is used to describe piecewise polynomial surfaces. The use of finite elements makes problems (9)–(10) and (8) numerically tractable and reduces them to linear systems.

Convenient domain partitions, in both the planar and non-planar settings, are provided by triangular meshes. Figure 1(b) and (c) and Figure 2 show examples of both non-planar and planar triangular meshes. In the triangulation, two adjacent triangles share either a vertex or a complete edge and the union of all the triangles provides an approximation for the whole domain. The boundary of the domain and any interior hole are represented by the polygon generated by the outer edges of the triangulation. In the simulations and applications considered in this paper the vertices of the non-planar mesh coincide with the data location points $x_1, \ldots, x_n$, so that the conformal flattening provides a planar mesh with n vertices coinciding with the corresponding data locations $u_i = X^{-1}(x_i)$ on the planar domain. Below we consider the general case where mesh vertices and data locations do not necessarily coincide.

Starting from a triangulation $\mathcal{T}$ of a non-planar or planar domain, we can introduce a locally supported finite element basis that spans the space of the piecewise polynomials over $\mathcal{T}$. This finite element space, denoted by $H^1_{\mathcal{T}}$, discretizes the infinite-dimensional space H^1. In this paper, we use linear finite elements where each basis function ψ_j is associated with a triangle vertex ξ_j $(j = 1, \ldots, N)$. The basis function ψ_j coincides with the jth hat function, a piecewise linear polynomial which assumes the value one at the vertex ξ_j and the value zero on all the other vertices of the mesh, i.e., $\psi_j(\xi_l) = \delta_{jl}$, where δ_{jl} denotes the Kronecker delta symbol. Figure 3 shows an example of a linear finite element basis function on non-planar and planar triangulations. The locally supported nature of the basis functions ψ_j is evident.

Now, let $\psi = (\psi_1, \ldots, \psi_N)^T$ be the column vector collecting the N basis functions associated with the N vertices ξ_j $(j = 1, \ldots, N)$. Then, each function h in the finite element space $H^1_{\mathcal{T}}$ can be represented as an expansion in terms of the basis function $\psi_1, \ldots, \psi_N$. In particular, $h(\cdot) = \sum_{j=1}^{N} h(\xi_j)\psi_j(\cdot) = \mathrm{h}^T\psi(\cdot)$, where $\mathrm{h} = \{h(\xi_1), \ldots, h(\xi_N)\}^T$ is the column vector of the

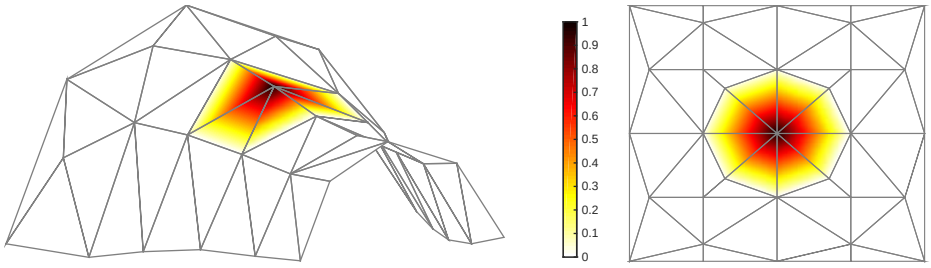

Fig. 3. Colormap of a linear finite element basis on a non-planar (left) and on a planar (right) triangulation.

evaluations of h on the N vertices of the mesh. Each function $h \in H^1_{\mathcal{T}}$ is thus uniquely identified by its evaluations on the mesh vertices.

The finite element space $H^1_{\mathcal{T}}$ is characterized by less regularity with respect to the spaces involved in (9)–(10) and (8) which are subsets of H^2. This will lead us to provide an equivalent formulation for problems (9)–(10) and (8) in a suitable subset of the space H^1, that only involves first order partial derivatives; see (11) and (12), respectively.

3·2. *The finite element computation of the flattening map*

To compute the conformal coordinates in (9)–(10) we approximate both coordinates $u_{[1]}$ and $u_{[2]}$ with functions $H^1_{\mathcal{T}}(\Gamma)$, which are globally continuous and coincide with a linear polynomial over each triangle of the non-planar mesh $\Gamma_{\mathcal{T}}$ approximating the manifold Γ. To obtain an equivalent formulation of problems (9)–(10) suited for this finite element discretization we use a classical result stating that $u_{[1]}$ in (9) is the minimizer of the energy functional

$$\mathcal{E}(u_{[1]}) = \frac{1}{2}\int_{\Gamma} \|\nabla_{\Gamma} u_{[1]}\|^2 d\Gamma \tag{11}$$

and the same result applies to $u_{[2]}$ in (10) (Haker et al., 2000; Dierkes et al., 2010). This reformulation reduces the regularity assumptions on $u_{[1]}$ and $u_{[2]}$ with respect to the original problems (9)–(10) requiring $u_{[1]}, u_{[2]} \in H^2(\Gamma)$. Now the problem is well defined in the space $H^1(\Gamma)$, which can be approximated with the finite element space $H^1_{\mathcal{T}}(\Gamma)$.

The conformal flattening method maps the non-planar triangular mesh $\Gamma_{\mathcal{T}}$, that approximates the surface domain Γ, to a planar triangulation $\Omega_{\mathcal{T}}$ of Ω, without flipping or breaking any of the triangles; see Figure 1, (b) to (c). Moreover, this map provides the terms $K(u)$ and $\mathcal{D}(u)$ associated with the map X. The computation of these terms, together with the planar mesh $\Omega_{\mathcal{T}}$, conveniently set up the planar estimation problem (8) to be solved via planar finite elements, as detailed in the following section. The Supplementary Material furnishes more details.

Remark 1. The conformal rectangular domain has two artificial boundaries, generated by the cut, denoted by C and C'; see Figure 2, bottom left. There are two alternative approaches to ensure the necessary periodicity of the estimate of f along this artificial cut. The first approach consists in repeating the planar triangulation and the corresponding data values along C and C', as shown in Figure 2, bottom left; due to the locally supported nature of a finite element basis, it is in fact sufficient to repeat only a small portion of the triangulated domain. The second approach consists in imposing periodic boundary conditions, enforcing $(f \circ X)|_C = (f \circ X)|_{C'}$ and $\nu^T \nabla(f \circ X)\big|_C = -\nu'^T \nabla(f \circ X)\big|_{C'}$, where ν and ν' are the unitary outward normal vectors to C and C', while we denote the restriction of the function $f \circ X$ to the generic set $A \subset \mathbb{R}^2$ by $(f \circ X)|_A$. The second approach leads to computational savings with respect to the first, but formally demands the introduction of a new function space to properly rewrite the estimation

problem as well as to deal with boundary terms. In Sections 3·3 and 4, for simplicity of exposition, we follow the former approach, but both methods are implemented and numerically assessed in Sections 5 and 6.

3·3. *Finite element solution to the estimation problem*

We now show how to solve the estimation problem (8) using the finite element space $H^1_{\mathcal{T}}(\Omega)$, associated with the planar mesh $\Omega_{\mathcal{T}}$ yielded in the previous step of the procedure. By introducing a proper auxiliary function g, as detailed in the Supplementary Material, it is possible to obtain the following reformulation of the estimation problem (8), suited for a finite element approximation in the space $H^1_{\mathcal{T}}(\Omega)$: find $(\hat{f} \circ X, g \circ X) \in \{H^1_{n0,K}(\Omega) \cap C^0(\bar{\Omega})\} \times H^1(\Omega)$ such that

$$\begin{aligned} \mathrm{q}_n^T \hat{\mathrm{f}}_n - \lambda \int_\Omega \nabla(q \circ X)^T K \nabla(g \circ X) d\Omega &= \mathrm{q}_n^T z \\ \int_\Omega (p \circ X)(g \circ X) \mathcal{D} d\Omega + \int_\Omega \nabla(p \circ X)^T K \nabla(\hat{f} \circ X) d\Omega &= 0 \end{aligned} \tag{12}$$

for any $(q \circ X, p \circ X) \in \{H^1_{n0,K}(\Omega) \cap C^0(\bar{\Omega})\} \times H^1(\Omega)$. This formulation requires less regularity on the functions involved with respect to formulation (8), defined in $H^2_{n0,K}(\Omega)$. Nevertheless it can be shown that the solution $\hat{f} \circ X$ to (12) belongs to $H^2_{n0,K}(\Omega)$. See, e.g., Lions & Magenes (1973).

The functions and integrals in (12) can be approximated using functions in the finite element space $H^1_{\mathcal{T}}(\Omega)$, so that problem (12) is approximated with its discrete counterpart: find $(\hat{f} \circ X, g \circ X) \in H^1_{\mathcal{T}}(\Omega) \times H^1_{\mathcal{T}}(\Omega)$ that satisfy (12) for any $(q \circ X, p \circ X) \in H^1_{\mathcal{T}}(\Omega) \times H^1_{\mathcal{T}}(\Omega)$, where the integrals are now computed over the triangulation $\Omega_{\mathcal{T}}$. Let $\psi = (\psi_1, \ldots, \psi_N)^T$ be the column vector of finite element basis functions that span $H^1_{\mathcal{T}}(\Omega)$, associated with the nodes $\xi_1, \ldots, \xi_N$ of $\Omega_{\mathcal{T}}$. Let $\{\Psi\}_{i,j} = \psi_j(u_i)$, $(i = 1, \ldots, n;\ j = 1, \ldots, N)$ be the $n \times N$ matrix of the evaluations of the N basis functions at the n data locations $u_1, \ldots, u_n$, where $u_i = X^{-1}(x_i)$. For any function $h \circ X \in H^1_{\mathcal{T}}(\Omega)$, denote $\mathrm{h}_N = \{h \circ X(\xi_1), \ldots, h \circ X(\xi_N)\}^T$, so that, in accordance with the notation introduced in (7), $\mathrm{h}_n = \{h \circ X(u_1), \ldots, h \circ X(u_n)\}^T = \Psi \mathrm{h}_N$. If the mesh vertices coincide with the data location points, as in the case of the simulations and applications considered in this paper, then $\mathrm{h}_N = \mathrm{h}_n$ and $\Psi = I_n$. Moreover, consider the $N \times N$ matrices

$$R_0 = \int_{\Omega_{\mathcal{T}}} \psi \psi^T \, \mathcal{D} d\Omega, \qquad R_1 = \int_{\Omega_{\mathcal{T}}} \nabla \psi^T K \nabla \psi \, d\Omega.$$

Thus, the integrals in (12) can be expressed as

$$\int_{\Omega_{\mathcal{T}}} \nabla(q \circ X)^T K \nabla(g \circ X) d\Omega = \mathrm{q}_N^T R_1 \mathrm{g}_N,$$

$$\int_{\Omega_{\mathcal{T}}} (p \circ X)(g \circ X) \mathcal{D} d\Omega = \mathrm{p}_N^T R_0 \mathrm{g}_N, \qquad \int_{\Omega_{\mathcal{T}}} \nabla(p \circ X)^T K \nabla(\hat{f} \circ X) d\Omega = \mathrm{p}_N^T R_1 \hat{\mathrm{f}}_N.$$

Let 0_N denote the column vector of length N of zero entries. The discrete counterpart of the estimation problem (12) thus reduces to solving a linear system, as stated in the following proposition.

PROPOSITION 2. *The estimator $\hat{f} \circ X \in H^1_{\mathcal{T}}(\Omega)$ that solves the discrete counterpart of the estimation problem (12) is given by $\hat{f} \circ X = \hat{\mathrm{f}}_N^T \psi$, such that $\hat{\mathrm{f}}_N$ satisfies*

$$\begin{pmatrix} -\Psi^T\Psi & \lambda R_1 \\ \lambda R_1 & \lambda R_0 \end{pmatrix} \begin{pmatrix} \hat{\mathrm{f}}_N \\ \mathrm{g}_N \end{pmatrix} = \begin{pmatrix} -\Psi^T z \\ 0_N \end{pmatrix}, \tag{13}$$

where g_N is the vector associated with the auxiliary function g in (12). Moreover $\hat{f} \circ X$ is uniquely determined.

Uniqueness of the solution to (13) is ensured by positive definiteness of the matrices R_0 and $\left(\Psi^T\Psi + \lambda R_1 R_0^{-1} R_1\right)$. Owing to the positive definiteness and thus invertibility of these matrices, Proposition 2 implies that $\hat{\mathrm{f}}_N = \left(\Psi^T\Psi + \lambda R_1 R_0^{-1} R_1\right)^{-1} \Psi^T z$. The estimator is thus linear in the observed data z and has the typical penalized regression form, with the roughness penalty matrix $R_1 R_0^{-1} R_1$ accounting for the domain deformation implied by the flattening map via the terms $\mathcal{D}$ and K in the matrices R_0 and R_1.

4. MODEL WITH COVARIATES

Space-varying covariates may be included in the model by considering a generalized additive framework, as done in Sangalli et al. (2013) for planar domains. Let $w_i = (w_{i[1]}, \ldots, w_{i[d]})^T$ be a d-vector of covariates associated with the variable of interest z_i observed at the location x_i. We modify model (1) as

$$z_i = w_i^T \beta + f(x_i) + \epsilon_i \qquad (i = 1, \ldots, n), \tag{14}$$

where $\beta \in \mathbb{R}^d$ is the vector of regression coefficients and the remaining terms are defined as in (1). This can be regarded as a partially linear model; see, e.g., Härdle et al. (2000). To estimate β and f in (14), we minimize the following penalized sum of squared errors functional

$$J_{\Gamma,\lambda}(\beta, f) = \sum_{i=1}^{n} \left\{ z_i - w_i^T \beta - f(x_i) \right\}^2 + \lambda \int_{\Gamma} \left\{ \Delta_{\Gamma} f(x) \right\}^2 \, d\Gamma.$$

Thus, using the map X, we can consider the following equivalent estimation problem over the planar domain Ω: find $\beta \in \mathbb{R}^d$ and $f \circ X \in H^2_{n0,K}(\Omega)$ that minimize

$$J_{\Omega,\lambda}(\beta, f \circ X) = \sum_{i=1}^{n} \left[z_i - w_i^T \beta - f\{X(u_i)\} \right]^2 + \lambda \int_{\Omega} \frac{1}{\mathcal{D}(u)} \Big(\mathrm{div}\big[K(u) \nabla f\{X(u)\} \big] \Big)^2 d\Omega. \tag{15}$$

Let W be the $n \times d$ matrix whose i-th row is the vector w_i^T of the covariates associated with the i-th data location. We assume W has full rank. Define $P = W(W^TW)^{-1}W^T$ to be the matrix that orthogonally projects on the subspace of $\mathbb{R}^n$ spanned by the columns of W. Let $Q = (I - P)$. Then the following result holds.

PROPOSITION 3. *The estimators $\hat{\beta}$ and $\hat{f} \circ X$ that minimize (15) over $\mathbb{R}^d$ and $H^2_{n0,K}(\Omega)$, are given respectively by $\hat{\beta} = (W^TW)^{-1}W^T(z - \hat{\mathrm{f}}_n)$ and $\hat{f}$ satisfying*

$$\mathrm{q}_n^T Q \hat{\mathrm{f}}_n + \lambda \int_{\Omega} \frac{1}{\mathcal{D}} \mathrm{div}\big\{ K \nabla (q \circ X) \big\} \, \mathrm{div}\big\{ K \nabla (\hat{f} \circ X) \big\} d\Omega = \mathrm{q}_n^T Q z \tag{16}$$

for any function q defined on Γ such that $q \circ X \in H^2_{n0,K}(\Omega)$. Moreover, the estimators $\hat{\beta}$ and $\hat{f} \circ X$ exist and are unique.

To prove this result, we first differentiate the functional $J_{\Omega,\lambda}(\beta, f \circ X)$ with respect to β, obtaining the minimizer $\hat{\beta} = \hat{\beta}(f)$. After plugging $\beta = \hat{\beta}$ into (15), the proof follows the one of Proposition 1.

Following the same arguments invoked for the model in Section 1·2, an auxiliary function g is introduced to derive the weak formulation of (16), analogously to (12). Hence the estimation problem is well suited for the discretization via a finite element space, and yields the following result.

PROPOSITION 4. *The estimators $\hat{\beta} \in \mathbb{R}^d$ and $\hat{f} \circ X \in H^1_{\mathcal{T}}(\Omega)$ that solve the discrete counterpart of the estimation problem with covariates are $\hat{\beta} = (W^T W)^{-1} W^T (z - \hat{\mathrm{f}}_n)$ and $\hat{f} \circ X = \hat{\mathrm{f}}_N^T \psi$, such that $\hat{\mathrm{f}}_N$ satisfies*

$$\begin{pmatrix} -\Psi^T Q \Psi & \lambda R_1 \\ \lambda R_1 & \lambda R_0 \end{pmatrix} \begin{pmatrix} \hat{\mathrm{f}}_N \\ \mathrm{g}_N \end{pmatrix} = \begin{pmatrix} -\Psi^T Q z \\ 0_N \end{pmatrix}, \tag{17}$$

where g_N is the vector associated with the auxiliary function g. Moreover $\hat{\beta}$ and $\hat{f} \circ X$ are uniquely determined.

From (17) we derive $\hat{\mathrm{f}}_N = (\Psi^T Q \Psi + \lambda R_1 R_0^{-1} R_1)^{-1} \Psi^T Q z$, and the estimators $\hat{\beta}$ and $\hat{\mathrm{f}}_N$ are thus linear in the observed data values z. Denote by S_f the $n \times n$ matrix $S_f = \Psi(\Psi^T Q \Psi + \lambda R_1 R_0^{-1} R_1)^{-1} \Psi^T Q$. Using this notation, $\hat{\mathrm{f}}_n = S_f\, z$ and $\hat{\beta} = (W^T W)^{-1} W^T \big(I - S_f\big)\, z$. Since $E(z) = W\beta + \mathrm{f}_n$ and $\mathrm{var}(z) = \sigma^2\, I_n$, with some algebra we derive

$$E(\hat{\mathrm{f}}_n) = S_f\, \mathrm{f}_n, \qquad E(\hat{\beta}) = \beta + (W^T W)^{-1} W^T \big(I_n - S_f\big) \mathrm{f}_n, \tag{18}$$

$$\mathrm{var}(\hat{\mathrm{f}}_n) = \sigma^2\, S_f\, S_f^T, \tag{19}$$

$$\mathrm{var}(\hat{\beta}) = \sigma^2 (W^T W)^{-1} + \sigma^2 (W^T W)^{-1} W^T \big(S_f\, S_f^T\big) W (W^T W)^{-1}. \tag{20}$$

Now let S denote the smoothing matrix $P + Q\, S_f$ and consider the vector $\hat{z}$ of fitted values at the n data locations, $\hat{z} = W\hat{\beta} + \hat{\mathrm{f}}_n = S\, z$. To measure the equivalent degrees of freedom for this linear estimator, we can use the trace of the smoothing matrix, $\mathrm{tr}(S)$. We can then estimate σ^2 by $\hat{\sigma}^2 = (z - \hat{z})^T (z - \hat{z}) / \{n - \mathrm{tr}(S)\}$. This estimate of σ, in combination with the expression in (20), can be used to derive approximate confidence intervals for β and for the pointwise evaluations of f. The Supplementary Material provides a simulation study in support of the asymptotic normality of these estimators. Moreover, the smoothing parameter λ can be selected by generalized cross validation, minimizing the quantity $\mathrm{GCV}(\lambda) = (z - \hat{z})^T (z - \hat{z}) / [n\{1 - \mathrm{tr}(S)/n\}^2]$. The predicted value of a new observation at x_{n+1} with associated covariates w_{n+1} is given by $\hat{z}_{n+1} = w_{n+1}^T \hat{\beta} + \hat{f}(u_{n+1}) = w_{n+1}^T \hat{\beta} + \hat{\mathrm{f}}_N^T \psi(u_{n+1})$, where $X(u_{n+1}) = x_{n+1}$. The mean and variance of $\hat{z}_{n+1}$ can be obtained from the expression above and approximate prediction intervals may be derived. Owing to the conformal map X between the Γ and Ω, the solution to the estimation problem and these distributional results can be obtained on the original non-planar domain; see Figure 1, (d) to (e).

The estimators have two sources of bias. One source is the penalty term, and this is typical of regression models involving a roughness penalty: unless the true function f is such that it annihilates the penalty term, this term will bias the estimator. This source of bias disappears as n increases, in the sense of infill asymptotics, if the smoothing parameter λ decreases with n. Letting λ decrease with n appears natural, since having more observations decreases the need for regularization. A second source of bias is the discretization, and is common to any model employing a basis expansion. This source of bias disappears as the mesh is appropriately re-

fined. These convergence results follow from Azzimonti et al. (2014a) and are formalized in the Supplementary Material. The bias always appeared negligible in our simulations.

5. Simulation Studies

In this section we illustrate the performance of the proposed technique on different non-planar domains. We compare three methods:

A: the proposed spatial regression model over non-planar domains;
B: the spatial regression model over planar domains introduced in Ramsay (2002) combined with the cylindrical flattening map;
C: the iterative heat kernel smoothing described in Chung et al. (2005).

Since iterative heat kernel smoothing is not designed for the inclusion of covariates, this simulation study does not include covariates. The three methods are compared on four tubular domains: three test domains and a real geometry of an internal carotid artery from the AneuRisk dataset; see Figure 4. Since the spatial regression model over planar domains is combined with the cylindrical flattening map, we consider an artery geometry that does not present an aneurysm; in fact the cylindrical map is not suited to deal with large aneurisms as multiple points on the wall of the carotid artery would be mapped to the same point in the plane.

To compute the cylindrical flattening map, we consider the centerline of the tubular domain, which for simplicity is assumed known, this being the case for the AneuRisk data. Then, each point $x = (x_{[1]}, x_{[2]}, x_{[3]})$ on the non-planar domain Γ is associated with the closest point on the centerline. Hence, it is possible to consider the cylindrical parameterization defined by (s, r, θ), where s is the curvilinear abscissa computed along the centerline, r is the artery radius, i.e., the distance between x and the associated centerline point, and θ is the angle in radians identified by x with respect to the curvilinear abscissa. The cylindrical flattening map thus changes Γ into the rectangle $(s, \theta\bar{R})$, where $\bar{R}$ is the average radius. Analogously to the spatial regression model over non-planar domains, the necessary $(2\pi\bar{R})$-periodicity of the estimate along the coordinate $\theta\bar{R}$ can be guaranteed either by repeating the data values at the same abscissas but with ordinates $(\theta + 2\pi)\bar{R}$ and $(\theta - 2\pi)\bar{R}$, or by imposing periodic boundary conditions. Both approaches are implemented, and the same is done for the spatial regression over non-planar domains. Finally, for these methods the optimal value of the smoothing parameter λ is selected at each simulation replicate and for each domain by generalized cross validation.

Iterative heat kernel smoothing has been developed for neuroimaging applications, to deal with very complex domain geometries such as the cortical surface of the brain, which is usually approximated by three-dimensional meshes with more than 10^6 nodes. This smoothing algorithm aims at reducing the computational burden associated with such geometries and works directly on the mesh without any flattening. To do this, the Laplace–Beltrami eigenvalue problem is solved directly on Γ, i.e., ordered eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \cdots$ and the corresponding eigenfunctions $\phi_0, \phi_1, \ldots$ are found by solving the eigenvalue problem $-\Delta_\Gamma \phi_j = \lambda \phi_j$ on Γ. Thus, the heat kernel with bandwidth t is constructed from the eigenvalue-eigenfunction pairs $\{(\lambda_j, \phi_j)\}$ as $\mathcal{K}_t(p, q) = \sum_{j=0}^{\infty} e^{-\lambda_j t} \phi_j(p) \phi_j(q)$, where p and q are two generic points on Γ. The heat kernel smoothing of z_i is given by $\mathcal{K}_t * z_i = \sum_{j=0}^{\infty} e^{-\lambda_j t} \beta_j(x_i) \phi_j(x_i)$, where $\beta_j(x_i) = \langle z_i, \phi_j \rangle$. In practice, only k eigenvalue-eigenfunction pairs are chosen via an iterative residual fitting algorithm. For our simulations, we made a heuristic choice by selecting the bandwidth with the best performance after some test runs. In particular, we set $t = 10^{-3}$. To determine the level of smoothing, the optimal number of iterations is selected for each simulation replicate and for each domain via the F-test criterion outlined in Chung et al. (2005).

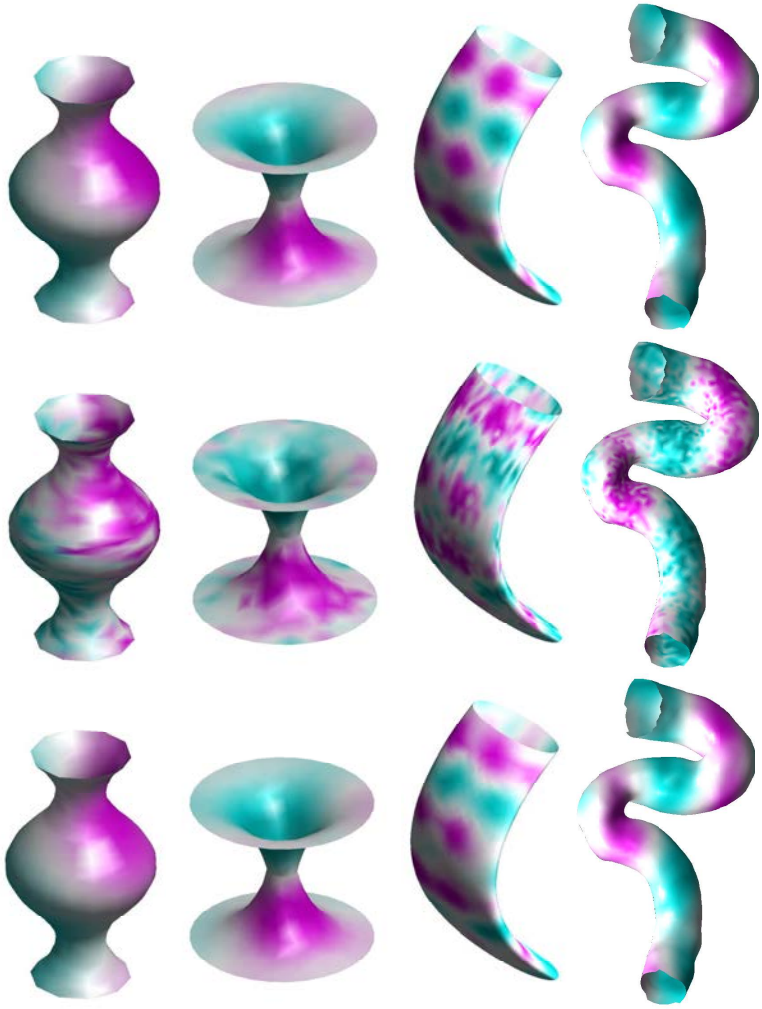

Fig. 4. Simulation study. Top: colormap of a test function generated by (21) on four tubular domains; the fourth one is a real geometry of an internal carotid artery from the AneuRisk dataset. Middle: data with noise. Bottom: estimates provided by the spatial regression model over non-planar domains. The colormap range is optimized for each geometry, magenta to green (min,max); Geometry 1: (-4.06,6.06); Geometry 2: (-0.79,2.79); Geometry 3: (-5.14,7.17); Geometry 4: (-2.09,4.10).

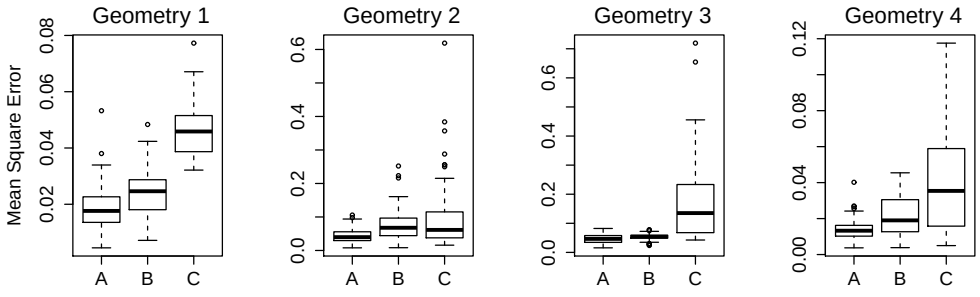

Fig. 5. Box plots of the mean square errors for the three estimators over fifty simulations. A: spatial regression over non-planar domains; B: spatial regression over planar domains, combined with cylindrical flattening; C: iterative heat kernel smoothing.

We generate data as follows. Over each non-planar test domain, we consider fifty test functions, having the form

$$f(x_{[1]}, x_{[2]}, x_{[3]}) = a_1 \sin(2\pi x_{[1]}) + a_2 \sin(2\pi x_{[2]}) + a_3 \sin(2\pi x_{[3]}) + 1, \qquad (21)$$

with coefficients a_j $(j = 1, 2, 3)$ randomly generated from independent normal distributions with mean and standard deviation equal to one. For each geometry, the data locations x_i, $(i = 1, \ldots, n)$ coincide with the nodes of the three-dimensional mesh used to approximate the domain; see, e.g., Figure 2, top left, for the triangular mesh approximating Geometry 2. Noisy data values are obtained, in accordance with the model (1), by adding independent normally distributed errors with mean zero and a standard deviation 0.5 at each of the data locations. For each geometry, the top row of Figure 4 shows an example of a test function generated by (21), the middle row shows the corresponding level of noise, and the bottom row shows the associated estimate by the proposed spatial regression model. The first evidence of the good performance of the proposed method is provided by a visual comparison between each test function and the corresponding estimate. The colormaps are obtained by linear interpolation of the data at x_i, using Matlab code of Chung et al. (2005).

For each simulation replicate and test domain, we compute the mean square error of the estimator associated with each of the three different methods. Figure 5 shows that the spatial regression model over non-planar domains yields better estimates with a smaller variance. Pairwise Wilcoxon tests confirm that the mean square errors of the estimators obtained by spatial regression model over non-planar domains are stochastically lower than those associated with the competitor methods, with p-values of the order 10^{-5} to 10^{-8} for the comparison with the spatial regression model over planar domains and p-values of the order 10^{-8} to 10^{-10} for the comparison with iterative heat kernel.

The results shown for the regression model over non-planar domains correspond to the approach imposing periodic boundary conditions. No statistically significant difference is found between these estimates and those obtained with the approach based on the repetition of the domain and data, and likewise for the regression models over planar domains. Nevertheless, the periodic boundary conditions lead to a computationally more efficient method. For instance, with reference to Geometry 3, characterized by 1648 nodes, the spatial regression model over non-planar domains takes a total of 12s to fit the test data in Figure 4 when using periodic boundary conditions, 11s for flattening the mesh and less than 1s for fitting the data, and 16s with data repetition. For the same geometry and data, spatial regression model over planar domains takes 2s with periodic boundary conditions, 1s for flattening and less than 1s for fitting, and 16s with data repetition. The shorter times taken by the model over planar domains do not account for the computation of the centerline, which for simplicity is here considered as already known. Finally, iterative heat kernel smoothing takes 44s for fitting the same data. The simulations were run with Matlab 7.12.0 on 2 GHz Intel Core i7 processor in a MacBook Pro with a 256GB Solid State hard drive.

The Supplemetary Material extends this comparative study to a technique based on tensor product of univariate smoothing splines, illustrating the advantages of the methods considered in this section with respect to a smoothing technique that is not constrained over the manifold. The Supplemetary Material also offers a simulation study with covariates.

6. Application to hemodynamic data

In this section, we apply the non-planar spatial regression model to the AneuRisk data. The AneuRisk project (`https://statistics.mox.polimi.it/aneurisk/` and

http://ecm2.mathcs.emory.edu/aneuriskweb/index) gathered researchers of scientific fields from neurosurgery and neuroradiology to statistics, numerical analysis and bioengineering, with the aim of studying the pathogenesis of cerebral aneurysms. Aneurysms are deformations of the vessel wall. The formation of an aneurysm is usually ascribed to the complex interplay between the biomechanical properties of artery walls and the effects of hemodynamic forces exerted on the vessel walls, such as wall shear stress and pressure. These forces in turn depend on the vessel morphology itself; exploring the relationships between the vessel morphology and the hemodynamics was one of the main goals of the AneuRisk project.

Here, we analyze hemodynamic data on the real internal carotid artery displayed in Figure 1(a), and we study its relationship with local vessel morphology. Iterative heat kernel smoothing cannot be implemented in this study, as it cannot yet account for space-varying covariates, while the cylindrical map for the planar spatial regression model cannot be implemented due to the large aneurysm. The hemodynamic quantities of interest, such as wall shear stress and pressure, are computed via computational fluid dynamics simulations over the real anatomy (Passerini et al., 2012). The inner carotid artery geometry is obtained via the reconstruction algorithm coded in the Vascular Modeling ToolKit, from three-dimensional angiographic images belonging to the AneuRisk data warehouse (Piccinelli et al., 2009). Figure 1(a) shows the modulus of the simulated wall shear stress (dyn/cm^2) at the systolic peak on the three-dimensional geometry. Figure 1(b) shows the three-dimensional triangular mesh approximating the artery. Each triangle vertex in this mesh coincides with a data location x_i, at which the wall shear stress and covariates describing the local morphology are measured. According to the approach proposed in this paper, the three-dimensional triangular mesh is flattened via the conformal map described in Section 3·2 to create the planar triangulation in Figure 1(c). The sides of the planar triangulation labeled inflow and outflow correspond to the open ends of the carotid artery. The sides indicated by dashed lines correspond to the longitudinal cut along the artery wall, connecting the open boundaries of the artery. To maintain the periodicity of the solution we impose periodic boundary conditions on these sides. The aneurysm and major curves of the artery are recognizable also in the flattened domain. In particular, the area of the mesh which is very fine and close to the outflow side corresponds to the aneurysmal sac.

To explore the relationship between the wall shear stress and the local morphology of the artery, we consider the following space-varying covariates: the local curvature of the vessel wall, the curvature of the artery centerline and the local radius of the vessel. The wall shear stress obtained via the numerical simulations depends on the complex three-dimensional geometry of the carotid, and thus also implicitly on the simple geometric features that we consider as covariates, although the form of such relationship is a priori unclear. Hence, we aim to explore the form of this relationship and to understand to what extent these geometric features can be considered a good summary of the three-dimensional artery morphology. The local curvature of the vessel wall is calculated from the three-dimensional mesh and varies between -20.63 cm^{-1} and 36.46 cm^{-1}. The curvature of the vessel centerline identifies the curvature of the whole vascular geometry. The artery centerline and its curvature are computed as described in Sangalli et al. (2009b). In particular, to measure the centerline curvature at each point on the vessel wall we refer to the curvature at the associated centerline point. The centerline curvature varies from 0.05 cm^{-1} to 4.64 cm^{-1}. Finally, the local radius of the vessel is measured as the distance from the artery wall to the associated centerline point, and ranges from 0.14 cm to 0.43 cm.

Figure 1(d) and (e) display the wall shear stress estimated via the non-planar spatial regression model including the covariates, for $\lambda = 10^{0.5}$. The planar view in Figure 1(d) allows us to see the entire geometry without rotating the figure, thus making areas of interest easier to highlight. We recognize large areas of high wall shear stress, in correspondence with the neck of the aneurysm

and along the first major bend of the carotid syphon, thus highlighting the sensitivity of the wall shear stress with respect to the complex geometry of the artery. Notice in particular the sharp spatial change in the wall shear stress values, from the very high values at the aneurysmal neck to the very low values in correspondence with the aneurysmal sack. These sharp changes in wall shear stress are conjectured to be very dangerous for the aneurysmal pathology and its evolution. The estimated covariates effects are $\hat{\beta}_1 = -0.99$ dyn cm with a standard error of 0.03 dyn cm for the local curvature of the vessel wall, $\hat{\beta}_2 = 0.35$ dyn cm with a standard error of 0.11 dyn cm for the curvature of the artery centerline, and $\hat{\beta}_3 = -7.834$ dyn/cm with a standard error of 0.13 dyn/cm for the local radius. Hence, the local curvature of the vessel wall and the local radius are negatively associated with the wall shear stress, while the artery centerline curvature is positively associated. This supports clinical intuition about the effect of the local vessel morphology, and is consistent with hemodynamical studies in simplified geometries. As expected, the most influential factor for the wall shear stress is the local radius. These preliminary findings will be further investigated in future work, with statistical analysis across patients. An important advantage of the non-planar spatial regression model approach is that patient-specific estimates can all be mapped into a common planar domain where, after suitable registration among patients, comparisons across patients are possible.

On the computational side, there is also a possibility to solve the estimation problem in (2) directly on the non-planar domain, without resorting to a flattening map. This would probably lead to computational savings, although mapping the estimates to a reference domain is still of interest, allowing for more direct comparisons across different geometries.

ACKNOWLEDGEMENT

This work was developed while B. Ettinger was postdoctoral fellow at MOX, Dipartimento di Matamatica, Politecnico di Milano. We are grateful to James Ramsay, Piercesare Secchi and John Aston, for their advice on this work, to Tiziano Passerini, for providing the data and for assistance with VMTK software, and to Alessandro Veneziani, P.I. of the AneuRisk project. We thank the Editor, Associate Editor and three anonymous Referees for their very constructive comments. This work has been funded by the research program Dote Ricercatore Politecnico di Milano - Regione Lombardia and by the Ministero dell'Istruzione dell'Università e della Ricerca, FIRB starting grant project SNAPLE http://mox.polimi.it/users/sangalli/firbSNAPLE.html.

SUPPLEMENTARY MATERIAL

Supplementary material available at *Biometrika* online includes proofs, technical details and further simulation studies.

REFERENCES

ALFELD, P., NEAMTU, M. & SCHUMAKER, L. L. (1996). Fitting scattered data on sphere-like surfaces using spherical splines. *J. Comput. Appl. Math.* **73**, 5–43.

ANGENENT, S., HAKER, S., TANNENBAUM, A. & KIKINIS, R. (1999). On the laplace-beltrami operator and brain surface flattening. *IEEE Trans. Med. Imag.* **18**, 700–711.

AZZIMONTI, L., NOBILE, F., SANGALLI, L. M. & SECCHI, P. (2014a). Mixed finite elements for spatial regression with PDE penalization. *SIAM/ASA J. Uncertain. Quantif.* **2**, 305–335.

AZZIMONTI, L., SANGALLI, L. M., SECCHI, P., DOMANIN, M. & NOBILE, F. (2014b). Blood flow velocity field estimation via spatial regression with PDE penalization. *J. Amer. Statist. Assoc.* DOI:10.1080/01621459.2014.946036.

BARAMIDZE, V., LAI, M. J. & SHUM, C. K. (2006). Spherical splines for data interpolation and fitting. *SIAM J. Sci. Comput.* **28**, 241–259.

CHUNG, M. K., ROBBINS, S. M., DALTON, K. M., DAVIDSON, R. J., ALEXANDER, A. L. & EVANS, A. C. (2005). Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage* **25**, 1256–1265.

DIERKES, U., HILDEBRANDT, S. & SAUVIGNY, F. (2010). *Minimal Surfaces*, vol. 1. Heidelberg: Springer, 2nd ed.

GNEITING, T. (2013). Strictly and non-strictly positive definite functions on spheres. *Bernoulli* **19**, 1087–1500.

GOCKENBACH, M. (2006). *Understanding and Implementing the Finite Element Method*. Philadelphia: Society for Industrial and Applied Mathematics.

GU, X. & YAU, S.-T. (2003). Global conformal surface parameterization. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, SGP '03. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association.

HAGLER, JR., D. J., SAYGIN, A. P. & SERENO, M. I. (2006). Smoothing and cluster thresholding for cortical surface-based group analysis of fMRI data. *NeuroImage* **33**, 1093–1103.

HAKER, S., ANGENENT, S., TANNENBAUM, A. & KIKINIS, R. (2000). Nondistorting flattening maps and the 3-D visualization of colon CT images. *IEEE Trans. Med. Imag.* **19**, 665–670.

HÄRDLE, W., LIANG, H. & GAO, J. (2000). *Partially linear models*. Physica-Verlag, Heidelberg.

HURDAL, M. K. & STEPHENSON, K. (2009). Discrete conformal methods for cortical brain flattening. *NeuroImage* **45**, S86 – S98.

JUN, M. (2011). Non-stationary cross-covariance models for multivariate processes on a globe. *Scandinavian Journal of Statistics* **38**, 726–747.

JUN, M. & STEIN, M. (2007). An approach to producing space-time covariance functions on spheres. *Technometrics* **49**, 468–479.

LINDGREN, F., RUE, H. & LINDSTRÖM, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **73**, 423–498. With discussion and a reply by the authors.

LIONS, J.-L. & MAGENES, E. (1973). *Non-Homogeneous Boundary Value Problems and Applications*, vol. III. New York: Springer-Verlag.

PASSERINI, T., SANGALLI, L. M., VANTINI, S., PICCINELLI, M., BACIGALUPPI, S., ANTIGA, L., BOCCARDI, E., SECCHI, P. & VENEZIANI, A. (2012). An integrated CFD-statistical investigation of parent vasculature of cerebral aneurysms. *Cardio. Eng. and Tech.* **3**, 26–40.

PICCINELLI, M., VENEZIANI, A., STEINMAN, D., REMUZZI, A. & ANTIGA, L. (2009). A framework for geometric analysis of 852 vascular structures: applications to cerebral aneurysms. *IEEE Trans. Med. Imag.* **28**, 1141–1155.

QUARTERONI, A. (2014). *Numerical models for differential problems*, vol. 8 of *Modeling, Simulation and Applications*. Springer, Milan, 2nd ed.

RAMSAY, T. O. (2002). Spline smoothing over difficult regions. *J. R. Statist. Soc. Ser. B* **64**, 307–319.

SANGALLI, L. M., RAMSAY, J. O. & RAMSAY, T. O. (2013). Spatial spline regression models. *J. R. Statist. Soc. Ser. B* **75**, 681–703.

SANGALLI, L. M., SECCHI, P., VANTINI, S. & VENEZIANI, A. (2009a). A case study in exploratory functional data analysis: geometrical features of the internal carotid artery. *J. Am. Stat. Assoc.* **104**, 37–48.

SANGALLI, L. M., SECCHI, P., VANTINI, S. & VENEZIANI, A. (2009b). Efficient estimation of three-dimensional curves and their derivatives by free-knot regression splines, applied to the analysis of inner carotid artery centrelines. *J. R. Statist. Soc. Ser. C* **58**, 285–306.

WAHBA, G. (1981). Spline interpolation and smoothing on the sphere. *SIAM J. Sci. Stat. Comput.* **2**, 5–16.

[*Received April* 2012. *Revised September* 2012]