

An elasticity-based covariance analysis of shapes

Martin Rumpf · Benedikt Wirth

Received: date / Accepted: date

Abstract We introduce the covariance of a number of given shapes if they are interpreted as boundary contours of elastic objects. Based on the notion of nonlinear elastic deformations from one shape to another, a suitable linearization of geometric shape variations is introduced. Once such a linearization is available, a principal component analysis can be investigated. This requires the definition of a covariance metric—an inner product on linearized shape variations. The resulting covariance operator robustly captures strongly nonlinear geometric variations in a physically meaningful way and allows to extract the dominant modes of shape variation. The underlying elasticity concept represents an alternative to Riemannian shape statistics. In this paper we compare a standard L^2 -type covariance metric with a metric based on the Hessian of the nonlinear elastic energy. Furthermore, we explore the dependence of the principal component analysis on the type of the underlying nonlinear elasticity. For the built-in pairwise elastic registration, a relaxed model formulation is employed which allows for a non-exact matching. Shape contours are approximated by single well phase fields, which enables an extension of the method to a covariance analysis of image morphologies. The model is implemented with multilinear finite elements embedded in a multi-scale approach. The characteristics of the approach are demonstrated on a number of illustrative and real world examples in 2D and 3D.

Keywords shape analysis · principal components · non-rigid registration · nonlinear elasticity · covariance

metric · phase field approximation · finite element discretization

1 Introduction

During the past decade, the problem of finding the average and the analysis of principal modes of variation on a given number of shapes has attracted much attention in anatomy research, where anatomical atlases are constructed from data sets of different subjects [38, 5, 25, 43, 30, 4]. In this paper we will investigate a notion of shape averages and principal modes of shape variation based on concepts from continuum mechanics, namely nonlinear and linearized elasticity. As shapes we consider object contours, encoded as edge sets in images. Compared to a classical principal component analysis in a vector space, where an average and a covariance tensor can be computed directly on the linear space itself, in the case of shapes we are dealing with highly nonlinear geometric variations. Thus, we aim for a suitable interpretation of actually nonlinear geometric variations of shapes in terms of vector fields in a linear tangent space to the manifold of shapes. Indeed the space of stresses on an averaged shape or an induced space of displacement fields turn out to be suitable and mechanically sound spaces for such an analysis of dominant modes of shape variation. As a first motivation for the model to be developed in this paper, let us mention a simple physical interpretation of the arithmetic mean and the covariance tensor for n points x_1, \dots, x_n in \mathbb{R}^d . Indeed, the arithmetic mean $x \in \mathbb{R}^d$ minimizes $\sum_{i=1, \dots, n} \alpha d(x, x_i)^2$, where $d(x, x_i)$ is the distance between x and x_i , and due to Hooke's law the stored elastic energy $\alpha d(x, x_i)^2$ in the spring connecting x_i and x is proportional to the squared distance (where α repre-

sents half the spring constant). Hence, the arithmetic mean minimizes the total elastic energy of the system of connected springs. Likewise, the covariance tensor ($\langle x_i - x, x_j - x \rangle$) can—up to the spring constant—be identified with the covariance tensor ($\langle \sigma_i, \sigma_j \rangle$) of the forces σ_i pulling at the mean x . Here, we will generalize this mechanical interpretation of point averaging and dominant modes of point variation to the case of shapes in 2D and 3D.

For the zero moment analysis, i.e. the definition of a suitable shape average, we ask for minimizers of a total elastic energy on a set of in general nonlinear elastic deformations from the input shapes onto a single image shape. At the energy minimum the corresponding image shape is defined as the shape average. We note that this zero moment analysis has been the subject of a previous publication [40] and is therefore only briefly summarized in Section 3.

Concerning the first moment analysis, we pick up the above-mentioned linear vector space approach and propose a mechanically sound linearization of shape variations in terms of stresses and induced displacement fields. As a next step we compare different notions of a metric on the linear vector space which we apply to define a covariance tensor. Indeed, each input shape is associated with an elastically optimal deformation onto the shape average. This deformation induces stresses on the shape average, which can be regarded as the imprint of the input shape. Modulating these stresses leads to displacements on the shape average, where the mapping from stresses to displacements is linear and locally well-defined (for larger stress modulations there might however exist multiple equilibrium displacements). Each of these displacements can be regarded as a linearization of the usually nonlinear elastic deformation from one of the image shapes onto the shape average. Despite being a vector field and thus by itself a linear object, both the stresses and the induced displacements encode nonlinear geometric variations of the average shape. To obtain a mechanically meaningful interpretation of these variations, it is important to take into account the fact that the average shape as an image of different elastic deformations of different input shapes is not in a stress-free state. This significantly influences displacement response to imposed stresses.

Given the set of resulting displacements, a covariance tensor can be computed for a prescribed inner product, i.e. a metric, on the space of displacements of the shape average. This tensor linearly encodes all modes of variation of the shape average induced by the set of input shapes. The metric, via which the covariance tensor is defined, can be chosen in various ways. A standard L^2 -metric pronounces shape variations with large displace-

ments even though they are energetically cheap (e.g. rotation of some structure around a joint such as the right arm in Fig. 1 rotating around the shoulder which only produces little deformation energy near the shoulder, cf. also Fig. 2), while the Hessian of the nonlinear elastic energy serves as the appropriate inner product so as to measure distances between displacements solely based on the associated change of elastic energy. Indeed, the Hessian represents an averaged linearized elasticity tensor at the deformed configuration. Thus, displacements in regions and directions which are significantly loaded are weighted strongly, which is mechanically sound. Finally, given the metric and the covariance tensor, we perform a principal component analysis (PCA), which allows to identify the (mechanically) dominant modes of variation of the input shapes.

2 Review of related work

Let us briefly review some of the main contributions on zero and first moment shape analysis. At first, shape analysis was mainly based on correspondences between landmark positions on different shapes [24, 36]. Principal component analysis (PCA) is a classical, by definition linear statistical tool. Chalmond and Girard [7] have proposed a PCA which incorporates also truly nonlinear geometric transformations. Dominant eigenmodes of landmark displacement on human organs have been investigated by Söhn et al. in [41]. A survey on the potential of shape analysis in brain imaging is given by Faugeras and coworkers in [19]. Another important application concerns ready-made clothing, where it would be favorable to know the shape of the average human body and its principal modes of variation to design clothes which sufficiently fit as many people as possible.

Conceptually, correlations of shapes have been studied on the basis of a general framework of a space of shapes and its intrinsic structure. The notion of a shape space was introduced by Kendall [27] already in 1984. Charpiat et al. [11, 10] discuss shape averaging and shape statistics based on the notion of the Hausdorff distance of sets. They propose to use smooth approximations of the Hausdorff distance based on the difference of signed distance functions. Applying a singular value decomposition of the covariance matrix in a suitable scalar product they have additionally analyzed the dominant modes of shape variations. Statistics on signed distance functions were also studied by Leventon et al. [29], whereas Dambreville et al. [17] used shape statistics based on characteristic functions to define a robust shape prior in image segmentation. Ker-

nel density estimation in feature space was introduced by Cremers et al. [15] to incorporate the probability of 2D silhouettes of 3D objects in 2D image segmentation. An overview on related kernel density methods is given by Rathi et al. [37]. An abstract distance measure between objects in (different) metric spaces is the Gromov–Hausdorff distance, which allows to compute an isometrically invariant distance measure between shapes. It measures globally and based on an L^∞ -type functional the lack of isometry between two different shapes. Mémoli and Sapiro [31] have proposed an efficient method to compute distances between point clouds. Bronstein et al. incorporate the Gromov–Hausdorff distance concept in various classification and modeling approaches in geometry processing [6]. In contrast to these global measures of the defect from an isometry, the nonlinear elastic energy functional involved in our approach measures this defect locally, and locally isometric deformations indeed minimize the corresponding local functional.

Understanding shape space as an infinite-dimensional Riemannian manifold has been studied extensively by Miller et al. [32,33]. On a Riemannian manifold \mathcal{M} , all infinitesimal variations of a point $x \in \mathcal{M}$ span the so-called tangent space $\mathcal{T}_x\mathcal{M}$ to \mathcal{M} at x . The set of all tangent spaces $\mathcal{T}_x\mathcal{M}$, $x \in \mathcal{M}$, is denoted the tangent bundle, and each tangent space $\mathcal{T}_x\mathcal{M}$ is equipped with an inner product $g_x(\cdot, \cdot)$, the so-called Riemannian metric. This Riemannian metric then induces a geodesic distance $d(\cdot, \cdot)$ on \mathcal{M} . The metric $d(\cdot, \cdot)$ on the shape space by Miller et al. is defined as the geodesic distance based on a Riemannian metric on the tangent bundle of transport vector fields generating deformations of shapes. Hence, the corresponding shape average is a generalization of the geometric mean for points on a Riemannian manifold proposed by Karcher [26]. Fuchs et al. [22] proposed a viscoelastic notion of the distance between shapes \mathcal{S} given as boundaries of physical objects \mathcal{O} . They define a metric on infinitesimal variations of a shape \mathcal{S} via the dissipation induced by the corresponding normal displacement on the object boundary. The method is applied to compute geodesics between two-dimensional shapes using finite elements on a given triangulation of one of the two objects bounded by the corresponding shape contour. Physically, they suppose that due to the viscous dissipation, the object material relaxes instantaneously and remains in a stress-free state. This is a major difference to our approach, where the objects in the deformed configurations bear elastic stresses and the averaged shape \mathcal{S} is characterized by a true balance of these stresses on the shape contour. The elasticity paradigm for shape analysis on which our approach is founded differs significantly from these metric

approaches to shape space (cf. Sec. 6 for a detailed discussion of the conceptual difference).

Our method inherently establishes a correspondence between points on different shapes, which is implicitly defined by energy-minimizing matching deformations between the shapes. This represents a substantial difference to approaches in which shapes contain or are composed of landmarks and in which there is thus an explicit point correspondence, as e. g. in [27,14,24,28]. Such point correspondences have to be obtained in an additional preprocessing step and might be desirable if an exact matching between manually placed landmarks is required. However, as this is typically not the case, the consistent placement of landmarks often involves a shape registration [36,41]. Then it would of course be more natural to use the full shape correspondence from the registration (as it is done in our approach) than to select only a discrete number of landmark correspondences and neglect the remaining information.

Furthermore, since shapes are often encoded in images or volume data, the averaging procedure is inherently associated with the detection or segmentation of structures. An enormous body of literature on segmentation refers to the seminal paper by Mumford and Shah [35] on a variational approach which states the segmentation problem as a free discontinuity problem for the edge set of an image. The theoretic examination of this variational problem is quite elaborate [34,18,16]. A very successful computational approach by Chan and Vese describes the segments by level sets [8,9]. Another widely used approach due to Ambrosio and Tortorelli [1] encodes the segment boundaries as a diffused edge set represented by a so-called phase field function.

In this paper, shapes are represented implicitly via a diffused phase field description. This description is naturally obtained from a Mumford–Shah segmentation of input images via the Ambrosio–Tortorelli approximation, and it will allow a natural extension of the method from shapes to general segmented images or image morphologies (cf. Sec. 8). This enables a robust and flexible application in two and three dimensions.

3 Elastic shape average revisited

In this section we briefly recall an elastic approach to shape averaging already presented in [40]. We consider shapes \mathcal{S}_i as the boundaries $\partial\mathcal{O}_i$ of sufficiently regular (e. g. Lipschitz) objects \mathcal{O}_i . These objects will be encoded as characteristic functions in binary images $I_i : \Omega \rightarrow \mathbb{R}$, defined on some domain $\Omega \subset \mathbb{R}^d$, so that the shapes \mathcal{S}_i can be regarded as the edge sets of these images.

Given a set of shapes, $\mathcal{S}_1, \dots, \mathcal{S}_n$, we seek an average shape \mathcal{S} that reflects the geometric characteristics of the given shapes in a physically intuitive manner. For that purpose it seems generic to interpret the different shapes \mathcal{S}_i and corresponding objects \mathcal{O}_i as deformed configurations of each other. Then, the average shape \mathcal{S} clearly can also be described as a deformed configuration of the input shapes, i.e. there are deformations $\phi_i : \mathcal{O}_i \rightarrow \mathbb{R}^d$, $i = 1, \dots, n$, with $\mathcal{S} = \phi_i(\mathcal{S}_i)$ (see Fig. 1). As corresponding average object we obtain $\mathcal{O} = \phi_i(\mathcal{O}_i)$. A natural choice for the definition of the shape average \mathcal{S} then is given by that particular shape which minimizes the total accumulated deformation energy of all deformations,

$$\mathcal{E}[\mathcal{S}, (\phi_i)_{i=1, \dots, n}] = \begin{cases} \frac{1}{n} \sum_{i=1}^n \mathcal{W}[\mathcal{O}_i, \phi_i] ; \phi_i(\mathcal{S}_i) = \mathcal{S}, \\ \infty & ; \text{else} \end{cases},$$

where $\mathcal{W}[\mathcal{O}_i, \phi_i]$ represents the stored deformation energy of the deformation ϕ_i . In principle, $\mathcal{W}[\mathcal{O}_i, \phi_i]$ can admit various forms such as e.g. linear or nonlinear elastic deformation energy.

The form of energy $\mathcal{E}[\mathcal{S}, (\phi_i)_{i=1, \dots, n}]$ might however not be sufficient to ensure a certain regularity of the minimizing shape \mathcal{S} . In particular, it is unclear whether we can expect \mathcal{S} to have a finite perimeter since standard regularity theory only allows to prove Lipschitz continuity of deformations ϕ_i under certain strong conditions [12]. Hence, we add a regularizing prior $\mathcal{L}[\mathcal{S}]$ to the energy. For the actual implementation, we consider the \mathcal{H}^{d-1} -measure of \mathcal{S} as the prior $\mathcal{L}[\mathcal{S}]$, i.e. $\mathcal{L}[\mathcal{S}] = \int_{\mathcal{S}} da$. Finally the shape average \mathcal{S} is defined as a minimizer of the energy $\mathcal{E}[\mathcal{S}, (\phi_i)_{i=1, \dots, n}] + \mu \mathcal{L}[\mathcal{S}]$ for some small $\mu > 0$. Note that the prior induces a bias towards smaller perimeters, however, since it has been introduced mainly for theoretical purposes, μ may be chosen very small.

As deformation energy $\mathcal{W}[\mathcal{O}_i, \phi_i]$ we will employ a nonlinear, hyperelastic energy $\mathcal{W}[\mathcal{O}, \phi] = \int_{\mathcal{O}} W(\mathcal{D}\phi) dx$. By the fundamental axiom of elasticity, the integrand only depends on the deformation gradient $\mathcal{D}\phi$. From the frame indifference principle and under an isotropy assumption, W can be rewritten as a function of only the three invariants

$$W(\mathcal{D}\phi) = \hat{W}(\mathcal{D}\phi, \text{cof}\mathcal{D}\phi, \det(\mathcal{D}\phi)) = \bar{W}(I_1, I_2, I_3)$$

with $(I_1, I_2, I_3) := (|\mathcal{D}\phi|_2^2, |\text{cof}\mathcal{D}\phi|_2^2, \det(\mathcal{D}\phi))$, where $\det(\mathcal{D}\phi)$ describes the local volume change, $|\mathcal{D}\phi|_2 := \sqrt{\text{tr}(\mathcal{D}\phi^T \mathcal{D}\phi)}$ the averaged local change of length, and $|\text{cof}(\mathcal{D}\phi)|_2$ the averaged local change of area. Here, for an invertible matrix $A \in \mathbb{R}^{d,d}$, $\text{cof}A := (\det A)A^{-T}$ denotes its cofactor matrix. The first Piola–Kirchhoff stress tensor, which describes force per unit area in the

reference configuration \mathcal{O} , is then recovered as $\sigma^{\text{ref}} = W_{,A}(\mathcal{D}\phi)$, where $W_{,A}(A) = \frac{\partial W(A)}{\partial A}$. The Cauchy stress, describing force per unit area in the deformed configuration $\phi(\mathcal{O})$, reads $\sigma = \sigma^{\text{ref}}(\text{cof}\mathcal{D}\phi)^{-1}$. Here, we consider so-called polyconvex energy functionals [13], where \hat{W} is supposed to be convex, and isometries, i.e. deformations with $\mathcal{D}\phi^T(x)\mathcal{D}\phi(x) = \mathbb{1}$, are local minimizers (cf. Fig. 2). Typical energy densities in this class are of the form $\bar{W}(I_1, I_2, I_3) = \alpha_1 I_1^{\frac{p}{2}} + \alpha_2 I_2^{\frac{q}{2}} + \Gamma(I_3)$ with $p > 0$, $q \geq 0$, $\alpha_1, \alpha_2 > 0$, and Γ convex with $\Gamma \xrightarrow{I_3 \rightarrow 0} \infty$. The penalization of volume shrinkage, $\bar{W} \xrightarrow{I_3 \rightarrow 0} \infty$, enables us to control local injectivity (cf. [3]). This type of energy has two major advantages: it allows to incorporate large deformations with strong material and geometric nonlinearities, and the form of the energy follows from first principles and allows to distinguish the physical effects of length, area, and volume distortion, which reflect the local distance from an isometry.

Let us emphasize that due to

$$\mathcal{W}[\mathcal{O}_i, \phi_i] = \mathcal{W}[\psi(\mathcal{O}_i), \phi_i \circ \psi^{-1}]$$

for any rigid body motion ψ , the average shape is invariant to rotation and translation of any of the input shapes.

Note that instead of minimizing the accumulated energy of the deformations ϕ_i from the objects \mathcal{O}_i onto \mathcal{O} , one might want to consider the inverse deformations $\psi_i = \phi_i^{-1}$ and define the shape average as the minimizer of $\frac{1}{n} \sum_{i=1}^n \mathcal{W}[\mathcal{O}, \psi_i]$ under the constraints $\psi_i(\mathcal{S}) = \mathcal{S}_i$. Which view is more appropriate certainly depends on the application. However, note that $\det \mathcal{D}\psi_i = \frac{1}{\det \mathcal{D}\phi_i}$, $\mathcal{D}\psi_i = \frac{\text{cof}\mathcal{D}\phi_i^T}{\det \mathcal{D}\phi_i}$, and $\text{cof}\mathcal{D}\psi_i = \frac{\mathcal{D}\phi_i^T}{\det \mathcal{D}\phi_i}$ so that $\mathcal{W}[\mathcal{O}, \psi_i] = \int_{\mathcal{O}} \hat{W}(\mathcal{D}\psi_i, \text{cof}\mathcal{D}\psi_i, \det \mathcal{D}\psi_i) dx$ can be rephrased as a particular elastic energy

$$\widetilde{\mathcal{W}}[\mathcal{O}_i, \phi_i] = \int_{\mathcal{O}_i} |\det \mathcal{D}\phi_i| \hat{W}\left(\frac{\text{cof}\mathcal{D}\phi_i^T}{\det \mathcal{D}\phi_i}, \frac{\mathcal{D}\phi_i^T}{\det \mathcal{D}\phi_i}, \frac{1}{\det \mathcal{D}\phi_i}\right) dx$$

of ϕ_i , which leads us back to the original approach. For details we refer to [40].

To simplify numerical treatment and to allow for slight topological differences between the shapes \mathcal{S}_i , which might e.g. result from noise in the shape acquisition, we relax the constraint $\phi_i(\mathcal{S}_i) = \mathcal{S}$, $i = 1, \dots, n$, and introduce a penalty functional

$$\mathcal{F}[\mathcal{S}_i, \phi_i, \mathcal{S}] = \mathcal{H}^{d-1}(\mathcal{S}_i \setminus \phi_i^{-1}(\mathcal{S}) \cup \phi_i^{-1}(\mathcal{S}) \setminus \mathcal{S}_i)$$

which measures the symmetric difference of the input shapes \mathcal{S}_i and the pull back $\phi_i^{-1}(\mathcal{S})$ of \mathcal{S} . Note that for properly chosen deformation energies, deformations ϕ_i with finite energy are homeomorphisms [3] so that ϕ_i^{-1} is well-defined. The numerical computation of the inverse deformation will however not be required since

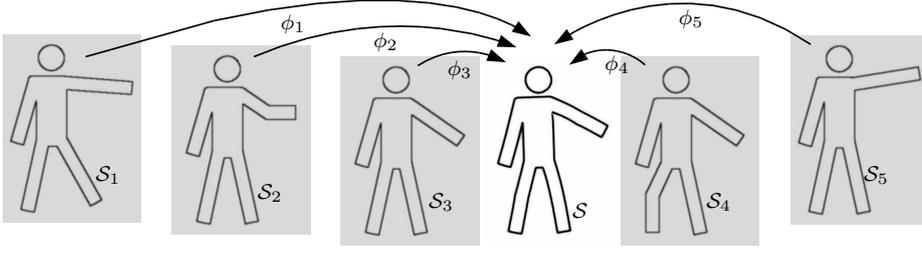


Fig. 1 Sketch of elastic shape averaging. The input shapes \mathcal{S}_i ($i = 1, \dots, 4$) are mapped onto a shape \mathcal{S} via elastic deformations ϕ_i . The shape \mathcal{S} which minimizes the elastic deformation energy is denoted the shape average.

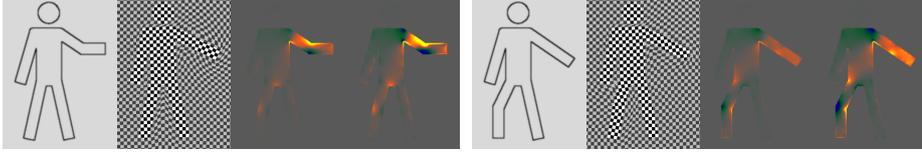


Fig. 2 For two input shapes from Fig. 1 the deformation (via a deformed checkerboard), the averaged local change of length $\frac{1}{\sqrt{2}} |\mathcal{D}\phi_i|_2$, and the local change of area $\det(\mathcal{D}\phi_i)$ are depicted (colors encode range $[0.95, 1.05]$).

$\phi_i^{-1}(\mathcal{S})$ will be represented by the concatenation of a phase field v (which describes \mathcal{S}) with ϕ_i (cf. Sec. 7). Our shape averaging model is thus based on the energy

$$\begin{aligned} \mathcal{E}^\gamma[\mathcal{S}, (\phi_i)_{i=1, \dots, n}] \\ = \frac{1}{n} \sum_{i=1}^n \left(\int_{\mathcal{O}_i} W(\mathcal{D}\phi_i) dx + \gamma \mathcal{F}[\mathcal{S}_i, \phi_i, \mathcal{S}] \right) + \mu \mathcal{L}[\mathcal{S}]. \end{aligned}$$

4 Linearization of shape variations

As outlined in the introduction, our first moment analysis on shapes is based on an analysis of stresses induced on the shape average by each individual input shape. Modulation of each of these stresses results in a certain displacement, and the proposed principal component analysis on shapes will be performed on these displacements. To comprehensively derive this model we proceed in several steps:

Boundary stresses as shape representations. Let us at first review the underlying physical concept of stress. By the Cauchy stress principle, each deformation $\phi_i : \mathcal{O}_i \rightarrow \mathcal{O}$ is characterized by pointwise boundary stresses on \mathcal{S} in the deformed configuration. The stress at some point x on \mathcal{S} is given by the application of the Cauchy stress tensor σ_i to the outer normal ν on \mathcal{S} . The resulting stress $\sigma_i \nu$ is a force density acting on a local surface element of \mathcal{S} . The shape \mathcal{S} is in an equilibrium configuration if the opposite force is applied as an external surface load. Otherwise, by the axiom of elasticity, releasing the object \mathcal{O} , the elastic body will snap back to the original reference configuration of the input object \mathcal{O}_i . Let us assume that the above relation

between the energetically favorable deformation and its induced stresses is one-to-one. Hence, the average shape can be described in terms of the input shape \mathcal{S}_i and the boundary stress $\sigma_i \nu$, and we write $\mathcal{S} = \mathcal{S}_i[\sigma_i \nu]$. If we now scale the stress with a weight $t \in [0, 1]$, we obtain a one-parameter family of shapes $\mathcal{S}(t) = \mathcal{S}_i[t\sigma_i \nu]$ connecting $\mathcal{S}_i = \mathcal{S}(0)$ with $\mathcal{S} = \mathcal{S}(1)$. Thus, we can regard $\sigma_i \nu$ as a representative of shape \mathcal{S}_i in the linear space of vector fields on \mathcal{S} .

Zero mean of shape representations. Let us denote by ν_i the outer normal on \mathcal{S}_i . From elasticity theory we know that the force $\sigma_i(x) \nu(x) da[\mathcal{S}](x)$ on the surface element $da[\mathcal{S}]$ at some point x in the deformed configuration equals the force $\sigma_i^{\text{ref}}(x^{\text{ref}}) \nu_i(x^{\text{ref}}) da[\mathcal{S}_i](x^{\text{ref}})$ at $x^{\text{ref}} = \phi_i^{-1}(x)$ in the reference configuration of the input shape, \mathcal{S}_i . Here, we use the above transformation rule for the Piola–Kirchhoff and the Cauchy stress and the purely geometric observation that $\nu da[\mathcal{S}] = \text{cof} \mathcal{D}\phi_i \nu_i da[\mathcal{S}_i]$. This will allow us to effectively treat the stress induced by a deformation on the given input shapes.

Due to the set of constraints $(\phi_i(\mathcal{S}_i) = \mathcal{S})_{i=1, \dots, n}$, the Euler Lagrange conditions on the different objects \mathcal{O}_i and shapes \mathcal{S}_i are interlinked: Consider a small variation of the average shape \mathcal{S} by some displacement u according to $(\mathbb{1} + \epsilon u)(\mathcal{S})$ and correspondingly varied deformations $(\mathbb{1} + \epsilon u) \circ \phi_i$. The optimality implies

$$\frac{d}{d\epsilon} \mathcal{E}[(\mathbb{1} + \epsilon u)(\mathcal{S}), ((\mathbb{1} + \epsilon u) \circ \phi_i)_{i=1, \dots, n}] = 0. \quad (1)$$

After a brief calculation (given in appendix A.1), this results not only in the classical system of partial differential equations of elasticity, $\text{div} W_{,A}(\mathcal{D}\phi_i) = 0$, which holds on every domain \mathcal{O}_i , but additionally in a stress

balance

$$0 = \sum_{i=1, \dots, n} \sigma_i(x) \nu(x) \quad (2)$$

for every $x \in \mathcal{S}$ (cf. Fig. 3). Here, we have assumed that \mathcal{S} is the outer boundary of the object domain \mathcal{O} and that there are no interior interfaces. A generalization to interior interfaces as shape components is straightforward, but requires the proper handling of jumps of stresses. The above stress balance (2) can be regarded as a pointwise physical interpretation of the shape average as the stable shape on which all Cauchy stresses $\sigma_i \nu$, induced by the deformations ϕ_i for $i = 1, \dots, n$, are balanced.

Displacements representing shape variations. Let us now study how the average shape \mathcal{S} varies if we increase the impact of a particular input shape \mathcal{S}_k for some $k \in \{1, \dots, n\}$. In fact, we intend to associate to every surface load $\sigma_k \nu$ a displacement on the averaged object domain \mathcal{O} via the solution operator of a suitable linearized elasticity problem. Here, the object \mathcal{O} actually is a deformed configuration of different original objects \mathcal{O}_i . Hence, we have to choose a proper elasticity tensor which reflects the compound stress configuration of the averaged domain \mathcal{O} . A simple isotropic linearized elasticity model would not take into account the nonlinear geometric nature of our zero and first order analysis. To achieve this, we apply the Cauchy stress $\sigma_k \nu$ to the average shape \mathcal{S} , scaled with a small constant δ . Based on our above discussion of stresses and due to the sketched equilibrium condition, this additional boundary stress $\delta \sigma_k \nu$ acts as a first Piola–Kirchhoff stress on the (reference) configuration \mathcal{S} . The elastic response is given by a correspondingly scaled displacement $u_k : \mathcal{O} \rightarrow \mathbb{R}^d$. These displacements u_k are considered as representatives of the variation of the average shape \mathcal{S} with respect to the input shape \mathcal{S}_k given as functions on the fixed compound object \mathcal{O} . The space of these displacements is a linear vector space, which will give rise to the actual covariance analysis on the set of displacements u_1, \dots, u_n in Section 5.

Computation of the u_k . To properly model the loaded configurations we concatenate this displacement with every nonlinear deformation ϕ_i and take into account the sum of the resulting elastic energies plus a term involving the given Cauchy stress in the following energy,

$$\mathcal{E}_k[\delta, u] = \frac{1}{n} \sum_{i=1}^n \mathcal{W}[\mathcal{O}_i, (\mathbb{1} + \delta u) \circ \phi_i] - \delta^2 \int_{\mathcal{S}} \sigma_k \nu \cdot u \, da \quad (3)$$

Now, the displacement u_k is obtained as a minimizer of this modulated energy for a fixed set of deformations $(\phi_i)_{i=1, \dots, n}$ under the constraints $\int_{\mathcal{O}} u_k \, dx = 0$

and $\int_{\mathcal{O}} x \times u_k \, dx = 0$, which encode zero average translation and rotation. Indeed, as the corresponding Euler–Lagrange condition for u_k we obtain $\operatorname{div} \sigma[\delta u_k] = 0$ on \mathcal{O} and $\sigma[\delta u_k] \nu = \delta \sigma_k \nu$ on \mathcal{S} (see derivation in appendix A.2), where we denote by

$$\sigma[\delta u_k] := \frac{1}{n} \sum_{i=1}^n W_{,A}((\mathbb{1} + \delta \mathcal{D}u_k) \mathcal{D}\phi_i \circ \phi_i^{-1}) \operatorname{cof} \mathcal{D}(\phi_i^{-1})$$

the first Piola–Kirchhoff stress tensor on the compound object \mathcal{O} , which effectively reflects an average of all stresses in the n deformed configurations $\phi_i(\mathcal{O}_i)$ for $i = 1, \dots, n$. This becomes obvious, noting that

$$\begin{aligned} & \frac{1}{\det \mathcal{D}(\mathbb{1} + \delta u_k)} \sigma[\delta u_k] \mathcal{D}(\mathbb{1} + \delta u_k)^T \\ &= \left(\frac{1}{n} \sum_{i=1}^n \frac{W_{,A}(\mathcal{D}((\mathbb{1} + \delta u_k) \circ \phi_i)) \mathcal{D}((\mathbb{1} + \delta u_k) \circ \phi_i)^T}{\det \mathcal{D}((\mathbb{1} + \delta u_k) \circ \phi_i)} \right) \circ \phi_i^{-1} \\ &= \left(\frac{1}{n} \sum_{i=1}^n \sigma_i[\delta u_k] \right) \circ (\mathbb{1} + \delta u_k) \quad , \end{aligned}$$

where the $\sigma_i[\delta u_k]$ are indeed the Cauchy stresses of the different objects \mathcal{O}_i when deformed into $(\mathbb{1} + \delta u_k)(\mathcal{O}) = ((\mathbb{1} + \delta u_k) \circ \phi_i)(\mathcal{O}_i)$. Let us remark that the boundary integral in (3) can be replaced by the volume integral $\int_{\mathcal{O}} \sigma_k : \mathcal{D}u \, dx$, which is more convenient with respect to a numerical discretization. To verify this, we use integration by parts and the fact that $\operatorname{div} \sigma_k = 0$ holds on \mathcal{O} .

Linearization for infinitesimal variations. As long as $A \mapsto W(A)$ is not quadratic in A , u_k still solves a nonlinear elastic problem. The advantage of this nonlinear variational formulation is that it is of the same type as the one for the zero moment analysis, and it encodes in a natural way the compound elasticity configuration of the averaged shape domain \mathcal{O} . As an obvious drawback we have to consider the sum of n nonlinear elastic energies for the computation of every displacement u_k , $k = 1, \dots, n$. To avoid this shortcoming, we consider the limit of the Euler–Lagrange equations for $\delta \rightarrow 0$, and—with a slight misuse of notation—obtain u_k as the solution of the actually linear elasticity problem

$$\operatorname{div}(\mathbf{C} \epsilon[u]) = 0 \text{ in } \mathcal{O}, \quad \mathbf{C} \epsilon[u] \nu = \sigma_k \nu \text{ on } \mathcal{S} \quad (4)$$

for the symmetric displacement gradient $\epsilon[u] = (\mathcal{D}u + \mathcal{D}u^T)/2$ under the constraints $\int_{\mathcal{O}} u \, dx = 0$ and $\int_{\mathcal{O}} x \times u \, dx = 0$. Here, the in general inhomogeneous and anisotropic elasticity tensor \mathbf{C} is derived in appendix A.3 as

$$\mathbf{C} = \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\det \mathcal{D}\phi_i} \mathcal{D}\phi_i W_{,AA}[\mathcal{D}\phi_i] \mathcal{D}\phi_i^T \right) \circ \phi_i^{-1} \quad ,$$

based on an appropriate transformation of the Hessian of the energy density W . This elasticity tensor takes

into account the loads of the compound configuration based on the combination of all deformations ϕ_i on the input objects \mathcal{O}_i for $i = 1, \dots, n$. For implementation, however, it is simpler to implement the bilinear operator

$$(u, v) \mapsto \int_{\mathcal{O}} \mathbf{C}\epsilon[u] : \epsilon[v] \, dx \\ = \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{O}_i} \langle W_{,AA}[\mathcal{D}\phi_i], \mathcal{D}(u \circ \phi_i), \mathcal{D}(v \circ \phi_i) \rangle \, dx$$

directly without first applying the transformation rule as detailed in Section 7. The resulting u_k will be used as linearizations of shape variations for the covariance analysis in the next section. Since the boundary stresses σ_k are independent of the original position of object \mathcal{O}_k , the covariance analysis is inherently invariant to rotation and translation of any input shape. However, scale invariance (as e. g. in the approach by Kendall [27]) cannot be incorporated and does not make much sense in an approach based on elastic deformations. However, if required, one might scale the input objects to have equal volume in a preprocessing step.

5 Covariance analysis on shapes

In order to perform a principal component analysis on a set of shapes $\mathcal{S}_1, \dots, \mathcal{S}_n$ we sought for representatives of the shapes in a linear vector space and chose the above displacements $u_k : \mathcal{O} \rightarrow \mathbb{R}^d$, $k = 1, \dots, n$, as these representatives. As explained above, these displacements reflect the variations of the average shape induced by a modulation of the stresses σ_k from the deformations ϕ_k of the input shapes \mathcal{S}_k into the average shape \mathcal{S} . Let us emphasize that the stresses already have zero mean due to (2). Hence, the same holds true for the displacements, i. e.

$$\sum_{k=1}^n u_k = 0.$$

Now, we will define a covariance operator, which requires to select a suitable inner product (metric) $g(u, \tilde{u})$ on displacements $u, \tilde{u} : \mathcal{O} \rightarrow \mathbb{R}^d$.

Let us remark, that g induces a metric $\tilde{g}(\sigma\nu, \tilde{\sigma}\nu) := g(u, \tilde{u})$ on the associated boundary stress. Hence, the covariance analysis presented here can be considered as a corresponding analysis directly on boundary stresses $\sigma_1\nu, \dots, \sigma_n\nu$. Indeed, we use that the solvability condition

$$\int_{\mathcal{O}} \operatorname{div}(\mathbf{C}\nabla u) \, dx = 0 = \int_{\mathcal{S}} \mathbf{C}\nabla u \nu \, da[\mathcal{S}]$$

is fulfilled. Thus, the solution u_k for given boundary stress $\sigma_k\nu = \mathbf{C}\nabla u_k$ is uniquely determined up to a

linearized rigid body motion (i. e. an affine displacement with skew-symmetric matrix representation), which is fixed by the conditions of zero mean displacement and angular momentum for u . Due to the linearity of the operator $\sigma\nu \mapsto u$ the metric \tilde{g} is bilinear and symmetric as well. Finally, the positive definiteness of \tilde{g} follows from the positive definiteness of g and the injectivity of the map $\sigma\nu \mapsto u$.

In this paper, we consider two different inner products:

- *The L^2 -product.* Given two square integrable displacements u, \tilde{u} we define

$$g(u, \tilde{u}) := \int_{\mathcal{O}} u \cdot \tilde{u} \, dx$$

This product weights local displacements equally on the whole compound object \mathcal{O} .

- *The Hessian of the energy as inner product.* Different from the L^2 -metric, we now measure displacement gradients in a non-homogeneous way. In fact, we define

$$g(u, \tilde{u}) := \int_{\mathcal{O}} \mathbf{C}\epsilon[u] : \epsilon[\tilde{u}] \, dx$$

for displacements u, \tilde{u} with square integrable gradients. Hence, the contribution to the inner product is larger in areas of the compound object which are in a significantly stressed configuration.

Once an inner product is given, we can define the covariance operator \mathbf{Cov} by

$$\mathbf{Cov} u := \frac{1}{n} \sum_{k=1}^n g(u, u_k) u_k.$$

Obviously, \mathbf{Cov} is positive definite on $\operatorname{span}(u_1, \dots, u_n)$. Hence, we can diagonalize \mathbf{Cov} on this finite-dimensional space and obtain a set of g -orthogonal eigenfunctions $w_k : \mathcal{O} \rightarrow \mathbb{R}^d$ —actually again displacements—and eigenvalues $\lambda_k > 0$ with

$$\mathbf{Cov} w_k = \lambda_k w_k.$$

These eigenfunctions can be considered as the principal modes of variation of the average object \mathcal{O} and hence of the average shape \mathcal{S} , given the n input shapes $\mathcal{S}_1, \dots, \mathcal{S}_n$. The eigenvalues encode the actual strength of these variations. The resulting modes of variation can easily be visualized via a scalar modulation of the δw_k for varying values of δ , which is associated with a corresponding modulation of the underlying stresses $\delta \mathbf{C}\nabla w_k \nu$. If an amplified visualization of the modes is required, it is preferable to depict displacements w_δ^k which are defined as minimizers of the nonlinear variational energy

$$\frac{1}{n} \sum_{i=1}^n \mathcal{W}[\mathcal{O}_i, (\mathbb{1} + w) \circ \phi_i] - \delta^2 \int_{\mathcal{S}} \mathbf{C}\nabla w_k \nu \cdot w \, da$$

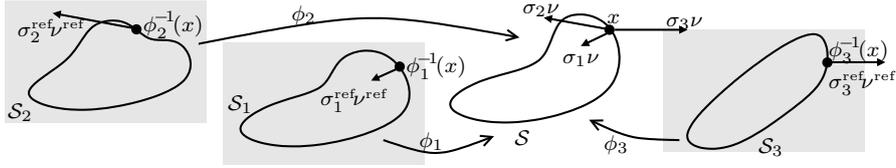


Fig. 3 Sketch of the pointwise stress balance relation on the averaged shape.



Fig. 4 The two dominant modes (right) for four different shapes (left) demonstrate that our principal component analysis properly captures strong geometric nonlinearities.

(cf. (3)). The diagonalization of \mathbf{Cov} can be performed by diagonalizing the symmetric matrix $\frac{1}{n} (g(u_i, u_j))_{ij} = \mathbf{O}\mathbf{A}\mathbf{O}^T$, where $\mathbf{A} = \text{diag}(\lambda_1, \lambda_2, \dots)$ and \mathbf{O} is orthogonal. The eigenfunctions are then obtained as $w_k = \frac{1}{\sqrt{\lambda_k}} \sum_{j=1}^n O_{jk} u_j$.

Let us underline that this covariance analysis properly takes into account the usually strong geometric nonlinearity in shape analysis via the transfer of geometric shape variation to elastic stresses on the average shape, based on paradigms from nonlinear elasticity. This is illustrated in Fig. 4 for the L^2 -metric as underlying inner product. The interpretation of displacements and stresses, respectively, is the proper linearization of shapes. In abstract terms, either the space of displacements or stresses can be considered as the tangent space of shape space at the average shape. The identification of displacements and stresses via (4) provides a suitable physical interpretation of stresses as modes of shape variation.

In Fig. 5 a larger set of 20 binary images “device7” from the MPEG7 shape database serves as input shapes. Apparently, the first principal component is given by a thickening or thinning of the leaves, accompanied by a change of indentation depth between them. The second mode obviously corresponds to bending the leaves, and the third mode represents local changes at the tips: A sharpening and orientation of neighboring tips towards each other, originating e. g. from the sixth or the second last input shape.

The impact of the chosen metric. Naturally, the modes of variation depend on the chosen inner product. We have already mentioned that in order to be physically meaningful, the inner product should be based on the induced displacements of the compound object (which is composed of all deformed input shapes). If instead we would apply the boundary stresses $\sigma_{i\nu}$ to an object which just looks like the average shape, but does not contain the information how strongly the input shapes

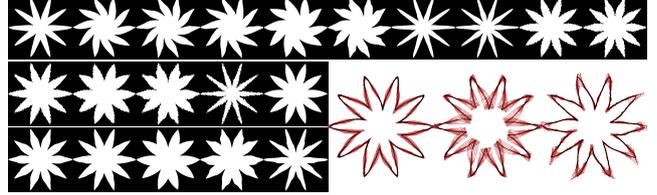


Fig. 5 Original shapes and their first three modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.20, and 0.05.

had to be deformed to arrive at the average, we obtain a different result as shown in Fig. 6. Regions which were more heavily deformed than others need higher stresses to be deformed even further. Therefore, these regions exhibit much stronger variations if the inner product is based on a non-compound shape (which is not already prestressed). Figure 6 furthermore exemplifies the impact of the underlying metric on the obtained displacements. In contrast to the L^2 -metric, the metric defined via the Hessian of the elastic energy captures differences in the deformation in exactly those regions where the deformation takes place. Furthermore, a clearer separation of mechanically separated regions is observed compared to the L^2 -metric.

The impact of the nonlinear elasticity model. Likewise, the particular choice of the nonlinear elastic energy density has a considerable effect on the average shape and its modes of variation. Figure 7 has been obtained using $\bar{W}(I_1, I_2, I_3) = \frac{\mu}{2} I_1 + \frac{\lambda}{4} I_3^2 - (\mu + \frac{\lambda}{2}) \log I_3 - \mu - \frac{\lambda}{4}$, where μ and λ are the factors of length and volume change penalization, respectively, for deformations near the identity. A low penalization of volume changes apparently leads to local compression or inflation at the dumbbell ends (left), while for deformations with a strong volume change penalization (right), material is squeezed from one end to the other, which becomes especially apparent in the second and third mode. Note that in the first mode, the volume of one dumbbell end shrinks while the other increases, whereas for the second mode both ends change equally. This is a very illustrative example of the orthogonality of the decomposition of the tangent space into subspaces according to the different modes of variation. Here, the underlying metric is the based on the Hessian of the energy.

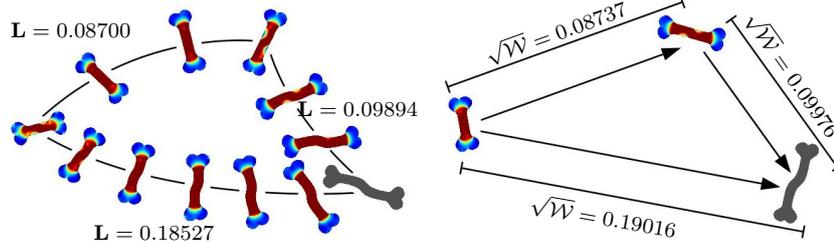


Fig. 8 Left: Viscosity-based geodesics between the shapes at the corners (the shapes are taken from [23], the geodesics are computed with the method from [44]). The triangle inequality holds. The colors encode local viscous dissipation, from blue (low) to red (high). Right: Elastic distances between the same shapes, where the arrows point from the reference to the deformed configuration. The triangle inequality does not hold. The colors encode the local deformation energy density.

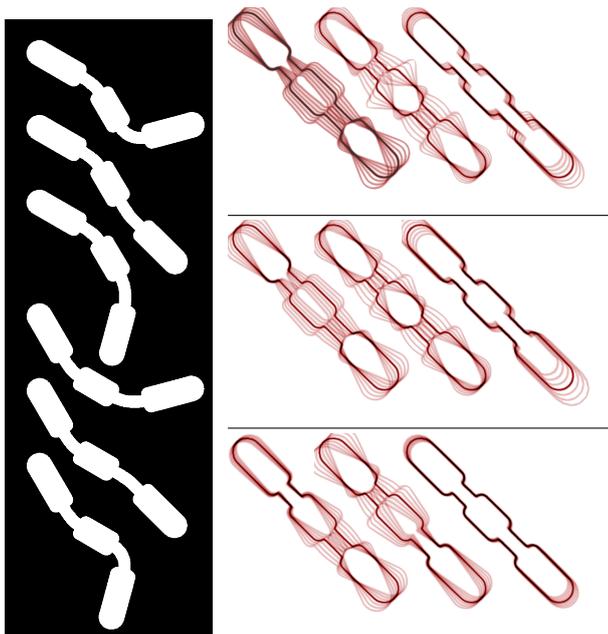


Fig. 6 The first three dominant modes of variation for the six input shapes on the left obviously depend on the employed metric: The top row depicts the modes belonging to the L^2 -metric on displacements of a non-prestressed object (with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.23, and 0.07), the middle row corresponds to the L^2 -metric on displacements of the proper compound object (with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.28, and 0.03), the bottom row represents the results for the energy Hessian based metric on displacements of the compound object (with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.61, and 0.24).

6 Elastic Versus Riemannian Shape Analysis

If one aims at a physical model of shape space in which the corresponding objects are made of a physical material, there is basically the choice between solid, elastic objects (as in our approach) or liquid, viscous objects. The latter case would result in a Riemannian shape space (as e.g. in [45]), and it does of course strongly depend on the application whether an elastic or a viscous-flow perspective is more appropriate. The elasticity paradigm, on which our zero and first order shape analysis are based, differs significantly from

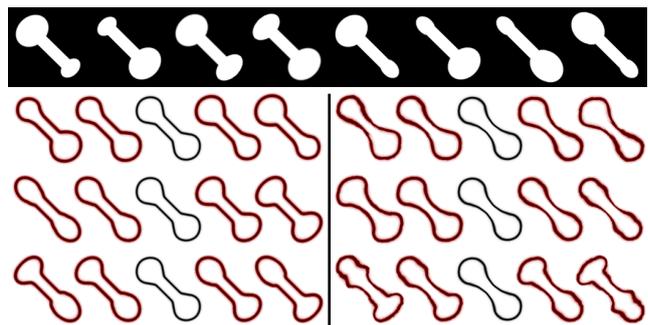


Fig. 7 The first three modes of variation for eight dumbbell shapes, left for a 100 times stronger penalization of length changes than of volume changes (with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.22, and 0.05), right for the reverse (with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.41, and 0.07).

a Riemannian approach to shape space as proposed for instance by Srivastava et al. [42]. Due to the axiom of elasticity, the energy at the deformed configuration \mathcal{S} is independent of the path from a shape $\hat{\mathcal{S}}$ to the shape \mathcal{S} along which the deformation is generated in time. Hence, there is no notion of shortest paths if we consider a purely elastic shape model. The visco-plastic model by Fuchs et al. [22] and the related model by Younes [45] define energies based on an integration of dissipation along transformation paths, where dissipation is understood as a Riemannian metric. This approach is not elastic in the classical axiomatic sense we consider here, and it particularly requires that the intermediate configurations are all stress-free at rest.

The above-mentioned conceptual differences are reflected in a different behavior. In particular, the non-vanishing stresses in the deformed state of an elastic object prevent the elastic energy (or rather its square root) from being a metric: Typically, neither symmetry (the energy to deform \mathcal{O}_A into \mathcal{O}_B is different from the energy to deform \mathcal{O}_B into \mathcal{O}_A) nor the triangle inequality hold (Fig. 8).

If we regard shapes from a flow-oriented perspective, then a visco-elastic approach would be more appropriate. However, the elastic approach is favorable for rather rigid, more stable shapes, since it prevents lo-



Fig. 9 A set of input shapes (cf. Fig. 1) and their modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.22, 0.15, and 0.06.

cally strong isometry violation. The strong local rigidity and isometry preservation of the elasticity concept becomes particularly evident in Fig. 4 and Fig. 9, where non-isometric deformations are concentrated only at joints. This holds true already in the case of an underlying L^2 -metric as inner product as it is taken into account here.

On a Riemannian manifold, the exponential map allows to describe geodesics from an averaged shape \mathcal{S} —in the sense of Karcher [26]—to the input shapes \mathcal{S}_k via $\mathcal{S}_k = \exp_{\mathcal{S}}(v_k)$ for some tangent vector v_k to shape space at \mathcal{S} . Hence, a covariance analysis will be performed on the tangent vectors v_1, \dots, v_n with respect to the Riemannian metric $g(\cdot, \cdot)$. In the strictly elastic setup, the shape space is in general not metrizable. Instead, the stresses σ_k play the role of the v_k , imprinting the impact of \mathcal{S}_k on the average shape \mathcal{S} in terms of an induced displacement u_k .

7 Phase Field Approximation and Finite Element Discretization

Since explicit treatment of an edge set is difficult in a variational setting, we consider a phase field model picking up the approach by Ambrosio and Tortorelli [2] for the discretization of the Mumford–Shah model [35]. Hence, a shape \mathcal{S} is encoded by a smooth phase field function $v : \Omega \rightarrow \mathbb{R}$, which is close to zero on \mathcal{S} and one in between and thus provides a diffuse interface representation. In our approach we construct such phase field functions v_i for the input shapes \mathcal{S}_i in advance. Usually, v_i can be computed based on the model in [2] applied to the input images u_i . The specific form of the phase field function v for the averaged shape \mathcal{S} is then directly determined via a phase field approximation of our variational model for the zero moment analysis. Given a phase field parameter ϵ , which will determine the width of the diffuse edge representation, we first define an approximate mismatch penalty $\mathcal{F}^\epsilon[v_i, \phi_i, v] = \frac{1}{\epsilon} \int_{\Omega} (v \circ \phi_i)^2 (1 - v_i)^2 + v_i^2 (1 - v \circ \phi_i)^2 dx$. Here, we suppose v to be extended by 1 outside the computational domain Ω . Next, we consider the energy $\mathcal{L}^\epsilon[v] = \frac{1}{2} \int_{\Omega} \epsilon |\nabla v|^2 + \frac{1}{4\epsilon} (v - 1)^2 dx$, which acts as an approximation of the prior $\mathcal{L}[\mathcal{S}]$. Furthermore, we simplify the later numerical implementation by assuming that the whole computational domain behaves

elastically with an elasticity several orders of magnitude softer outside the object domains \mathcal{O}_i on the complement set $\Omega \setminus \mathcal{O}_i$. Thus, given a smooth approximation $\chi_{\mathcal{O}_i}^\epsilon$ of the characteristic function of the object domain \mathcal{O}_i , we define an approximate elastic energy $\mathcal{W}^\epsilon[\mathcal{O}_i, \phi_i] = \int_{\Omega} \left((1 - \eta) \chi_{\mathcal{O}_i}^\epsilon + \eta \right) W(\mathcal{D}\phi_i) dx$, where in our applications $\eta = 10^{-4}$. Finally, the resulting approximation of the total energy functional for the variational description of the average shape reads

$$\mathcal{E}^{\gamma, \epsilon}[v, (\phi_i)_{i=1, \dots, n}] = \frac{1}{n} \sum_{i=1}^n (\mathcal{W}^\epsilon[\mathcal{O}_i, \phi_i] + \gamma \mathcal{F}^\epsilon[v_i, \phi_i, v]) + \mu \mathcal{L}^\epsilon[v]. \quad (5)$$

Analogously, a phase field approximation $\mathcal{E}_k^{\gamma, \epsilon}$ of the energy \mathcal{E}_k can be constructed. In the approximation, \mathcal{F}^ϵ acts as a penalty with $\gamma \gg 1$ and \mathcal{L}^ϵ ensures a mild regularization of the averaged shape with $\mu \ll 1$. The results typically are not very sensitive to variations of γ and μ , and $\gamma = 0.2$, $\mu = 2 \cdot 10^{-10}$ represent typical parameter values. Integration is performed only in regions where all integrands are defined.

The actual spatial discretization is based on finite elements. We consider the phase fields v , v_i and deformations ϕ_i as being represented by continuous, piecewise multilinear (trilinear in 3D and bilinear in 2D) finite element functions on an image domain $\Omega = [0, 1]^d$. The energy is computed using Gaussian quadrature of third order on each element. The energy is minimized via a gradient descent scheme, in which we alternate between a number of gradient descent steps for the deformations ϕ_i and the solution of the (linear) Euler–Lagrange equation for the average phase field v for fixed deformations ϕ_i . As stepsize control we employ Armijo’s rule. Robustness is ensured by a cascadic multi-scale approach which helps to avoid local minima: The energy is first minimized on a coarse scale, using the input images at a coarse resolution. Then, the results are repeatedly prolonged to a higher resolution and refined on that resolution, until the final resolution is achieved. For this purpose we utilize a sequence of dyadic grids with $2^l + 1$ nodes in each coordinate direction, where the resolution level l is increased in steps of 1. For further details both on the phase field approximation and the numerical discretization we refer to [40].

Concerning the covariance analysis, the finite element approximation of the elastic covariance metric is given by the matrix

$$\left(\frac{1}{n} \sum_{k=1}^n \int_{\mathcal{O}_k} \langle W_{,AA}[\mathcal{D}\phi_k], \mathcal{D}(\theta_i \circ \phi_k), \mathcal{D}(\theta_j \circ \phi_k) \rangle dx \right)_{ij},$$

where θ_i, θ_j represent the vector-valued finite element basis functions. This stiffness matrix is assembled by

first computing $\phi_k(x)$ at each quadrature point x . Then, all those basis functions θ_i are identified whose support contains $\phi_k(x)$. Points which are displaced outside the computational domain are projected back on its boundary. The corresponding evaluation of $\theta_i \circ \phi_k$ at x contributes to the assembly of the stiffness matrix. After assembly, a standard conjugate gradient iteration is applied to solve (4) for the displacement u_k , given the stress σ_k from the zero moment analysis. The eigenvalues and eigenvectors of the correlation matrix $\frac{1}{n}(g(u_i, u_j))_{ij}$ are then obtained by first tridiagonalizing the matrix via Householder's method and then applying a QR iteration with implicit shifts.

Concerning computation time, the solution of the linear equation (4) as well as the diagonalization of the correlation matrix are negligible (of the order of seconds) compared to the minimization of (5) for the average phase field and the matching deformations, which on resolution level $l = 9$ in 2D needs almost ten minutes per input shape on a Pentium IV PC at 1.8 GHz running under Linux (using a non-optimized code). Via parallelization on m processors, computation time can be reduced by almost a factor m : For this purpose, the value and the derivative of m summands in (5) are evaluated simultaneously on different processors.

8 Applications

We have applied our shape analysis approach to various collections of 2D and 3D shapes. These shapes were obtained via segmentation from images, where the input images were aligned a priori (so that the numerical algorithm did not get stuck in local minima). In principle, this alignment could of course also be done as a preprocessing step, finding translations and rotations such that the center of mass and the principal axes of the objects \mathcal{O}_k are aligned.

2D shapes. Some results of shape averages and corresponding dominant modes of shape variation for shapes of 2D objects are already depicted in Figs. 1 to ?? as first illustrative examples. Figure 1 shows the average of five human silhouettes. The corresponding deformations ϕ_i and local deformation invariants are displayed in Fig. 2 for two of the input shapes. Particularly the deformed checkerboard patterns show that—due to the invariance properties of the energy—isometries are locally preserved. Also, the indicators of length and area variation only peak locally at the person's joints. The corresponding principal components are given in Fig. 9. The average shape is represented by the dark line, whereas the light red lines signify deformations of the shape

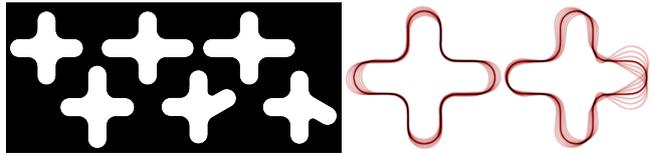


Fig. 10 Six input shapes and their first two modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1 and 0.34.

along the principal components. Here, we see the bending of the arm and the leg basically decoupled as the first two dominant modes of variation. The silhouette variations of raising the arm or the leg can only be obtained as linear combinations of the first and fourth or of the second and third mode of variation, respectively. This coupling is not too surprising, noting that the average has a slightly bent leg and arm so that the influence of all input shapes on the average also incorporates a straightening or bending of these limbs.

In the following examples the covariance analysis is performed based on the metric induced by the Hessian of the elastic energy. The decoupling of shape variations becomes even more obvious in Fig. 10, where we have more input shapes but fewer variation among them. The first mode describes the shortening of the horizontal and stretching of the vertical axis (or vice versa), whereas the second mode corresponds to bending the right branch of the cross-shape. In order to depict the principal component deformations of the average shape we deducted global rotations which for the second mode produces the tilt of the deformed shapes. The complete decoupling of bending and stretching is here achieved by including the cross with bent branch as well as its symmetric counterpart as input shapes. The average shape then only has straight branches so that the bending is invisible to the stretching modes. The results from Fig. 11 are shown for comparison with [11]. Interestingly, the obtained modes of variation differ slightly: In both cases, the first mode of variation is some kind of height variation of the mean fish (though locally, the variation looks different). While our second mode of variation is more or less an overall variation in fish length, especially pronounced at the tailfin, they obtain a combination of different local variations like tailfin thickness, pectoral fin length, and chest shape. Such a type of eigenmode in our computation only occurs as the third mode of variation. A statistical analysis of the hand shapes in Fig. 12 has also been performed in [14] and [21], where the shapes are represented as vectors of landmark positions. The average and the modes of variation are quite similar, representing different kinds of spreading the fingers.

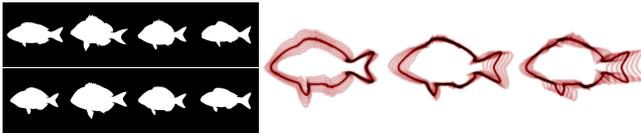


Fig. 11 First three modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.49, and 0.26 for 8 fish silhouettes from the shape database at the Centre for Vision, Speech, and Signal Processing, University of Surrey.

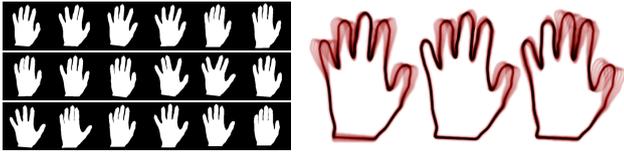


Fig. 12 First three modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.88, and 0.42 for 18 hand silhouettes from [14].

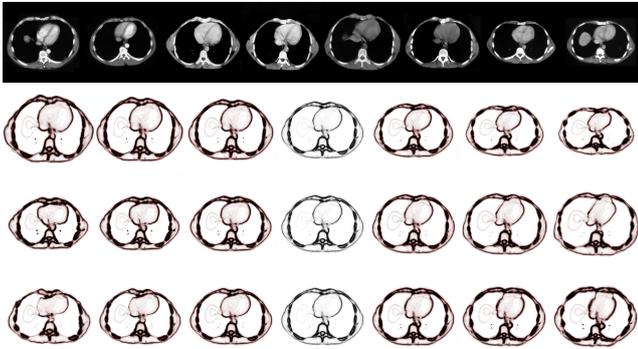


Fig. 13 8 thorax CT scans from different patients and their first three modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.12, and 0.07. Note that the thin lines which can be seen left of the heart correspond to contours of the liver, which are only visible in the first and last input image.

Images. Figure 13, using thorax CT scans as input images, shows that the approach also works for image morphologies instead of shapes. For images, the edge set is considered as the corresponding shape. Hence, these shapes are usually significantly more complex and characterized by nested shape contours. In our example, the first mode of variation represents a variation of the chest size, the next mode corresponds to a change of heart and scapula shape, while the third mode mostly concerns the rib position. For this example, the input images were not segmented in advance, but simultaneously to the averaging procedure to exploit the stronger robustness of joint segmentation and registration. In this way, artifacts have been avoided that would appear otherwise due to the visibility of liver contours in only some of the input images. Note that the local shape variation at the sides of the chest in the second and third mode of variation originates from the visibility of the scapula in some input images.

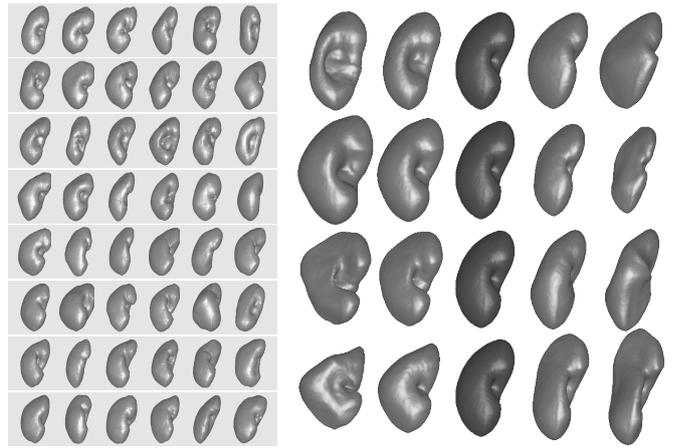


Fig. 14 48 input kidneys and their first four modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.72, 0.37, and 0.31.

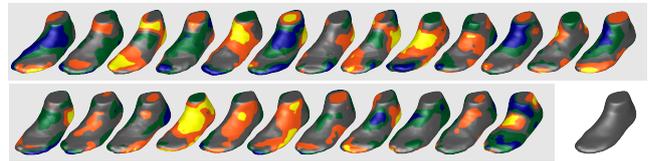


Fig. 15 24 given foot shapes, textured with the distance to the surface of the average foot (bottom right). The range $[-6 \text{ mm}, 6 \text{ mm}]$ is color-coded as .

3D shapes. Next, let us investigate the dominant modes of variation of shapes in \mathbb{R}^3 , where the computation is based on the L^2 -metric for the first example and on the elasticity tensor-based inner product for the second example. In our first 3D example we compute the first four modes of variation for a given set of 48 kidney shapes as depicted in Fig. 14. For all modes we show the average in the middle and its configurations after deformation according to the principal components. Local structures seem to be quite well represented and preserved during the averaging process and the subsequent covariance analysis compared to e.g. the PCA on kidney shapes in [20] where a medial representation is used. The second example takes 24 foot-shapes as input (which were originally provided as triangulated surfaces and then converted to characteristic functions on the unit cube). The average shape is shown along with the original shapes in Fig. 15, where the input feet are color-coded according to their local distance to the surface of the average foot. For that purpose the input shapes were optimally aligned with the average. It is doubtlessly difficult to analyze the shape variation on this basis: We see modest variation at the toes and the heel as well as on the instep, but any correlation between these variations is difficult to determine. The corresponding modes of variation in Fig. 16, however, are quite intuitive. The first mode apparently rep-

resents changing foot volume, the second mode shows width and length changes, the third and fourth mode belong to different variants of toe positions, and the fifth and sixth mode correspond to variations in relative heel position and instep shape. Note that in contrast to [39], where dominant modes from a PCA are shown for the same example, this time the results were obtained using the elasticity tensor of the compound object as the covariance metric. As a further example,

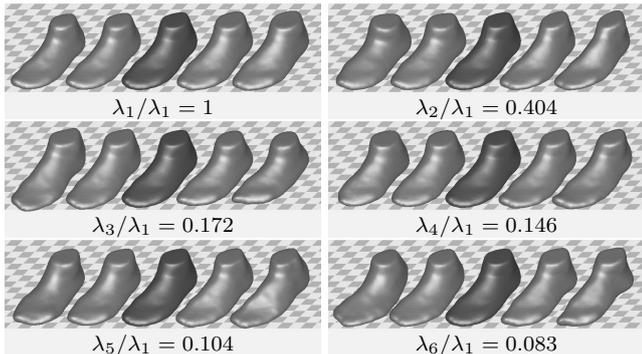


Fig. 16 The first six dominant modes of variation for the feet from Fig. 15.

let us now use these modes to evaluate how well a new, additional foot shape $\hat{S} = \partial\hat{O}$ fits into the given distribution of foot shapes. To this end, we first compute a linear representative of \hat{S} : We find an optimal matching deformation $\hat{\phi}$ that deforms \hat{O} into the average object \mathcal{O} , calculate the corresponding boundary stresses $\hat{\sigma}\nu$, and finally compute the corresponding displacement \hat{u} via (4). Then we can reconstruct \hat{u} using the different dominant modes of variation. Figure 17 shows that a reconstruction with only the first dominant mode yields roughly the correct volume variation, but the actual shape is not well reconstructed. The first three modes already serve to reconstruct the correct width to length ratio, and using all 23 modes of variation, \hat{u} can almost be reconstructed. Only the bunion or hallux valgus type position of the big toe cannot be modeled. The Mahalanobis distance of the new foot to the average (neglecting those components of shape variation which are not present in the original set of foot shapes) is computed as $d_M(\hat{S}, \mathcal{S}) = \sqrt{\frac{1}{n} \sum_{i=1}^{23} \frac{g(\hat{u}, w_i)^2}{\lambda_i}} = 1.23$, where a value of 1 would correspond to the standard deviation.

9 Conclusion

We have developed an elasticity-based notion of shape variation. Since the shape space of elastically deformable objects inherently does not possess a Riemannian structure, we utilized an alternative shape space structure, in

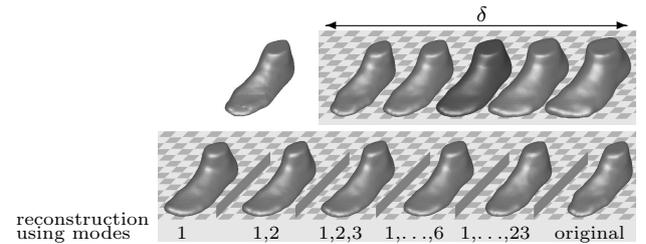


Fig. 17 For a new, additional foot (top left), a linear representative \hat{u} is computed via the method from Sec. 4, and it is visualized as variation of the average foot \mathcal{S} via $(\mathbb{1} + \delta\hat{u})(\mathcal{S})$ (top right). This variation (the bottom right shape is a copy of the leftmost shape in the visualization of \hat{u} from the first line) is reconstructed based on the PCA, using only one, two, three, six, and all 23 dominant modes of variation (bottom).

which distance is replaced by elastic deformation energy and boundary stresses play the role of linear representations of shapes. Such an approach imposes a physically and mathematically sound structure on spaces of elastic objects. Its computational feasibility has been proven by application to sets of 2D and 3D shapes. The interpretation of shapes as boundaries of objects with volume (where the covariance analysis takes the whole volume in account) has the effect of robustness with respect to noise in the shape contour: The (large-scale) object variation takes place mostly inside the volume and thus is rather insensitive to perturbations at the boundary. Also, due to the underlying elastic model, the object positioning has no influence on the result, which approximately also holds in the numeric computations due to a multi-scale approach.

Acknowledgements The authors thank Guillermo Sapiro for pointing them to the issue of an elastic principal component analysis. We are grateful to Heiko Schlarb from *adidas*, Herzogenaurach, Germany, for providing 3D scans of feet. Furthermore, we acknowledge support by the German Science Foundation via the Hausdorff Center for Mathematics and the SPP 1253. Benedikt Wirth has been supported by the Bonn International Graduate School.

A Appendix

A.1 Stress balance on the average shape. Applying the transformation rule and integration by parts we deduce from (1) that

$$\begin{aligned} 0 &= \sum_{i=1}^n \int_{\mathcal{O}_i} W_{,A}(\mathcal{D}\phi_i) : \mathcal{D}(u \circ \phi_i) \, dx \\ &= \sum_{i=1}^n \int_{\mathcal{O}} \left(\frac{W_{,A}(\mathcal{D}\phi_i) \mathcal{D}\phi_i^T}{\det \mathcal{D}\phi_i} \right) \circ \phi_i^{-1} : \mathcal{D}u \, dx \\ &= \int_{\mathcal{S}} \sum_{i=1}^n \sigma_i : u \otimes \nu[\mathcal{S}] \, da[\mathcal{S}] - \int_{\mathcal{O}} \sum_{i=1}^n \operatorname{div}(\sigma_i) \cdot u \, dx, \end{aligned}$$

where $\nu[\mathcal{S}]$ is the outer normal in \mathcal{S} and “ \otimes ” denotes the rank-1 product $v \otimes w = v w^T$ for two vectors $v, w \in \mathbb{R}^d$. Hence, since (1)

holds for any u , we obtain $\operatorname{div} W_{,A}(\mathcal{D}\phi_i) = 0$ on every domain \mathcal{O}_i and the stress balance $0 = \sum_{i=1,\dots,n} \sigma_i(x)\nu(x)$ on \mathcal{S} .

A.2 Euler–Lagrange condition for offset displacements. Abbreviating $\psi_k := \mathbb{1} + \delta u_k$, the optimality condition for (3) reads

$$\begin{aligned} 0 &= \langle \partial_{\psi_k} \mathcal{E}_k, \psi \rangle \\ &= \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{O}_i} W_{,A}(\mathcal{D}(\psi_k \circ \phi_i)) : \mathcal{D}\psi \circ \phi_i \mathcal{D}\phi_i \, dx - \delta \int_{\mathcal{S}} (\sigma_k \nu) \cdot \psi \, da[\mathcal{S}] \\ &= \int_{\mathcal{O}} \sigma[\delta u_k] : \mathcal{D}\psi \, dx - \delta \int_{\mathcal{S}} (\sigma_k \nu) \cdot \psi \, da[\mathcal{S}] \\ &= \int_{\mathcal{S}} ((\sigma[\delta u_k] - \delta \sigma_k) \nu) \cdot \psi \, da[\mathcal{S}] - \int_{\mathcal{O}} \operatorname{div} \sigma[\delta u_k] \cdot \psi \, dx \end{aligned}$$

for all test functions ψ , where we denote by

$$\sigma[\delta u_k] := \frac{1}{n} \sum_{i=1}^n W_{,A}((\mathbb{1} + \delta \mathcal{D}u_k) \mathcal{D}\phi_i \circ \phi_i^{-1}) \operatorname{cof} \mathcal{D}(\phi_i^{-1})$$

the first Piola–Kirchhoff stress tensor on the compound object \mathcal{O} . Hence, as Euler–Lagrange condition for u_k we obtain $\operatorname{div} \sigma[\delta u_k] = 0$ on \mathcal{O} and $\sigma[\delta u_k] \nu = \delta \sigma_k \nu$ on \mathcal{S} .

A.3 Derivation of the averaged elasticity tensor. Linearization of (3) for small δ yields

$$\begin{aligned} \mathcal{E}_k[\delta, u] &\doteq \frac{1}{n} \sum_{i=1}^n \mathcal{W}[\mathcal{O}_i, \phi_i] + \delta \int_{\mathcal{O}_i} W_{,A}[\mathcal{D}\phi_i] : \mathcal{D}(u \circ \phi_i) \, dx \\ &\quad + \frac{\delta^2}{2} \int_{\mathcal{O}_i} \langle W_{,AA}[\mathcal{D}\phi_i], \mathcal{D}(u \circ \phi_i), \mathcal{D}(u \circ \phi_i) \rangle \, dx - \delta^2 \int_{\mathcal{O}} \sigma_k : \mathcal{D}u \, dx \\ &= C + \delta \int_{\mathcal{O}} \left(\frac{1}{n} \sum_{i=1}^n \sigma_i \right) : \mathcal{D}u \, dx - \delta^2 \int_{\mathcal{O}} \sigma_k : \mathcal{D}u \, dx \\ &\quad + \frac{\delta^2}{2} \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{O}_i} \langle W_{,AA}[\mathcal{D}\phi_i], \mathcal{D}(u \circ \phi_i), \mathcal{D}(u \circ \phi_i) \rangle \, dx \\ &= C + \delta \int_{\mathcal{S}} \left(\frac{1}{n} \sum_{i=1}^n \sigma_i \nu \right) \cdot u \, da[\mathcal{S}] \\ &\quad - \delta \int_{\mathcal{O}} \left(\frac{1}{n} \sum_{i=1}^n \operatorname{div} \sigma_i \right) \cdot u \, dx - \delta^2 \int_{\mathcal{O}} \sigma_k : \mathcal{D}u \, dx \\ &\quad + \frac{\delta^2}{2} \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{O}_i} \langle W_{,AA}[\mathcal{D}\phi_i], \mathcal{D}(u \circ \phi_i), \mathcal{D}(u \circ \phi_i) \rangle \, dx \\ &= C + \delta^2 \int_{\mathcal{O}} \frac{1}{2} (\mathbf{C}, \mathcal{D}u, \mathcal{D}u) - \sigma_k : \mathcal{D}u \, dx. \end{aligned}$$

for the tensor

$$\mathbf{C} = \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\det \mathcal{D}\phi_i} \mathcal{D}\phi_i W_{,AA}[\mathcal{D}\phi_i] \mathcal{D}\phi_i^T \right) \circ \phi_i^{-1}.$$

References

1. L. Ambrosio and V. M. Tortorelli. Approximation of functionals depending on jumps by elliptic functionals via Γ -convergence. *Comm. Pure Appl. Math.*, 43:999–1036, 1990.
2. L. Ambrosio and V. M. Tortorelli. On the approximation of free discontinuity problems. *Bollettino dell’Unione Matematica Italiana, Sezione B*, 6(7):105–123, 1992.
3. J.M. Ball. Global invertibility of Sobolev functions and the interpenetration of matter. *Proceedings of the royal Society of Edinburgh*, 88A:315–328, 1981.
4. K. K. Bhatia, J. V. Hajnal, A. Hammers, and D. Rueckert. Similarity metrics for groupwise non-rigid registration. In N. Ayache, S. Ourselin, and A. Maeder, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2007*, volume 4792 of *LNCS*, pages 544–552, 2007.
5. K. K. Bhatia, J. V. Hajnal, B. K. Puri, A. D. Edwards, and D. Rueckert. Consistent groupwise non-rigid registration for atlas construction. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro*, volume 1, pages 908–911, 2004.
6. A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Monographs in Computer Science. Springer, 2008.
7. Bernard Chalmond and Stéphane C. Girard. Nonlinear modeling of scattered multivariate data and its application to shape change. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):422–432, 1999.
8. Tony F. Chan and Luminita A. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.
9. Tony F. Chan and Luminita A. Vese. A level set algorithm for minimizing the Mumford-Shah functional in image processing. In *IEEE/Computer Society Proceedings of the 1st IEEE Workshop on Variational and Level Set Methods in Computer Vision*, pages 161–168, 2001.
10. G. Charpiat, O. Faugeras, R. Keriven, and P. Maurel. Distance-based shape statistics. In *Acoustics, Speech and Signal Processing, 2006 (ICASSP 2006)*, volume 5, 2006.
11. Guillaume Charpiat, Olivier Faugeras, and Renaud Keriven. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics*, 5(1):1–58, 2005.
12. M. Chipot and L. C. Evans. Linearization at infinity and lipschitz estimates in the calculus of variations. *Proceedings of the Royal Society of Edinburgh A*, 102(3–4):291–303, 1986.
13. P. G. Ciarlet. *Three-dimensional elasticity*. Elsevier Science Publishers B. V., 1988.
14. T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
15. Daniel Cremers, Timo Kohlberger, and Christoph Schnörr. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36:1929–1943, 2003.
16. G. Dal Maso, J.M. Morel, and S. Solimini. A variational method in image segmentation: existence and approximation results. *Acta Math.*, 168(1-2):89–151, 1992.
17. Samuel Dambreville, Yogesh Rathi, and Allen Tannenbaum. A shape-based approach to robust image segmentation. In A. Campilho and M. Kamel, editors, *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 4141 of *LNCS*, pages 173–183, 2006.
18. E. De Giorgi, M. Carriero, and A. Leaci. Existence theorem for a minimum problem with free discontinuity set. *Arch. Rat. Mech. and Anal.*, 108:195–218, 1989.
19. Olivier Faugeras, Geoffroy Adde, Guillaume Charpiat, Christophe Chef’d’Hotel, Maureen Clerc, Thomas Deneux, Rachid Deriche, Gerardo Hermosillo, Renaud Keriven, Pierre Kornprobst, Jan Kybic, Christophe Lenglet, Lucero Lopez-Perez, Théo Papadopoulos, Jean-Philippe Pons, Florent Segonne, Bertrand Thirion, David Tschumperlé, Thierry Viéville, and Nicolas Wotawa. Variational, geometric, and statistical methods for modeling brain anatomy and function. *NeuroImage*, 23:S46–S55, 2004.
20. P. Thomas Fletcher, Conglin Lu, and Sarang Joshi. Statistics of shape via principal geodesic analysis on Lie groups. In *IEEE Computer Society Conference on Computer Vision*

- and *Pattern Recognition CVPR*, volume 1, pages 95–101, 2003.
21. Tom Fletcher, Suresh Venkatasubramanian, and Sarang Joshi. Robust statistics on Riemannian manifolds via the geometric median. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
 22. M. Fuchs, B. Jüttler, O. Scherzer, and H. Yang. Shape metrics based on elastic deformations. Technical report, FWF, Joint Research Program of Industrial Geometry, 2008.
 23. Matthias Fuchs, Bert Jüttler, Otmar Scherzer, and Huaiping Yang. Shape metrics based on elastic deformations. *J. Math. Imaging Vis.*, 35(1):86–102, 2009.
 24. B.J. Hafner, S.G. Zachariah, and J.E. Sanders. Characterisation of three-dimensional anatomic shapes using principal components: application to the proximal tibia. *Med. Biol. Eng. Comput.*, 38:9–16, 2000.
 25. S. Joshi, B. Davis, M. Jomier, and G. Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23:151–160, 2004. Supplement 1.
 26. H. Karcher. Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30(5):509–541, 1977.
 27. D. G. Kendall. Shape manifolds, procrustean metrics, and complex projective spaces. *Bull. London Math. Soc.*, 16:81–121, 1984.
 28. M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. In *ACM Transactions on Graphics*, volume 26, pages #64, 1–8, 2007.
 29. Michael E. Leventon, W. Eric L. Grimson, and Olivier Faugeras. Statistical shape influence in geodesic active contours. In *5th IEEE EMBS International Summer School on Biomedical Imaging, 2002.*, 2002.
 30. S. Marsland, C. J. Twining, and C. J. Taylor. Group-wise non-rigid registration using polyharmonic clamped-plate splines. In R. E. Ellis and T. M. Peters, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI*, volume 2879 of *LNCS*, pages 771–779, 2003.
 31. F. Méoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5:313–347, 2005.
 32. M. I. Miller and L. Younes. Group actions, homeomorphisms, and matching: a general framework. *International Journal of Computer Vision*, 41(1–2):61–84, 2001.
 33. M.I. Miller, A. Trouvé, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering*, 4:375–405, 2002.
 34. J.-M. Morel and S. Solimini. Segmentation of images by variational methods: a constructive approach. *Revista Matemática de la Universidad Complutense de Madrid*, 1(1):169–182, 1988.
 35. David Mumford and Jayant Shah. Optimal approximation by piecewise smooth functions and associated variational problems. *Communications on Pure Applied Mathematics*, 42:577–685, 1989.
 36. Dimitrios Perperidis, Raad Mohiaddin, and Daniel Rueckert. Construction of a 4d statistical atlas of the cardiac anatomy and its use in classification. In J. Duncan and G. Gerig, editors, *Medical Image Computing and Computer Assisted Intervention*, volume 3750 of *LNCS*, pages 402–410, 2005.
 37. Yogesh Rathi, Samuel Dambreville, and Allen Tannenbaum. Comparative analysis of kernel methods for statistical shape learning. In R.R. Beichel and M. Sonka, editors, *computer vision approaches to medical image analysis*, volume 4241 of *LNCS*, pages 96–107, 2006.
 38. D. Rueckert, A. F. Frangi, and J. A. Schnabel. Automatic construction of 3D statistical deformation models using non-rigid registration. In W. Niessen and M. Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI*, volume 2208 of *LNCS*, pages 77–84, 2001.
 39. Martin Rumpf and Benedikt Wirth. An elasticity approach to principal modes of shape variation. In *Proceedings of the Second International Conference on Scale Space Methods and Variational Methods in Computer Vision (SSVM 2009)*, volume 5567 of *Lecture Notes in Computer Science*, pages 709–720, 2009.
 40. Martin Rumpf and Benedikt Wirth. A nonlinear elastic shape averaging approach. *SIAM Journal on Imaging Sciences*, 2(3):800–833, 2009.
 41. M. Söhn, M. Birkner, D. Yan, and M. Alber. Modelling individual geometric variation based on dominant eigenmodes of organ deformation: implementation and evaluation. *Phys. Med. Biol.*, 50:5893–5908, 2005.
 42. Anuj Srivastava, Aastha Jain, Shantanu Joshi, and David Kaziska. Statistical shape models using elastic-string representations. In P.J. Narayanan, editor, *Asian Conference on Computer Vision*, volume 3851 of *LNCS*, pages 612–621, 2006.
 43. C. Studholme. Simultaneous population based image alignment for template free spatial normalisation of brain anatomy. In J. C. Gee, J. B. A. Maintz, and M. W. Vannier, editors, *Second International Workshop, WBIR, Biomedical Image Registration*, volume 2717 of *LNCS*, pages 81–90, 2003.
 44. Benedikt Wirth, Leah Bar, Martin Rumpf, and Guillermo Sapiro. A continuum mechanical approach to geodesics in shape space. *submitted to IJCV*, 2010.
 45. Laurent Younes. Computable elastic distances between shapes. *SIAM J. Appl. Math.*, 58(2):565–586, April 1998.