

An Elasticity Approach to Principal Modes of Shape Variation

Martin Rumpf¹ and Benedikt Wirth¹

Bonn University, 53113 Bonn, Germany,
{martin.rumpf,benedikt.wirth}@ins.uni-bonn.de,
WWW home page: <http://www.ins.uni-bonn.de>

Abstract. Concepts from elasticity are applied to analyze modes of variation on shapes in two and three dimensions. This approach represents a physically motivated alternative to shape statistics on a Riemannian shape space, and it robustly treats strong nonlinear geometric variations of the input shapes.

To compute a shape average, all input shapes are elastically deformed into the same configuration. That configuration which minimizes the total elastic deformation energy is defined as the average shape. Each of the deformations from one of the shapes onto the shape average induces a boundary stress. Small amplitude stimulation of these stresses leads to displacements which reflect the impact of every single input shape on the average. To extract the dominant modes of variation, a PCA is performed on this set of displacements.

To make the approach computationally tractable, a relaxed formulation is proposed, and sharp contours are approximated via phase fields. For the spatial discretization of the resulting model, piecewise multilinear finite elements are applied. Applications in 2D and in 3D demonstrate the qualitative properties of the presented approach.

1 Introduction

This paper is concerned with the 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. Hence, for the zero moment analysis – i. e. the definition of a suitable shape average – the total elastic energy stored in a set of deformations from the input shapes onto a single image shape is minimized. At the energy minimum the corresponding image shape is defined as the shape average.

Concerning a first moment analysis, we propose a physically sound linearization of shape variations which allows to define a covariance tensor. Each deformation from an input onto the average shape 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 well-defined. 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. Thus, a covariance tensor can be computed based on these displacements of the shape average. It linearly encodes the modes of variation of the shape average induced by the set of input shapes, even though the underlying deformations are usually large and nonlinear. Finally, we perform a principal component analysis based on this covariance tensor, which allows to identify the dominant modes of variation of the input shapes.

Our model is related to the 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 . 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. 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 .

At first, shape analysis was mainly based on correspondences between landmark positions on different shapes as in the influential work by Cootes et al. [1]. Principal component analysis (PCA) is a classical, by definition linear statistical tool. Chalmond and Girard [2] have proposed a PCA which incorporates also truly nonlinear geometric transformations. A survey on the potential of shape analysis in brain imaging is given by Faugeras and coworkers in [3]. 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 shape space was introduced by Kendall [4] already in 1984. Charpiat et al. [5] discuss shape averaging and shape statistics based on the Hausdorff distance of sets. Statistics on signed distance functions was also studied by Leventon et al. [6], whereas Dambreville et al. [7] used shape statistics based on characteristic functions to define a robust shape prior in image segmentation. Kernel density estimation in feature space was introduced by Cremers et al. [8] to incorporate the probability of 2D silhouettes of 3D objects in image segmentation. An overview on related kernel density methods is given by Rathi et al. [9]. Mémoli and Sapiro [10] have investigated the Gromov–Hausdorff distance as a global measure for the lack of isometry in shape analysis. In contrast to such a global measure for 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. [11,12]. Fuchs et al. [13] proposed a vis-




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.

coelastic notion of the distance between shapes \mathcal{S} given as boundaries of physical objects \mathcal{O} . The elasticity paradigm for shape analysis on which our approach is founded differs significantly from these metric approaches to shape space (cf. Sect. 4 for a detailed discussion of the conceptual difference).

In this paper, shapes are represented implicitly via a diffused phase field description. This in particular enables a robust and flexible application in two and three dimensions.

2 Zero Moment Analysis

In this section we briefly recall an elastic approach to shape averaging already presented in [14]. We consider shapes \mathcal{S}_i as the boundaries $\partial\mathcal{O}_i$ of sufficiently regular objects \mathcal{O}_i .

Given n 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 physical manner. For that purpose we assume that the average shape \mathcal{S} can 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). A natural choice for the shape average \mathcal{S} is that particular shape which minimizes the total accumulated deformation energy of all deformations, $\mathcal{E}[\mathcal{S}, (\phi_i)_{i=1, \dots, n}] = \frac{1}{n} \sum_{i=1}^n \mathcal{W}[\mathcal{O}_i, \phi_i]$, where $\mathcal{W}[\mathcal{O}_i, \phi_i]$ represents the stored deformation energy of the deformation ϕ_i .

To ensure existence of a minimizing shape \mathcal{S} , we add a regularizing prior $\mathcal{L}[\mathcal{S}]$ to the energy. Here, we consider the \mathcal{H}^{d-1} -measure of \mathcal{S} , i. e. $\mathcal{L}[\mathcal{S}] = \int_{\mathcal{S}} da$, and the shape average \mathcal{S} is defined as a minimizer of the energy $\mathcal{E}[\mathcal{S}, (\phi_i)_{i=1, \dots, n}] + \mu\mathcal{L}[\mathcal{S}]$. As deformation energy $\mathcal{W}[\mathcal{O}_i, \phi_i]$ we will employ a nonlinear, hyperelastic energy $\mathcal{W}[\mathcal{O}, \phi] = \int_{\mathcal{O}} W(\mathcal{D}\phi) dx$, whose integrand 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 $(I_1, I_2, I_3) := (|\mathcal{D}\phi|_2^2, |\text{cof}\mathcal{D}\phi|_2^2, \det(\mathcal{D}\phi))$. $|\mathcal{D}\phi|_2 := \sqrt{\text{tr}(\mathcal{D}\phi^T \mathcal{D}\phi)}$, $|\text{cof}(\mathcal{D}\phi)|_2$, and $\det(\mathcal{D}\phi)$ describe the averaged local change of length, area, and volume, respectively. We consider polyconvex energy functionals [15], where \hat{W} is convex and isometries, i. e. deformations with $\mathcal{D}\phi^T \mathcal{D}\phi = \mathbb{1}$, are local minimizers (cf. Fig. 2). Typical energy densities are of the form $\bar{W}(I_1, I_2, I_3) = \alpha_1 I_1^{\frac{p}{2}} + \alpha_2 I_2^{\frac{q}{2}} + \alpha_3 I_3^{-s} + \alpha_4 I_3^r$ with $\alpha_1, \dots, \alpha_4 > 0$, where the penalization of volume





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]).

shrinkage, i. e. $\bar{W} \xrightarrow{I_3 \rightarrow 0} \infty$, enables us to control local injectivity (cf. [16]).

This type of energy has two major advantages: it allows to incorporate large deformations with strong material and geometric nonlinearities, and its form 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. The first Piola–Kirchhoff stress tensor, which describes force per unit area in the reference configuration \mathcal{O} , is then recovered as $\sigma^{\text{ref}}[\phi] = W_{,A}(\mathcal{D}\phi) := \frac{\partial W(A)}{\partial A}$. The Cauchy (real) stress, describing the force per unit area in the deformed configuration $\phi(\mathcal{O})$, reads $\sigma[\phi] = \sigma^{\text{ref}}[\phi](\text{cof}\mathcal{D}\phi)^{-1}$.

To simplify the numerical treatment and to allow for slight topological differences between the shapes \mathcal{S}_i 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} . Our shape averaging model is thus based on the energy

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

3 First Moment Analysis

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:

Encoding nonlinear deformations via stresses on a linear vector space.

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, which try to restore the undeformed configuration \mathcal{O}_i . The stress at some point x on \mathcal{S} is given by the application of the Cauchy stress tensor $\sigma_i = \sigma[\phi_i]$ 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} . 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} .

Modeling the impact of an input shape on the average shape. 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 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 $u_k : \mathcal{O} \rightarrow \mathbb{R}^d$. 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,

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

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. Let us remark that the boundary integral 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} . As Euler Lagrange condition for u_k we obtain $\operatorname{div} \sigma_k[\delta u_k] = 0$ on \mathcal{O} and $\sigma[\delta u_k]\nu = \delta\sigma_k\nu$ on \mathcal{S} after a tedious but straightforward computation. Here,

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

is 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 $i = 1, \dots, n$. As long as $A \mapsto W(A)$ is not quadratic in A , u_k still solves




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

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 , $k = 1, \dots, n$. In the limit for $\delta \rightarrow 0$, we would obtain u_k as the solution of the actually linear elasticity problem

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

for the symmetric displacement gradient $\epsilon[u] = (\mathcal{D}u + \mathcal{D}u^T)/2$ under the constraint $\int_{\mathcal{O}} u \, dx = 0$. Here, the in general inhomogeneous and anisotropic elasticity tensor \mathbf{C} is defined by

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

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 $i = 1, \dots, n$. In our current implementation, we avoid the evaluation of \mathbf{C} and consider the above nonlinear approximation, which is simpler to implement but computationally more expensive.

The actual covariance analysis based on the derived displacements.

Now, we have a set of displacements $u_k : \mathcal{O} \rightarrow \mathbb{R}^d$ at hand which represent 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} . On this space of displacements, we consider the standard L^2 -product $(u, \tilde{u})_2 := \int_{\mathcal{O}} u \cdot \tilde{u} \, dx$ and define the covariance operator

$$\mathbf{Cov} : L^2(\mathcal{O}) \rightarrow L^2(\mathcal{O}); \quad u \mapsto \mathbf{Cov}u := \frac{1}{n} \sum_{k=1, \dots, n} (u, u_k)_2 u_k.$$

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 L^2 -orthogonal eigenfunctions $w_k : \mathcal{O} \rightarrow \mathbb{R}^d$ – actually displacements – and eigenvalues $\lambda_k > 0$ with

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



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

These eigenfunctions can be considered as principal modes of variation of the average object \mathcal{O} and hence of the average shape \mathcal{S} , given the n input shapes. The eigenvalues encode the actual strength of these variations. 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 (cf. Fig. 4). These stresses lie in a linear vector space and thus allow for a covariance analysis, which is by definition linear. The interpretation of stresses in terms of displacements can be regarded as a proper choice of a scalar metric $g(\cdot, \cdot)$ on the space of stresses interpreted as a tangent space of the shape space at the average shape: we define $g(\sigma\nu, \tilde{\sigma}\nu) := (u, \tilde{u})_2$, given the above identification of stresses $\sigma\nu, \tilde{\sigma}\nu$ with induced displacements u, \tilde{u} via the proper compound elasticity problem. Finally, this identification provides a suitable physical interpretation of stresses as modes of shape variation.

4 Elastic Versus Riemannian Shape Analysis

The elasticity paradigm, on which our zero and first order shape analysis are based, differs significantly from a Riemannian approach to shape space as proposed for instance by Srivastava et al. [17]. Due to the axiom of elasticity, the energy at the deformed configuration \mathcal{S} is independent of the path from a shape $\tilde{\mathcal{S}}$ 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. [13] and the related model by Younes [18] 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 at rest the intermediate configurations are all stress-free.

The above-mentioned conceptual differences are reflected in a different behavior. 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 locally strong isometry violation. An example is provided in Fig. 5: The input shapes are regarded as two versions of an object that may have none, one, or two pins at more or less stable positions. Both pins are apparently not interpreted as shifted versions of each other since a shifting deformation would cost too much energy. However, if the material was visco-plastic, a horizontal shift of each pin would be easier and



Fig. 5. Average and variation (right) for two shapes with pins at different positions (left). The pins are not interpreted as shifted versions of each other.

result in an average shape with just one centered pin and its variation being a sideward movement. This corresponds to a completely different perception of the input shapes. The strong local rigidity and isometry preservation of the elasticity concept becomes particularly evident in Fig. 4 and Fig. 6, where non-isometric deformations are concentrated only at joints.

On a Riemannian manifold, the exponential map allows to describe geodesics from an averaged shape \mathcal{S} – in the sense of Karcher [19] – to the input shapes \mathcal{S}_k via $\mathcal{S}_k = \exp_{\mathcal{S}}(v_k)$ for some tangent vector v_k at the shape \mathcal{S} in shape space. 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 .

5 Finite Element Phase Field Approximation

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 [20] for the discretization of the Mumford–Shah model [21]. 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. 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 [20] 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. Given a phase field parameter ϵ , which will determine the width of the phase field, 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 $\mathcal{L}^\epsilon[v] = \int_{\Omega} \epsilon |\nabla v|^2 + \frac{1}{4\epsilon} (v - 1)^2 dx$, 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




Fig. 6. 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.

energy functional for the variational description of the average shape reads

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

In analogy, a phase field approximation $\mathcal{E}_k^{\gamma, \epsilon}$ of the energy \mathcal{E}_k can be constructed. In these approximations, \mathcal{F}^ϵ acts as a penalty with $\gamma \gg 1$ and \mathcal{L}^ϵ ensures a mild regularization of the averaged shape with $\mu \ll 1$. 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 $\Omega = [0, 1]^d$. A cascadic multi scale approach is applied for the relaxation of the energy. For details both on the phase field approximation and the numerical discretization we refer to [14].

6 2D and 3D Applications

We have applied our shape analysis approach to various collections of 2D and 3D shapes. The computed average and dominant variations for sets of 2D shapes are depicted in Figs. 1 to 7 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. 6. The average shape is represented by the dark line, whereas the light red lines signify deformations of the shape 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. A larger set of shapes is treated in Fig. 7, where 20 binary images “device7” from the MPEG7 shape database serve 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




Fig. 7. Original shapes and their first three modes of variation with ratios $\frac{\lambda_i}{\lambda_1}$ of 1, 0.20, and 0.05.

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 final example uses 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. 8, where the input feet are color-coded according to their local distance to the surface of the average foot. 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. 9, however, are quite intuitive.





Fig. 8. 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 .

For all modes we show the average in the middle and its configurations after deformation according to the principal components. The first mode apparently represents changing foot lengths, the second and third mode belong to different variants of combined width and length variation, and the fourth to sixth mode correspond to variations in relative heel position, ankle thickness, and instep height.

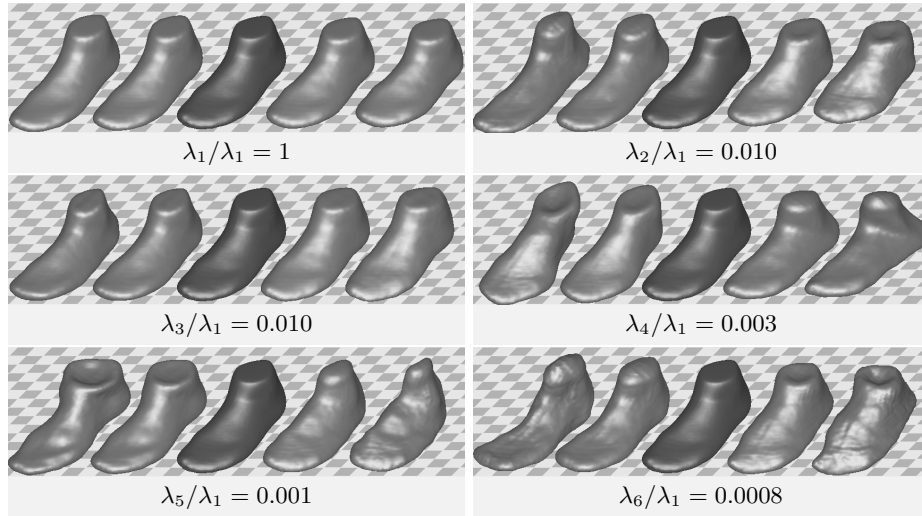


Fig. 9. The first six dominant modes of variation for the feet from Fig. 8.

7 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 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.

Acknowledgment. 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 Hausdorff Center for Mathematics. Benedikt Wirth has been supported by the Bonn International Graduate School.

References

1. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models—their training and application. *Computer Vision and Image Understanding* **61**(1) (1995) 38–59
2. Chalmond, B., Girard, S.C.: Nonlinear modeling of scattered multivariate data and its application to shape change. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21**(5) (1999) 422–432
3. Faugeras, O., Adde, G., Charpiat, G., Chefd’Hotel, C., Clerc, M., Deneux, T., Deriche, R., Hermosillo, G., Keriven, R., Kornprobst, P., Kybic, J., Lenglet, C., Lopez-Perez, L., Papadopoulou, T., Pons, J.P., Segonne, F., Thirion, B., Tschumperlé, D.,

- Viéville, T., Wotawa, N.: Variational, geometric, and statistical methods for modeling brain anatomy and function. *NeuroImage* **23** (2004) S46–S55
4. Kendall, D.G.: Shape manifolds, procrustean metrics, and complex projective spaces. *Bull. London Math. Soc.* **16** (1984) 81–121
 5. Charpiat, G., Faugeras, O., Keriven, R.: Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics* **5**(1) (2005) 1–58
 6. Leventon, M., Grimson, W., Faugeras, O.: Statistical shape influence in geodesic active contours. In: 5th IEEE EMBS International Summer School on Biomedical Imaging, 2002. (2002)
 7. Dambreville, S., Rathi, Y., Tannenbaum, A.: A shape-based approach to robust image segmentation. In Campilho, A., Kamel, M., eds.: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Volume 4141 of LNCS. (2006) 173–183
 8. Cremers, D., Kohlberger, T., Schnörr, C.: Shape statistics in kernel space for variational image segmentation. *Pattern Recognition* **36** (2003) 1929–1943
 9. Rathi, Y., Dambreville, S., Tannenbaum, A.: Comparative analysis of kernel methods for statistical shape learning. In Beichel, R., Sonka, M., eds.: *computer vision approaches to medical image analysis*. Volume 4241 of LNCS. (2006) 96–107
 10. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics* **Volume 5** (2005) 313–347
 11. Miller, M.I., Younes, L.: Group actions, homeomorphisms and matching: a general framework. *International Journal of Computer Vision* **41**(1-2) (2001) 61–84
 12. Miller, M., Trounev, A., Younes, L.: On the metrics and euler-lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering* **4** (2002) 375–405
 13. Fuchs, M., Jüttler, B., Scherzer, O., Yang, H.: Shape metrics based on elastic deformations. *Forschungsschwerpunkt S92, Industrial Geometry 71, Universität Innsbruck* (2008)
 14. Rumpf, M., Wirth, B.: A nonlinear elastic shape averaging approach. Submitted to *SIAM Journal on Imaging Sciences* (2008)
 15. Ciarlet, P.G.: *Three-dimensional elasticity*. Elsevier Science Publishers B. V. (1988)
 16. Baker, T.: Three dimensional mesh generation by triangulation of arbitrary point sets. In: *Computational Fluid Dynamics Conference, 8th, Honolulu, HI, June9-11, 1987*. Volume 1124-CP. (1987) 255–271
 17. Srivastava, A., Jain, A., Joshi, S., Kaziska, D.: Statistical shape models using elastic-string representations. In Narayanan, P., ed.: *Asian Conference on Computer Vision*. Volume 3851 of LNCS. (2006) 612–621
 18. Younes, L.: Computable elastic distances between shapes. *SIAM J. Appl. Math* **58** (1998) 565–586
 19. Karcher, H.: Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics* **30**(5) (1977) 509–541
 20. Ambrosio, L., Tortorelli, V.M.: On the approximation of free discontinuity problems. *Bollettino dell’Unione Matematica Italiana, Sezione B* **6**(7) (1992) 105–123
 21. Mumford, D., Shah, J.: Optimal approximation by piecewise smooth functions and associated variational problems. *Communications on Pure Applied Mathematics* **42** (1989) 577–685