

Statistical analysis of multivariate planar curves and applications to X-ray classification

Moindjié, Issam-Ali^{1,*}, Descary, Marie-Hélène¹, and Beaulac, Cédric¹

¹Department of mathematics, Université du Québec à Montréal

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

Abstract

Recent developments in computer vision have enabled the availability of segmented images across various domains, such as medicine, where segmented radiography images play an important role in diagnosis-making. As prediction problems are common in medical image analysis, this work explores the use of segmented images (through the associated contours they highlight) as predictors in a supervised classification context. Consequently, we develop a new approach for image analysis that takes into account the shape of objects within images. For this aim, we introduce a new formalism that extends the study of single random planar curves to the joint analysis of multiple planar curves — referred to here as multivariate planar curves. In this framework, we propose a solution to the alignment issue in statistical shape analysis. The obtained multivariate shape variables are then used in functional classification methods through tangent projections. Detection of cardiomegaly in segmented X-rays and numerical experiments on synthetic data demonstrate the appeal and robustness of the proposed method.

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

1 Introduction

Image segmentation has long been a central research area in pattern recognition and image analysis, aiming to divide images into homogeneous regions that share common characteristics such as texture, intensity, or gray level. Since the development of computer graphics, several methods have been investigated (see Fu and Mui (1981) for a survey). A recent breakthrough has been the development of deep learning approaches for image segmentation (see, e.g., Minaee et al. (2021)). These methods demonstrate precise recovery properties in practical cases, especially in medicine, where segmentation greatly aids practitioners in scan-based diagnosis—making such techniques perhaps the most widely used in this context (Wang et al., 2022).

With constant advancements in data collection technologies and deep learning methods, the number of segmented images to analyze is expected to grow, motivating the development of specialized statistical methodologies for image analysis. This paper contributes in this direction. Rather than viewing images as collections of pixels, we aim to interpret them as collections of objects defined by their shape, color, and texture. In this manuscript, we explore the concept of multivariate shape analysis by extracting and analyzing shapes from segmented images. Essentially, we propose a novel representation of images as data.

In particular, we focus on classification: a common statistical problem in fields such as medicine. Our motivating example illustrates a classical situation where, given segmented chest X-ray images showing the two lungs and the heart, the task is to detect a morphological anomaly. Specifically, the application concerns cardiomegaly, an anomaly characterized by an abnormally large (relative) heart size. Detecting this condition has been the subject of numerous studies; see, e.g., Lee et al. (2024) and Bougias et al. (2021). However, to our knowledge, most of the methods investigated to date have been pixel-based.

We argue that this belongs to a new paradigm of statistical classification problems, one in which the shapes of objects in images, along with their location, size, and orientation, are used as explanatory variables.

Figure 1 illustrates the image representation used in the proposed image analysis procedure. Panel (a) shows the original X-ray as a collection of pixel intensity values stored as a matrix, while panel (b) displays the masks obtained through image segmentation (Gaggion et al., 2024). Panels (c) and (d) present the data we aim to analyze: the contours of the three objects of interest, denoted C_1 , C_2 and C_3 , which can be represented by two coordinate functions, i.e., $C_j = (X_j, Y_j)^\top$, $j = 1, 2, 3$. Together, these contours form a multivariate planar curve $\mathbf{C} = (C_1^\top, \dots, C_p^\top)^\top$ (with $p = 3$), which constitutes the parsimonious image representation studied in this work.

Statistical shape analysis and functional data analysis underlie this formalism. While the first is generally used to study single random planar curves ($p = 1$), the second focuses on various types of random functions.

In particular, foundational works in statistical shape analysis (see e.g. Kendall (1989), Dryden and Mardia (2016)) analyze the single or univariate planar curve \mathbf{C} by focusing on observations of \mathbf{C} at k specific points called landmarks. However, in general image analysis applications, landmarks are not provided, and choosing them appears to be crucial, as different landmarks might lead to different results. Modern works (see, e.g., Srivastava et al. (2010), Younes (1998)) propose an alternative setting that considers the entire random planar curve from a theoretical aspect, borrowing tools from differential geometry. The authors in Moindjié et al. (2025a) leverage the large amount of image data and proposed a fully functional data analysis framework to study contours of objects in images. Their work introduces an adaptation of (linear) functional principal analysis (Happ and Greven (2018), Ramsey and Silverman (2005)) for random planar curves. Indeed, linear functional data analysis methods cannot be directly applied to the space of random shapes due to their complex geometry. Even though \mathbf{C} (for $p \geq 1$) can be viewed as a multivariate functional variable in the sense of Happ and Greven (2018), i.e., a vector of random functions, adaptations of the methods developed for classification and regression (see e.g. Moindjié et al. (2025b), Febrero-Bande et al. (2017)) are mandatory for the same reasons as in the univariate case. This justifies the study of \mathbf{C} for $p \geq 2$, which, to our knowledge, has not been considered in

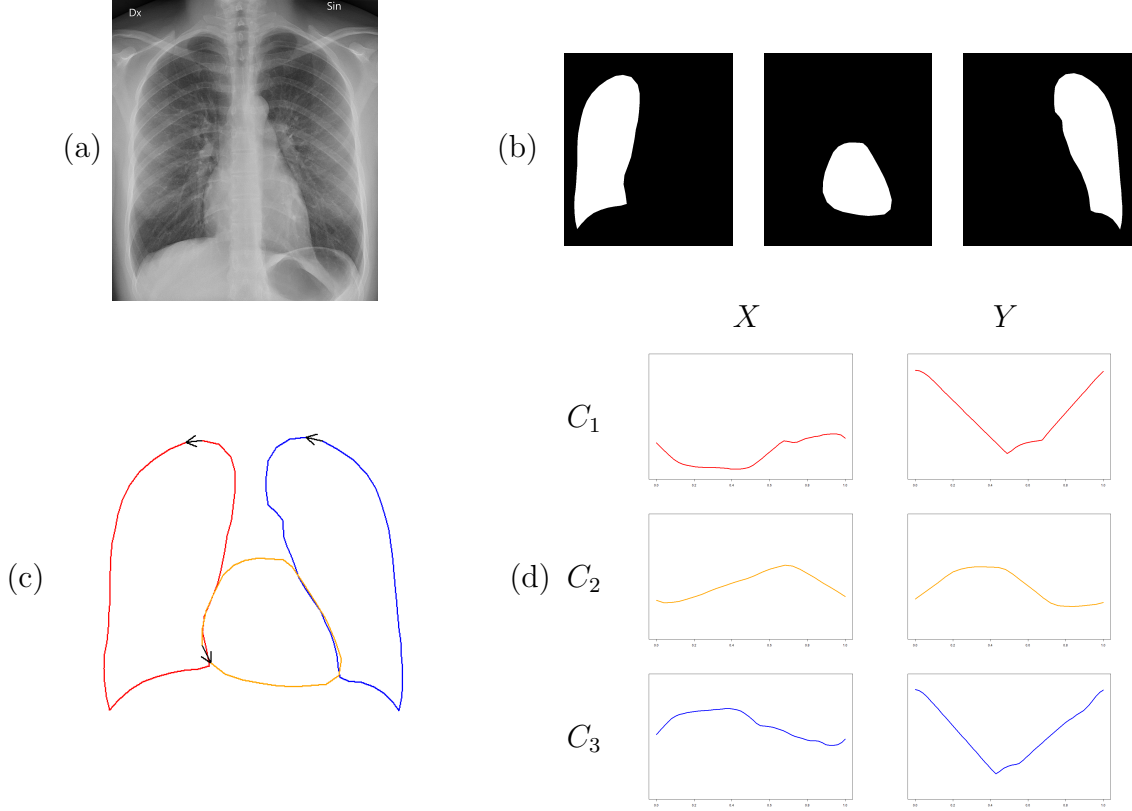


Figure 1: Representation of a segmented chest X-ray (Gaggion et al., 2024) as a multivariate planar curve. (a) Original X-ray image. (b) Segmentation masks of the three organs of interest: right lung, heart, and left lung. (c) Contours that form the multivariate planar curve \mathbf{C} . (d) Coordinate 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.

the literature, neither in a supervised nor in an unsupervised context. Indeed, although the case of several objects has been explored in the statistical shape analysis community, the focus seems to have been 3D objects (see e.g. Gorczowski et al. (2007), Khan et al. (2023)) rather than on 2D objects in images.

In this paper, we propose a modeling framework for \mathbf{C} when $p \geq 2$, by relying on the single planar curve modeling framework (Srivastava et al. (2010), Dryden and Mardia (2016)). This leads to a formal definition of the shape variable and the associated deformation variables. We also provide methods for estimating these parameters from random multivariate planar curves. Finally, we study the classification problem when the predictor is a multivariate planar curve. We show that, by relying on the tangent space transformations, this problem can be solved using (linear) functional data analysis methods (see e.g. Moindjié et al. (2024), Aguilera et al. (2008), Febrero-Bande et al. (2017)).

In summary, the main contributions of this paper are:

- A generalization of the functional shape model previously proposed in Moindjié et al. (2025a) to incorporate contour information from multiple objects simultaneously;

- An iterative procedure for the alignment of such multivariate curves;
- An application of the proposed model to classify observations from a chest X-ray dataset.

The rest of the paper is organized as follows. In Section 2, we provide a formal definition of multivariate random planar closed curves and describe the statistical framework for their analysis. In section 3, we discuss the estimation of the parameters involved in the analysis of multivariate curves. Section 4 presents an application of the proposed model on the chest X-ray data set shown in Figure 1. Section 5 contains a short discussion on future work.

2 Random planar variables

In this section, we introduce our proposed model for the multivariate random planar curve \mathbf{C} and define the functional space used to analyze observations of the shape associated with this variable. Since our construction builds upon the univariate model introduced in Moindjié et al. (2025a), when $p = 1$, we begin by reviewing their model.

2.1 Univariate random planar closed curves

In Moindjié et al. (2025a), 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. 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 $\mathcal{C} \in \mathbf{S}^\infty$, where

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

and where $\mathbf{S}^\infty = \left\{ \mathbf{f} \in \mathcal{H} : \int_0^1 f_1(t)dt = \int_0^1 f_2(t)dt = 0, \|\mathbf{f}\|_{\mathcal{H}} = 1 \right\}$ is 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. The function γ_δ is a reparametrization function belonging to

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

which parametrizes the starting location of the path traversed along the contour $C(t)$. The latent variable \mathcal{C} represents the shape of C , that is, the geometric information remaining once scale, rotation, and translation effects are removed (Dryden and Mardia, 2016).

2.2 Multivariate random planar closed curves

Let $\mathbf{C} = \text{Vec}(C_1, \dots, C_p) = (C_1^\top, \dots, C_p^\top)^\top$ be a multivariate planar curve, where each $C_j = (X_j, Y_j)^\top \in \mathcal{H}$ is a univariate planar closed curve as described above. Thus, $\mathbf{C} = (X_1, Y_1, \dots, X_p, Y_p)^\top \in \mathcal{H}^p$, where \mathcal{H}^p is the product Hilbert space, i.e., $\mathcal{H}^p = \mathcal{H} \times \dots \times \mathcal{H}$ (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 corresponding norm induced by this inner product.

A straightforward extension of the univariate model (1) would be to treat each component C_j independently, by setting

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

with \mathcal{C}_j the latent shape of C_j . However, this independent modeling ignores the internal structure of \mathbf{C} as a whole. Specifically, it fails to capture inter-component information such as relative positioning, scaling, and orientation. For instance, in our motivating example from Section 1, this would imply that the size relationships between organs (heart and lungs) are unimportant, thereby preventing the detection of relative size variations, like cardiomegaly anomalies.

To overcome this limitation, we propose the following joint model:

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

where the deformation parameters ρ , θ , and \mathbf{T} are shared across all components and where each curve is allowed its own reparameterization function γ_{δ_j} , reflecting independent shape traversal. Consequently, given that all curves are parametrized with a single position, scaling, and orientation parameter, the individual shapes \tilde{C}_j will carry the relative information.

In this model, the global shape of \mathbf{C} is encoded in the latent multivariate planar curve $\tilde{\mathbf{C}} = \text{Vec}(\tilde{C}_1, \dots, \tilde{C}_p) \in \mathbf{S}_p^\infty$ with \mathbf{S}_p^∞ being the centered unit p -dimensional Hilbert sphere:

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

Note that model (3) can be rewritten in a 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 a p -dimensional column vector of ones, and $\boldsymbol{\gamma} = (\gamma_{\delta_1}, \dots, \gamma_{\delta_p}) \in \Gamma^p$ is the vector of reparameterization functions. The composition $\tilde{\mathbf{C}} \circ \boldsymbol{\gamma}$ is understood component-wise:

$$\tilde{\mathbf{C}} \circ \boldsymbol{\gamma} = \begin{pmatrix} \tilde{C}_1 \circ \gamma_{\delta_1} \\ \vdots \\ \tilde{C}_p \circ \gamma_{\delta_p} \end{pmatrix}.$$

As in the univariate case, (Γ^p, \circ) is a commutative group and we have the following isometry property

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

The proof follows directly from Proposition 2.1 in Moindjié et al. (2025a).

Note that, from model (4) and the fact that $\tilde{\mathbf{C}}$ is centered and has unit norm, the translation and scaling parameters are defined as follows

$$\mathbf{T} = \frac{1}{p} \sum_{j=1}^p \int_0^1 C_j(t) dt = \left(\frac{1}{p} \sum_{j=1}^p \int_0^1 X_j(t) dt, \frac{1}{p} \sum_{j=1}^p \int_0^1 Y_j(t) dt \right)^\top, \quad \rho = \|\mathbf{C} - \mathbf{1}_p \otimes \mathbf{T}\|_{\mathcal{H}^p}. \quad (5)$$

2.3 Statistical analysis of shape variables

The space \mathbf{S}_p^∞ is called the *pre-shape* space, as it contains the shape $\tilde{\mathbf{C}}$ as well as any functions obtained by applying rotation and/or reparametrization deformations, i.e.,

$$(\mathbf{I}_p \otimes \mathbf{O}_\theta) \tilde{\mathbf{C}} \circ \gamma \in \mathbf{S}_p^\infty \quad \text{for all } \theta \in [0, 2\pi] \text{ and } \gamma \in \Gamma^p.$$

The shape space is then defined as the quotient space $\mathbf{S}_p^\infty / \sim$, where \sim is the equivalence relation:

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

for $\mathbf{f}, \mathbf{g} \in \mathbf{S}_p^\infty$. This relation simply states that $\mathbf{f} \sim \mathbf{g}$ if they share the same shape.

The statistical analysis of a random sample $\tilde{\mathbf{C}}_1, \dots, \tilde{\mathbf{C}}_n$ from $\tilde{\mathbf{C}}$ is challenging due to the nonlinear and complex geometry of the shape space $\mathbf{S}_p^\infty / \sim$. Classical statistical tools—especially those relying on linear structures—are generally inapplicable. Existing methods to analyze such data fall into two main categories: (i) direct inference on the nonlinear manifold \mathbf{S}_p^∞ , and (ii) inference based on a transformed version of the data that linearizes the structure while preserving essential geometric features.

The first category of methods attempts to work directly on the shape manifold. This includes generalizations of directional statistics to infinite-dimensional spaces, extending ideas used for data on spheres or tori (Mardia, 1975). While conceptually appealing, such approaches are limited by the lack of a well-defined probability density and by the difficulty of constructing tractable models in infinite dimensions (Delaigle and Hall, 2010). These limitations make direct inference on \mathbf{S}_p^∞ impractical in most applications.

The second approach, which we adopt here, addresses these issues by mapping the data into a more manageable linear space. In the classical literature on statistical shape analysis, a widely used strategy consists in projecting the data onto the tangent space \mathcal{T}_μ at a reference point $\mu \in \mathbf{S}_p^\infty$ (see e.g. Dryden and Mardia (2016), Srivastava et al. (2005)). The tangent space is defined as

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

Unlike \mathbf{S}_p^∞ , this space has a linear structure: it is closed under linear combinations, meaning that for any $\mathbf{f}, \mathbf{g} \in \mathcal{T}_\mu$ and $u, v \in \mathbb{R}$, we have

$$u\mathbf{f} + v\mathbf{g} \in \mathcal{T}_\mu.$$

This property enables the use of classical (functional) linear statistical methods in \mathcal{T}_μ . Details on statistical inference theory within this tangent space framework can be found in Dai (2022).

In the next subsections, we introduce a mapping to determine the projection onto \mathcal{T}_μ , then we discuss the choice of the reference point μ , which is crucial for ensuring a projection that approximately preserves the local geometry of the shape space. Finally, we briefly present how we can use the tangent space to address classification problems when the explanatory variable is a multivariate shape.

2.3.1 The tangent space \mathcal{T}_μ

Several techniques can be used to project elements from \mathbf{S}_p^∞ to \mathcal{T}_μ , such as orthogonal projections. Here, we use the exponential and logarithm mappings because of their well-established framework in both practical and theoretical settings (see e.g. Dryden and Mardia (2016), Dai (2022)).

The logarithm mapping is defined as $\mathcal{L}: \mathbf{S}_p^\infty / \{-\mu\} \rightarrow \mathcal{T}_\mu$, with

$$\mathcal{L}(\mathbf{f}) = \frac{\omega}{\sin(\omega)} (\mathbf{f} - \cos(\omega)\mu), \text{ where } \omega = d_{\mathbf{S}_p^\infty}(\mathbf{f}, \mu).$$

Here $d_{\mathbf{S}_p^\infty}: \mathbf{S}_p^\infty \times \mathbf{S}_p^\infty \rightarrow \mathbb{R}$ is the Riemann metric on the Hilbert sphere defined as

$$d_{\mathbf{S}_p^\infty}(\mathbf{f}, \mathbf{g}) = \cos^{-1}(\langle \mathbf{f}, \mathbf{g} \rangle_{\mathcal{H}^p}). \quad (6)$$

One interesting propriety of \mathcal{L} is that when $\mathbf{f} \rightarrow \mu$, we have $\mathcal{L}(\mathbf{f}) \rightarrow 0_{\mathcal{H}^p}$; since $\omega \rightarrow 0$ implies $\omega/\sin(\omega) \rightarrow 1$ and $\cos(\omega) \rightarrow 1$. Moreover, we have that $\|\mathcal{L}(\mathbf{f})\|_{\mathcal{H}^p} = d_{\mathbf{S}_p^\infty}(\mathbf{f}, \mu)$.

The inverse transformation of \mathcal{L} is the exponential mapping, defined as $\mathcal{E}: \mathcal{T}_\mu \rightarrow \mathbf{S}_p^\infty$, with

$$\mathcal{E}(\mathbf{v}) = \cos(\|\mathbf{v}\|_{\mathcal{H}^p})\mu + \|\mathbf{v}\|_{\mathcal{H}^p}^{-1} \sin(\|\mathbf{v}\|_{\mathcal{H}^p})\mathbf{v}.$$

Although the logarithm and exponential maps satisfy $\mathcal{L}(\mathcal{E}(\mathbf{v})) = \mathbf{v}$ for any $\mathbf{v} \in \mathcal{T}_\mu$, they are not bijective except in a neighborhood of μ .

Our approach to analyze a sample drawn from $\tilde{\mathbf{C}}$ is thus to first project each observation onto the tangent space via the logarithm map, perform the statistical analysis in this linear space, and then map the results back onto the sphere using the exponential map. For this procedure to be accurate, it is important that the observations lie sufficiently close to the reference point μ , ensuring that both projections remain valid and stable.

This consideration naturally motivates the choice of the Fréchet mean of $\tilde{\mathbf{C}}$ as the reference point μ , as we explain in the following subsection.

2.3.2 Functional mean on the Hilbert sphere

Inspired by classical shape analysis (see Dryden and Mardia (2016)), two types of mean for $\tilde{\mathbf{C}}$ can be defined on \mathbf{S}_p^∞ : extrinsic or intrinsic. Both can be written in the form of a Fréchet mean:

$$\mu = \arg \min_{\mathbf{f} \in \mathbf{S}_p^\infty} \mathbb{E}(d_e^2(\tilde{\mathbf{C}}, \mathbf{f})), \quad (7)$$

where d_e is a chosen distance.

The *extrinsic mean* relies on a distance d_e induced by embedding the shape space into a higher-dimensional linear space (e.g., using $\|\cdot\|_{\mathcal{H}^p}$). This approach involves applying a mapping $\Phi : \mathbf{S}_p^\infty \rightarrow \mathcal{H}^p$, computing the Euclidean mean in \mathcal{H}^p , and then mapping it back to \mathbf{S}_p^∞ via Φ^{-1} (see Dryden and Mardia, 2016). While computationally appealing, this method depends heavily on the choice of Φ , which is often nontrivial and can reduce interpretability.

A more natural choice is the *intrinsic mean*, where $\boldsymbol{\mu}$ is a point on the sphere which minimizes the distance-based variance for a selected intrinsic metric in \mathbf{S}_p^∞ . Here, we define d_e in (7) as the Riemann metric $d_{\mathbf{S}_p^\infty}$ given in (6).

In this intrinsic setting, no closed-form expression exists for $\boldsymbol{\mu}$ in (7), so it must be obtained via optimisation. Several authors have addressed this problem (see e.g. Srivastava et al. (2010), Dryden and Mardia (2016)). Inspired by this literature, we present our estimation procedure in Section 3.2.

2.3.3 Classification of multivariate shapes

Working in the tangent space greatly simplifies the analysis of $\tilde{\mathbf{C}}$. In particular, for classification tasks, it enables the use of linear methods when the explanatory variable is a random multivariate shape, as commonly done in classical shape analysis (see Dryden and Mardia (2016), Chap 13).

Formally, consider a binary response variable Y associated with a multivariate planar curve \mathbf{C} . Using the projection $\mathcal{L}(\tilde{\mathbf{C}})$ as the predictor embeds the problem in the classical framework of functional classification with multivariate functional predictors, which is a well-studied topic in the functional data analysis literature (see e.g. Godwin (2013), Moindjié et al. (2024), Febrero-Bande et al. (2017)). A typical model in this context is the functional linear model

$$Y = \psi \left(\beta_0 + \langle \mathcal{L}_\mu(\tilde{\mathbf{C}}), \boldsymbol{\beta} \rangle_{\mathcal{H}^p} \right) \quad (8)$$

where the parameters $\beta_0 \in \mathbb{R}$ and $\boldsymbol{\beta} \in \mathcal{H}^p$ are estimated from the data, and the link function $\psi : \mathbb{R} \rightarrow \{0, 1\}$ is chosen depending on the classification framework: linear discriminant analysis or logistic regression.

3 Functional estimation

We aim to analyze a dataset consisting of multivariate planar curves $\mathbf{C}_1, \dots, \mathbf{C}_n$ observed on a finite grid determined by the resolution of the images. The objective of this section is to present the complete pipeline for analyzing such a dataset. To do this, we begin by describing how to estimate the deformation parameters in model (3), which allows us to extract the shape associated with each multivariate planar curve \mathbf{C}_i . We then explain how to estimate the Fréchet mean from a set of shapes. Finally, we present the full procedure for analyzing a dataset, which includes smoothing the raw data and obtaining the shapes and the Fréchet mean through an interlaced estimation procedure.

In the following, we adopt a classical approach in FDA which assumes that the variable $\mathbf{C} = \text{Vec}(C_1, \dots, C_p)$ can be sufficiently well approximated by a finite Fourier basis expansion. In our setting, this means we assume that

$$C_j(t) = \begin{pmatrix} X_j(t) \\ Y_j(t) \end{pmatrix} = \begin{pmatrix} B_{j1} \\ B_{j2} \end{pmatrix} + \sum_{m=1}^M \begin{pmatrix} A_{jm1} \\ A_{jm2} \end{pmatrix} \phi_m(t) = B_j + A_j \boldsymbol{\phi}(t), \quad j = 1, \dots, p, \quad (9)$$

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

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

Note that (9) can be rewritten in the following compact form :

$$\mathbf{C} = \mathbf{B} + \mathbf{A}\boldsymbol{\phi}, \quad (10)$$

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 .

3.1 Estimation of the shape variable

We now describe how to estimate the shape variable $\tilde{\mathbf{C}}$ from a given multivariate planar curve \mathbf{C} , assuming that its representation in the form of (10) is known — that is, the coefficient matrices \mathbf{A} and \mathbf{B} are given.

According to model (3), each component \tilde{C}_j of the shape vector $\tilde{\mathbf{C}}$ can be expressed as

$$\tilde{C}_j = \mathbf{O}_\theta^\top \left(\frac{1}{\rho} (C_j - \mathbf{T}) \right) \circ \gamma_{1-\delta_j} = \mathbf{O}_\theta^\top (C_j^*) \circ \gamma_{1-\delta_j}, \quad j = 1, \dots, p, \quad (11)$$

where $\gamma_{1-\delta_j}$ is the inverse function of γ_{δ_j} (Moindjié et al., 2025a) and $\mathbf{C}^* = \text{Vec}(C_1^*, \dots, C_p^*) \in \mathbf{S}_p^\infty$ represents the pre-shape of \mathbf{C} .

Section 3.1.1 presents the direct method for computing the pre-shape \mathbf{C}^* , while Section 3.1.2 introduces a new algorithm for estimating the shape $\tilde{\mathbf{C}}$ from \mathbf{C}^* .

3.1.1 Translation and scaling parameters

Translation and scaling are computed independently for each curve of the dataset, as defined in (5). Using the representation (10), and the fact that the Fourier functions ϕ_m are orthonormal and satisfy $\int_0^1 \phi_m(t) dt = 0$, we obtain the following expressions:

$$\mathbf{T} = \left(\frac{1}{p} \sum_{j=1}^p B_{j1}, \frac{1}{p} \sum_{j=1}^p B_{j2} \right)^\top = \frac{1}{p} \sum_{j=1}^p B_j \quad \text{and} \quad \rho = \sqrt{\|\mathbf{A}\|_F^2 + \|\mathbf{B} - \mathbf{1}_p \otimes \mathbf{T}\|_F^2},$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

Substituting into (11), we can write the pre-shape as

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

3.1.2 Alignment by Iterative Closest Function

Extracting the shape from a pre-shape is essentially an alignment problem. The goal is to estimate the rotation and reparametrization associated with each pre-shape \mathbf{C}_i^* , so as to obtain well-aligned shapes $\tilde{\mathbf{C}}_1, \dots, \tilde{\mathbf{C}}_n$. The strategy is to align the preshapes to a common template $\bar{\mathbf{C}} = \text{Vec}(\bar{C}_1, \dots, \bar{C}_p) \in \mathbf{S}_p^\infty$, chosen to be a good representative of the dataset. For example, $\bar{\mathbf{C}}$ can be a Fréchet-mean-type function or simply one of the observed pre-shapes.

Aligning a pre-shape \mathbf{C}^* to a given template $\bar{\mathbf{C}}$ involves estimating the rotation matrix \mathbf{O}_θ and the vector of reparametrization functions $\gamma = (\gamma_{\delta_1}, \dots, \gamma_{\delta_p})$, by solving the following optimization problem:

$$\min_{\boldsymbol{\delta} \in [0,1]^p, \theta \in [0,2\pi]} \sum_{j=1}^p \|\mathbf{O}_\theta \bar{C}_j \circ \gamma_{\delta_j} - C_j^*\|_{\mathcal{H}}^2, \quad (12)$$

where $\boldsymbol{\delta} = (\delta_1, \dots, \delta_p)$.

In the univariate case ($p = 1$), this problem has been addressed using grid search techniques (Moindjié et al., 2025a). However, for $p \geq 2$, the dimensionality of the search space makes such methods computationally prohibitive.

To overcome this issue, we propose a novel alternating optimization algorithm inspired by the iterative closest point algorithm (Zhang, 2021), which operates in a fully functional setting – without requiring discretization of the curves. We refer to this algorithm as *Iterative Closest Function* (ICF). It proceeds by alternating between two steps: (i) rotation estimation assuming the reparametrization functions are fixed, and (ii) reparametrization estimation assuming the rotation angle is fixed. Each of these steps is described in detail below.

(i) Estimation of θ for a given γ

In this step, the estimator of θ is obtained by solving a Procrustes orthogonal problem:

$$\hat{\theta} = \arg \min_{\theta \in [0,2\pi]} \sum_{j=1}^p \|\mathbf{O}_\theta \bar{C}_j \circ \gamma_{\delta_j} - C_j^*\|_{\mathcal{H}}^2. \quad (13)$$

The multivariate version of this problem has been extensively studied (see, e.g., Schönemann (1966)). In the functional setting, by arguments analogous to those in the classical case, one can show that $\hat{\theta}$ belongs to the set $\{\theta_1, \theta_2\}$, where θ_1 and θ_2 are the two solutions to the following equation:

$$\tan(\theta_k) = \left(\sum_{j=1}^p \langle C_j^*, \bar{C}_j \circ \gamma_{\delta_j} \rangle_{\mathcal{H}} \right)^{-1} \sum_{j=1}^p \{ \langle X_j^*, \bar{X}_j \circ \gamma_{\delta_j} \rangle_{L_2} - \langle Y_j^*, \bar{Y}_j \circ \gamma_{\delta_j} \rangle_{L_2} \}, k = 1, 2,$$

where $\bar{C}_j = (\bar{X}_j, \bar{Y}_j)$ and $C_j^* = (X_j^*, Y_j^*)$ for $j = 1, \dots, p$. Once θ_1 and θ_2 have been computed, we evaluate the objective function in (13) at both values and select $\hat{\theta}$ as the one yielding the smallest value.

(ii) Estimation of γ for a given θ

When the rotation angle θ is known, solving problem (12) reduces to solving p independent optimization problems:

$$\hat{\delta}_j = \arg \min_{\delta \in [0,1]} \|\mathbf{O}_\theta \bar{C}_j \circ \gamma_\delta - C_j^*\|_{\mathcal{H}}^2, \quad j = 1, \dots, p. \quad (14)$$

Unlike step (i), this problem does not generally admit a closed-form solution. Instead, we solve it numerically using the following result:

Proposition 3.1. *Suppose that $\bar{\mathbf{C}}$ admits a Fourier basis expansion such that:*

$$\mathbf{O}_\theta \bar{\mathbf{C}}_j = \bar{\mathbf{B}}_j^\theta + \bar{\mathbf{A}}_j^\theta \boldsymbol{\phi}, \quad j = 1, \dots, p,$$

where $\bar{\mathbf{B}}_j^\theta \in \mathbb{R}^2$ and $\bar{\mathbf{A}}_j^\theta \in \mathbb{R}^{2 \times M}$ are coefficient matrices. Then, for $j \in \{1, \dots, p\}$, the solution $\hat{\delta}_j$ of problem (14) satisfies the equation

$$\sum_{k \in \{1, 3, \dots, M-1\}} w_{j,k}^{1,\theta} \sin((k+1)\pi \hat{\delta}_j) = \sum_{k \in \{1, 3, \dots, M-1\}} w_{j,k}^{2,\theta} \cos((k+1)\pi \hat{\delta}_j), \quad (15)$$

where the coefficients $w_{j,k}^{1,\theta}$ and $w_{j,k}^{2,\theta}$ are derived from the matrix $(\bar{\mathbf{A}}_j^\theta)^\top \mathbf{A}_j^*$, with $\mathbf{A}_j^* = \rho^{-1} \mathbf{A}_j$.

For clarity, the exact expressions of the coefficients $(w_{j,k}^{l,\theta})_{j,k,l}$ and the proof of the proposition are provided in Appendix A.

We solve equation (15) with respect to $\hat{\delta}_j$ using a bisection method. This approach yields accurate results, as demonstrated in the numerical experiments of Section 4.1.

Remark. *Note that the performance of the proposed ICF algorithm depends on the initialization of $\boldsymbol{\delta}$ in step (i). We therefore recommend running the algorithm with multiple initial values and selecting the pair $(\hat{\theta}, \hat{\boldsymbol{\delta}})$ that minimizes the objective function in (12). In our numerical experiments, we found that using five different initializations was sufficient to obtain good results.*

3.2 Estimation of the Fréchet mean

Recall that in Section 2.3.2, we defined the intrinsic Fréchet mean as

$$\boldsymbol{\mu} = \arg \min_{\mathbf{f} \in \mathbb{S}_p^\infty} \mathbb{E} \left[\left\{ \cos^{-1} \left(\langle \tilde{\mathbf{C}}, \mathbf{f} \rangle_{\mathcal{H}^p} \right) \right\}^2 \right]. \quad (16)$$

We now describe how to estimate this mean from a sample of shapes $\tilde{\mathbf{C}}_1, \dots, \tilde{\mathbf{C}}_n$. We assume that each shape admits a Fourier basis expansion of the form:

$$\tilde{\mathbf{C}}_i = \tilde{\mathbf{B}}_i + \tilde{\mathbf{A}}_i \boldsymbol{\phi}, \quad i = 1, \dots, n,$$

where $\tilde{\mathbf{B}}_i \in \mathbb{R}^{2p}$ and $\tilde{\mathbf{A}}_i \in \mathbb{R}^{2p \times M}$ are coefficient matrices.

We define the estimator $\hat{\boldsymbol{\mu}}$ of the Fréchet mean as

$$\hat{\boldsymbol{\mu}} = \hat{\mathbf{F}} + \hat{\mathbf{E}} \boldsymbol{\phi},$$

where

$$(\hat{\mathbf{E}}, \hat{\mathbf{F}}) = \arg \min_{\mathbf{E} \in \mathbb{R}^{2p \times M}, \mathbf{F} \in \mathbb{R}^{2p}} \frac{1}{n} \sum_{i=1}^n \left\{ \cos^{-1} \left(\langle \tilde{\mathbf{A}}_i, \mathbf{E} \rangle_F + \langle \tilde{\mathbf{B}}_i, \mathbf{F} \rangle_F \right) \right\}^2, \quad (17)$$

and \langle, \rangle_F denotes the Frobenius inner product.

Note that the minimization problem in (17) is non-convex. It is structurally similar to the Fréchet mean estimation problem in classical statistical shape analysis, where a numerical solution was proposed in Dryden (2023). We rely on their implementation to solve (17), using the generic non-linear minimization function *nlm* from the **stats** package in R (R Core Team, 2023). Further details of the algorithm can be found in our publicly available code (<https://github.com/imoindjie/Shape-MFDA>).

3.3 Iterative estimation of shapes and of the Fréchet mean

In this work, we use the intrinsic Fréchet mean defined in (16) as the template for the alignment procedure described in Section 3.1.2, i.e., we set $\bar{\mathbf{C}} = \boldsymbol{\mu}$. There is a strong interdependence between the estimation of the shapes and the estimation of $\boldsymbol{\mu}$: we need $\boldsymbol{\mu}$ to estimate the aligned shapes, and we need the aligned shapes to estimate $\boldsymbol{\mu}$. A poor estimate of one can thus negatively impact the other. To make the estimation procedure robust, we propose an iterative algorithm that alternates between shape alignment and Fréchet mean estimation.

The complete estimation pipeline, from raw data to aligned shapes and Fréchet mean, is summarized below:

- Input:
 - $\{(\mathbf{C}_i(t_{i1}), \dots, \mathbf{C}_i(t_{iK_i}))\}_{i=1}^n$: the observed curves on a discrete grid
 - M : the number of Fourier basis functions
 - ξ : a stopping threshold for convergence
- Algorithm:
 1. **Smoothing**
 For each $i \in \{1, \dots, n\}$, estimate the coefficient matrices \mathbf{A}_i and \mathbf{B}_i (defined in (10)) by minimizing a least squares criterion using the discretized curve $(\mathbf{C}_i(t_{i1}), \dots, \mathbf{C}_i(t_{iK_i}))$. Denote the resulting smoothed curves by $\hat{\mathbf{C}}_i = \hat{\mathbf{B}}_i + \hat{\mathbf{A}}_i \boldsymbol{\phi}$.
 2. **Centering and normalizing**
 For each $i \in \{1, \dots, n\}$, estimate the translation vector \mathbf{T}_i and the scaling factor ρ_i (as defined in (3.1.1)) from $\hat{\mathbf{A}}_i$ and $\hat{\mathbf{B}}_i$. Define the estimated pre-shape as $\hat{\mathbf{C}}_i^* = \hat{\rho}_i^{-1}(\hat{\mathbf{C}}_i - \mathbf{1}_p \otimes \hat{\mathbf{T}}_i)$.
 3. **Iterative alignment and Fréchet mean estimation**
 Initialize $\hat{\boldsymbol{\mu}}$ by randomly choosing one pre-shape from $\hat{\mathbf{C}}_1^*, \dots, \hat{\mathbf{C}}_n^*$, then repeat the following two steps until convergence, i.e., until $\eta \leq \xi$:
 - a. For each $i \in \{1, \dots, n\}$, estimate the rotation angle θ_i and the reparametrization function $\boldsymbol{\gamma}_i$ using the ICF algorithm with $\hat{\boldsymbol{\mu}}$ as the template. Define the estimated shape as $\hat{\mathbf{C}}_i = (\mathbf{I}_p \otimes \mathbf{O}_{\hat{\theta}}^\top) \hat{\mathbf{C}}_i^* \circ \boldsymbol{\gamma}_{1-\hat{\delta}}$, where $\boldsymbol{\gamma}_{1-\hat{\delta}} = (\gamma_{1-\hat{\delta}_1}, \dots, \gamma_{1-\hat{\delta}_p})$.

- b. Using the shapes $\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_n$, estimate the Fréchet mean $\boldsymbol{\mu}$ as described in Section 3.2. Then compute $\eta = \frac{1}{n} \sum_{i=1}^n \left(\cos^{-1} \left(\langle \hat{\mathbf{C}}_i, \hat{\boldsymbol{\mu}} \rangle_{\mathcal{H}^p} \right) \right)^2$, the distance-based variance used in the stopping criterion.

4 Numerical experiments

In this section, we present numerical experiments to assess the performance of our proposed method. One of our objectives is to evaluate whether our approach performs well in a classification setting. To this end, we apply it to the real-world task of detecting cardiomegaly from chest X-ray images. However, before turning to this application, we first demonstrate the near-perfect accuracy of our alignment procedure using a synthetic dataset, as classification of shapes heavily depends on the accurate estimation of deformation parameters.

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 3.1. Since translation and scaling parameters are known to be accurately estimated in practice (Moindjié et al., 2025a), we focus exclusively on the estimation of rotation and reparametrization. That is, we evaluate the estimation of shapes from preshapes using the ICF algorithm in a controlled setting.

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 ($p = 3$) representing the contours of the heart and both lungs. We randomly select one curve, smooth it using $M = 22$ Fourier basis functions, and center it. This results in a function, denoted by \mathbf{c}^0 , which, following Equation (9), has the representation:

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

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.

We then generate preshapes $\mathbf{c}_1^*, \dots, \mathbf{c}_n^*$ by applying 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:

$$\mathbf{c}_{ij}^* = \kappa_i \mathbf{O}_{\theta_i} (b_j^0 + \tilde{a}_{ij} \phi \circ \gamma_{\delta_{ij}}), \quad j = 1, 2, 3, \quad (18)$$

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. By construction, the function $\mathbf{c}^0 / \|\mathbf{c}^0\|_{\mathcal{H}^p}$ is the Fréchet mean of the simulated dataset.

We simulate $n = 500$ curves for three different values of σ , namely 0.1, 0.5, and 1.0, to generate preshapes that range from very similar to substantially different compared to the template \mathbf{c}^0 . For each value of σ , Figure 3 illustrates three examples of simulated preshapes.

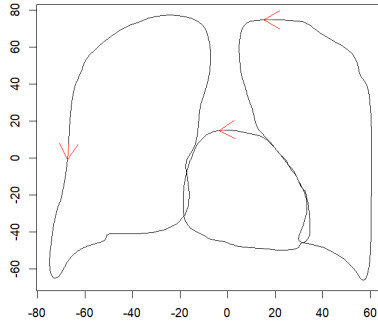


Figure 2: The function \mathbf{c}^0

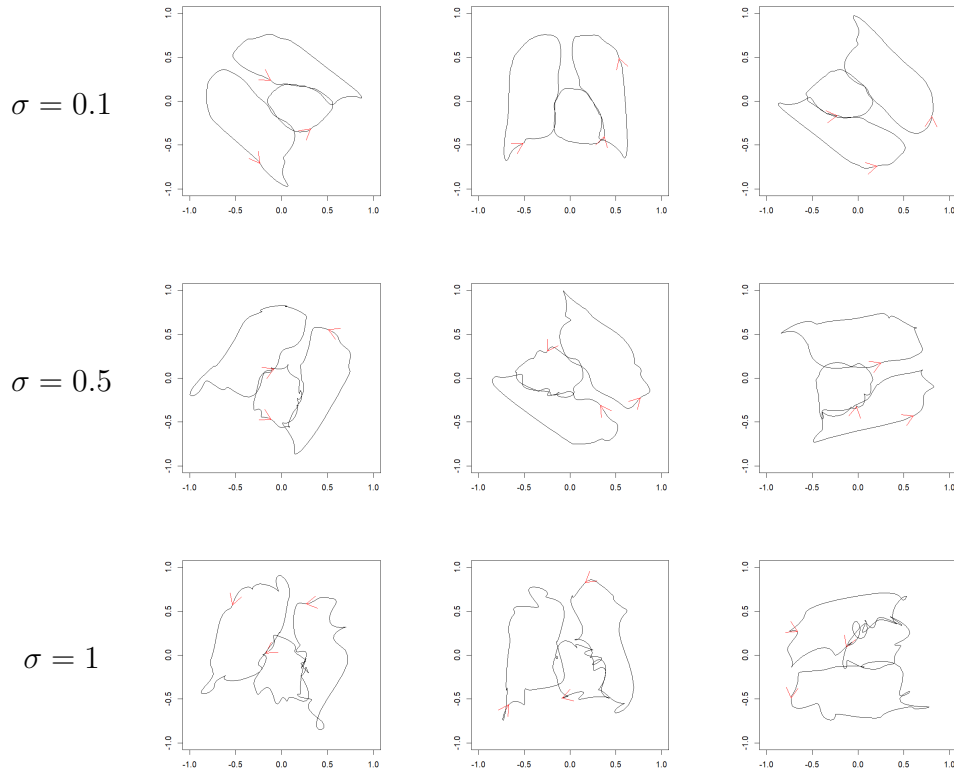


Figure 3: Simulated preshapes generated from the template \mathbf{c}^0 under three levels of deformation noise σ . Each deformation involves random rotation and reparametrization, and the degree of shape distortion increases with σ .

For each of the three datasets, we estimate the deformation parameters θ_i and $\boldsymbol{\delta}_i = (\delta_{i1}, \delta_{i2}, \delta_{i3})$ for every preshape by applying only steps 1 and 3.a. of the algorithm described in Section 3.3, using $M = 22$ and $\boldsymbol{\mu} = \mathbf{c}^0 / \|\mathbf{c}^0\|_{\mathcal{H}^p}$. Within the ICF algorithm, we initialize the reparametrization vector $\boldsymbol{\delta}_i$ with 5 randomly selected values to improve robustness.

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), defined as:

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

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 proposed 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 of estimated deformation parameters for different levels of noise σ .

Figure 4 further illustrates this result by displaying a subsample of simulated preshapes 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

In this section, we illustrate the practical relevance of our methodology by applying it to a real-world classification problem: detecting cardiomegaly from chest X-ray data. 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 (see Figure 5). This condition is clinically significant, as it may indicate underlying heart diseases such as heart failure or arrhythmia, which can lead to severe complications including sudden cardiac arrest. Early and accurate detection is therefore essential.

We base our analysis on the chexmask dataset (Gaggion et al., 2024), which contains 657,566 segmented masks derived from several large-scale public databases (ChestX-ray8, CheXpert, MIMIC-CXR-JPG, PadChest, and VinDr-CXR). In this study, we focus on samples from the ChestX-ray8 subset (Wang et al., 2017), which includes associated pathology labels, enabling supervised learning.

Each segmented image in the dataset consists of a multivariate planar curve, with $p = 3$ components corresponding to the contours of the heart and the left and right lungs. This structure makes the dataset a natural setting for multivariate shape analysis.

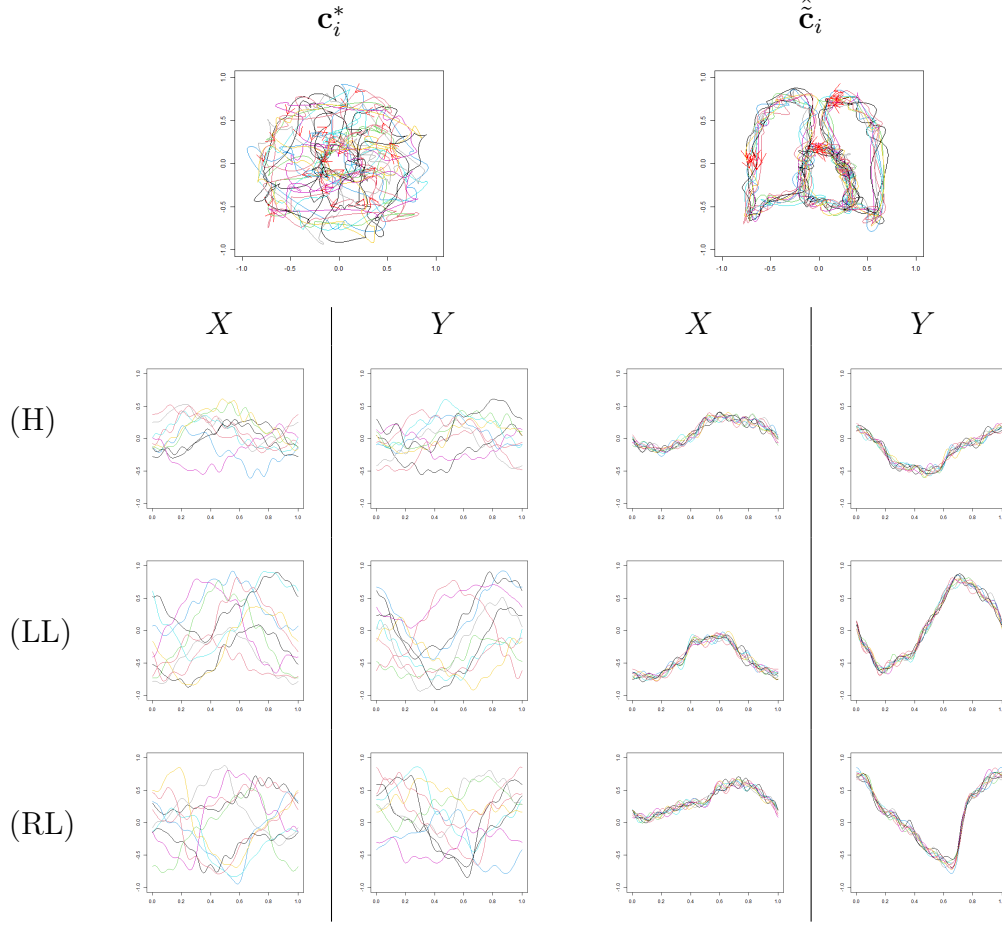


Figure 4: Left panel: a subsample of simulated preshapes for $\sigma = 1$. Right panel: aligned shapes obtained with the proposed ICF method. The three components of each multivariate curve are the heart (H), the left lung (LL) and the right lung (RL).

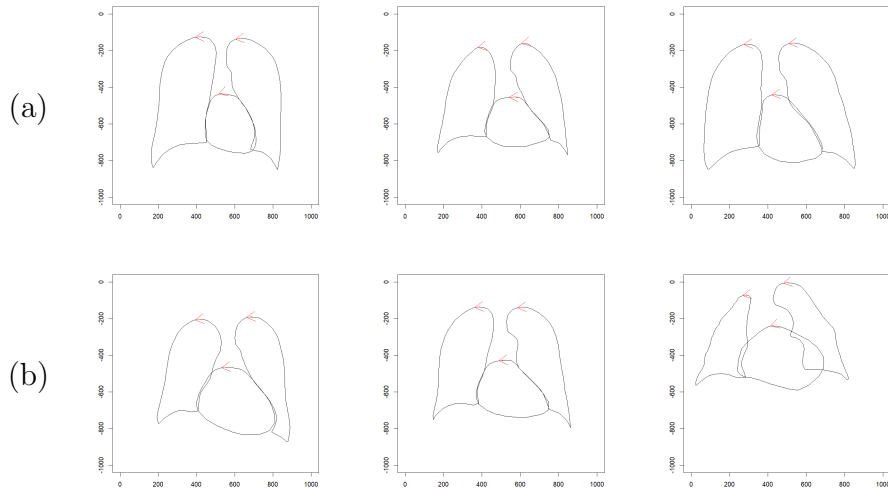


Figure 5: Examples of multivariate planar curves corresponding to healthy patients (a) and patients diagnosed with cardiomegaly (b). Each image shows the contours of the heart and both lungs.

As illustrated in Figure 5, the distinction between healthy and cardiomegaly patients can be subtle: there is high variability in both the shape and the deformations of the anatomical structures, making automated classification challenging.

4.2.1 Dataset

From the original dataset, we extract a subset of $n = 300$ images, with 150 labeled as cardiomegaly and 150 as healthy (labeled “no finding” in the ChestX-ray8 dataset). This subset allows us to evaluate classification performance in a controlled and balanced setting. The resulting dataset, denoted $\mathbf{c}_1, \dots, \mathbf{c}_n$, is illustrated on Figure 6, and can be found in our online code repository (<https://github.com/imoindjie/Shape-MFDA>).

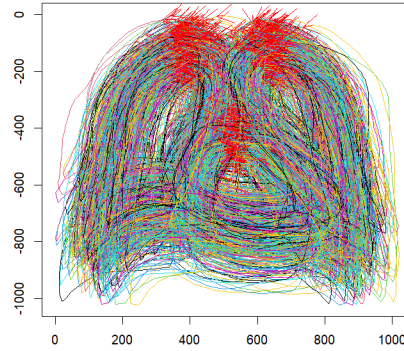


Figure 6: The 300 selected multivariate planar curves \mathbf{c}_i

Importantly, note that the curves provided in Gaggion et al. (2024) have already been preprocessed to correct for translation, rotation, and scale variability. This preprocessing is achieved using standard image processing techniques applied to the outputs of the HybridGNet segmentation algorithm (Gaggion et al., 2021).

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

Scenario 1 - Aligned data (ideal case) We directly use the pre-aligned curves $\mathbf{c}_i, i = 1, \dots, n$, as they are provided in the dataset.

Scenario 2 - Unaligned data (realistic case) To simulate a more realistic setting where pre-alignment is not performed, we introduce artificial deformations to each observation. We set $\mathbf{c}_i, i = 1, \dots, n$ as

$$\mathbf{c}_i \leftarrow (\mathbf{I}_3 \otimes \mathbf{O}_{2\pi\zeta_i})\mathbf{c}_i \circ \gamma_{\boldsymbol{\delta}_i}, \text{ with } \boldsymbol{\delta}_i = (\delta_{i1}, \delta_{i2}, \delta_{i3}) \in [0, 1]^3, \quad (19)$$

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 our method handles raw, misaligned input data.

4.2.2 Classification models

As introduced in Section 2.3.3, our 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}_μ 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 3.3. This method jointly models the three anatomical components.

- *UNI* - univariate approach: each anatomical component is analyzed separately, as in (2), and the shapes are estimated following the univariate procedure of Moindjié et al. (2025a). We use

$$(\mathcal{L}_{\hat{\mu}_1}(\hat{\mathbf{C}}_{i1}), \mathcal{L}_{\hat{\mu}_2}(\hat{\mathbf{C}}_{i2}), \mathcal{L}_{\hat{\mu}_3}(\hat{\mathbf{C}}_{i3}))^\top, i = 1, \dots, n,$$

as predictors, where \mathcal{L} is defined for $p = 1$ and $\hat{\mu}_j$ is the estimated 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 (8).

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, from the set $\{0.96^l \lambda_{\max}, l = 0, 1, \dots, 148\} \cup \{0\}$, where λ_{\max} is the smallest value for which the penalty term vanishes, as the value that minimizes classification error.

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.

4.2.3 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 as it treats the heart and each lung as independent univariate curves, their relative spatial relationships is ignored. Nevertheless, its accuracy remains relatively high, underlining the strong discriminative power of shape information alone.

In Scenario 2 (unaligned curves), a clear accuracy gap emerges between the shape-based methods (*OUR* and *UNI*) and the classical *CLA* approach. The *CLA* accuracies drop to around 50–57%, close to random guessing, highlighting the detrimental impact of ignoring alignment in multivariate shape classification. In contrast, both *OUR* and *UNI* maintain accuracies similar to those obtained in Scenario 1, demonstrating their robustness to rotational and reparametrization variability.

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

		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: Accuracies (%) for 10-fold cross-validation

5 Conclusion

This work extends statistical shape analysis for multiple objects by adopting a functional data analysis perspective. Similarly to the univariate case, the observed planar curve is viewed as a deformed version of an unobserved latent variable of interest: the shape. The novelty of our work lies in considering planar curves as multivariate, allowing the study of multiple objects simultaneously. After defining a suitable and intuitive model in this setting, we introduce estimation procedures for the deformation and shape variables. In particular, the proposed procedure borrows concepts from the iterative closest point approach to estimate the reparametrization and rotation effects, leading to the ICF procedure. Once obtained, the shape variables are estimated through inverse transformations. In the final part of our

work, we discuss the use of these variables in prediction models through projections onto the tangent space.

The numerical experiments demonstrate the proficiency of our method in estimating the deformation parameters and recovering the shape across diverse scenarios. They also show that using multivariate shapes in a classification task improves precision compared to functional linear methods and univariate shape analysis-based procedures, particularly in realistic scenarios where images are not necessarily aligned. The proposed representation for multivariate planar curves is 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

- Aguilera, A. M., Escabias, M., and Valderrama, M. J. (2008). Discussion of different logistic models with functional data. application to systemic lupus erythematosus. *Computational Statistics & Data Analysis*, 53(1):151–163.
- Bougias, H., Georgiadou, E., Malamateniou, C., and Stogiannos, N. (2021). Identifying cardiomegaly in chest x-rays: a cross-sectional study of evaluation and comparison between different transfer learning methods. *Acta Radiologica*, 62(12):1601–1609.
- Dai, X. (2022). Statistical inference on the hilbert sphere with application to random densities. *Electronic Journal of Statistics*, 16(1):700–736.
- Delaigle, A. and Hall, P. (2010). Defining probability density for a distribution of random functions. *The Annals of Statistics*, pages 1171–1193.
- Dryden, I. L. (2023). *shapes: Statistical Shape Analysis*. R package version 1.2.7.
- 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., Mansilla, L., Milone, D. H., and Ferrante, E. (2021). Hybrid graph convolutional neural networks for landmark-based anatomical segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, pages 600–610. Springer.
- 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.
- Gorczowski, K., Styner, M., Jeong, J.-Y., Marron, J., Piven, J., Hazlett, H. C., Pizer, S. M., and Gerig, G. (2007). Statistical shape analysis of multi-object complexes. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE.

- 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.
- Happ, C. and Greven, S. (2018). Multivariate functional principal component analysis for data observed on different (dimensional) domains. *Journal of the American Statistical Association*, 113(522):649–659.
- 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.
- Kendall, D. G. (1989). A survey of the statistical theory of shape. *Statistical Science*, 4(2):87–99.
- Khan, N., Peterson, A. C., Aubert, B., Morris, A., Atkins, P. R., Lenz, A. L., Anderson, A. E., and Elhabian, S. Y. (2023). Statistical multi-level shape models for scalable modeling of multi-organ anatomies. *Frontiers in Bioengineering and Biotechnology*, 11:1089113.
- Lee, K.-H., Choi, J.-W., Park, C.-O., Han, D.-H., and Kang, M.-S. (2024). A development and validation of an ai model for cardiomegaly detection in chest x-rays. *Applied Sciences*, 14(17):7465.
- Mardia, K. V. (1975). Statistics of directional data. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 37(3):349–371.
- 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. (2025a). A functional approach for 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.
- Moindjié, I.-A., Preda, C., and Dabo-Niang, S. (2025b). Fusion regression methods with repeated functional data. *Computational Statistics & Data Analysis*, 203:108069.
- Preda, C., Saporta, G., and Lévêder, C. (2007). Pls classification of functional data. *Computational Statistics*, 22(2):223–235.
- R Core Team (2023). *The R Stats Package*. R package version 4.3.0.

- Ramsey, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer New York, 2nd edition.
- Schönemann, P. H. (1966). A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10.
- 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.
- 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.

A Proof of Proposition 3.1

Using Lemma 2.3 in Moindjié et al. (2025a), we have, for $j \in \{1, \dots, p\}$,

$$\begin{aligned} \|\mathbf{O}_\theta \bar{C}_j \circ \gamma_\delta - C_j^*\|_{\mathcal{H}}^2 &= \|\bar{A}_j^\theta \phi \circ \gamma_\delta - A_j^* \phi + \bar{B}_j^\theta - B_j^*\|_{\mathcal{H}}^2 \\ &= \|\bar{A}_j^\theta P_\delta - A_j^*\|_F^2 + \|\bar{B}_j^\theta - B_j^*\|_F^2, \end{aligned}$$

where

$$P_\delta = \begin{pmatrix} \mathbf{O}_{2\pi\delta} & 0 & \dots & 0 \\ 0 & \mathbf{O}_{4\pi\delta} & \dots & 0 \\ \vdots & \dots & & \vdots \\ 0 & 0 & \dots & \mathbf{O}_{M\pi\delta} \end{pmatrix}.$$

Since the second term on the right-hand side does not depend on δ , (14) is equivalent to

$$\hat{\delta}_j = \arg \min_{\delta \in [0,1]} \|\bar{A}_j^\theta P_\delta - A_j^*\|_F^2.$$

This can be seen as a constrained Procrustes problem, since the orthogonal matrix P_δ has a fixed form. Therefore, we have

$$\hat{\delta}_j = \arg \min_{\delta \in [0,1]} \|P_\delta - (\bar{A}_j^\theta)^\top A_j^*\|_F^2. \quad (20)$$

Since P_δ is sparse orthogonal, Equation (20) reduces to

$$\hat{\delta}_j = \arg \min_{\delta \in [0,1]} \left(\sum_{k \in \{1,3,\dots,M-1\}} \|\Sigma_{j,k,\theta} - \mathbf{O}_{\pi\delta(k+1)}\|^2 \right), \quad (21)$$

where

$$((\bar{A}_j^\theta)^\top A_j^*) \odot \begin{pmatrix} \mathbf{I}_2 & 0 & \dots & 0 \\ 0 & \mathbf{I}_2 & \dots & 0 \\ \vdots & \dots & & \vdots \\ 0 & 0 & \dots & \mathbf{I}_2 \end{pmatrix} = \begin{pmatrix} \Sigma_{j,1,\theta} & 0 & \dots & 0 \\ 0 & \Sigma_{j,3,\theta} & \dots & 0 \\ \vdots & \dots & & \vdots \\ 0 & 0 & \dots & \Sigma_{j,M-1,\theta} \end{pmatrix},$$

and \odot denotes the Hadamard product.

Regarding the resolution of (21), note that

$$\begin{aligned} \frac{\partial}{\partial \delta} \|\Sigma_{j,k,\theta} - \mathbf{O}_{\pi\delta(k+1)}\|^2 &= \pi(k+1) (\text{Tr}(\Sigma_{j,k,\theta}) \sin(\pi(k+1)\delta) \\ &\quad + \text{Tr}(\Sigma_{j,k,\theta} \mathbf{O}_{\pi/2}) \cos(\pi(k+1)\delta)) \end{aligned}$$

which leads to

$$\begin{aligned} \sum_{k \in \{1,3,\dots,M-1\}} \frac{\partial}{\partial \delta} \|\Sigma_{j,k,\theta} - \mathbf{O}_{\pi\delta(k+1)}\|^2 &= \pi \sum_{k \in \{1,3,\dots,M-1\}} \left(w_{j,k}^{1,\theta} \sin(\pi(k+1)\delta) \right. \\ &\quad \left. - w_{j,k}^{2,\theta} \cos(\pi(k+1)\delta) \right), \end{aligned}$$

where $w_{j,k}^{1,\theta} = \text{Tr}((k+1)\mathbf{\Sigma}_{j,k,\theta})$, $w_{j,k}^{2,\theta} = \text{Tr}(-(k+1)\mathbf{\Sigma}_{j,k,\theta}\mathbf{O}_{\pi/2})$.

Therefore, the solution $\hat{\delta}_j$ of (14) belongs to \mathcal{S}_M , where

$$\mathcal{S}_M = \left\{ \delta \in [0, 1], \quad \sum_{k \in \{1,3,\dots,M-1\}} w_{j,k}^{1,\theta} \sin(\pi(k+1)\delta) = \sum_{k \in \{1,3,\dots,M-1\}} w_{j,k}^{2,\theta} \cos(\pi(k+1)\delta) \right\}.$$

This concludes the proof.