Statistical Science
2019, Vol. 34, No. 2, 236–252
https://doi.org/10.1214/18-STS675

A Kernel Regression Procedure in the 3D Shape Space with an Application to Online Sales of Children's Wear

Gregorio Quintana-Ortí and Amelia Simó

Abstract. This paper is focused on kernel regression when the response variable is the shape of a 3D object represented by a configuration matrix of landmarks. Regression methods on this shape space are not trivial because this space has a complex finite-dimensional Riemannian manifold structure (non-Euclidean). Papers about it are scarce in the literature, the majority of them are restricted to the case of a single explanatory variable, and many of them are based on the approximated tangent space. In this paper, there are several methodological innovations. The first one is the adaptation of the general method for kernel regression analysis in manifold-valued data to the three-dimensional case of Kendall's shape space. The second one is its generalization to the multivariate case and the addressing of the curse-of-dimensionality problem. Finally, we propose bootstrap confidence intervals for prediction. A simulation study is carried out to check the goodness of the procedure, and a comparison with a current approach is performed. Then, it is applied to a 3D database obtained from an anthropometric survey of the Spanish child population with a potential application to online sales of children's wear.

Key words and phrases: Shape space, statistical shape analysis, Kernel regression, Fréchet mean, children's wear.

1. INTRODUCTION

Many problems in medical imaging analysis and computer vision involve predicting the shape of an object as a function of a set of numerical covariates (age, dose, time, etc.) and therefore new statistical methodologies are needed to address this interesting problem. With the expression "shape of an object" we mean all the information of the object that remains invariant under Euclidean similarity transformations, that is, translations, rotations and scale changes.

Gregorio Quintana-Ortí is Associate Professor in Computer Science, Department of Computer Science and Engineering, Universitat Jaume I, Avda. de Sos Baynat, s/n. 12071-Castellón de la Plana, Spain (e-mail: gquintan@uji.es). Amelia Simó is Professor of Statistics, Department of Mathematics-IMAC, Universitat Jaume I, Avda. de Sos Baynat, s/n. 12071-Castellón de la Plana, Spain (e-mail: simo@uji.es).

There is a large number and variety of mathematical tools to represent the shape of an object. Depending on the type of mathematical representation chosen, the statistical methodology to be used will be different since in most cases the spaces in which they are embedded are different. Moreover, the task to be carried out is also important when choosing the representation to use: shape summary, shape modelling, shape deformation, etc. The four major mathematical approaches to shape analysis are described below.

The first approach considers objects as subsets of $\mathbb{R}^m$. When the aim is the description, ideas of set geometry and mathematical morphology can be used (Serra, 1984). When the aim is modelling, they are considered as a realization of a random compact set X (Matheron, 1975, Stoyan and Stoyan, 1994, Baddeley and Molchanov, 1998, Simó, de Ves and Ayala, 2004, Molchanov, 2005), that is, a random variable taking values in the space of all compact sets in $\mathbb{R}^m$ with the myopic topology. Because of the difficulty of this

space, the random distribution of X is summarized using a set of random variables taking values in a Euclidean space.

The second approach uses functions to represent the contours of the objects (Prince and Willsky, 1990, Del Bimbo et al., 1998, Loncaric, 1998, Kindratenko, 2003, Gual-Arnau, Herold-García and Simó, 2013). This approach transforms the objects into elements of a function space, and tools developed for analyzing the functions are used for analyzing the shape. Usually, the contour is represented on a particular functional basis, thus reducing the representation to a small number of parameters. This approach includes methods such as the cross section function, the radius-vector function, the support function, the tangent-angle function or the Fourier descriptors and it has been mainly used for summarizing and classification purposes.

The third and more recent approach is called elastic shape analysis (Azencott, 1994, Younes, 1996, Younes, 1998, Klassen et al., 2004, Younes et al., 2008, Srivastava et al., 2011, Gual-Arnau, Herold-García and Simó, 2015). It proposes the representation of planar shapes by closed simple curves defined by the boundaries of objects (Younes, 1998) or alternatively by functions that characterize these curves, such as the square-root velocity (Srivastava et al., 2011). After the identification of elements invariant under translations, rotations, scale changes and reparameterizations, the resulting shape space is an infinite dimensional Riemannian manifold and its Riemannian metric is used to compare shapes. The focuses of this approach are deformations, computation of intrinsic statistics and statistical modelling. Thus, the shape representation and Riemannian metric chosen are critically important. The corresponding theory for the shape space of surfaces was generalized by Bauer, Harms and Michor (2012) and Jermyn et al. (2017); however, its practical application is very challenging because of the infinite-dimensional nature of the space.

Finally, the majority of the research has been restricted to landmark-based analysis, where objects are represented using k labelled points in the Euclidean space $\mathbb{R}^m$. These landmarks are required to appear in each data object, and to correspond to each other in a physical sense. They are given by certain geometrical or anatomical properties. Seminal papers on this topic are Bookstein (1978), Kendall (1984) and Goodall (1991). The main references are Dryden and Mardia (1998, 2016) and Kendall et al. (1999). This paper will focus on this approach, which will be explained in detail in the next section.

Recall that shape is defined as the geometrical information about the object that is invariant under Euclidean similarity transformations. When the landmark-based approach is used, the shape space is the resulting quotient space, and it is a finite-dimensional Riemannian manifold. Therefore, standard statistical methodologies on linear spaces based on Euclidean distance cannot be used.

There are several difficulties in generalizing probability distributions and statistical procedures to measurements in a nonvectorial space like a Riemannian manifold, but fortunately there has been a significant amount of research and activity in this area over the recent years. An excellent review can be found in Pennec (2006). Our aim is to take advantage of and extend this recent significant research on statistics on Riemannian manifolds. Hence, our approach to the statistical shape analysis will be mainly focused from the point of view of Riemannian geometry.

To begin with one important difficulty, the concept of expectation of a random element in a manifold cannot be generalized, since it would be an integral with values in the manifold. In a Euclidean space, there is a clear and unique concept of mean, which corresponds to the arithmetic average of realizations. In Riemannian manifolds different kinds of means have been defined as Fréchet parameters associated with different types of distances (Bhattacharya and Patrangenaru, 2002, 2003, Kobayashi and Nomizu, 1969). Since a mean in a manifold is the result of a minimization, its existence is not ensured. Karcher (1977) and Kendall (1990) established conditions on a manifold to ensure the existence and uniqueness of the mean, and a gradient descent algorithm in the manifold can be found in Woods (2003).

Although statistical analysis of manifold-valued data has gained a great deal of attention in recent years, there is not much literature on regression analyses on manifolds. Early papers were developed for directional data (Jupp and Kent, 1987, Mardia and Jupp, 2000). In regression of directional data, parametric distributions, such as the Von Mises distribution, are commonly assumed. However, it is very challenging to assume useful parametric distributions for other manifold-valued data. Hence, nonparametric regression has been very common until now. Local constant regressions have been developed for manifold-valued data defined with respect to the Fréchet mean in Davis et al. (2007). On the other hand, Shi et al. (2009) developed a semiparametric regression model that uses a link function to map from the Euclidean space of covariates to the Riemannian manifold. However, the drawback of this

approach is that the link function must be known, and this is not common in usual applications.

When dealing with data in a manifold, linear methods are not applicable because the scalar product and the addition operations are not defined there. In the particular case of a scalar covariable, some authors have tried to use a geodesic curve as the natural generalization of a "straight line" on a manifold, and they have defined "linear models" as "geodesic models." Fletcher (2011) introduced a regression method for modelling the relationship between a manifold-valued random variable and a real-valued independent parameter based on a geodesic curve, parameterized by the independent parameter. The multivariate case using multiple geodesic basis on the manifold and a variational algorithm is treated in Kim et al. (2014). Recently, a regression parametric model based on a normal probability distribution has been introduced in Fletcher and Zhang (2016).

This paper was motivated by an important current application: a 3D anthropometric study of the child population in Spain developed by the Biomechanics Institute of Valencia. The aim of this study was to generate anthropometric data to help and inform decision makers (parents/relatives/children) during the size selection process, focusing on online shopping for children's wear. For this purpose, 739 randomly-selected Spanish children from 3 to 12 years old were scanned using a 3D body scanner.

Several new technologies and online services addressing the selection of proper garment sizes or models for the consumer have been developed in recent years. These applications can be classified into two groups. The first group uses neural network algorithms to match the garment being considered to other clothes worn by the user (see, for instance, www.whatfitsme.com). This method requires an initial database (the user's virtual closet) for training algorithms. The second group predicts the size and fit of the garment using user's anthropometric measurements and their relationship with the dimensions of the garment (see, for instance, www.fits.me). In this paper, instead of correlating children's anthropometric measurements with the dimensions of the garment, we propose to use them to predict the children's body shape represented by landmarks. In order to achieve this, we propose a new methodology in the shape space.

There are several methodological novelties in our work. The first novelty is the adaptation of the general method for kernel regression analysis in manifold-valued data of Davis et al. (2007) to the landmark-based 3D shape space. Although this method has been used in previous works, it has never been applied to 3D landmark data before because of its analytical and computational challenges. The second contribution of this work is the generalization of this method to multiple explanatory variables introducing a multivariate kernel and the addressing of the curse-of-dimensionality problem. Finally, we propose bootstrap confidence intervals for analyzing the predictions. A simulation study with simple and predictable objects is performed to check the validity of our approach, and then a performance comparison with a current approach is conducted on this dataset.

Moreover, this new methodology has been applied to the aforementioned children database with excellent results. The resulting predicted shapes can then be used to choose the most suitable size for the selected garments.

Vinué, Simó and Alemany (2016) used women's body shapes represented by landmarks to define a new sizing system by adapting clustering algorithms to the shape space. The 3D database used was very similar to the one used in this paper, and it was obtained from an anthropometric survey of the Spanish female population.

The R language (R Development Core Team, 2014) was employed in our implementations. We used the `shapes` package by Ian Dryden (Dryden, 2012). This is a very powerful and complete package for the statistical analysis of shapes. As its efficiency for medium and large data sets is limited, we rewrote some parts to accelerate it and to run our codes in a shorter time.

The code developed in this work is available as an open-source software in the https://github.com/gquintanaorti/kernel_regression_of_3d_shapes public code repository. This repository contains both the code and the simulated dataset employed in Section 5.

The article is organized as follows. Section 2 describes the basic concepts of statistical shape analysis. Section 3 shows the kernel regression for shape analysis. Some important details regarding our implementations are described in Section 4. A simulation study is conducted in Section 5. The application for regression to children's body shapes is detailed in Section 6. Finally, conclusions are discussed in Section 7.

2. SHAPE SPACE

Because our goal is to apply statistical methods on manifold-valued data to analyze shapes, first we introduce here a brief summary of the shape space as a Riemannian manifold and the main concepts that will be needed later in our approach.

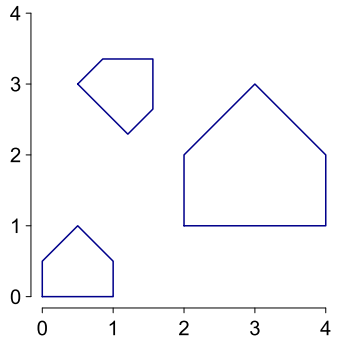

FIG. 1. *Three planar objects with the same shape.*

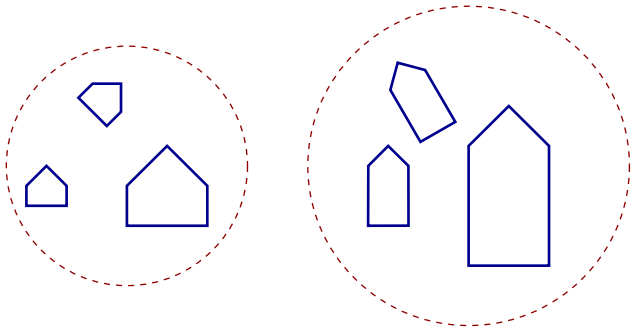

FIG. 2. *Two classes of equivalence each with three geometrical planar objects. Each class corresponds to a different shape.*

The word "shape" is very common in everyday language, usually referring to the appearance of a geometric object. Shape can be defined as the geometrical information of the object that is invariant under a Euclidean similarity transformation, that is, location, orientation and scale (Kendall, 1977, Dryden and Mardia, 1998, 2016). Thus, two objects have the same shape if we can find a transformation (translations, rotations and/or scale changes) that makes the two objects fit perfectly. For instance, the three objects in Figure 1 have the same shape. Invariance under reflections is not considered.

In this work, a geometrical m-dimensional object (usually $m = 2, 3$) is determined by a finite number of $k > m$ coordinate points, known as landmark points. Each object is then described by a $k \times m$ configuration matrix Y containing the m Cartesian coordinates of its k landmarks. For instance, the left-most object in Figure 1 is represented by the following configuration matrix:

$$Y^T = \begin{bmatrix} 0 & 0 & 0.5 & 1 & 1 \\ 0 & 0.5 & 1.5 & 0.5 & 0 \end{bmatrix}.$$

In this way, an object could be regarded as an element of the Ecuclidean space $\mathbb{R}^{mk}$.

However, a configuration matrix Y is not a proper shape descriptor because it is not invariant under similarity transformations. For any similarity transformation, that is, for any translation vector $t \in \mathbb{R}^m$, scale parameter $s \in \mathbb{R}^+$ and rotation matrix R in the special orthogonal group $SO(m)$, the configuration matrix given by $sYR + 1_k t^T$ (where 1_k is the $k \times 1$ vector of ones) describes the same shape as Y.

For instance, the central object in Figure 1 is obtained by applying the above transformation to the left object with:

$$s = 1, \quad R = \begin{bmatrix} \cos 45^\circ & -\sin 45^\circ \\ \sin 45^\circ & \cos 45^\circ \end{bmatrix}, \quad t = \begin{bmatrix} 0.5 \\ 3 \end{bmatrix}.$$

As was previously said, all the objects in Figure 1 correspond to the same shape, that is, they are equivalent. Hence, from a theoretical point of view we can define the shape space as:

DEFINITION 1. The shape space Σ_m^k is the set of equivalence classes $[Y]$ of $k \times m$ configuration matrices $Y \in R^{k \times m}$ under the action of Euclidean similarity transformations.

For example, Figure 2 shows two different equivalence classes.

The configuration matrix could be modified in many different ways to effectively correspond to a shape descriptor, that is, to obtain a representative of each equivalence class $[Y]$. Among popular shape coordinates are, for instance, Bookstein coordinates (Bookstein, 1986) and Kendall coordinates (Kendall, 1984). We will use the most suitable one for our goal of using statistical tools on Riemannian manifolds. This is described below.

Let Y be a configuration matrix, we remove the similarity transformations one at a time. A way to remove the location effect consists of multiplying it by the Helmert submatrix H, that is, $Y_H = HY$.

The Helmert submatrix H is obtained by removing the first row in the Helmert matrix. The Helmert matrix is an $h \times h$ orthogonal matrix with its first row of elements equal to $1/\sqrt{h}$ and the remaining rows orthogonal to the first row. The j-th row of the Helmert submatrix H is given by the number $-1/\sqrt{j(j+1)}$ repeated j times, followed by $-j/\sqrt{j(j+1)}$, and then $(h - j - 1)$ zeros.

Note that after removing the location, the object could be regarded again as an element of the Euclidean space, but now the Euclidean space is $\mathbb{R}^{mk-m}$.

To filter scale, we can divide Y_H by the centroid size, which is given by

$$S(Y) = \|Y_H\| = \|HY\| = \sqrt{\operatorname{trace}\big((HY)^T(HY)\big)} = \|CY\|.$$

The quotient $Z = Y_H/\|Y_H\|$ is called the pre-shape of the configuration matrix Y because all information about location and scale has been removed, but rotation information still remains.

DEFINITION 2. The pre-shape space S_m^k is the set of all possible pre-shapes.

It is important to note now that after removing scale, our matrix Z can no longer be represented by an element on a Euclidean space because we have the restriction of size one. Therefore, S_m^k is a hypersphere of unit radius in $\mathbb{R}^{mk-m}$.

S_m^k is a Riemannian manifold that has been widely studied and therefore well known. Σ_m^k is the quotient space of S_m^k under rotations, and a shape $[Y]$ is an orbit associated with the action of the rotation group $SO(m)$ on the pre-shape. The complexity of the Riemannian structure of the shape space depends on k and m. For example, the quotient space Σ_2^k is isometric with the complex projective space $\mathbb{CP}_{k-2}$, a familiar and well-known Riemannian manifold. Even more, the particular case of Σ_2^3, the shape space of planar triangles, is the sphere in three dimensions with radius 1/2. An illustrative description can be found in Dryden and Mardia (2016).

For $m > 2$, which is the case of our application, Σ_m^k is not a familiar space, and it has singularities, that is, points at which the differential structure is not defined.

From now on, in order to simplify the notation, we will use Y to denote both, a configuration matrix and its shape, provided that its meaning is understood in its context.

2.1 Riemannian Manifolds

Before continuing with the description of the Riemannian structure of Σ_m^k, let us review some important concepts on Riemannian manifolds which will be needed later.

A differential manifold immersed in $\mathbb{R}^m$ is a topological space which is locally homeomorphic with a Euclidean space. The local homeomorphisms can be smoothly patched together. The locally homeomorphic vector space at a point Y is called the *tangent space* and it is denoted by T_Y. This space contains all the possible directions in which one can tangentially pass through Y.

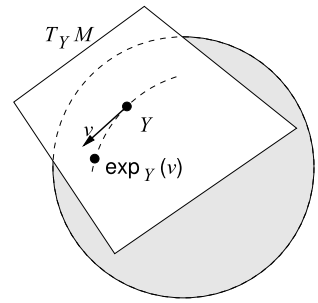

FIG. 3. *Tangent space at point Y on a sphere.*

A Riemannian manifold is a differential manifold equipped with an inner product on the tangent space that varies smoothly from point to point. As a result, it induces a metric on the manifold by means of the projection of points in the tangent. This metric is called the *Riemannian metric* and allows distances and angles on the manifold to be measured. The Riemannian distance from a point in the manifold to Y is equal to the Euclidean distance between its projections in the tangent space T_Y.

The locally shortest constant-velocity curves between two points in a manifold are called *geodesics*.

The map that allows us to move from the the tangent space to the manifold is called the *exponential map* (see Figure 3). This is an important expression that we will need in our work. Given a vector in the tangent T_Y, the exponential map is defined by the geodesic starting at Y and goes in that direction for a unit time. The inverse of the exponential map is called the *logarithmic map*.

2.2 Riemannian Structure of the Shape Space

Moving on to the shape space, remember that the case $m > 2$, which is the case of our application, is a complicated space with singularities. Since it is easier to work with the pre-shape hypersphere, we must make use of the theory of Riemannian submersions to obtain the Riemannian structure of the nonsingular part of the shape space.

Let π be the map to the quotient space that assigns the corresponding element on the shape space to each preshape Z:

$$\pi : S_m^k \to \Sigma_m^k = S_m^k/SO(m)$$
$$Z \mapsto \pi(Z).$$

The orbit $\pi^{-1}(\pi(Z)) = \{XR : R \in SO(m)\}$ is the set of all the pre-shapes that have the same shape as Z, and is also referred to as the fibre.

The theory of Riemmanian submersions tell us that $S_m^k/SO(m)$ is a Riemannian manifold with the quotient

metric if the action of the group $SO(m)$ is free, that is, $R, S \in SO(m)$, $R \neq S$ implies $ZR \neq ZS$.

For $m > 2$ this action is not everywhere free because if Z is a preshape with rank $m-2$ or less and we consider the matrix in $SO(m)$ defined as $R_m \text{diag}(I_{m-2}, R_2) R_m^T$ with $R_2 \in SO(2)$ and $R_m \in SO(m)$, then $ZR_m \text{diag}(I_{m-2}, R_2) R_m^T = Z$. Hence, the images of preshapes with rank $m-2$ or less are singularities in the shape space. Outside the singularity set, the shape space inherits the Riemannian structure from the pre-shape sphere (Le and Kendall, 1993, Kendall et al., 1999).

It is important to note that when we work with real-world applications, we can usually assume that our data are almost sure in the nonsingular part of the shape space. For instance, in our application we can obviously assume probability zero of finding a body of a child with all the landmarks of his pre-shape lying on a line.

Using this previous result, the Riemannian distance and the exponential and logarithmic maps in Σ_m^k can be computed (Le and Kendall, 1993, Dryden and Mardia, 1998, Kendall et al., 1999).

The inherited Riemannian distance in the shape space is the quotient metric, and it is obtained rotating the representative of one of the shapes to be as close as possible with respect to the great circle distance to the other one (see Figure 5). This distance is also called the Procrustes distance.

DEFINITION 3. Given two configuration matrices Y_1 and Y_2, the Procrustes distance between their respective shapes is:

$$\rho(Y_1, Y_2) = \min_{R \in SO(m)} d(Z_1, Z_2 R),$$

where $Z_j = HY_j / \|HY_j\|$, $j = 1, 2$, and $d(Z_1, Z_2 R) = 2 \arcsin(\|Z_1 - Z_2 R\|/2)$ is the Riemannian distance (great circle distance) in the hypersphere S_m^k.

The solution to this optimization problem is

$$\rho(Y_1, Y_2) = \arccos\left(\sum_{i=1}^{m} \lambda_i\right),$$

where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{m-1} \geq |\lambda_m|$ are the square roots of the eigenvalues of $Z_1^T Z_2 Z_2^T Z_1$, and the smallest value λ_m is the negative square root if and only if $\det(Z_1^T Z_2) < 0$ (Dryden and Mardia, 2016). By definition, note that the range of this distance is $[0, \pi/2]$.

Evidence of the increasing complexity in space when passing from $m = 2$ to $m = 3$ is described next. As was commented before, the space of the triangles in $\mathbb{R}^2$ (Σ_2^3) has a well-known Riemannian structure and the Riemannian distance can be analytically obtained in an easy way. However, in the case of the simplest 3D shape space Σ_3^4, an explicit simple expression for the Procrustes distance between two tetrahedrons does not exist.

Figure 4(a) shows three tetrahedrons with three different shapes. To get a visual idea of the Procrustes distance in this space, a multidimensional scaling (MDS) has been applied to the Procrustes distance matrix between them. In this way, they are displayed as points on a Euclidean 2D space with the corresponding distances (Cox and Cox, 2000) in Figure 4(b).

Let us see now how the Riemannian submersion theory allows us to calculate the exponential and logarithmic maps. Given a point in the pre-shape sphere Z, we can find two kinds of tangent vectors at Z, horizontal and vertical (see Figure 5). The horizontal vectors are invariant under rotations. A Riemannian submersion maps the horizontal subspace of the tangent space

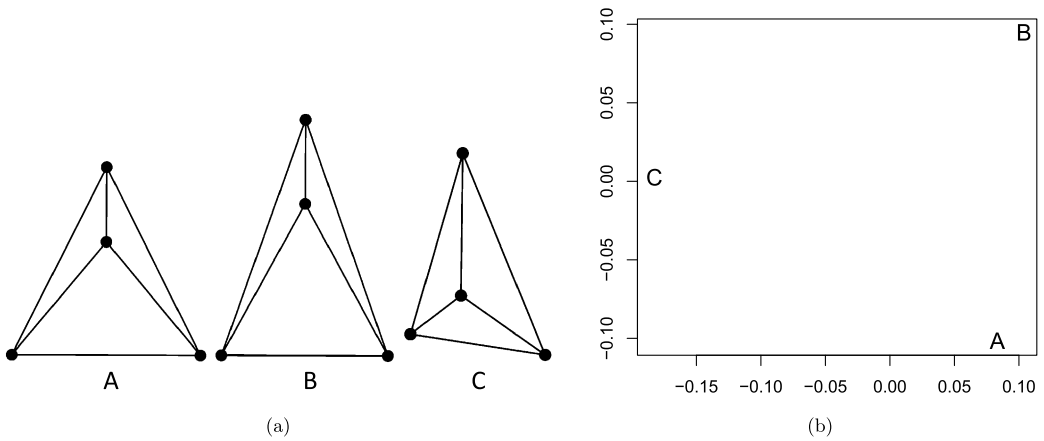

FIG. 4. *Three tetrahedrons and MDS applied on its Procrustes distance matrix.*

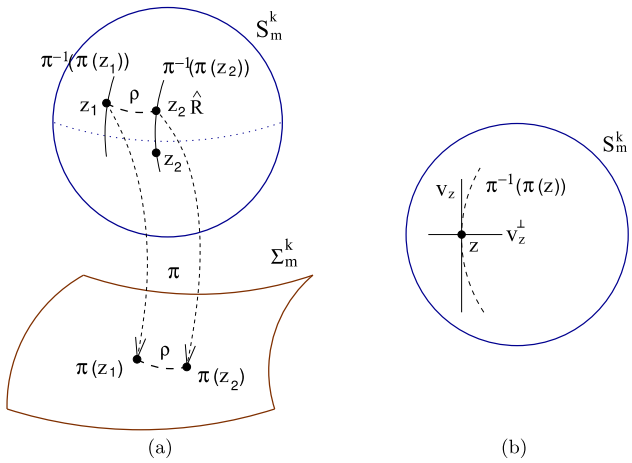

FIG. 5. *A graphical representation of the Riemannian submersion of the shape space in the pre-shape hypersphere*: (*a*) *induced distance between* $\pi(Z_1)$ *and* $\pi(Z_2)$. (*b*) *Horizontal and vertical tangent spaces at* Z.

to the pre-shape sphere isometrically onto the tangent space to the shape space at $\pi(Z)$.

As a result, if Y is the representative of a point in Σ^k_m, to obtain the expression of the projection onto the tangent plane at Y of a pre-shape Z, Z is rotated to be as close as possible again with respect to the great circle distance to its preshape. We write the rotated preshape as $Z\hat{R}$. The expression of $\hat{R}$ can be found on page 61 of Dryden and Mardia (1998):

$$\hat{R} = UV^T,$$

where $U, V \in SO(m)$ are respectively, the right and left matrices of the singular value decomposition of $Z_Y^T Z$. Then (Le and Kendall, 1993, Dryden and Mardia, 2016):

$$\log_Y(Z) = \left(I_{km-m} - \operatorname{vec}(Z_Y)\operatorname{vec}(Z_Y)^T\right) \times \operatorname{vec}(Z_Y\hat{R})\frac{\rho}{\sin(\rho)}, \tag{1}$$

where I_{km-m} is the $(km-m)\times(km-m)$ identity matrix, $\rho = \rho(\pi(Z_Y), \pi(Z))$, and the vectorizing operator of an $l \times m$ matrix A with columns $a_1, a_2, \ldots, a_m$ is defined as: $\operatorname{vec}(A) = (a_1^T, a_2^T, \ldots, a_m^T)^T$.

Given v in the tangent space at Y,

$$\exp_Y(v) = \operatorname{vec}^{-1}\left(\cos\left((v^T v)^{1/2}\right)\operatorname{vec}(Z_Y) + \frac{\sin((v^T v)^{1/2})}{(v^T v)^{1/2}} v\right)\hat{R}^T. \tag{2}$$

See Dryden and Mardia (1998, 2016) and Small (1996) for a more complete discussion of the tangent space.

Because of the non-Euclidean structure of the shape space, to introduce the concept of mean shape of a given set of shape realizations, a Fréchet-type mean (Fréchet, 1948) must be used, that is, one that minimizes the sum of squared distances from any shape in the set. We have the following definition.

DEFINITION 4. Given a set of configuration matrices $Y_1, \ldots, Y_n$, the empirical Fréchet mean in Σ^k_m is given by $\widehat{\mu}$, where

$$\widehat{\mu} = \arg\min_{\mu\in\Sigma^k_m} \sum_{i=1}^n \rho^2(Y_i, \mu). \tag{3}$$

An explicit analytical solution for the optimization problem 3 was found for two-dimensional data (Kent, 1994, Le, 2001). This solution is based on the eigenvectors of a complex sum of squares and matrix products (see page 44 in Dryden and Mardia, 1998). For $m = 3$ and higher-dimensional data, the solution cannot be found analytically, and an iterative procedure based on a gradient descent algorithm must be used.

In Pennec (2006), we can find this algorithm for a general Riemannian manifold $\mathcal{M}$. To characterize a local minimum of a twice differentiable function, we just have to require a null gradient and a positive definite Hessian matrix.

Given a point $z \in \mathcal{M}$, the gradient of the function

$$h_z(y) = \rho^2(y, z), \quad y \in \mathcal{M},$$

is, according to Pennec (2006),

$$(\operatorname{grad} h_z)(y) = -2\log_y(z),$$

where $\log_y(z)$ denotes the projection of z onto the tangent plane at y, that is, the inverse of the exponential map.

Therefore, given a set of points $\{x_1, \ldots, x_n\} \in \mathcal{M}$, if we consider the function $f : \mathcal{M} \longrightarrow \mathbb{R}$ defined as

$$f(y) = \frac{1}{n}\sum_{i=1}^n \rho^2(y, x_i),$$

where ρ denotes the Riemannian distance in $\mathcal{M}$ and we suppose that the points x_i are away from any singularity, we have

$$(\operatorname{grad} f)(y) = \frac{1}{n}\sum_{i=1}^n (\operatorname{grad} h_{x_i})(y) = -\frac{2}{n}\sum_{i=1}^n \log_y(x_i). \tag{4}$$

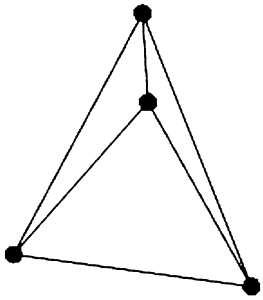

FIG. 6. *Empirical Fréchet mean of the three tetrahedrons of Figure 4.*

The gradient descent algorithm is

$$y_{t+1} = \exp_{y_t}\left(\frac{\sum_{i=1}^{n} \log_{y_t}(x_i)}{n}\right) \tag{5}$$

This constant step-size gradient descent algorithm is the most popular and easiest version of the gradient descent method on Riemannian manifolds (Afsari, Tron and Vidal, 2013).

A modification of this algorithm will be used to obtain our nonparametric regression procedure in Σ_m^k.

The empirical Fréchet mean of the three tetrahedrons of Figure 4 can be seen in Figure 6.

It is worth noting at this point that if the data are fairly concentrated around the mean, the Euclidean distance in the tangent space around the mean shape is a good approximation to ρ, that is, the tangent space is the linearized version of the shape space in the vicinity of the mean, and so we can perform standard multivariate statistical techniques in this space. This is an approach to inference on shape space that is widely used in many applications. A comparison of this approach and our methodology will be shown in Section 5.1.

3. KERNEL REGRESSION ALGORITHM IN THE SHAPE SPACE

In this section, we describe our new methodological proposals.

Let us consider the regression problem with response in the shape space and Euclidean p-dimensional covariates, that is, given a sample $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$, where Y_i, $i = 1, \ldots, n$ are random configuration matrices and X_i, $i = 1, \ldots, n$ are real valued p-dimensional vectors (random or not). Our aim is to estimate the regression function, $\mu(X)$, to predict the shape of an object given the covariates $X \in \mathbb{R}^p$.

As was commented before, classical regression methods are not applicable in this setting because they rely on the vector space structure of the observations.

Davis et al. (2007) used the notion of Fréchet expectation $\mu(X) = E(Y|X)$ to generalize the Euclidean case regression to a general Riemannian manifold $\mathcal{M}$. They proposed a method that generalizes Nadaraya-Watson kernel regression (Nadaraya, 1964) in order to predict manifold-valued data from (t_i, p_i), where t_i are drawn from a univariate random variable and p_i are points in the manifold. They define a manifold kernel regression estimator using the Fréchet empirical mean estimator:

$$m_h(t) = \arg\min_{q \in \mathcal{M}} \left(\frac{\sum_{i=1}^{n} K_h(t - t_i)\rho^2(q, p_i)}{\sum_{i=1}^{n} K_h(t - t_i)}\right), \tag{6}$$

where K_h is a univariate kernel function with bandwidth h.

They used this method to study spatio-temporal change in a random design database consisting of three-dimensional MR (Magnetic Resonance) images of healthy adults to compute representative images over time. This work is completely different from ours for two reasons. First, as they are interested in the study of changes in anatomy, its response variable takes values in the Riemannian manifold of diffeomorphic transformations and not in the landmark-based 3D shape space. Second, their regression is a function of just one scalar variable: the age.

There are many situations, in particular in our application, where there are many explanatory variables that determine the shape of an object, instead of just one. We propose to extend the Davis et al. (2007) estimator for general manifold-valued data to multiple explanatory variables by using a multivariate kernel (Härdle et al., 2004). So

$$m_H(X) = \arg\min_{Z \in \Sigma_m^k} \left(\frac{\sum_{i=1}^{n} K_H(X - X_i)\rho^2(Z, Y_i)}{\sum_{i=1}^{n} K_H(X - X_i)}\right) \tag{7}$$

where $K_H(X) = |H|^{-1/2} K(H^{-1/2}X)$, H is the $p \times p$ matrix of smoothing parameters, symmetric and positive definite and $K : \mathbb{R}^p \to [0, \infty)$ is a multivariate probability density.

As it is well known, there are a great number of possible kernel choices, but the difference between two functions K is almost negligible. The choice of the bandwidth matrix H is the most important factor affecting the accuracy of the estimator.

In our application we have chosen a multivariate Gaussian kernel because it is the easiest way to incorporate the correlation among covariates. In this way, we can put more emphasis on regions with more data

Algorithm 1 Kernel regression algorithm to predict the shape of a 3D object

Given a sample $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$, where $Y_i, i = 1, \ldots, n$ are configuration matrices and $X_i \in \mathbb{R}^p$, $i = 1, \ldots, n$. Let X_0 be a vector of covariate values. The algorithm to predict the shape corresponding to X_0 is:

Initialize $m_0 = Y_i$ (i at random), $\delta \in (0, 1)$, $d = 2\delta$, $j = 0$, h;
Compute the preshapes of $Y_1, \ldots, Y_n \to Z_1, \ldots, Z_n$.
while $(d > \delta)$ and $(j <$ maxSteps$)$ **do**
  Compute the preshape of m_j.
  for $i = 1, \ldots, n$ **do**
    Compute the singular value decomposition of $m_j^T Z_i$, and let u and v be the left and right matrices of this decomposition.
    $\phi = vu^T$
    Compute $\log_{m_j}(Z_i)$ from Eq. (1)
    $k_i = K_H(X_0 - X_i)$
  end for
  $v = \sum_i k_i \log_{m_j}(Z_i) / \sum_i k_i$
  $m_{j+1} = \exp_{m_j}(v)$ (Eq. (2))
  $d = \rho(m_j, m_{j+1})$
  $j = j + 1$
end while
Return m_j

and assign less weight to observations in regions with sparse data. With respect to the choice of the bandwidth matrix H, we propose to use $H = hS_X$, where S_X is the sample covariance matrix of $\{X_1, \ldots, X_n\}$, and then choosing the positive constant h by cross validation.

The algorithm that we propose for solving our regression problem is based on applying the concepts introduced in the previous section to Eq. (7). The distance ρ is given in Definition 3, and a modification of the algorithm stated in Eq. (5) is used for solving the minimization problem.

Taking into account all these considerations, we propose Algorithm 1.

As mentioned in Section 1, this algorithm will be used to predict the body shape of a child given a number of features such age, height, waist circumference, etc.

With respect to theoretical results about convergence and consistency, we would like to note that the rate of convergence of Newton-type algorithms in Riemannian manifolds is studied in Dedieu, Priouret and Malajovich (2003) and Alvarez, Bolte and Munier (2008). On the other hand, Bhattacharya and Patrangenaru (2002, 2003) focus on the asymptotic consistency properties of the intrinsic means and variances for large sample sizes on Riemannian manifolds. Here, our concern is more applied and our aim is just to find at least some good approximations. For this reason, these properties will be checked empirically in the experimental sections.

3.1 Confidence Regions

It is also of interest for the apparel industry to generalize confidence intervals, which are widely used in statistics, to build a region wherein the predictions lie with a given confidence level.

Our approach follows the ideas stated by González-Rodríguez, Trutschnig and Colubi (2009) for obtaining confidence regions for the mean of a fuzzy random variable. It is well known that given X, a real-valued random variable with mean μ and finite variance, an $(1-\alpha) \times 100\%$ confidence interval for μ can be determined as $\mathrm{CI} = [\bar{X} - \delta, \bar{X} + \delta]$, where $\bar{X}$ is the sample mean of a random sample of n independent variables $X_1, \ldots, X_n$, with the same distribution as X, and where $\delta = \delta(X_1, \ldots, X_n)$ is such that $P(\mu \in \mathrm{CI}) = 1 - \alpha$. Therefore, conventional confidence intervals for the mean μ can equivalently be seen as balls with respect to the Euclidean distance, centered in the sample mean $\bar{X}$, and with a suitable radius δ.

Applying these ideas to our regression context, we can define the confidence ball for the mean $\mu(X_0) = E(Y|X_0)$, with level of confidence $1 - \alpha$, $\mathrm{CB}_{1-\alpha}$, as

$$\mathrm{CB}_{1-\alpha} = \left\{Y \in \Sigma_m^k : \rho\big(Y, m(X_0)\big) \leq \delta\right\} : \tag{8}$$
$$P\big(\mu(X_0) \in \mathrm{CB}_{1-\alpha}\big) = 1 - \alpha$$

As in many other statistical problems, no procedure to calculate δ is available other than bootstrap methods. In particular, we propose to use pairwise resampling bootstrap; see Mammen (2000).

Given the sample $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$, and given $\alpha \in (0, 1)$, the chosen significance level, the procedure for building the confidence region can be schematized as follows:

1. Let $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$ be a random sample where Y_i is a shape and X_i a vector of real covariates. Let X_0 be the vector of covariate values whose shape is to be predicted, and let $m(X_0)$ be the mean estimated with this random sample.
2. Obtain B bootstrap sample sets $\{(X_1, Y_1)^{b^*}, \ldots, (X_n, Y_n)^{b^*}\}$ (where $b^* = 1, \ldots, B$) from the original

random sample $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$. For each resample, compute its corresponding mean, and let this be $m(X_0)^{b^*}$.

3. Compute the distances between the sample mean and each bootstrap sample mean, that is, calculate

$$d_b^* = \rho\big(m(X_0)^{b^*}, m(X_0)\big), \quad \text{for } b = 1, \ldots, B.$$

4. Choose δ as one of the $(1-\alpha)$ quantiles of the sample $(d_1^*, \ldots, d_B^*)$.

3.2 The Curse-of-Dimensionality Problem

It is well known (Stone, 1985, Hastie, Tibshirani and Friedman, 2001, Györfi et al., 2002, Klemelä, 2014) that kernel regression models can suffer from the "curse of dimensionality" if the number of explanatory variables is large. This is the case of our application in Section 6. It is impossible to simultaneously maintain localness (low bias) and a sizeable sample in the neighborhood (low variance) as the dimension increases, without the total sample size increasing exponentially in the dimension of the data (Hastie, Tibshirani and Friedman, 2001).

Several solutions can be found in the literature (Hastie, Tibshirani and Friedman, 2001, Klemelä, 2014) to deal with the curse of dimensionality. The usual approaches include single-index models, additive models, partially linear models and several extensions and combinations of them. We must note that many of these approaches cannot be directly used when the response variable is in a manifold because the addition is not defined.

Among the different possibilities, a simple approach is the single-index model (Klemelä, 2014). The single-index model is a semiparametric model that assumes that the regression function satisfies

$$m(X) = g\big(X^T\theta\big), \quad X \in \mathbb{R}^p,$$

where $g : \mathbb{R} \to \mathbb{R}$ is an unknown link function, and $\theta \in \mathbb{R}^p$ is the unknown index vector with $\|\theta\| = 1$. This approach is well suited for our application because the covariates may be highly correlated.

The estimation of θ and g can be made iteratively using the so called M-Estimation approach. For a given $\theta \in \mathbb{R}^p$, g can be estimated using a univariate kernel type estimator, then the estimator of θ is obtained by minimizing $\sum_{i=1}^n \rho^2(Y_i - \hat{g}_\theta(X_i^T\theta))$.

In our case, as our response variable is on a Riemanian manifold, we will use the generalization of the Nadaraya–Watson kernel regression to manifold-valued data given in equation (6):

$$\hat{g}_\theta\big(X_i^T\theta\big) = \arg\min_{Z \in \Sigma_m^k} \left(\frac{\sum_{i=1}^n K_h((X - X_i)^T\beta)\rho^2(Z, Y_i)}{\sum_{i=1}^n K_h((X - X_i)^T\beta)} \right),$$

where K_h is a univariate kernel function with bandwidth h.

We have to note that this estimator has a very large computational complexity.

4. IMPLEMENTATIONS

In our implementations, we have used the R language and the `shapes` package by Ian Dryden. This package provides many useful tools for the statistical analysis of shapes that allowed us to reduce the time spent on the implementation. It works very fast for small data sets, but its speed is somewhat limited for medium and large data sets.

Hence, we have rewritten some parts to accelerate it and enable us to run our codes in a shorter time. Specifically, we have improved routines `preshape` and `centroid.size` since they were the most time-consuming parts in our application. We computed a performance profile of both routines, and in our case they had the same bottleneck: Their main cost was the explicit building of the Helmert matrix and then the product of that matrix by the input argument (our dataset). We have improved the code so that the Helmert matrix is not explicitly built and therefore it is implicitly applied to the input argument.

In our case, the input argument (our dataset) was a matrix with dimensions 3075×3. The original routine `preshape` took an average of 49.13 seconds with these data. The new implementation takes an average of 0.056 seconds. Hence, the new code is about 877 times faster.

The original routine `centroid.size` took an average of 24.54 seconds. The new implementation takes an average of 0.028 seconds. Hence, the new code is about 876 times faster.

These improvements in speed have made the full procedure much faster and its overall time length is now more reasonable.

5. SIMULATION STUDY

As an illustration of the performance of the methodology, we have carried out a simulation study.

First, we generate a compact geometric figure Z described by k landmarks (see Figure 7). Then, we

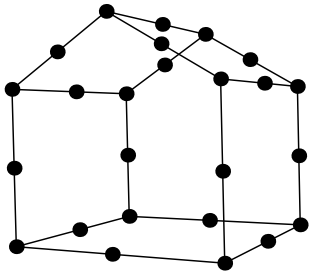

FIG. 7. *Landmarks corresponding to the original geometric figure.*

generate eight new geometric figures modifying the shape of Z by means of three scalar covariates $X = (X^{(1)}, X^{(2)}, X^{(3)})$. Each covariate takes two different values $\{10, 20\}$, and each new figure, that we denote by $\mu(X)$, is obtained by scaling the j-th coordinate of the Z landmarks by $X^{(j)}$, $j = 1, 2, 3$. Figure 8 shows the landmarks of these eight objects. Finally, a random geometric object Y given X is defined by a multivariate normal distribution with $3k$-dimensional mean vector $\mu(X)$ and a $3k \times 3k$ covariance matrix Σ, that is,

$$\operatorname{vec}(Y|X) \sim N_{3k}\big(\operatorname{vec}\big(\mu(X)\big), \Sigma\big).$$

Given X, the distribution of the shape of Y is called normal offset. In Dryden and Mardia (1998) p. 130, the density with respect to the uniform measure in the shape space is given for the 2-dimensional isotropic case, i.d., when the covariance matrix Σ is a multiple of the identity. In this case, the mean shape $m(X)$, calculated by means of the general algorithm, is a consistent estimate of $\mu(X)$.

Two random samples of sizes 50 and 25, respectively, of Y given X are obtained for each combination of X-values resulting in random samples of size $n = 400$ and $n = 200$: $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$. We take $\Sigma = \sigma I_{k\times 3}$, and two values for σ $(0.01, 0.05)$ are selected in such a way that the data are more or less dispersed.

Figure 9 shows a simulated shape Y_i given X, its prediction, and $\mu(X)$ for $X = (10, 10, 20)$ and both σ-values. We have to note that in the case of the simulated random house, landmarks do not always lie on lines. However, we have added these lines as an artefact to better visualize the 3D shape.

In order to do a quantitative analysis and to choose the optimum value of h, the Procrustes distances between the real and the predicted shape for each one of the eight sets of covariates are computed. As an illustration, each cell of Table 1 shows the mean of these values for $\sigma = 0.05$, different values of h, and different

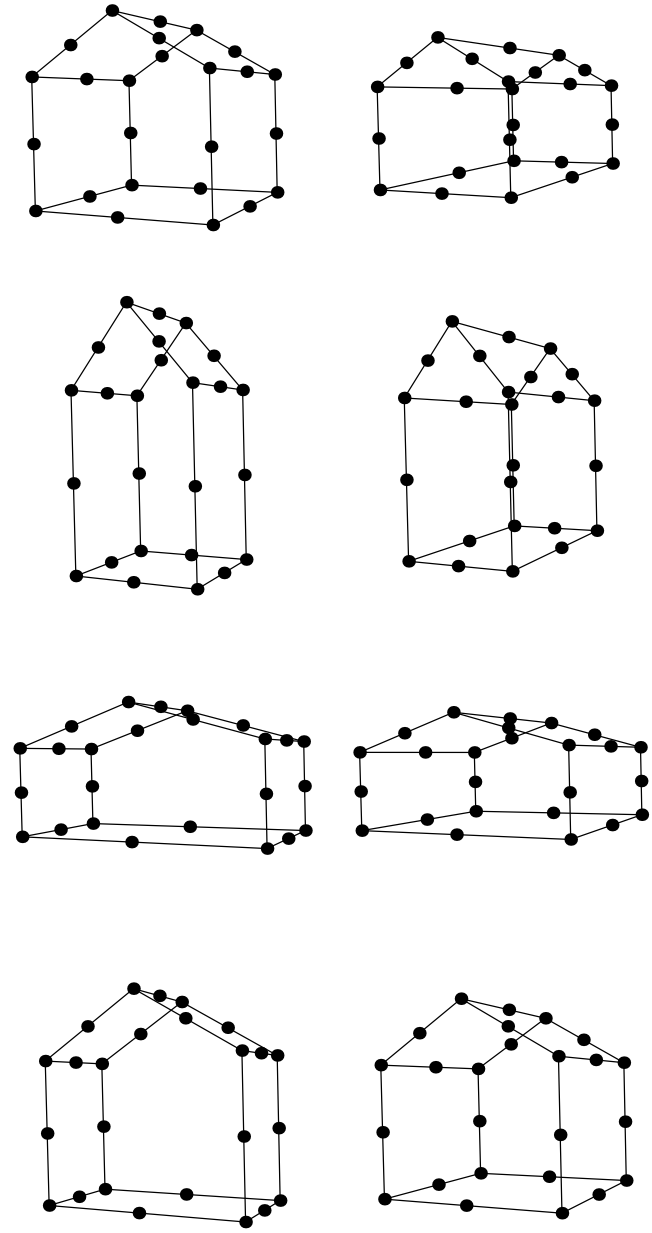

FIG. 8. *Mean shapes for different combinations of X-values.*

numbers of iterations. We can see that they reach the minimum values and become stable after around 2000 iterations. In addition, they are quite robust for small values of h reaching the minimum for $h = 0.25$. These obtained distances are much smaller than the average of pairwise distances in the simulated set (0.3809). We think that these good results validate the proposed method.

With respect to the confidence regions, there are theoretical consistency results that justify bootstrap confidence intervals in Euclidean spaces, but these results are not available in our context; hence, simulation stud-

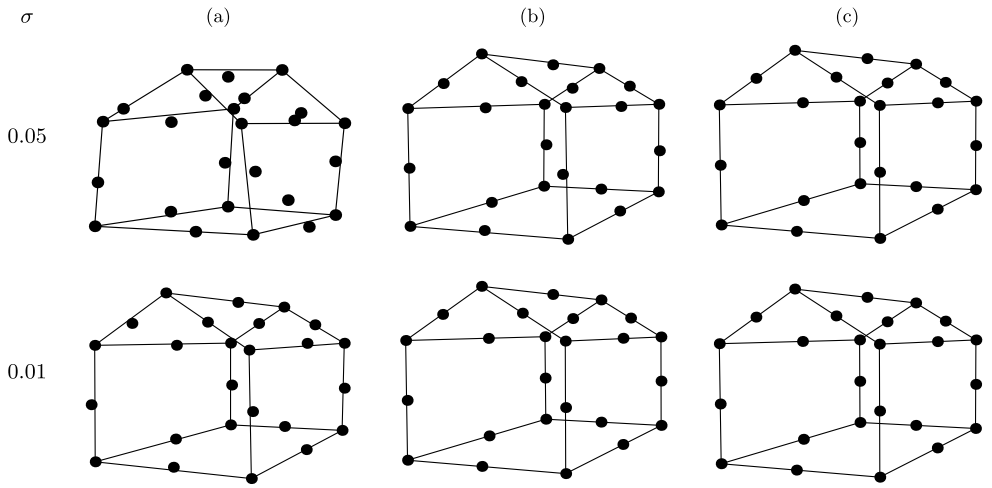

FIG. 9. *(a) A simulated shape Y_i given X, (b) predicted mean $m(X)$, (c) theoretical mean $\mu(X)$ for $X = (10, 10, 20)$ for both σ values (top row for $\sigma = 0.05$ and bottom row for $\sigma = 0.01$).*

TABLE 1
Procrustes distances between the real and the predicted shape

	Number of iterations					
h	100	250	500	1000	2000	3000
0.1	0.209	0.145	0.080	0.029	0.017	0.017
0.25	0.209	0.145	0.080	0.029	0.017	0.017
0.5	0.211	0.149	0.085	0.035	0.020	0.019
1.0	0.220	0.169	0.116	0.074	0.059	0.058
1.5	0.229	0.187	0.145	0.111	0.099	0.098

ies are, at this moment, the only way to assess its performance.

In order to evaluate the actual performance of the bootstrap confidence sets, a total number of 100 original samples of size 400 and 200 are generated, $S_i = \{(X_{i1}, Y_{i1}), \ldots, (X_{in}, Y_{in})\}$ for $i = 1, \ldots, 100$, and the corresponding prediction means are obtained, $m(X)_1, \ldots, m(X)_{100}$.

Bootstrap samples with $B = 100$ are taken from each sample S_i and the corresponding bootstrap confidence sets at a 95% confidence level (nominal coverage) are constructed: $\mathrm{CB}^1_{0.95}, \ldots, \mathrm{CB}^{100}_{0.95}$, or in other words, the radii $\delta_1, \ldots, \delta_{100}$ are obtained. The observed coverage proportion of the theoretical prediction in such confidence regions is calculated as

(9)
$$\hat{p}_i = \frac{\mathrm{card}\{\mathrm{CB}^i_{0.95} : m(X) \in \mathrm{CB}^i_{0.95}, i = 1, \ldots, 100\}}{100}.$$

The results of the simulation study show that the method achieves good observed coverage proportions.

TABLE 2
Simulation results showing observed coverage proportions for a nominal coverage of 95%

Sample size	σ	Coverage proportion
$n = 200$	0.01	100%
	0.05	100%
$n = 400$	0.01	89%
	0.05	98%

Table 2 summarizes the numerical outputs of our simulation study.

5.1 Comparison with the Multivariate Linear Model in the Approximated Tangent Space

Because regression is a topic not often addressed in shape analysis, there is no baseline model that works in the shape space and completely suits our application. In order to assess the goodness of our method, we compare our results with the result of applying the linear regression model to our data projected in the Euclidean tangent space. As was previously commented, this is a very common approximation in the shape analysis literature. The comparison has been reported in Table 3. This table shows that the average and maximum of the prediction errors are clearly lower using our new methodology.

In addition to this error assessment, we have also compared the speed of both methods since the time performance is an important factor in our application and in many other cases. Table 4 shows the time spent performing one prediction of a house by using the kernel

TABLE 3
Minimum, average, and maximum prediction error for both the kernel regression in the shape space ($h = 0.25$, 2000 iterations) and the linear model in the tangent space

	Kernel regression in the shape space	Linear regression in the tangent space
Minimum	0.015	0.012
Average	0.017	0.029
Maximum	0.019	0.045

regression in the shape space ($h = 0.25$ and 2000 iterations) and the linear regression in the tangent space. The performances have been assessed with an increasing number of landmarks. Each row of the table contains the results obtained on a different dataset with an increasing number of landmarks per house, where the first row of the table is the dataset described above. As it can be seen, the times required by the kernel regression are much smaller than the times required by the linear regression in the tangent space. The kernel regression is about 10 times faster for the smallest case, and about 30 times faster for the largest case, thus showing a smaller dependence on the number of landmarks. As the number of landmarks in the children dataset is much larger (3075), we believe that our new approach can be much faster on that dataset.

6. APPLICATION TO CHILDREN'S BODY SHAPES

The aim of this section is to show how the aforementioned algorithm can be used to predict the body shape of a child based on a small number of his or her anthropometric measurements and his or her age. The predicted shape could then be used to choose the most suitable size in a potential online sales application. There are multiple ways to perform this last step and all of them depend on the manufacturer. For instance, one possibility would be to calculate the Procrustes distance between the predicted shape and the shapes of the mannequins for each size.

TABLE 4
Time in seconds for both the kernel regression in the shape space ($h = 0.25$, 2000 iterations) and the linear model in the tangent space

Landmarks per object	Landmarks per edge	Kernel regression in the shape space	Linear regression in the tangent space
25	3	0.19	1.94
55	5	0.23	5.64
100	8	0.25	7.27
250	18	0.37	11.76
505	35	0.67	19.49

6.1 Children Dataset

Our database consists on the 3D full body scans of a randomly selected sample of 739 Spanish children aged from 3 to 12 and several anthropometric measurements.

They were scanned using a Vitus Smart 3D body scanner from Human Solutions, a nonintrusive laser system which performs a sweep of the whole body. First, several cameras captured images. Second, associated software provided by the scanner manufacturer detected the brightest points and used them to make a triangulation that provides information about the 3D spatial location of points on the body surface.

With the raw data provided by the body scanner, the following preprocessing steps were applied: First, the data were cleaned to remove redundancies, to fill holes and to remove noise. Then, the posture-harmonised homologous models were obtained from the previously obtained 3D scan data. This harmonization process consisted of two steps for each body scan: The first one was to fit a template with an anatomical skeleton to the scan surface to obtain a homologous structured representation. The second step was to use the skeleton to make the homologous representation adopt the same standard posture. For a detailed description, see Ballester et al. (2014). Finally, a database of individual 3D homologous avatars with anatomical one-to-one vertex correspondence among them was created. In the final dataset the 3D scan of the body of each child comprises 3075 landmarks.

From the 3D mesh, several anthropometric measurements were calculated semi-automatically by combining automatic measurements based on geometric characteristic points with a manual review. The anthropometric measurements obtained for each child are, among others, the following: height, bust circumference, waist circumference, hip circumference, right leg length, left leg length, right arm length and left arm length. In addition to those measurements other sociodemographic variables were available for each child, such as gender and age.

Figure 10 shows two examples contained inside our dataset. The left example is a boy and the right example is a girl. Both of them are 12 years old. Each example consists of 3075 landmarks, generated with the previously described preprocessing procedure.

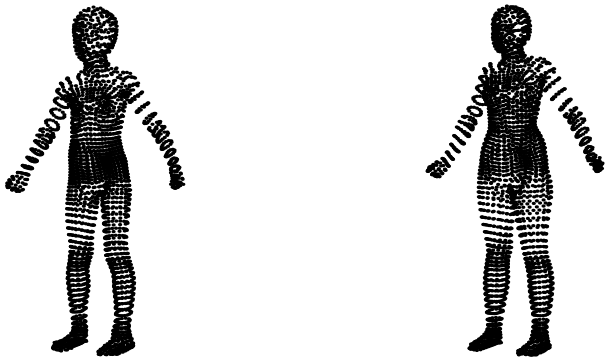

FIG. 10. *Two examples of 3D body scans from our dataset (a boy in the left and a girl in the right).*

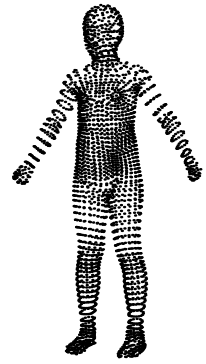

FIG. 11. *Shape predicted for a boy with the following covariates* $X_0 = (12, 1508.5, 784.75, 694.25, 845.25, 964, 955.75, 500.5, 502.25)$.

6.2 Procedure

In order to illustrate our procedure, two subsamples of our dataset have been chosen. For both samples, children over 7 years old were selected. The first sample consisted of 244 boys; the second consisted of 251 girls. The body shape of each child in our data-set was represented by 3075 3D landmarks, that is, by a 3075×3 configuration matrix.

Nine covariates were chosen in order to predict the shape of a child: age, height, bust circumference, waist circumference, hip circumference, right leg length, left leg length, right arm length and left arm length. We chose these covariates because they are the most common covariates asked for in online sales of wear, they are easy to measure in a child, and they are well known for everybody.

As an illustration of the developed application, Figure 11 shows the prediction that was obtained when algorithm 1 was applied to predict the shape of a boy with the covariates X_0. The following values for the covariates were employed: age = 12 years, height = 1508.5 mm, bust circumference = 784.75 mm, waist circumference = 694.25 mm, hip circumference = 845.25 mm, right leg length = 964 mm, left leg length = 955.75 mm, right arm length = 500.5 mm and left arm length = 502.25 mm.

In this particular application, it would be desirable to have, in addition, a prediction of the children's size. Although a kernel regression in the size and shape manifold could be applied, we can employ a rather simpler approach. Because shape and size could be considered independent, we can conduct separately a kernel regression for the shape and a univariate kernel regression for the size, given the above set of covariates. Then, we combine both predictions.

In order to perform a quantitative analysis of the effectiveness of the method, a 3-fold cross validation study was conducted. The full sample was randomly partitioned into three equal-sized subsamples. At each step of this study, one of the three subsamples was predicted from the covariates and the landmark data of the other two subsamples. Next, the Procrustes distance between the real shape and the predicted shape was calculated. The means of these prediction errors for different values of h and different numbers of iterations are shown in Table 5. We can see that these prediction errors are larger for small numbers of iterations, and they reach the minimum and become stable after around 1000 iterations. In contrast, the prediction errors are quite robust against h-values, reaching the minimum for $h = 1.0$. In general, these errors are considered acceptable in our application, especially taking into account that just 8 anthropometrical measures plus the age are used to predict the shape. In addition, unlike the simulated dataset in the previous section, we must consider that all the shapes in this dataset belong to children and, therefore, they show very similar shapes.

TABLE 5
Distances between real and predicted shapes for different values of h and number of iterations

	Number of iterations				
h	100	250	500	1000	2000
0.5	0.0447	0.0376	0.0351	0.0350	0.0350
1.0	0.0448	0.0373	0.0341	0.0335	0.0334
1.5	0.0450	0.0375	0.0342	0.0334	0.0334
2.0	0.0451	0.0377	0.0344	0.0336	0.0335
2.5	0.0452	0.0379	0.0346	0.0338	0.0338
3.0	0.0453	0.0381	0.0348	0.0341	0.0340
∞ (all children)	0.0459	0.0394	0.0367	0.0362	0.0362

The cross validation study is also used to test how reasonable the bootstrap confidence regions are. At each one of the 24 steps of the cross-validation study, the confidence interval (determined by the corresponding δ-value) is obtained. Table 6 shows these values and the distances between the real and predicted shapes. As we can see, distances are always smaller than δ-values.

As we have commented previously, although the results seem acceptable for our application, the fact that we are using nine covariates might introduce problems with the dimensionality. To get some idea on how big the sample is in the neighborhood of a child, Table 7 reports the average number of observations located at a distance lesser or equal to h. As we can see, the optimum value of h chosen in each case provides a poor size of the sample.

The 3-fold cross validation study was repeated using the single-index model reported in Section 3.2, and the results are shown in Table 8. Note that in this case using the single-index model does not improve the results, and they are even a bit worse than those of the previous proposal. We think that the problem is a combination of the optimization process used to find out the optimal value of the parameter of the model and the high computational cost of a prediction in our application. Recall that alternative methods (i.e. additive models) are not applicable because the addition is not defined in a differential manifold. Even though some other method could be applied in this case, the very high computational cost of our dataset (a large number of landmarks and a rather large number of objects) would make its practical application very difficult.

TABLE 6
δ-values of the 95% CI and distances between the real and predicted shape in 24 cases of the cross-validation study

δ-value	Procrustes distance
0.037	0.036
0.032	0.029
0.035	0.034
0.029	0.027
0.056	0.053
0.033	0.030
0.037	0.036
0.035	0.032
0.029	0.027
0.045	0.043
0.046	0.042
0.025	0.024
0.034	0.032
0.046	0.044
0.029	0.027
0.039	0.036
0.036	0.034
0.036	0.034
0.029	0.027
0.057	0.055
0.025	0.024
0.029	0.027
0.033	0.031
0.024	0.023

TABLE 7
Average number of observations at a distance less than or equal to 2h of each child

	Avg. no. of children	
h	Boys	Girls
0.5	1.02	1.02
1.0	3.38	4.33
1.5	34.04	39.16
2.0	116.84	127.30
2.5	197.78	205.61
3.0	233.65	236.58

7. DISCUSSION

In this paper, we have proposed an approach that represents a novelty in terms of integrating concepts from statistical shape analysis and regression procedures. Although it is an important and common problem in real applications, papers on this subject are scarce in the literature. The main novelty of this work is to show how to apply a general nonparametric regression method in manifold-valued data to the shape space based on landmarks. The second novelty is the generalization of the previous procedure to the case of

TABLE 8
Distances between real and predicted shapes for different values of h and 1000 iterations using the single-index model

h	1000
0.5	0.0487
1.0	0.0485
1.5	0.0490
2.0	0.0483
2.5	0.0484
3.0	0.0489

multiple covariates. Moreover, we propose bootstrap interval confidences for the predictions. A single-index model has been proposed for handling the curse-of-dimensionality problem when the number of covariates is great with respect to the sample size. Furthermore, a simulation study with simple objects was conducted to validate the procedure, and a successful comparison with a current approach was performed.

To illustrate our new methodology, it has been applied to a children body dataset with hundreds of subjects in order to predict the shape of a child given a small set of quantitative measures. The results obtained with this dataset avail the feasibility of our new method. This regression method could be useful for the implementation of an online sales application.

We used the R language and the `shapes` package in our implementations. Due to the large size and large dimensionality of our dataset, the overall computational cost was too large. Therefore, we improved the speed of some parts of certain routines of the aforementioned package to reduce the computational cost of the procedure. The new implementations were significantly faster than the original ones. Moreover, we implemented the most compute-intensive part of the code by using the C programming language and a very efficient library such as BLAS (Basic Linear Algebra Subroutines) to accelerate it even further.

ACKNOWLEDGMENTS

This paper has been partially supported by the grants DPI2013-47279-C2-1-R, DPI2017-87333-R and RTI2018-098156-B-C54 from the Spanish Ministry of Science, Innovation and Universities co-financed by FEDER funds. We would like to thank the Biomechanics Institute of Valencia (Spain) for providing us with the dataset. We also thank Francisco D. Igual-Peña from the Universidad Complutense de Madrid (Spain) for granting access to the computing platform where some of the experiments were conducted.

REFERENCES

AFSARI, B., TRON, R. and VIDAL, R. (2013). On the convergence of gradient descent for finding the Riemannian center of mass. *SIAM J. Control Optim.* **51** 2230–2260. MR3057324

ALVAREZ, F., BOLTE, J. and MUNIER, J. (2008). A unifying local convergence result for Newton's method in Riemannian manifolds. *Found. Comput. Math.* **8** 197–226. MR2407031

AZENCOTT, R. (1994). Deterministic and random deformations; applications to shape recognition. In *Conference at HSSS Workshop in Cortona, Italy*.

AZENCOTT, R., COLDEFY, F. and YOUNES, L. (1996). A distance for elastic matching in object recognition. In *Proceedings of* 12*th ICPR*.

BADDELEY, A. and MOLCHANOV, I. (1998). Averaging of random sets based on their distance functions. *J. Math. Imaging Vision* **8** 79–92. MR1612209

BALLESTER, A., PARRILLA, E., URIEL, J., PIEROLA, A., ALEMANY, S., NACHER, B., GONZALEZ, J. and GONZALEZ, J. C. (2014). 3D-based resources fostering the analysis, use, and exploitation of available body anthropometric data. In 5*th International Conference on* 3*D Body Scanning Technologies*.

BAUER, M., HARMS, P. and MICHOR, P. W. (2012). Sobolev metrics on shape space, II: Weighted Sobolev metrics and almost local metrics. *J. Geom. Mech.* **4** 365–383. MR3011892

BHATTACHARYA, R. and PATRANGENARU, V. (2002). Nonparametic estimation of location and dispersion on Riemannian manifolds. *J. Statist. Plann. Inference* **108** 23–35. MR1947389

BHATTACHARYA, R. and PATRANGENARU, V. (2003). Large sample theory of intrinsic and extrinsic sample means on manifolds. I. *Ann. Statist.* **31** 1–29. MR1962498

BOOKSTEIN, F. L. (1978). Lecture notes in biomathematics. In *The Measurement of Biological Shape and Shape Change* Springer, Berlin.

BOOKSTEIN, F. L. (1986). Size and shape spaces for landmark data in two dimensions. *Statist. Sci.* **2** 181–222.

COX, T. F. and COX, M. A. (2000). *Multidimensional Scaling*. CRC Press, Boca Raton, FL.

DAVIS, B. C., FLETCHER, P. T., BULLITT, E. and JOSHI, S. (2007). Population shape regression from random design data. In *Computer Vision*, 2007. *ICCV* 2007. *IEEE* 11*th International Conference on* 1–7. IEEE, New York.

DEDIEU, J.-P., PRIOURET, P. and MALAJOVICH, G. (2003). Newton's method on Riemannian manifolds: Convariant alpha theory. *IMA J. Numer. Anal.* **23** 395–419. MR1987937

DEL BIMBO, A., DE MARSICO, M., LEVIALDI, S. and PERITORE, G. (1998). Query by dialog: An interactive approach to pictorial querying. *Image Vis. Comput.* **16** 557–569.

DRYDEN, I. L. (2012). `shapes` package. R Foundation for Statistical Computing, Vienna, Austria. Contributed package.

DRYDEN, I. L. and MARDIA, K. V. (1998). *Statistical Shape Analysis*. *Wiley Series in Probability and Statistics*: *Probability and Statistics*. Wiley, Chichester. MR1646114

DRYDEN, I. L. and MARDIA, K. V. (2016). *Statistical Shape Analysis*: *With Applications in R*. Wiley, Chichester. MR3559734

FLETCHER, T. (2011). Geodesic regression on Riemannian manifolds. In *Proceedings of the Third International Workshop on Mathematical Foundations of Computational Anatomy-Geometrical and Statistical Methods for Modelling Biological Shape Variability* 75–86.

FLETCHER, P. T. and ZHANG, M. (2016). Probabilistic geodesic models for regression and dimensionality reduction on Riemannian manifolds. In *Riemannian Computing in Computer Vision* 101–121. Springer, Cham. MR3444348

FRÉCHET, M. (1948). Les éléments aléatoires de nature quelconque dans un espace distancié. *Ann. Inst. Henri Poincaré* **10** 215–310. MR0027464

GONZÁLEZ-RODRÍGUEZ, G., TRUTSCHNIG, W. and COLUBI, A. (2009). Confidence regions for the mean of a fuzzy random variable. In *IFSA-EUSFLAT*, *Lisbon*, *Portugal* (J. P. Carvalho, D. Dubois, U. Kaymak and J. M. da Costa Sousa, eds.) 1433–1438.

GOODALL, C. (1991). Procrustes methods in the statistical analysis of shape. *J. Roy. Statist. Soc. Ser. B* **53** 285–339. MR1108330

GUAL-ARNAU, X., HEROLD-GARCÍA, S. and SIMÓ, A. (2013). Shape description from generalized support functions. *Pattern Recogn. Lett.* **34** 619–626.

GUAL-ARNAU, X., HEROLD-GARCÍA, S. and SIMÓ, A. (2015). Geometric analysis of planar shapes with applications to cell deformations. *Image Anal. Stereol.* **34** 171–182. MR3513425

GYÖRFI, L., KOHLER, M., KRZYŻAK, A. and WALK, H. (2002). *A Distribution-Free Theory of Nonparametric Regression. Springer Series in Statistics*. Springer, New York. MR1920390

HÄRDLE, W., MÜLLER, M., SPERLICH, S. and WERWATZ, A. (2004). *Nonparametric and Semiparametric Models. Springer Series in Statistics*. Springer, New York. MR2061786

HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning. Data Mining, Inference, and Prediction. Springer Series in Statistics*. Springer, New York. MR1851606

JERMYN, I. H., KURTEK, S., LAGA, H. and SRIVASTAVA, A. (2017). *Elastic Shape Analysis of Three-Dimensional Objects. Synthesis Lectures on Computer Vision* **12** 1–185.

JUPP, P. E. and KENT, J. T. (1987). Fitting smooth paths to spherical data. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **36** 34–46. MR0887825

KARCHER, H. (1977). Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.* **30** 509–541. MR0442975

KENDALL, D. G. (1977). The diffusion of shape. *Adv. in Appl. Probab.* **9** 428–430.

KENDALL, D. G. (1984). Shape manifolds, Procrustean metrics, and complex projective spaces. *Bull. Lond. Math. Soc.* **16** 81–121. MR0737237

KENDALL, W. S. (1990). Probability, convexity, and harmonic maps with small image. I. Uniqueness and fine existence. *Proc. Lond. Math. Soc.* (3) **61** 371–406. MR1063050

KENDALL, D. G., BARDEN, D., CARNE, T. K. and LE, H. (1999). *Shape and Shape Theory. Wiley Series in Probability and Statistics*. Wiley, Chichester. MR1891212

KENT, J. T. (1994). The complex Bingham distribution and shape analysis. *J. Roy. Statist. Soc. Ser. B* **56** 285–299. MR1281934

KIM, H., ADLURU, N., COLLINS, M., CHUNG, M., BENDLIN, B., JOHNSON, S., DAVIDSON, R. and SINGH, V. (2014). Multivariate general linear models (MGLM) on Riemannian manifolds with applications to statistical analysis of diffusion weighted images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 2705–2712.

KINDRATENKO, V. V. (2003). On using functions to describe the shape. *J. Math. Imaging Vision* **18** 225–245. MR1971180

KLASSEN, E., SRIVASTAVA, A., MIO, M. and JOSHI, S. H. (2004). Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **26** 372–383.

KLEMELÄ, J. (2014). *Multivariate Nonparametric Regression and Visualization. Wiley Series in Computational Statistics*. Wiley, Hoboken, NJ. With R and applications to finance. MR3222314

KOBAYASHI, S. and NOMIZU, K. (1969). *Foundations of Differential Geometry. Vol. II. Interscience Tracts in Pure and Applied Mathematics, No.* 15 *Vol. II*. Interscience Publishers Wiley, New York. MR0238225

LE, H. (2001). Locating Fréchet means with application to shape spaces. *Adv. in Appl. Probab.* **33** 324–338. MR1842295

LE, H. L. and KENDALL, D. G. (1993). The Riemannian structure of Euclidean shape spaces: A novel environment for statistics. *Ann. Statist.* **21** 1225–1271. MR1241266

LONCARIC, S. (1998). A survey of shape analysis techniques. *Pattern Recognit.* **31** 983–1001.

MAMMEN, E. (2000). Resampling methods for nonparametric regression. In *Smoothing and Regression: Approaches, Computation, and Application* 425–450. Wiley, New York.

MARDIA, K. V. and JUPP, P. E. (2000). *Directional Statistics. Wiley Series in Probability and Statistics*. Wiley, Chichester. Revised reprint of ıt Statistics of directional data by Mardia [MR0336854 (49 #1627)]. MR1828667

MATHERON, G. (1975). *Random Sets and Integral Geometry*. Wiley, New York. MR0385969

MOLCHANOV, I. (2005). *Theory of Random Sets. Probability and Its Applications* (*New York*). Springer, London. MR2132405

NADARAYA, E. A. (1964). On estimating regression. *Theory Probab. Appl.* **9** 141–142.

PENNEC, X. (2006). Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *J. Math. Imaging Vision* **25** 127–154. MR2254442

PRINCE, J. L. and WILLSKY, A. S. (1990). Reconstructing convex sets from support line measurements. *IEEE Trans. Pattern Anal. Mach. Intell.* **12** 377–389.

R DEVELOPMENT CORE TEAM (2014). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

SERRA, J. (1984). *Image Analysis and Mathematical Morphology*. Academic Press, London. English version revised by Noel Cressie. MR0753649

SHI, X., STYNER, M., LIEBERMAN, J., IBRAHIM, J. G., LIN, W. and ZHU, H. (2009). Intrinsic regression models for manifold-valued data. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI* 2009 192–199.

SIMÓ, A., DE VES, E. and AYALA, G. (2004). Resuming shapes with applications. *J. Math. Imaging Vision* **20** 209–222. MR2060144

SMALL, C. G. (1996). *The Statistical Theory of Shape. Springer Series in Statistics*. Springer, New York. MR1418639

SRIVASTAVA, A., KLASSEN, E., JOSHI, S. H. and JERMYN, I. H. (2011). Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **33** 1415–1428.

STONE, C. J. (1985). Additive regression and other nonparametric models. *Ann. Statist.* **13** 689–705. MR0790566

STOYAN, D. and STOYAN, H. (1994). *Fractals, Random Shapes and Point Fields*. Wiley, Chichester. Translated from the 1992 German original by N. Bamber and R. B. Johnson. MR1297125

VINUÉ, G., SIMÓ, A. and ALEMANY, S. (2016). The k-means algorithm for 3D shapes with an application to apparel design. *Adv. Data Anal. Classif.* **10** 103–132. MR3464302

WOODS, R. P. (2003). Characterizing volume and surface deformations in an atlas framework: Theory, applications, and implementation. *NeuroImage* **18** 769–788.

YOUNES, L. (1998). Computable elastic distances between shapes. *SIAM J. Appl. Math.* **58** 565–586. MR1617630

YOUNES, L., MICHOR, P. W., SHAH, J. and MUMFORD, D. (2008). A metric on shape space with explicit geodesics. *Atti Accad. Naz. Lincei Cl. Sci. Fis. Mat. Natur. Rend. Lincei* (9) *Mat. Appl.* **19** 25–57. MR2383560