

Multivariate Planar Curves: A Statistical Framework for Shape Analysis in Images

Issam-Ali Moindjié^{*1,2}, Cédric Beaulac¹, and Marie-Hélène Descary¹

¹Department of mathematics, Université du Québec à Montréal, Canada

²LAMPS, Université de Perpignan Via Domitia, France

^{*}Corresponding author: issam-ali.moindjie@univ-perp.fr

Abstract

Recent developments in computer vision have made segmented images widely available across many domains, such as medicine, where segmented radiographs play an important role in diagnosis. As prediction problems are common in image analysis, this work explores the use of the object contours highlighted by such images as predictors in a supervised classification context. To this end, we develop a new statistical learning framework that accounts for the joint shape of the multiple objects contained in an image. We introduce a formalism that extends the study of a single random planar curve to the joint analysis of several planar curves, referred to as a multivariate planar curve. Modeling the contours jointly, rather than separately, preserves the inter-component information, such as their relative position, scale, and orientation, which is often essential to the analysis. Based on this model, we propose a joint alignment procedure and we extend core inferential tools to multivariate shapes: shape dissimilarity, Fréchet mean estimation, and tangent-space representation. These tangent coordinates are then used as predictors in standard functional classification models. A simulation study shows accurate recovery of deformation parameters over increasing noise levels. Then, through a cardiomegaly detection problem on segmented chest X-rays, we show that jointly modeling the contours is robust to misalignment and improves classification accuracy over both a contour-wise univariate analysis and a naive approach based on the raw curves.

Keywords: alignment, functional data analysis, multivariate functional data, multiple objects, image data, statistical shape analysis, supervised classification.

1 Introduction

Images constitute an important source of information in many scientific and industrial applications, including biomedical research Lu et al. (2019) and road safety Gu et al. (2020), to name just a few. Due to acquisition technologies, images are naturally represented as

matrices of pixels, where each pixel records information such as grayscale intensity or color values.

In many applications, the primary object of interest is not the collection of pixels itself, but rather the geometric structures represented in the image. In particular, the shape of objects often carries meaningful and interpretable information that may be difficult to extract from pixel-level representations alone. When the goal is to study geometry, a shape-based representation can provide a substantially more parsimonious description of the data while preserving the structural characteristics of the objects under investigation.

A common strategy for isolating objects and their shapes is image segmentation, whose goal is to partition an image into homogeneous regions sharing common characteristics such as texture, intensity, or gray level. Image segmentation has long been a central research area in pattern recognition and image analysis. Since the early developments of computer vision, numerous segmentation techniques have been proposed (see Fu and Mui (1981) for a survey). More recently, deep learning approaches have significantly improved segmentation performance in a wide range of applications (see, e.g., Minaee et al. (2021)). These methods have demonstrated remarkable accuracy in practical settings, particularly in medicine, where segmentation plays an important role in image-based diagnosis and clinical decision-making (Wang et al., 2022).

The output of a segmentation algorithm is typically a collection of masks identifying the objects of interest within an image. While such masks provide a natural pixel-based representation of shape, they remain discrete approximations of inherently continuous geometrical structures. Consequently, we argue that segmented images should not be analyzed solely through their masks, but rather through continuous representations of the shapes they contain. Srivastava and Klassen (2016) advocate for a functional representation of shape that preserves the continuous nature of these objects, provides a parsimonious representation, and enables their analysis through the lens of Functional Data Analysis (FDA).

Building upon this idea, we propose to represent the multiple objects contained in an image as a multivariate planar curve. Each object is encoded by its contour, modeled as a planar closed curve described by two coordinate functions. An image containing multiple objects is therefore represented as a vector of planar curves, forming a multivariate functional observation.

As a motivating example, we consider segmented chest X-ray images containing the right lung, the heart, and the left lung, where the statistical problem is to detect cardiomegaly, a condition characterized by an abnormally large heart relative to the lungs. The relevant information is not contained in the shape of a single organ alone, but rather in the joint configuration of multiple anatomical structures. Figure 1 illustrates the proposed representation. Panel (a) shows the original chest X-ray, represented as a matrix of pixel intensities, while panel (b) displays the segmentation masks obtained using the method of Gaggion et al. (2024). Panels (c) and (d) present the objects that form the basis of our analysis. The contours of the right lung, heart, and left lung, denoted C_1 , C_2 , and C_3 , are extracted from the masks and represented as planar closed curves. By traversing the boundary of each organ from the starting point indicated by the black arrow, each contour can be represented by two coordinate functions, namely $C_j^\top = (X_j, Y_j)$, $j = 1, 2, 3$. Collectively, these contours form the multivariate planar curve $\mathbf{C} = (C_1^\top, \dots, C_p^\top)^\top$, with $p = 3$, which constitutes the parsimonious image representation studied throughout this paper.




Figure 1: Representation of a segmented chest X-ray (Gaggion et al., 2024) as a multivariate planar curve. (a) Original X-ray image. (b) Segmentation masks of the three organs of interest: right lung, heart, and left lung. (c) Contours that form the multivariate planar curve $\mathbf{C}$. (d) Coordinate functions (X_j, Y_j) for each contour C_j ($j = 1, 2, 3$), obtained by traversing the contour in (c) from the starting point indicated by the black arrow.

Statistical shape analysis has traditionally focused on objects represented by a single contour. In this setting, a shape is represented by a planar curve C , and a rich literature provides methodologies for registration, shape representation, similarity measurement, representative shape estimation, variability modeling, and classification (see, e.g., Dryden and Mardia (2016), Srivastava et al. (2010), Younes (1998), Srivastava and Klassen (2016)). The motivating example of Figure 1, however, departs from this single-contour setting, as the observation of interest is a collection of several contours rather than a single one. Analyzing such data calls for treating these contours jointly, as a single multivariate planar curve $\mathbf{C} = (C_1^\top, \dots, C_p^\top)^\top$.

In this paper, we introduce a statistical framework for multivariate planar curves $\mathbf{C}$ with $p \geq 2$, extending classical shape analysis methodologies developed for single contours (Srivastava et al., 2010; Dryden and Mardia, 2016). More precisely, we directly extend the work of Moindjié et al. (2025) to obtain a formal definition of multivariate shape together with associated deformation variables accounting for translation, rotation, scaling, and reparametrization effects.

Building upon this definition, we extend several fundamental tasks of statistical shape analysis identified by Srivastava and Klassen (2016) to multivariate planar curves. In particular, we develop methodologies for: (i) alignment, (ii) quantifying similarities and dissimilarities, (iii) estimating an intrinsic mean and (iv) supervised classification.

In summary, the main contributions of this work are threefold:

- We introduce the notion of a multivariate planar curve and provide a formal definition of multivariate shape.
- We develop a unified framework that extends several fundamental problems of statistical shape analysis from single contours to multivariate planar curves.
- Through an application to the motivating dataset, we demonstrate the practical benefits of jointly modeling multiple interacting contours and show that this approach outperforms a naive extension of single-contour methodologies.

The rest of the paper is organized as follows. Section 2 recalls the univariate model of Moindjié et al. (2025), introduces our multivariate model, and details the estimation of the shape variables, including task (i), the joint alignment of the multiple contours. Section 3 then develops the statistical analysis of multivariate shapes: it equips the shape space with a dissimilarity measure between shapes (task (ii)), summarizes a sample by its representative Fréchet mean (task (iii)), and uses the resulting shape variables as predictors in shape-based classification (task (iv)). Section 4 evaluates the methodology through numerical experiments on synthetic data and an application to the cardiomegaly detection problem, and Section 5 concludes with a discussion of future works.

2 Multivariate planar curves: model and estimation

In this section, we introduce our model for the multivariate random planar curve $\mathbf{C}$, together with the functional space it belongs to, and we describe the estimation of the deformation and shape variables. Since our construction builds upon the univariate case ($p = 1$) of Moindjié et al. (2025), we begin by briefly reviewing it.

2.1 Univariate random planar closed curves

In Moindjié et al. (2025), a univariate random planar closed curve $C = (X, Y)^\top$ is defined as a bivariate function such that $C(t) = (X(t), Y(t))^\top$ for $t \in [0, 1]$, with $C(0) = C(1)$. The functions X and Y represent the coordinates along the first and second axes, respectively, as the curve is traversed from a given starting location. Formally, C is assumed to belong to the Hilbert space $\mathcal{H} = L_2([0, 1]) \times L_2([0, 1])$, equipped with the inner product

$$\langle \mathbf{f}, \mathbf{g} \rangle_{\mathcal{H}} = \int_0^1 f_1(t)g_1(t)dt + \int_0^1 f_2(t)g_2(t)dt, \quad \text{for } \mathbf{f} = (f_1, f_2)^\top, \mathbf{g} = (g_1, g_2)^\top \in \mathcal{H},$$

and we denote the induced norm by $\|\cdot\|_{\mathcal{H}}$.

The closed curve C is modeled as a deformed version of a latent variable $\mathcal{C} \in \mathbf{S}^\infty$ according to

$$C(t) = \rho \mathbf{O}_\theta (\mathcal{C} \circ \gamma_\delta)(t) + \mathbf{T}, \quad t \in [0, 1], \quad (1)$$

where $\mathbf{S}^\infty = \left\{ \mathbf{f} \in \mathcal{H} : \int_0^1 f_1(t)dt = \int_0^1 f_2(t)dt = 0, \|\mathbf{f}\|_{\mathcal{H}} = 1 \right\}$ denotes the centered unit Hilbert sphere. In this model, $\rho \in \mathbb{R}^+$ is a scaling factor, $\mathbf{O}_\theta \in SO(2)$ is a rotation matrix of angle $\theta \in [0, 2\pi]$, and $\mathbf{T} \in \mathbb{R}^2$ is a translation vector.

In practice, the parametrization of a contour depends on the chosen starting point used to traverse the boundary. For instance, in Figure 1(c), the black arrows indicate the starting points used to generate the coordinate functions shown in Figure 1(d). Changing this starting point results in a different parametrization of the same geometric curve. In model (1), this effect is captured through the reparametrization function γ_δ , which shifts the starting point along the contour. Indeed, we assume that γ_δ belongs to

$$\Gamma = \{ \gamma_\delta : [0, 1] \rightarrow [0, 1], \gamma_\delta(t) = \text{mod}(t - \delta, 1), \delta \in [0, 1] \},$$

where $\text{mod}(\cdot, 1)$ is the modulo 1 function. The function γ_δ can therefore be viewed as an extension of the classical time-shift warping function used in functional data registration to the case of planar closed curves, where the modulo operator accounts for the cyclic nature of closed curves. The latent variable $\mathcal{C}$ represents the shape of the curve C , that is, the geometric information that remains once translation, rotation and scaling effects have been removed (Dryden and Mardia, 2016).

2.2 Multivariate random planar closed curves

We now extend the univariate model of Section 2.1 to the case where several contours are observed and analyzed jointly. Such an observation is modeled as a multivariate planar curve $\mathbf{C} = \text{Vec}(C_1, \dots, C_p) = (C_1^\top, \dots, C_p^\top)^\top$, where each component $C_j = (X_j, Y_j)^\top \in \mathcal{H}$ is a univariate planar closed curve. Thus, $\mathbf{C} = (X_1, Y_1, \dots, X_p, Y_p)^\top \in \mathcal{H}^p$, where $\mathcal{H}^p$ denotes the product Hilbert space $\mathcal{H} \times \dots \times \mathcal{H}$ (p times), equipped with the inner product

$$\langle \mathbf{f}, \mathbf{g} \rangle_{\mathcal{H}^p} = \sum_{j=1}^p \langle \mathbf{f}_j, \mathbf{g}_j \rangle_{\mathcal{H}}, \quad \text{for } \mathbf{f} = \begin{pmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_p \end{pmatrix}, \mathbf{g} = \begin{pmatrix} \mathbf{g}_1 \\ \vdots \\ \mathbf{g}_p \end{pmatrix} \in \mathcal{H}^p,$$

with $\mathbf{f}_j = (f_{j1}, f_{j2})^\top$ and $\mathbf{g}_j = (g_{j1}, g_{j2})^\top$. We denote by $\|\cdot\|_{\mathcal{H}^p}$ the norm induced by this inner product.

In order to establish a reasonable model, the key modeling question is to determine which geometric variations should be regarded as deformations caused by the image acquisition process and which should instead be considered meaningful characteristics of the multivariate shape. Since all contours originate from the same image, translation, rotation, and scaling should be viewed as global deformations affecting the image as a whole rather than each contour individually. Conversely, the relative position, orientation, and size of the objects are intrinsic characteristics of the multivariate shape itself and should therefore be encoded in the latent shape variable.

Motivated by this observation, we propose the following model:

$$C_j = \rho \mathbf{O}_\theta \tilde{C}_j \circ \gamma_{\delta_j} + \mathbf{T}, \quad j = 1, \dots, p, \quad (2)$$

where the scaling factor ρ , rotation angle θ , and translation vector $\mathbf{T}$ are global deformations shared across all contours, while each component retains its own reparametrization function γ_{δ_j} . Allowing contour-specific reparametrizations is natural since the starting point used to parameterize each contour depends on the contour extraction procedure and is independent across objects. In contrast, the remaining deformation parameters arise from the acquisition process and therefore affect every contour simultaneously.

Under model (2), the multivariate shape is not defined as a collection of independent shapes, but rather as the geometric structure formed by all contours after removing only the deformations that act globally on the image. As a result, relative variations between contours, such as the size, orientation, and position of one object with respect to another, are encoded directly in the multivariate shape variable $\tilde{\mathbf{C}} = \text{Vec}(\tilde{C}_1, \dots, \tilde{C}_p)$.

For comparison, a straightforward extension of the univariate model (1) would consist in modeling each contour independently by defining

$$C_j = \rho_j \mathbf{O}_{\theta_j} \mathcal{C}_j \circ \gamma_{\delta_j} + \mathbf{T}_j, \quad j = 1, \dots, p, \quad (3)$$

where each contour has its own deformation parameters. Although this formulation is a natural extension of the univariate setting, it implicitly removes the relative position, orientation, and scale of the objects from the multivariate shape by capturing them into contour-specific deformation parameters. Returning to the motivating example of Section 1, such a model would treat the heart and lungs independently. Consequently, the relative size of the heart with respect to the lungs would be treated as a deformation rather than a statistically meaningful characteristic of the multivariate shape. As a result, some important geometric patterns would no longer be encoded in the latent shape variable.

The global shape of $\mathbf{C}$ is therefore represented by the latent multivariate planar curve $\tilde{\mathbf{C}} = \text{Vec}(\tilde{C}_1, \dots, \tilde{C}_p) \in \mathbf{S}_p^\infty$, as defined in (2), where $\mathbf{S}_p^\infty$ denotes the centered unit Hilbert sphere in $\mathcal{H}^p$:

$$\mathbf{S}_p^\infty = \left\{ \mathbf{f} \in \mathcal{H}^p : \sum_{j=1}^p \int_0^1 f_{j1}(t) dt = \sum_{j=1}^p \int_0^1 f_{j2}(t) dt = 0, \|\mathbf{f}\|_{\mathcal{H}^p} = 1 \right\}.$$

Model (2) can be written in the following compact vectorized form:

$$\mathbf{C} = \rho(\mathbf{I}_p \otimes \mathbf{O}_\theta) \tilde{\mathbf{C}} \circ \boldsymbol{\gamma} + (\mathbf{1}_p \otimes \mathbf{T}), \quad (4)$$

where $\otimes$ denotes the Kronecker product, $\mathbf{I}_p$ is the identity matrix of dimension $p \times p$, $\mathbf{1}_p$ is the p -dimensional vector of ones, and $\boldsymbol{\gamma} = (\gamma_{\delta_1}, \dots, \gamma_{\delta_p}) \in \Gamma^p$ is the vector of reparametrization functions, the composition being defined component-wise as $\tilde{\mathbf{C}} \circ \boldsymbol{\gamma} = (\tilde{C}_1 \circ \gamma_{\delta_1}, \dots, \tilde{C}_p \circ \gamma_{\delta_p})^\top$. As in the univariate setting, $(\Gamma^p, \circ)$ is a commutative group and isometry property

$$\|\mathbf{f}\|_{\mathcal{H}^p} = \|(\mathbf{I}_p \otimes \mathbf{O}_\theta) \mathbf{f} \circ \boldsymbol{\gamma}\|_{\mathcal{H}^p}, \quad \text{for } \mathbf{f} \in \mathcal{H}^p \text{ and } \boldsymbol{\gamma} \in \Gamma^p,$$

holds; this result follows directly from Proposition 2.1 in Moindjié et al. (2025).

2.3 Functional estimation

We now describe how to recover the shape $\tilde{\mathbf{C}}$ from a curve $\mathbf{C}$. Inverting the deformation model (2), each component of the shape can be written as

$$\tilde{C}_j = \mathbf{O}_\theta^\top \left(\frac{1}{\rho} (C_j - \mathbf{T}) \right) \circ \gamma_{1-\delta_j} = \mathbf{O}_\theta^\top (C_j^*) \circ \gamma_{1-\delta_j}, \quad j = 1, \dots, p, \quad (5)$$

where $\gamma_{1-\delta_j}$ is the inverse of γ_{δ_j} (Moindjié et al., 2025) and $\mathbf{C}^* = \text{Vec}(C_1^*, \dots, C_p^*) \in \mathbf{S}_p^\infty$ is the *pre-shape* of $\mathbf{C}$, that is, the curve once translation and scaling have been removed. This inversion is purely algebraic, so recovering the shape amounts to estimating the deformation parameters $\mathbf{T}$, ρ , θ and $\boldsymbol{\delta} = (\delta_1, \dots, \delta_p)$.

To carry out this estimation in practice, we follow the standard functional data analysis approach and represent the curve in a finite Fourier basis. Each component admits the representation

$$C_j(t) = \begin{pmatrix} X_j(t) \\ Y_j(t) \end{pmatrix} = \begin{pmatrix} B_{j1} \\ B_{j2} \end{pmatrix} + \sum_{m=1}^M \begin{pmatrix} A_{jm1} \\ A_{jm2} \end{pmatrix} \phi_m(t) = B_j + A_j \boldsymbol{\phi}(t), \quad j = 1, \dots, p, \quad (6)$$

where $B_j \in \mathbb{R}^2$, $A_j \in \mathbb{R}^{2 \times M}$ are coefficient matrices and $\boldsymbol{\phi}(t) = (\phi_1(t), \dots, \phi_M(t))^\top$ contains the first M Fourier basis functions (M even)

$$\phi_k(t) = \begin{cases} \sqrt{2} \sin((k+1)\pi t) & \text{if } k \text{ is odd,} \\ \sqrt{2} \cos(k\pi t) & \text{if } k \text{ is even.} \end{cases}$$

In practice, a curve $\mathbf{C}$ is not observed as a continuous function but only on a finite grid of points, typically dictated by the image resolution. The coefficient matrices B_j and A_j are therefore unknown and are estimated from these observations by smoothing, for instance by minimizing a least-squares criterion between the expansion (6) and the observed values. Stacking these coefficients, we write the multivariate curve compactly as

$$\mathbf{C} = \mathbf{B} + \mathbf{A}\boldsymbol{\phi}, \quad (7)$$

where $\mathbf{B} = \text{Vec}(B_1, \dots, B_p) \in \mathbb{R}^{2p}$ and $\mathbf{A} \in \mathbb{R}^{2p \times M}$ is the matrix formed by vertically concatenating $A_1, \dots, A_p$. Building on this representation, Section 2.3.1 computes the translation and scaling parameters (hence the pre-shape $\mathbf{C}^*$) and Section 2.3.2 estimates the rotation and reparametrization that map $\mathbf{C}^*$ to the shape $\tilde{\mathbf{C}}$.

2.3.1 Translation and scaling parameters

Since $\tilde{\mathbf{C}}$ is centered and has unit norm, the translation and scaling factors are identified directly from model (4), in closed form, as

$$\mathbf{T} = \frac{1}{p} \sum_{j=1}^p \int_0^1 C_j(t) dt = \left(\frac{1}{p} \sum_{j=1}^p \int_0^1 X_j(t) dt, \frac{1}{p} \sum_{j=1}^p \int_0^1 Y_j(t) dt \right)^\top, \quad \rho = \|\mathbf{C} - \mathbf{1}_p \otimes \mathbf{T}\|_{\mathcal{H}^p}. \quad (8)$$

The translation $\mathbf{T}$ is thus the centroid of the multivariate curve (its average position in the image plane) while the scaling ρ is its overall size, measured as the distance from $\mathbf{C}$ to that centroid.

In practice, these two parameters are evaluated through the Fourier representation (7). Since the basis functions are orthonormal and satisfy $\int_0^1 \phi_m(t) dt = 0$, the integral and the norm in (8) reduce to the coefficient expressions

$$\mathbf{T} = \frac{1}{p} \sum_{j=1}^p B_j \quad \text{and} \quad \rho = \sqrt{\|\mathbf{A}\|_F^2 + \|\mathbf{B} - \mathbf{1}_p \otimes \mathbf{T}\|_F^2},$$

where $\|\cdot\|_F$ is the Frobenius norm. Removing these two effects from $\mathbf{C}$ yields the pre-shape $\mathbf{C}^*$ of (5), which in the Fourier basis reads

$$\mathbf{C}^* = \mathbf{B}^* + \mathbf{A}^* \phi, \quad \text{with} \quad \mathbf{A}^* = \frac{1}{\rho} \mathbf{A} \quad \text{and} \quad \mathbf{B}^* = \frac{1}{\rho} (\mathbf{B} - \mathbf{1}_p \otimes \mathbf{T}).$$

The pre-shape $\mathbf{C}^*$ lies in $\mathbf{S}_p^\infty$ but still carries the rotation and reparametrization effects, which are estimated in the next section.

2.3.2 Multivariate contour alignment

Recovering the shape from the pre-shape is an alignment problem: we estimate the rotation and reparametrization that best align a pre-shape $\mathbf{C}^*$ to a template $\bar{\mathbf{C}} = \text{Vec}(\bar{C}_1, \dots, \bar{C}_p) \in \mathbf{S}_p^\infty$. The choice of this template is discussed in Section 3; here it is taken as given. Aligning $\mathbf{C}^*$ to $\bar{\mathbf{C}}$ amounts to estimating the rotation parameter θ and the reparametrization parameter δ that solve

$$\min_{\delta \in [0,1]^p, \theta \in [0,2\pi]} \sum_{j=1}^p \left\| \mathbf{O}_\theta \bar{C}_j \circ \gamma_{\delta_j} - C_j^* \right\|_{\mathcal{H}}^2. \quad (9)$$

For a single contour ($p = 1$), Moindjié et al. (2025) solve (9) with the *Iterative Closest Function* (ICF) algorithm, an alternating optimization procedure, inspired by the iterative closest point (ICP) algorithm (Zhang, 2021), that operates directly in the functional space without discretizing the curves. Since the dimension of the search space in (9) grows with the number of contours p , optimizing jointly over $(\theta, \delta_1, \dots, \delta_p)$ is impractical. We therefore extend the ICF algorithm to the multivariate setting: starting from an initial value of δ , it alternates the two steps below until convergence. Both are the direct multivariate counterparts of the steps of Moindjié et al. (2025); the only structural difference is that the shared rotation θ is now estimated jointly from all p components.

- (i) **Estimation of θ for a fixed δ .** The shared rotation angle solves the Procrustes problem

$$\hat{\theta} = \arg \min_{\theta \in [0,2\pi]} \sum_{j=1}^p \left\| \mathbf{O}_\theta \bar{C}_j \circ \gamma_{\delta_j} - C_j^* \right\|_{\mathcal{H}}^2.$$

As in the univariate case (Moindjié et al., 2025), the minimizer belongs to $\{\theta_1, \theta_1 + \pi\}$, where

$$\tan(\theta_1) = \frac{\sum_{j=1}^p \left\{ \langle Y_j^*, \bar{X}_j \circ \gamma_{\delta_j} \rangle_{L_2} - \langle X_j^*, \bar{Y}_j \circ \gamma_{\delta_j} \rangle_{L_2} \right\}}{\sum_{j=1}^p \left\{ \langle X_j^*, \bar{X}_j \circ \gamma_{\delta_j} \rangle_{L_2} + \langle Y_j^*, \bar{Y}_j \circ \gamma_{\delta_j} \rangle_{L_2} \right\}},$$

with $\bar{C}_j = (\bar{X}_j, \bar{Y}_j)^\top$, $C_j^* = (X_j^*, Y_j^*)^\top$, and $\langle \cdot, \cdot \rangle_{L_2}$ the inner product of $L_2([0, 1])$. The two candidates correspond to opposite orientations; we evaluate the objective at both and keep the one with the smaller value.

- (ii) **Estimation of δ for a fixed θ .** For a fixed rotation, the objective in (9) is a sum of terms each depending on a single δ_j . The reparametrization parameters therefore decouple across components, and each is obtained by solving

$$\hat{\delta}_j = \arg \min_{\delta_j \in [0, 1]} \|\mathbf{O}_\theta \bar{C}_j \circ \gamma_{\delta_j} - C_j^*\|_{\mathcal{H}}^2, \quad j = 1, \dots, p. \quad (10)$$

Each subproblem (10) is exactly the univariate alignment problem of Moindjié et al. (2025); as there, it has no closed-form solution and each $\hat{\delta}_j$ is computed numerically by a bisection method.

Remark. *As in the univariate case, the multivariate ICF algorithm is sensitive to the initialization of δ . In practice, we run it from several initial values and retain the pair $(\hat{\theta}, \hat{\delta})$ that minimizes the objective in (9).*

Once $(\hat{\theta}, \hat{\delta})$ have been obtained, the shape $\tilde{\mathbf{C}}$ is recovered through the inverse transformations (5).

3 Statistical analysis of multivariate shapes

This section develops the statistical analysis of multivariate shapes and their use as predictors in a classification task. Throughout, we consider a sample of n multivariate planar curves $\mathbf{C}_1, \dots, \mathbf{C}_n$, for instance the n segmented images of a dataset, from which the pre-shapes $\mathbf{C}_1^*, \dots, \mathbf{C}_n^*$ are obtained as described in Section 2.3.

By construction, each pre-shape $\mathbf{C}_i^*$ lies on the sphere $\mathbf{S}_p^\infty$, but still differs from the shape $\tilde{\mathbf{C}}_i$ of $\mathbf{C}_i$ by a rotation and a reparametrization. The shape $\tilde{\mathbf{C}}_i$ is what remains once these last two deformations are removed as well; equivalently, it is identified with the equivalence class of $\mathbf{C}_i^*$ under the relation

$$\mathbf{f} \sim \mathbf{g} \iff \exists (\theta, \gamma) \in [0, 2\pi] \times \Gamma^p \text{ such that } \mathbf{f} = (\mathbf{I}_p \otimes \mathbf{O}_\theta) \mathbf{g} \circ \gamma, \quad \mathbf{f}, \mathbf{g} \in \mathbf{S}_p^\infty,$$

which holds when $\mathbf{f}$ and $\mathbf{g}$ coincide up to a rotation and a reparametrization. The shapes thus belong to the quotient *shape space* $\mathbf{S}_p^\infty / \sim$, while $\mathbf{S}_p^\infty$ itself is called the *pre-shape space*, as it contains each shape together with all its rotated and reparametrized versions. This shape space is nonlinear, so classical statistical tools, which rely on a linear structure, cannot be applied to the shapes directly.

We address this difficulty through the standard strategy of statistical shape analysis (see e.g. Dryden and Mardia (2016), Srivastava et al. (2005)), which proceeds in three steps. First, we summarize the sample by an intrinsic mean shape, the Fréchet mean $\boldsymbol{\mu}$ (Section 3.1). Second, we linearize the shape space in the neighborhood of $\boldsymbol{\mu}$ by projecting the shapes onto the tangent space at $\boldsymbol{\mu}$ (Section 3.2). Finally, we perform the statistical analysis, here the shape-based classification, in this linear space (Section 3.3).

3.1 The Fréchet mean of multivariate shapes

To compare two shapes $\tilde{\mathbf{C}}_1$ and $\tilde{\mathbf{C}}_2$, that is, the equivalence classes of their pre-shapes $\mathbf{C}_1^*$ and $\mathbf{C}_2^*$ under the relation introduced above, the distance $d_{\mathcal{H}^p}$ (the distance induced by $\|\cdot\|_{\mathcal{H}^p}$) is minimized over the rotation and reparametrization that align one pre-shape to the other:

$$d_{\mathcal{S}}(\tilde{\mathbf{C}}_1, \tilde{\mathbf{C}}_2) = \min_{\theta \in [0, 2\pi], \gamma \in \Gamma^p} d_{\mathcal{H}^p}(\mathbf{C}_1^*, (\mathbf{I}_p \otimes \mathbf{O}_\theta) \mathbf{C}_2^* \circ \gamma). \quad (11)$$

The minimization in (11) is the alignment problem (9) applied to two pre-shapes (the distance minimized in (9) and the distance $d_{\mathcal{H}^p}$ in (11) share the same minimizer on $\mathbf{S}_p^\infty$), so the multivariate ICF algorithm of Section 2.3.2 computes $d_{\mathcal{S}}$. By construction, $d_{\mathcal{S}}$ is a dissimilarity measure between multivariate shapes that is invariant to translation, scaling, rotation and reparametrization: it is nonnegative and symmetric, and vanishes if and only if the two curves share the same shape.

Following the classical approach in statistical shape analysis (Dryden and Mardia, 2016), we summarize the shape variable $\tilde{\mathbf{C}}$ by its Fréchet mean

$$\boldsymbol{\mu} = \arg \min_{\mathbf{f} \in \mathbf{S}_p^\infty} \mathbb{E} \left[d_{\mathcal{H}^p}^2(\tilde{\mathbf{C}}, \mathbf{f}) \right] \in \mathbf{S}_p^\infty, \quad (12)$$

which is the shape that is, on average, closest to the random shape $\tilde{\mathbf{C}}$, and thus serves as a natural representative of the population. It also settles the template choice left open in Section 2.3.2: as is standard in statistical shape analysis (Dryden and Mardia, 2016), each pre-shape is aligned to the mean shape itself, which places the whole sample in a common frame.

The Fréchet mean (12) has no closed-form expression and must be estimated from the sample. This estimation is itself circular, since the template above is the very mean we seek: $\boldsymbol{\mu}$ is needed to align the shapes, while the aligned shapes are needed to compute $\boldsymbol{\mu}$. Following Moindjié et al. (2025), we break this dependency with an iterative procedure that, starting from an initial template, repeats the two steps

- (a) *alignment*: for $i = 1, \dots, n$, align the pre-shape $\mathbf{C}_i^*$ to the current template with the ICF algorithm of Section 2.3.2, producing the shapes $\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_n$;
- (b) *update*: replace the template by the empirical Fréchet mean of $\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_n$, that is, the sample counterpart of (12),

until convergence. The procedure then returns the estimated Fréchet mean $\hat{\boldsymbol{\mu}}$ together with the aligned shapes $\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_n$. Both the empirical Fréchet mean of step (b) and the overall procedure are the multivariate counterparts of those in Moindjié et al. (2025), to which we refer for the numerical details.

3.2 The tangent space

The classical tangent projection (Dryden and Mardia (2016), Srivastava and Klassen (2016), Dai (2022)) can be extended in our case. We recall that the tangent space to $\mathbf{S}_p^\infty$ at $\boldsymbol{\mu}$ is the linear subspace

$$\mathcal{T}_{\boldsymbol{\mu}} = \{ \mathbf{f} \in \mathcal{H}^p : \langle \mathbf{f}, \boldsymbol{\mu} \rangle_{\mathcal{H}^p} = 0 \}.$$

Hence, equipped with the Fréchet mean, we can linearize the shape space in its neighborhood. Indeed, unlike $\mathbf{S}_p^\infty$, $\mathcal{T}_\mu$ is closed under linear combinations and therefore amenable to classical (functional) linear methods. Shapes are mapped to $\mathcal{T}_\mu$, and back, through the logarithm and exponential maps, whose well-established properties make them a standard choice in this context (see e.g. Dryden and Mardia (2016), Dai (2022)). The logarithm map $\mathcal{L}_\mu : \mathbf{S}_p^\infty \setminus \{-\mu\} \rightarrow \mathcal{T}_\mu$ projects a shape onto the tangent space,

$$\mathcal{L}_\mu(\mathbf{f}) = \frac{\omega}{\sin(\omega)} (\mathbf{f} - \cos(\omega) \mu), \quad \omega = d_{\mathbf{S}_p^\infty}(\mathbf{f}, \mu), \quad (13)$$

whereas the exponential map $\text{Exp}_\mu : \mathcal{T}_\mu \rightarrow \mathbf{S}_p^\infty$ maps a tangent vector back onto the sphere,

$$\text{Exp}_\mu(\mathbf{v}) = \cos(\|\mathbf{v}\|_{\mathcal{H}^p}) \mu + \sin(\|\mathbf{v}\|_{\mathcal{H}^p}) \frac{\mathbf{v}}{\|\mathbf{v}\|_{\mathcal{H}^p}}, \quad \mathbf{v} \in \mathcal{T}_\mu. \quad (14)$$

The two maps are mutual inverses in a neighborhood of μ , so that $\text{Exp}_\mu(\mathcal{L}_\mu(\mathbf{f})) = \mathbf{f}$: one can therefore carry out the statistical analysis on the tangent coordinates $\mathcal{L}_\mu(\mathbf{f})$ and, when needed, map the results back onto the sphere through Exp_μ .

In practice, each sample shape is represented by its tangent coordinate $\mathcal{L}_\mu(\hat{\mathbf{C}}_i)$, $i = 1, \dots, n$. This linearization is accurate as long as the shapes lie close to the reference point: choosing μ as the Fréchet mean, which minimizes the average squared distance (12), keeps the shapes as close as possible to μ and thus makes the projection as faithful as possible.

3.3 Classification of multivariate shapes

Because the tangent space $\mathcal{T}_\mu$ is linear, the projected shapes can serve as predictors in standard functional classification models, as is common in statistical shape analysis (Dryden and Mardia, 2016, Chap. 13). Formally, let Y be a binary response associated with a multivariate planar curve $\mathbf{C}$. Using the tangent coordinate $\mathcal{L}_\mu(\tilde{\mathbf{C}})$ as a predictor casts the problem into the well-studied framework of functional classification with multivariate functional predictors (see e.g. Godwin (2013), Moindjié et al. (2024), Febrero-Bande et al. (2017)). A typical instance is the functional linear model

$$Y = \psi\left(\beta_0 + \langle \mathcal{L}_\mu(\tilde{\mathbf{C}}), \boldsymbol{\beta} \rangle_{\mathcal{H}^p}\right), \quad (15)$$

where $\beta_0 \in \mathbb{R}$ and $\boldsymbol{\beta} \in \mathcal{H}^p$ are estimated from the data and $\psi : \mathbb{R} \rightarrow \{0, 1\}$ is a link function determined by the classification method, such as linear discriminant analysis or logistic regression. The specific methods used in our experiments are presented in Section 4.2.

4 Numerical experiments

In this section, we present numerical experiments designed to assess the performance of the proposed methodology. Our primary objective is to evaluate its effectiveness in a classification setting. To this end, we apply it to the real-world task of detecting cardiomegaly from chest X-ray images. Before considering this application, we first study the accuracy of the proposed alignment procedure on a synthetic dataset. This preliminary analysis is important because classification of shapes strongly depends on the accurate estimation of the deformation parameters, in particular the rotation and reparametrization components.

4.1 Performance study of the alignment procedure

We conduct a simulation study to numerically assess the accuracy of the deformation estimation described in Section 2.3. Since translation and scaling parameters are known to be accurately estimated in practice (Moindjié et al., 2025), this performance study focuses exclusively on the estimation of the rotation and reparametrization parameters. In other words, we evaluate the ability of our proposed alignment procedure to recover the underlying shapes from their corresponding pre-shapes.

To this end, we construct a synthetic dataset by deforming a single observation from the CheXmask dataset (Gaggion et al., 2024), which provides multivariate planar curves ($p = 3$) corresponding to the contours of the heart and both lungs. We randomly select one observation, smooth it using $M = 22$ Fourier basis functions, and center it. This produces a reference curve, denoted $\mathbf{c}^0$, which serves as the template for generating the synthetic data. The curve $\mathbf{c}^0$ can be written as

$$\mathbf{c}^0 = \text{Vec}(c_1^0, c_2^0, c_3^0), \quad \text{with} \quad c_j^0 = b_j^0 + a_j^0 \phi, \quad j = 1, 2, 3,$$

where $b_j^0 \in \mathbb{R}^2$ and $a_j^0 \in \mathbb{R}^{2 \times 22}$ are coefficient matrices. The resulting function $\mathbf{c}^0$ is illustrated in Figure 2. In particular, the red arrows indicate the starting point ($t = 0$) used to traverse each contour.

We then generate pre-shapes $\mathbf{c}_1^*, \dots, \mathbf{c}_n^*$ by applying random rotation and reparametrization deformations to perturbed versions of $\mathbf{c}^0$. Specifically, we define $\mathbf{c}_i^* = \text{Vec}(\mathbf{c}_{i1}^*, \mathbf{c}_{i2}^*, \mathbf{c}_{i3}^*)$ with:

$$\mathbf{c}_{ij}^* = \kappa_i \mathbf{O}_{\theta_i} (b_j^0 + \tilde{a}_{ij} \phi \circ \gamma_{\delta_{ij}}), \quad j = 1, 2, 3, \quad (16)$$

where $\theta_i \sim \mathcal{U}([0, 2\pi])$, $\delta_{ij} \sim \mathcal{U}([0, 1])$, $\text{Vec}(\tilde{a}_{ij}) \sim \mathcal{N}(\text{Vec}(a_j^0), \sigma^2 \mathbf{I}_{44})$, and where $\kappa_i = \left[\sum_{j=1}^3 (\|b_j^0\|^2 + \|\tilde{a}_{ij}\|^2) \right]^{-1/2}$ ensures normalization. In this construction, the function $\mathbf{c}^0 / \|\mathbf{c}^0\|_{\mathcal{H}^p}$ corresponds to the Fréchet mean of the simulated dataset.




Figure 2: The multivariate ($p = 3$) curve $\mathbf{c}^0$ used as the template for generating the synthetic data. Each closed curve represents the contour of an organ, and the red arrow indicates the starting point ($t = 0$) used to traverse the contour and define the coordinate functions.

To assess the robustness of the alignment procedures under various realistic conditions, we generate $n = 500$ curves under different levels of the perturbation parameter $\sigma \in$

$\{0.1, 0.5, 1.0\}$. Increasing σ produces pre-shapes that deviate progressively from the template $\mathbf{c}^0$, as illustrated in Figure 3.




Figure 3: Examples of three pre-shapes generated by applying random rotations and reparametrizations to perturbed versions of the template $\mathbf{c}^0$ under three levels of the perturbation parameter σ . The rotations change the global orientation of the curves, while the reparametrizations shift the starting points along the contours, illustrated by the red arrows indicating the starting point ($t = 0$) used to traverse the curves.

For each of the three datasets, we estimate the deformation parameters θ_i and $\boldsymbol{\delta}_i = (\delta_{i1}, \delta_{i2}, \delta_{i3})$ associated with the pre-shape $\mathbf{c}_i^*$ by applying the multivariate ICF alignment, described in Section 2.3.2, with the template $\bar{\mathbf{C}}$ set to $\boldsymbol{\mu} = \mathbf{c}^0 / \|\mathbf{c}^0\|_{\mathcal{H}^p}$, and we use $M = 22$ for the Fourier representation. To improve robustness, we run the ICF algorithm from 5 random initializations of $\boldsymbol{\delta}_i$ and retain the best one.

To quantitatively assess the accuracy of the estimated deformation parameters, while accounting for the periodic nature of rotations and reparametrizations, we compute the *cyclic mean squared error* (cMSE). For the rotation, it is defined as

$$\text{cMSE}_\theta = \frac{1}{500} \sum_{i=1}^{500} \left\| \begin{pmatrix} \cos(\theta_i) \\ \sin(\theta_i) \end{pmatrix} - \begin{pmatrix} \cos(\hat{\theta}_i) \\ \sin(\hat{\theta}_i) \end{pmatrix} \right\|_2^2,$$

while, separately for each component $j = 1, 2, 3$, the reparametrization cMSE is

$$\text{cMSE}_{\delta_j} = \frac{1}{500} \sum_{i=1}^{500} \left\| \begin{pmatrix} \cos(2\pi\delta_{ij}) \\ \sin(2\pi\delta_{ij}) \end{pmatrix} - \begin{pmatrix} \cos(2\pi\hat{\delta}_{ij}) \\ \sin(2\pi\hat{\delta}_{ij}) \end{pmatrix} \right\|_2^2.$$

The results are summarized in Table 1. The low cMSE values across all values of σ demonstrate the high accuracy of our alignment procedure, even when the deformation noise is substantial. This supports the reliability of the ICF method for estimating deformation parameters.

σ	cMSE $_{\theta}$	cMSE $_{\delta_1}$	cMSE $_{\delta_2}$	cMSE $_{\delta_3}$
0.1	8.40e-07	3.62e-04	3.40e-04	3.45e-04
0.5	1.78e-05	5.25e-04	4.17e-04	3.98e-04
1.0	6.55e-05	1.31e-03	6.14e-04	5.69e-04

Table 1: Cyclic mean squared error (cMSE) of the estimated rotation and reparametrization parameters, for different levels of the perturbation parameter σ .

Figure 4 further illustrates this result by displaying a subsample of simulated pre-shapes for $\sigma = 1$ (left panel) along with their corresponding aligned shapes (right panel). A visual inspection confirms the successful recovery of the aligned shapes.

4.2 Classification of multivariate shapes: a cardiomegaly detection problem

We now illustrate the practical relevance of the proposed methodology by applying it to a real-world classification problem: detecting cardiomegaly from chest X-ray images. Cardiomegaly refers to an abnormal enlargement of the heart, which manifests on radiographic images as the heart occupying more than half of the thoracic width. This condition is clinically important, as it may indicate underlying heart diseases such as heart failure or arrhythmia.

Our analysis is based on the CheXmask dataset (Gaggion et al., 2024), which contains segmented masks derived from several large-scale chest X-ray databases. We focus on the ChestX-ray8 subset (Wang et al., 2017), which provides pathology labels that allow supervised learning.

Each segmented image is represented as a multivariate planar curve with $p = 3$ components corresponding to the contours of the heart and the left and right lungs. This representation naturally fits the framework developed in this paper. Figure 5 illustrates examples of contours extracted from chest X-ray images of healthy patients and patients diagnosed with cardiomegaly. Although cardiomegaly is characterized by an enlargement of the heart relative to the lungs, the visual distinction between the two classes can be subtle. As shown in the figure, there is substantial variability in both the shapes of the organs and their relative deformations across individuals, which makes automated classification a challenging task.




Figure 4: Left panel: a subsample of simulated pre-shapes for $\sigma = 1$. Right panel: aligned shapes obtained with the ICF method. The three components of each multivariate curve are the heart (H), the left lung (LL) and the right lung (RL).




Figure 5: Examples of multivariate planar curves extracted from chest X-ray images: (a) healthy patients and (b) patients diagnosed with cardiomegaly. Each panel shows the contours of the heart and both lungs for three representative observations. The red arrows indicate the starting point $t = 0$.

From the original dataset, we extract a balanced subset of $n = 300$ images, consisting of the 150 X-rays labeled as “cardiomegaly” and 150 randomly selected healthy X-rays labeled as “no finding”. The resulting dataset, denoted $\mathbf{c}_1, \dots, \mathbf{c}_n$, is illustrated in Figure 6 and can be found in our online code repository (<https://github.com/imoindjie/Shape-MFDA>).

It is important to note that the curves provided in Gaggion et al. (2024) have already been preprocessed to reduce translation, rotation, and scale variability using standard image registration techniques. However, visual inspection of Figure 6 reveals that residual misalignments remain, particularly in the location of the starting points used to traverse the contours.




Figure 6: The considered subsample of 300 multivariate planar curves $\mathbf{c}_i$ from the chest X-ray dataset (Gaggion et al., 2024). Different observations are identified using different colors and the red arrows represent the starting point ($t = 0$) for each coordinate function.

To study the robustness of our classification pipeline, we consider two experimental scenarios:

Scenario 1 - Pre-aligned data We directly use the pre-aligned curves $\mathbf{c}_i, i = 1, \dots, n$, as provided in the dataset.

Scenario 2 - Misaligned data To assess the benefit of alignment in a controlled setting, we introduce artificial deformations to each observation. We set $\mathbf{c}_i, i = 1, \dots, n$, as

$$\mathbf{c}_i \leftarrow (\mathbf{I}_3 \otimes \mathbf{O}_{2\pi\zeta_i})\mathbf{c}_i \circ \gamma_{\delta_i}, \text{ with } \delta_i = (\delta_{i1}, \delta_{i2}, \delta_{i3}) \in [0, 1]^3, \quad (17)$$

where $\zeta_i \sim \mathcal{U}([0, 1])$ and each reparametrization component $\delta_{ij} \sim \mathcal{U}([0, 1])$, for $j = 1, 2, 3$. This setting allows us to evaluate how well the alignment methods handle misaligned input data.

4.2.1 Classification models

As introduced in Section 3.3, the classification framework uses a predictor of the form $\mathcal{L}_\mu(\tilde{\mathbf{C}})$, which, we recall, corresponds to the projection of the multivariate shape $\tilde{\mathbf{C}}$ on the tangent space $\mathcal{T}_\mu$ where μ is the Fréchet mean of $\tilde{\mathbf{C}}$.

To assess the benefits of our proposed multivariate approach, we compare three ways of defining the predictors:

- *OUR* - our proposed approach: we use $\mathcal{L}_{\hat{\mu}}(\hat{\mathbf{c}}_i)$, $i = 1, \dots, n$, as predictors, where the shapes and Fréchet mean are estimated using the algorithm described in Section 3.1. This method jointly models the three anatomical components.
- *UNI* - univariate approach: each anatomical component is analyzed separately, as in equation (3), and the shapes are estimated following the univariate procedure of Moindjié et al. (2025). We use $(\mathcal{L}_{\hat{\mu}_1}(\hat{\mathcal{C}}_{i1}), \mathcal{L}_{\hat{\mu}_2}(\hat{\mathcal{C}}_{i2}), \mathcal{L}_{\hat{\mu}_3}(\hat{\mathcal{C}}_{i3}))^\top$, $i = 1, \dots, n$, as predictors, where $\mathcal{L}$ is defined for $p = 1$ and $\hat{\mu}_j$ is the estimated empirical Fréchet mean of $\mathcal{C}_{1j}, \dots, \mathcal{C}_{nj}$, for $j = 1, 2, 3$. This approach ignores the global multivariate structure and treats each organ independently.
- *CLA* - classical naïve approach: we use the raw observed curves $\mathbf{c}_i$ as predictors, without any alignment.

To ensure that performance differences are not tied to a specific classifier, we evaluate these three predictor sets with four different functional linear classification methods, corresponding to different choices of the link function ψ in (15).

Two methods are based on penalized group logistic regression. Following the group lasso framework (Meier et al., 2008) (see also Godwin (2013) for its functional extension), functional predictors are partitioned into predefined groups before applying the penalty:

GL1: Three groups, each corresponding to one anatomical component $c_j = (x_j, y_j)^\top$, $j = 1, 2, 3$. The penalty can identify whether the heart or one of the lungs has no discriminative power in the model.

GL2: Six groups, each corresponding to a single coordinate function x_j or y_j , for $j = 1, 2, 3$. This allows for finer selection, identifying specific coordinates that are non-informative.

The penalty parameter λ is selected, independently for the two methods, as the value that minimizes the classification error over the set $\{0.96^l \lambda_{\max}, l = 0, 1, \dots, 148\} \cup \{0\}$, where $\lambda_{\max}$ is the smallest value for which the penalty term vanishes.

The two other considered methods rely on linear discriminant analysis where the discriminant coefficient function is estimated via partial least squares (PLS) and principal component regression (PCR) (Preda et al. (2007), Moindjié et al. (2024)). For both, the number of components is chosen from the set $\{1, 2, \dots, 6 \times 22 - 1\}$ to minimize classification error.

4.2.2 Results

We evaluate the performance of each classification strategy using a 10-fold cross-validation procedure. The results for both scenarios are reported in Table 2.

In Scenario 1 (aligned curves), the results show that although the four classification procedures (GL1, GL2, PLS, PCR) lead to varying accuracy levels, the *OUR* and *CLA* approaches achieve almost identical performance across all methods. This confirms that, when

the multivariate curves are already aligned, the advantage of a shape-based preprocessing step is minimal. The *UNI* approach performs slightly worse, likely because it treats the heart and each lung as independent univariate curves, so that their relative spatial relationships are ignored. Nevertheless, its accuracy remains relatively high, underlining the strong discriminative power of shape information alone.

In Scenario 2 (misaligned curves), the situation changes markedly. The classical *CLA* accuracies collapse to 53–57%, close to random guessing, which highlights the detrimental impact of ignoring alignment in multivariate shape classification. The two shape-based methods, by contrast, remain as accurate as in Scenario 1, confirming their robustness to rotation and reparametrization. Among them, our proposed multivariate approach *OUR* is the best for every classifier and outperforms the univariate *UNI* by 3 to 11 accuracy points. This gap reflects the core benefit of our framework: modeling the three contours jointly preserves the inter-component information, relative position, scale and orientation, that a separate, univariate treatment discards, yet that is discriminative for cardiomegaly.

Overall, these results emphasize two main points: (i) when data are aligned, our approach can match the performance of classical methods, and (ii) in settings with misaligned data, shape-based methods, particularly our proposed multivariate framework, are more reliable.

Scenario	Method	GL1	GL2	PLS	PCR
Scenario 1	<i>OUR</i>	79.65	77.34	85.22	82.94
	<i>UNI</i>	71.73	71.94	78.83	73.36
	<i>CLA</i>	80.00	78.27	85.25	80.60
Scenario 2	<i>OUR</i>	80.50	78.39	85.99	82.77
	<i>UNI</i>	77.37	74.11	77.92	71.67
	<i>CLA</i>	57.11	54.94	55.15	53.13

Table 2: Classification accuracies (%) obtained by 10-fold cross-validation, for the three predictor sets (*OUR*, *UNI*, *CLA*) and the four classifiers (GL1, GL2, PLS, PCR), under the two scenarios.

5 Conclusion

This work extends statistical shape analysis to multiple objects by adopting a functional data analysis perspective. As in the univariate case, the observed multivariate planar curve is viewed as a deformed version of an unobserved latent variable of interest: its shape. The novelty of our work lies in treating the multiple contours as a single multivariate planar curve, which allows several objects to be studied jointly while preserving the geometric information they share. After defining a suitable and intuitive model in this setting, we introduce estimation procedures for the deformation and shape variables. In particular, we extend the Iterative Closest Function (ICF) procedure of Moindjié et al. (2025) to jointly estimate the rotation and reparametrization effects of the multiple contours, from which the shapes are recovered through inverse transformations. We then describe how to summarize and analyze these shapes, through their Fréchet mean and a projection onto the associated tangent space, so that they can serve as predictors in functional classification models.

The numerical experiments demonstrate the proficiency of our method in estimating the deformation parameters and recovering the shapes across diverse scenarios. In the cardiomegaly detection problem, they further show that classifying multivariate shapes improves accuracy over both a univariate shape analysis, which treats each contour separately, and a classical approach that feeds the raw, unaligned curves to the same classifiers, particularly in the setting where images are not pre-aligned. This advantage stems from modeling the contours jointly, which preserves the inter-component information that is discriminative for the task. The proposed representation for multivariate planar curves is moreover parsimonious, and we believe that the low computational cost and transparency of the proposed shape-based methods, compared to pixel-based approaches, make them appealing alternatives for image analysis when objects are well identified.

Indeed, the availability of segmented images makes our approach readily applicable; however, the quality of current image segmentation models remains a limitation for analyzing shapes in more complex images. As image segmentation and contour detection constitute active research domains, future work should focus on the challenge of extracting multivariate planar curves directly from colored images. For this aim, we could rely on the rich literature within the computer vision community (see e.g., Kass et al. (1988), Ibrahim and El-kenawy (2020)).

Moreover, the underlying idea of this work, which defines an image solely through the objects’ contours, can be limiting in many application cases. Indeed, in addition to shapes, colors in images are strongly informative in the classification task (see e.g., (Gowda and Yuan, 2018), Funt and Zhu (2018)). Hence, investigating a fully functional approach to account for these variations is of interest. An intuitive idea is to shift the perspective from shapes to surfaces, where pixels represent a discrete sample over such surfaces. Future work should focus on keeping the procedure interpretable and the image representation as parsimonious as possible, to avoid requiring more computing resources than pixel-based methods. The resulting methods could then be used to fairly assess the advantages of using pixels versus shapes and colors as inputs in prediction problems.

References

- Dai, X. (2022). Statistical inference on the hilbert sphere with application to random densities. Electronic Journal of Statistics, 16(1):700–736.
- Dryden, I. L. and Mardia, K. V. (2016). Statistical shape analysis: with applications in R. John Wiley & Sons.
- Febrero-Bande, M., Galeano, P., and González-Manteiga, W. (2017). Functional principal component regression and functional partial least-squares regression: An overview and a comparative study. International Statistical Review, 85(1):61–83.
- Fu, K.-S. and Mui, J. (1981). A survey on image segmentation. Pattern recognition, 13(1):3–16.
- Funt, B. and Zhu, L. (2018). Does colour really matter? evaluation via object classification. In Color and Imaging Conference, volume 26, pages 268–271. Society for Imaging Science and Technology.
- Gaggion, N., Mosquera, C., Mansilla, L., Saidman, J. M., Aineseder, M., Milone, D. H., and Ferrante, E. (2024). Chexmask: a large-scale dataset of anatomical segmentation masks for multi-center chest x-ray images. Scientific Data, 11(1):511.
- Godwin, J. (2013). Group lasso for functional logistic regression. Master’s thesis, Auburn University.
- Gowda, S. N. and Yuan, C. (2018). Colornet: Investigating the importance of color spaces for image classification. In Asian conference on computer vision, pages 581–596. Springer.
- Gu, L., Wu, M., and Kwon, T. J. (2020). An enhanced spatial statistical method for continuous monitoring of winter road surface conditions. Canadian Journal of Civil Engineering, 47(10):1154–1165.
- Ibrahim, A. and El-kenawy, E.-S. M. (2020). Image segmentation methods based on superpixel techniques: A survey. Journal of Computer Science and Information Systems, 15(3):1–11.
- Kass, M., Witkin, A., and Terzopoulos, D. (1988). Snakes: Active contour models. International journal of computer vision, 1(4):321–331.
- Lu, D., Heisler, M., Lee, S., Ding, G. W., Navajas, E., Sarunic, M. V., and Beg, M. F. (2019). Deep-learning based multiclass retinal fluid segmentation and detection in optical coherence tomography images using a fully convolutional neural network. Medical image analysis, 54:100–110.
- Meier, L., Van De Geer, S., and Bühlmann, P. (2008). The group lasso for logistic regression. Journal of the Royal Statistical Society Series B: Statistical Methodology, 70(1):53–71.

- Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., and Terzopoulos, D. (2021). Image segmentation using deep learning: A survey. IEEE transactions on pattern analysis and machine intelligence, 44(7):3523–3542.
- Moindjié, I.-A., Beaulac, C., and Descary, M.-H. (2025). A functional approach to curve alignment and shape analysis. arXiv preprint arXiv:2503.05632.
- Moindjié, I.-A., Dabo-Niang, S., and Preda, C. (2024). Classification of multivariate functional data on different domains with partial least squares approaches. Statistics and Computing, 34(1):5.
- Preda, C., Saporta, G., and Lévêder, C. (2007). Pls classification of functional data. Computational Statistics, 22(2):223–235.
- Srivastava, A., Joshi, S. H., Mio, W., and Liu, X. (2005). Statistical shape analysis: Clustering, learning, and testing. IEEE Transactions on pattern analysis and machine intelligence, 27(4):590–602.
- Srivastava, A., Klassen, E., Joshi, S. H., and Jermyn, I. H. (2010). Shape analysis of elastic curves in euclidean spaces. IEEE transactions on pattern analysis and machine intelligence, 33(7):1415–1428.
- Srivastava, A. and Klassen, E. P. (2016). Functional and shape data analysis, volume 1. Springer.
- Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., and Nandi, A. K. (2022). Medical image segmentation using deep learning: A survey. IET image processing, 16(5):1243–1267.
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., and Summers, R. M. (2017). Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2097–2106.
- Younes, L. (1998). Computable elastic distances between shapes. SIAM Journal on Applied Mathematics, 58(2):565–586.
- Zhang, Z. (2021). Iterative closest point (icp). In Computer vision: a reference guide, pages 718–720. Springer.