

Regularized multi-structural shape modeling of the knee complex based on deep functional maps

Konstantinos Filip*, Evangelia I. Zacharaki, Konstantinos Moustakas

Department of Electrical and Computer Engineering, University of Patras, Greece

ARTICLE INFO

Keywords:

Multi-structure statistical shape model
Missing structure estimation
Shape correspondence
Deep functional maps
Shape reconstruction
Knee modeling

ABSTRACT

The incorporation of a-priori knowledge on the shape of anatomical structures and their variation through *Statistical Shape Models (SSMs)* has shown to be very effective in guiding highly uncertain image segmentation problems. In this paper, we construct multiple-structure SSMs of purely geometric nature, that describe the relationship between adjacent anatomical components through Canonical Correlation Analysis. Shape inference is then conducted based on a regularization term on the shape likelihood providing more reliable structure representations. A fundamental prerequisite for performing statistical shape analysis on a set of objects is the identification of corresponding points on their associated surfaces. We address the correspondence problem using the recently proposed *Functional Maps* framework, which is a generalization of point-to-point correspondence to manifolds. Additionally, we show that, by incorporating techniques from the *deep learning* theory into this framework, we can further enhance the ability of SSMs to better capture the shape variation in a given dataset. The efficiency of our approach is illustrated through the creation of 3D models of the human knee complex in two application scenarios: incomplete or noisy shape reconstruction and missing structure estimation.

1. Introduction

The reconstruction of geometric shapes plays an important role in many fields such as computer vision, augmented and virtual reality, personalized computer-aided intervention, robotic mapping, and other. If 3D scans are available, as in the case of medical imaging, automated image segmentation is usually performed to localize and extract the object of interest. However, in real case scenarios with noise and intrinsic artifacts, such as pathologies, most of the automated algorithms produce invalid geometric representations if they rely solely on the image content. Image segmentation methods are also especially sensitive to missing or incomplete data producing unrealistic shapes.

One way to reduce uncertainty in estimation is through the incorporation of prior information on the shape variability of anatomical structures, known as *statistical shape analysis* (Dryden and Mardia, 1998; Goodall, 1991). The use of *Statistical Shape Models (SSMs)* in image analysis has been well established more than a decade ago, with the most popular variants being the Active Shape Models and Active Appearance Models that, in addition to the expected shape, represent also the texture (complete appearance) of the volumetric object

(Heimann and Meinzer, 2009). Despite their success, the main disadvantage of these methods is the excessive memory usage in case of high texture resolution that often requires to be scaled down radically. On the other hand, purely geometrical shape models, often referred to as Point Distribution Models (PDMs) (Cootes et al., 1995) are more intuitive, easy to implement, reasonably robust, and fast.

Lately, there has been a growing interest in the construction of models that consist of multiple anatomical structures (Cerrolaza et al., 2015, 2016, 2019; Saito et al., 2013). The key concept in such methods is to overcome the limitations of the small availability of training data, by considering the inter-relation between neighboring structures, i.e., they describe how the shape variation of one structure affects the shape of the other, and vice versa. This can potentially lead to more efficient and accurate shape representation, while also enabling us to conduct shape inference about adjacent structures.

In this work we employ such methods for building statistical shape models of purely geometric nature and demonstrate their application in volumetric data, such as medical images. In particular, we are interested in constructing SSMs that capture the shape variation of multiple anatomical components, while also encoding the correlation between

* Corresponding author. Visualization and Virtual Reality Group, Department of Electrical and Computer Engineering, University of Patras, Rio Achaia, 26504, Greece. Tel.: +302610969809; fax: +302610991855.

E-mail addresses: filip.k@ece.upatras.gr (K. Filip), ezachar@upatras.gr (E.I. Zacharaki), moustakas@upatras.gr (K. Moustakas).

<https://doi.org/10.1016/j.compmedimag.2021.101890>

Received 16 September 2020; Received in revised form 13 February 2021; Accepted 19 February 2021

Available online 25 February 2021

0895-6111/© 2021 Elsevier Ltd. All rights reserved.

neighboring structures. Moreover, by using more sophisticated methods like the manifold representation and deep learning techniques, we are able to better capture the variation of an object's class.

1.1. Shape correspondence

The fundamental requirement in constructing an accurate shape model, is to first establish correspondence among the elements of the training shapes. The *correspondence problem*, or *shape matching*, is the most crucial task in order to capture true geometric variation, as false identification of point pairs may lead to unnatural shapes.

Solving the problem of surface matching is an ambitious goal and is still not adequately addressed, although, there are many techniques that tackle effectively this problem (Biasotti et al., 2016; Van Kaick et al., 2011; Sahillioğlu, 2020). For the construction of SSMs, the most popular shape matching method is the Iterative Closest Point (ICP) (Rusinkiewicz and Levoy, 2001) which is based on repetitive rigid transformations and assignment of closest points, until convergence to a local minimum. However, proximity-based methods, are often insufficient to correctly identify corresponding points. Therefore, the focus of shape matching methods is on computing correct correspondences from an invariant and robust point of view.

One promising framework was proposed by Ovsjanikov et al. called *Functional Maps (FM)* (Ovsjanikov et al., 2012), which is a generalization of the notion of point-to-point shape correspondence. The framework describes how mappings act on real-valued functions defined on manifolds and allows the computation of point correspondences between two shapes in a non-rigid fashion and independently of their spatial orientation, assuming that the shapes undergo (near-) isometric deformations. This method has achieved state-of-the-art performance in shape correspondence benchmarks for both full and partial shapes (Rodolà et al., 2017; Kovnatsky et al., 2015, 2013; Rodolà et al., 2017; Litany et al., 2016, 2017; Ren et al., 2018), but it has rarely been employed for estimation and refinement of anatomical structures in challenging situations including image noise, artifacts, and partial information.

Recently, there has been a growing interest in techniques that attempt to generalize deep neural networks to non-Euclidean domains like manifolds, which are collectively referred to as *geometric deep learning* (Bronstein et al., 2017). For shape matching, it has been shown that the extracted information from shape data can be used in order to compute more accurate point-to-point correspondences (Monti et al., 2017; Masci et al., 2015; Boscaini et al., 2016; Litany et al., 2017a; Roufousse et al., 2019).

1.2. Goals and contributions

The aim of this work is to investigate whether image-guided shape reconstruction methods can be further improved if SSMs are incorporated for refinement of the reconstructed shape, where we propose a regularized multi-structure statistical shape modeling approach for estimation of partial or degenerate 3D data. Specifically, we address the issue of topological alterations (noise, artifacts, missing parts) in the segmentation outcome by formulating an optimization function that balances between the original subject-specific shape and a prior likelihood term. The advantage of this approach is two-fold:

- We address cases where the full multi-shape model is fitted to new shapes (obtained from automatic image segmentation methods) that do not represent anatomically correct structures due to inaccuracies in image segmentation, such as in regions with low intensity contrast, inhomogeneity or imaging artifacts, resulting in missing parts or protrusions in the segmented image. A regularization term on the solution space can potentially prevent overfitting of the SSMs to such defective shapes.

- We address cases where whole structures are missing (not detected) due to the low sensitivity of some imaging modalities to depict certain tissues. For example, the relationship of a tumor to adjacent normal structures, including joints and neurovascular structures, is better assessed with MRI, whereas CT is superior in visualizing calcific deposits and pathologic fractures (Zimmer et al., 1985). Multi-structure SSMs, once constructed, have the potential to infer unknown or missing structures (due to availability of a single modality) through shape correlation analysis, by exploiting the observed shape conformations of one structure to approximate the shape of another highly correlated neighboring structure.

We examine the application of our approach in data of the human knee complex. While there have been various approaches on the construction of statistical shape models of the knee joint (Rao et al., 2013; Fitzpatrick et al., 2011; Baldwin et al., 2010; Williams et al., 2010; Bredbenner et al., 2010), to our best knowledge, it has never been carried out in such way.

To summarize, this work makes the following contributions to the field of 3D shape modeling:

- We employ the recently proposed framework of Functional Maps for solving the isometric shape matching problem, and we also incorporate deep learning techniques to enhance the ability of SSMs to capture shape variation.
- We construct multiple-structure shape models that encode the relationship between neighboring structures, and we jointly optimize reconstruction performance and shape likelihood to improve the quality of automated segmentation of the knee complex.
- Finally, we exploit the knowledge captured by the model regarding the inter-dependence of related structures through Canonical Correlation Analysis, to conduct inference about unknown or missing structures.

2. Methods

An overview of our approach is presented in Fig. 1, which consists of the training and testing phase. The training phase involves the calculation of the parameters of (i) the Deep Functional Maps network that finds point correspondences between two shapes, (ii) the SSM model that can be used to produce new realistic shapes for multiple structures, and (iii) the correlation matrix expressing the relationship between adjacent structures. The testing phase involves the search for the parameters \mathbf{b} of the built SSM and a global rigid transformation T that best describe a new shape \mathcal{S}_S acquired from a segmentation mask \mathcal{M}_S . Then inference of the new shape is performed by minimizing the following cost function:

$$(T, \mathbf{b}) = \arg \min_{T, \mathbf{b}} [d(\mathcal{S}_S, T \circ \mathcal{S}_M(\mathbf{b}|\mathbf{P})) + w\mathcal{R}(\mathbf{b})] \quad (1)$$

where \mathcal{S}_M denotes the reconstructed shape that depends on the parameters \mathbf{b} of the shape model, given a set of orthogonal basis \mathbf{P} that describe the shape variation; $d(\cdot)$ denotes a dissimilarity metric between two shapes (smaller values indicate higher similarity); w controls the weight of a regularization term \mathcal{R} on the shape variation. Next, we discuss separately the methods implemented for the shape matching problem and those involved with the construction of SSMs and how they are applied to new input data.

2.1. Estimation of shape correspondence

Manifolds. We model shapes as two-dimensional Riemannian manifolds \mathcal{X} possibly with boundary $\partial\mathcal{X}$. Given scalar functions $f, g : \mathcal{X} \rightarrow \mathbb{R}$ defined on a manifold \mathcal{X} , the standard inner product is defined as $\langle f, g \rangle_{\mathcal{L}^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$, where dx denotes the area element induced by

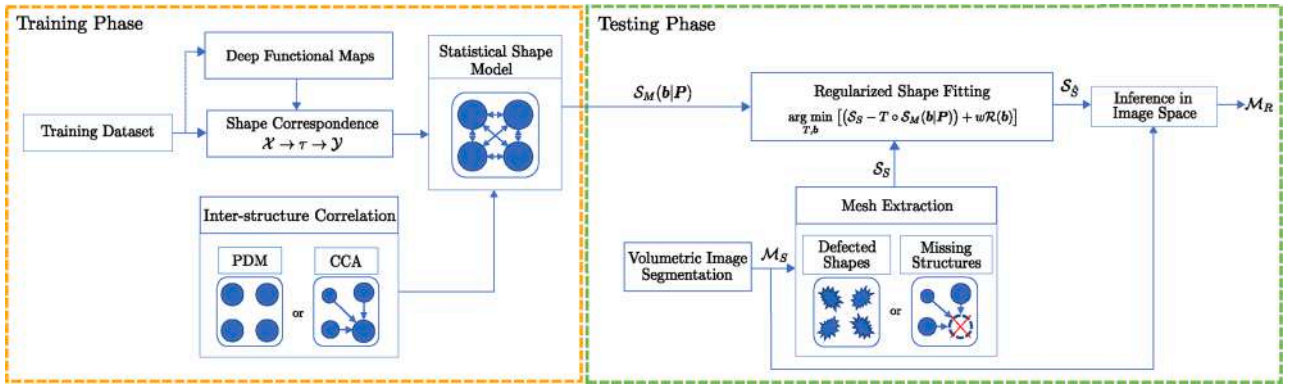


Fig. 1. Schematic illustration of the statistical modelling methodology. Training phase: Given a population of shapes, point correspondence across shapes is calculated based on the Deep Functional Maps framework. The network learns correspondences from training examples. Then, a multi-structure SSM is constructed, parameterized to mimic other approaches based on the inter-correlation of adjacent structures. Testing phase: The SSM is fitted on potentially defected or partial shapes S_S acquired from 3D image segmentation masks M_S using a regularization scheme, and shape inference is then conducted in image space resulting in a redefined image M_R .

the Riemannian metric. We denote with $\mathcal{L}^2(\mathcal{X}) = \{f : \mathcal{X} \rightarrow \mathbb{R} \mid \langle f, f \rangle_{\mathcal{L}^2(\mathcal{X})} < \infty\}$ the space of square-integrable functions on \mathcal{X} .

For shape representation we will utilize the positive semi-definite Laplace-Beltrami (LB) operator Δ , which generalizes the Laplacian to manifolds. It admits to an eigendecomposition problem $\Delta\phi_i = \lambda_i\phi_i$, where the eigenvalues λ_i form a discrete spectrum ordered from “low frequency” to “higher frequency” ($\lambda_0 = 0 \leq \lambda_1 \leq \lambda_2 \leq \dots$), and the eigenfunctions ϕ_0, ϕ_1, \dots form an orthonormal basis on $\mathcal{L}^2(\mathcal{X})$ (i.e., $\langle \phi_i, \phi_j \rangle_{\mathcal{L}^2(\mathcal{X})} = \delta_{ij}$), allowing any function $f \in \mathcal{L}^2(\mathcal{X})$ to be expanded into a *Fourier series* as:

$$f = \sum_{i \geq 0} \langle f, \phi_i \rangle_{\mathcal{L}^2(\mathcal{X})} \phi_i \quad (2)$$

Functional Maps. FM (Ovsjanikov et al., 2012) are a generalization of the notion of point-to-point shape correspondence, as they describe how mappings act on real-valued functions defined on shapes. Let $\tau : \mathcal{X} \rightarrow \mathcal{Y}$ be a map between manifolds \mathcal{X} and \mathcal{Y} . The map τ induces a *linear* transformation $\tau_F : \mathcal{L}^2(\mathcal{X}) \rightarrow \mathcal{L}^2(\mathcal{Y})$ on the functions $f : \mathcal{X} \rightarrow \mathbb{R}$ and $g : \mathcal{Y} \rightarrow \mathbb{R}$ defined on manifolds \mathcal{X} and \mathcal{Y} , respectively, such that $\tau_F(f) = g$.

Assuming the two orthogonal sets of basis $\{\phi_i\}$ and $\{\psi_j\}$ on $\mathcal{L}^2(\mathcal{X})$ and $\mathcal{L}^2(\mathcal{Y})$, respectively, to be given, the functional map τ_F can be represented as a matrix C with coefficients $c_{ij} = \langle \tau_F(\phi_i), \psi_j \rangle_{\mathcal{L}^2(\mathcal{Y})}$ leading to an expansion:

$$\tau_F(f) = \sum_{i,j \geq 0} \langle f, \phi_i \rangle_{\mathcal{L}^2(\mathcal{X})} c_{ij} \psi_j \quad (3)$$

Truncating the Fourier series at the first k coefficients, functional correspondence becomes the problem of finding a matrix C of size $k \times k$ which is a rank- k approximation of the original map. Given a set of q corresponding functions $f_i \in \mathcal{L}^2(\mathcal{X})$ and $g_i \in \mathcal{L}^2(\mathcal{Y})$, with spectral representations $\mathbf{a}_i = (\langle f_i, \phi_0 \rangle_{\mathcal{L}^2(\mathcal{X})}, \dots, \langle f_i, \phi_k \rangle_{\mathcal{L}^2(\mathcal{X})})$, $\mathbf{b}_i = (\langle g_i, \psi_0 \rangle_{\mathcal{L}^2(\mathcal{Y})}, \dots, \langle g_i, \psi_k \rangle_{\mathcal{L}^2(\mathcal{Y})})$ stored as columns in matrices A and B , respectively, the problem of functional correspondence boils down to a simple linear system:

$$C = \arg \min_C \|CA - B\|_F^2 \quad (4)$$

In practice, f_i and g_i are q -dimensional shape descriptors (e.g., HKS (Sun et al., 2009), SHOT (Tombari et al., 2010)) that capture some shape properties around the neighborhood of a point of interest. In order for the system in Eq. (4) to be (over)-determined, we must select $k \leq q$.

Discretization. In the discrete setting, where the manifolds are represented as triangle meshes with n vertices, the discrete LB operator can be expressed as $\Delta = S^{-1}W$ using the classical *cotangent weights*

scheme (Reuter et al., 2009); W is the $n \times n$ *stiffness matrix* of the edge weights, and $S = \text{diag}(s_1, \dots, s_n)$ is the diagonal *mass matrix* of vertex area weights. The first k eigenfunctions are found solving the eigendecomposition $\Delta\Phi = \Lambda\Phi$, where Φ is a $n \times k$ matrix containing as columns the first k eigenfunctions and $\Lambda = \text{diag}(\lambda_0, \dots, \lambda_{k-1})$ is the diagonal matrix of the corresponding eigenvalues. A function on the manifold is represented by an n -dimensional vector $f = (f(x_1) \dots f(x_n))$. The Fourier coefficients are then computed as $\mathbf{a} = \Phi^T S f$.

2.1.1. Estimating the functional map

Assuming that the shapes undergo (near)-isometric deformations, the underlying functional map is volume preserving, i.e., the matrix C is orthogonal $C^T C = I$. Then:

$$C = \arg \min_C \|CA - B\|_F^2, \quad \text{s.t. } C^T C = I \quad (5)$$

This problem is referred to as the *orthogonal Procrustes* problem (Kovnatsky et al., 2013) which has a closed form solution obtained through Singular Value Decomposition (SVD) of $M = BA^T$ as $C = U_M V_M^T$, where the columns of U_M and V_M are the left-singular vectors and right-singular vectors, respectively.

There are several regularizations and extensions that can help to improve this estimation significantly. For example, since the eigenfunctions diagonalize the Laplacian, $\langle \phi_i, \Delta\phi_j \rangle_{\mathcal{L}^2(\mathcal{X})} = \lambda_i \delta_{ij}$, the Dirichlet energy can be expressed as: $\text{tr}(\langle \tau_F(\phi_i), \Delta\tau_F(\phi_j) \rangle_{\mathcal{L}^2(\mathcal{X})})$ which is written in matrix form (for all eigenfunctions) as $\text{tr}(C^T \Lambda^X C)$ (Ovsjanikov et al., 2017). By adding this term in Eq. (5), smoothness on the functional map is enforced:

$$C = \arg \min_{C \in S(k,k)} \|CA - B\|_F^2 + \alpha \text{tr}(C^T \Lambda^X C) \quad (6)$$

where Λ^X is the diagonal matrix of the eigenvalues of LB operator of manifold \mathcal{X} , α is a scalar weight parameter, and $S(k, k) = \{X \in \mathbb{R}^{k \times k} : X^T X = I_k\}$ denotes the *Stiefel manifold* of $k \times k$ orthogonal matrices. The above problem is an instance of *manifold optimization* and can be solved using efficient numerical techniques (Boumal et al., 2014).

In both cases, the resulting matrix C can be interpreted as an alignment of the spectral representations of manifolds \mathcal{X} and \mathcal{Y} , and it can be further processed using the refinement technique described in Ovsjanikov et al. (2012). Point-to-point correspondence is then obtained by nearest neighbor assignment in the embedded functional space.

Deep Functional Maps. In cases where the shapes undergo deformations that are not strictly isometric, the methods described by the Eq. (5) or (6) do not necessarily result in the best solution that leads to

correct shape correspondence. By incorporating deep learning techniques, it is possible to extract features from the functions defined on the manifolds which, when used in conjunction with the functional map framework, will define more accurate correspondences.

In this work, we make use of the neural network presented by Litany et al. (Litany et al., 2017) called *FMNet*. Manually designed point-wise descriptors F and G from a pair of shapes are passed through a Siamese network with Residual blocks (He et al., 2016), resulting in refined descriptors \hat{F} , \hat{G} . These, in turn, are projected onto the Laplacian eigenbasis to produce the spectral representations \hat{A} , \hat{B} , which are used to solve Eq. (5). *FMNet* is trained using a specialized soft error loss, which punishes geodesically distant predictions stronger than predictions near the correct node.

2.2. Reconstruction of partial or inconsistent shapes by regularized multi-structure SSMs

Statistical Shape Models. SSMs are discrete, parametric models that capture the most prominent geometric variation in a given set of shapes. Each shape \mathcal{X}_i in the dataset is represented in terms of a vector of landmarks x_i . The fundamental prerequisite of building SSMs is the identification of corresponding landmarks on each surface. Shape analysis is then performed jointly by first aligning the shapes using *Generalized Procrustes Analysis* (Gower, 1975) (Goodall, 1991) to remove rigid or similarity transformations, and then using *Principal Component Analysis* (PCA) to extract the main modes that best describe the observed variation. Alternatively, *Robust PCA* (RPCA) Candès et al. (2011) can recover the principal modes of variation even if a fraction of computed correspondences is arbitrarily corrupted.

Multi-structure SSMs. The previous procedure can be extended in the case of multiple structures by taking into consideration the individual shapes' correlation. Let $x_i \in \mathbb{R}^{3n_i \times 1}$ represent the vector form of a single structure i in a set of m structures, with n_i landmarks. Then, the multi-structure shape can take the vector form

$$X = \{x_1, \dots, x_m\}, \quad X \in \mathbb{R}^{3n \times 1}$$

through the concatenation of the m structure vectors, with $n = \sum_{i=1}^m n_i$ the total number of points. If N such multi-structure shapes (observations) exist, let $L \in \mathbb{R}^{N \times 3n}$ represent the centered data matrix (i.e., with zero mean) containing the vectors X as rows.

In order to account for structure inter-dependence in the constructed multiple-structure SSMs, we employ the *Soft Multi-Organ Shape models* (SOMOS) framework (Cerroloza et al., 2016), which uses the weighted

variant of the generalized SVD (G-SVD) (Greenacre, 2021) of L to impose constraints on the left and right singular vectors. The matrix L can be written as $U_L \Sigma_L V_L^T$, where Σ_L is a diagonal matrix, and U_L, V_L are (not necessarily) orthonormal matrices that satisfy $U_L^T W U_L = V_L^T \Omega V_L = I$, with W and Ω being specified positive-definite symmetric matrices. G-SVD solves the regular SVD of:

$$\tilde{L} = W^{1/2} L \Omega^{1/2} \quad (7)$$

This framework allows the creation of a statistical model for each of the m structures individually (instead of a single global model), while also taking into account the relationship (e.g., correlation) between the structures. We describe below the procedure for the SSM creation of each individual structure by omitting the index of this structure for simplicity. For the examined structure, we can select the matrix Ω to be diagonal with entries $w_i \in [0, 1]$ representing the correlation of the examined structure with structure x_i . To allow multiplication with the data matrix, the values w_i are replicated for all points in the structure x_i . Moreover, assuming that all observations are equally important, identical weights are assigned such that $W = I$. Then, Eq. (7) becomes $\tilde{L} = L \Omega^{1/2}$, which corresponds to a weighted variant of the regular PCA. Fig. 2 shows a visual representation of matrices L and Ω in order to aid reader's understanding. In case of correlated structures (i.e., $w_i > 0, \forall i$), the matrix L can then be expressed as:

$$L = \tilde{L} \Omega^{-1/2} \Rightarrow U_L \Sigma_L V_L^T = U_L \Sigma_L V_L^T \Omega^{-1/2} \quad (8)$$

where Σ_L is a diagonal matrix whose diagonal elements are the singular values of the eigendecomposition of \tilde{L} , and U_L, V_L are unitary matrices whose columns are the left-singular and right-singular vectors, respectively. Thus:

$$V_L = \Omega^{-1/2} V_L^T, \quad \Sigma_L = \Sigma_L, \quad U_L = U_L^T \quad (9)$$

Like in the classic PCA-based PDM, each shape can be modeled as the best approximation of Y by a linear combination of the matrix P_c , which is sub-matrix of V_L with rank c that describes the desired (e.g., 98%) shape variation,

$$Y = (y_1; \dots; y_m)^T \approx \tilde{Y} = \bar{X} + P_c b \quad (10)$$

where the vector b is obtained as:

$$b = P_c^T \Omega^{1/2} (Y - \bar{X}) \quad (11)$$

This approach allows us also to replicate other approaches simply by




Fig. 2. Visual representation of matrices L and Ω . The columns of matrix L consist of the vector form of each structure in the training dataset with size N . Thus, each row in the matrix represents the vector form of each of the m structures comprising the multi-structure shape. In order to form the matrix \tilde{L} , the columns of matrix L are multiplied with the square root of w_i , which represents the correlation between the selected structure and structure i .

defining different inter-relation weights between the structures. For example, if the weight between structure i and j is $w_{ij} = 1, \forall i, j$ the problem becomes equivalent to the original PDM for multiple shapes, whereas if the weight is 1 only for $i = j$ and 0 otherwise, the problem becomes equivalent to independent modeling of each structure.

Canonical Correlation Analysis. The correlation factors between two structures are calculated using *Canonical Correlation Analysis (CCA)*. CCA determines the linear combination of the components in structures x_i and x_j that are maximally correlated with each other. The overall inter-structure correlation factor, w_{ij} , is defined as the average correlation coefficient over all calculated canonical modes (Cerrolaza et al., 2016).

Prior to CCA, it is required to perform PCA on the point coordinates for each individual structure across the training set to reduce the dimensionality of the data, ensuring that a percentage (95%) of the variation is retained. The dimensionality reduction minimizes the computational memory burden and also eliminates co-linearity in the data which can cause instability in the calculation of CCA (Rao et al., 2008).

Regularization of Estimated Shapes. In case of missing structures or partial shapes, the method described by Eq. (11) for approximating a new shape can not be used, since P_c forms an orthonormal basis only for the entire shape Y . That is, if we denote by $Y_i = (y_1; \dots; y_{i-1}; y_{i+1}; \dots; y_m)^T$ a new shape containing the coordinates of the $m - 1$ out of m structures, then projection on the truncated eigenvectors P_c^i will not provide the correct coefficients vector, since P_c^i does not necessarily constitute an orthogonal basis.

The problem of finding the vector b that represents an expansion of Y_i in terms of P_c^i can be formulated as an optimization problem, as proposed in a different setting (Mohamed et al., 2001), in which b is sought so that it yields a shape that fits Y_i but has also high likelihood expressed by:

$$g(b) = \frac{1}{\sqrt{(2\pi)^c \det(\Sigma_c)}} \exp\left(-\frac{b^T \Sigma_c^{-1} b}{2}\right) \quad (12)$$

which describes the probability density function of the coefficient vector b , assuming that the shapes x_i follow the Gaussian distribution. The matrix Σ_c is the c -rank submatrix of Σ containing the eigenvalues of the corresponding modes of variation (Eq. (9)). A solution is then given by minimizing the cost function:

$$\begin{aligned} \mathcal{E}(b) &= \|\hat{X}_i - Y_i\|^2 + w \frac{1}{g(b)} \\ &= \|\bar{X}_i + P_c^i b - Y_i\|^2 + w \frac{1}{g(b)} \end{aligned} \quad (13)$$

where $1/g(b)$ expresses the regularization term in Eq. (1), and w is the relative weighting factor. The first term in Eq. (13) seeks vectors that get as close as possible to the subject's observed shape Y_i , whereas the second term favors shape representations that are according to what has been observed in the training samples. The nonlinear Levenberg-Marquardt (Marquardt, 1963) optimization scheme is used to minimize this cost function. Since the parameters b determine also the global shape model described by Eq. (10), they can be used to provide an approximation of the missing structure y_i .

The same optimization scheme can also be used in cases of defected or partial shapes Y that contain all structures, in order to avoid unnatural representations due to overfitting.

Shape Fitting. Suppose we wish to find the best pose and shape parameters to match a shape model to a new set of points Y . This is achieved by minimizing the sum of square distances between the model and the target shape:

$$\|Y - T(\bar{X} + P_c b)\| \quad (14)$$

where T is a rigid transformation defining the translation and orientation of the estimated model. An efficient iterative approach for achieving this, is described by the algorithm presented by Cootes et al. in Cootes et al. (2004) that uses repetitive transformations and matching of closest points, thereby mimicking the ICP (Rusinkiewicz and Levoy, 2001) registration method. We modify the algorithm by replacing the step of finding the values of the vector b describing the target shape y , with that in Eq. (13) (*regularized shape fitting*) or Eq. (11) (*exact shape fitting*) if $w = 0$.

The incorporation of the FM framework to the iterative process of shape fitting has a higher potential in the identification of correspondences. It is however less efficient computationally, since it requires, not only the solution of the linear algebraic system and the refinement of the resulting functional map, but also the additional calculation of the shape descriptors and LB eigenfunctions in every iteration.

3. Implementation

The previously described methods find several application domains, however our main focus was the generation of robust simulation models for osteoarthritis (a degenerative disease of the articular cartilage) for personalized treatment design. In particular, we discuss the application of the multi-structure SSMS for improving the results of automatic segmentation of the human knee extracted from MRI data. The knee complex consists of many different biological tissues, which are difficult to depict accurately in a single image sequence. Especially the regions representing the cartilage and connective tissue are usually difficult to distinguish due to low intensity contrast. Thus, the incorporation of a good shape prior can be crucial for conducting shape inference in ambiguous regions of the segmented mask, while also being able to transfer across different image modalities.

Datasets. For the evaluation of the methods, we selected two separate datasets of annotated MRI of the human knee complex.

- 1 OAI ZIB¹: The dataset consists of MRI sequences from the Osteoarthritis Initiative² database for which 507 manual segmentations of five anatomical structures of the right knee – Femur Bone (FB), Tibia Bone (TB), Femoral Cartilage (FC), Medial Tibial Cartilage (TCM) and Lateral Tibial Cartilage (TCL) – were carried out by experienced users at Zuse Institute Berlin. The dataset is split into 253 training and 254 testing subjects, as previously performed by Ambellan et al. (2019).
- 2 OpenKnee³: The dataset contains knee MRIs from a representative sample of young and elderly, male and female, and healthy and osteoarthritic subjects. MR images and corresponding segmentations from only 7 subjects were complete, thus this dataset was used just for evaluation purposes (and not for training), focusing on the same (5) anatomical structures as for the OAI ZIB dataset.

Mesh Generation and Preprocessing. We construct triangle meshes from the given manual segmentation masks using the Marching Cubes algorithm (Lorensen and Cline, 1987). The meshes represent the different components of the knee structure. In order to alleviate the requirement for large computation power and memory footprint, the meshes are simplified (downsampled) using the *Quadric Edge Collapse Decimation* algorithm provided in Meshlab⁴. The target number of resulting faces for each mesh was kept constant ($\sim 8K$ faces, which is equivalent to $\sim 4K$ vertices).

FMNet Implementation. The FMNet implementation and

¹ <https://opus4.kobv.de/opus4-zib/frontdoor/index/index/docId/6995>

² <https://nda.nih.gov/oai/>

³ <https://simtk.org/projects/openknee>

⁴ <http://www.meshlab.net>

parameters are as described in Litany et al. (2017). The network extracts features from 352-dimensional SHOT (Tombari et al., 2010) descriptors defined on meshes, and uses those to find the mapping between a pair of shapes as described in Section (2.1.1). As orthonormal basis for the functional space of the surfaces we consider the eigenfunctions of the discrete LB operator derived from the Linear FEM scheme with the lumped mass matrix (Reuter et al., 2009). The first 35 eigenfunctions for each structure are kept, which proved sufficient for the functional map C to represent the ground truth map to a quality that is very close to the ground-truth point-to-point map.

In every training iteration, the network takes as input concatenated pairs of $\sim 2K$ randomly chosen ground-truth matching descriptors (1K from each shape). Training is performed separately for all the anatomical structures. Thus in our case, five different versions of the neural network are constructed, one for each structure of the knee complex.

The training session lasted 10 – 20 minutes (200–300 iterations). The average prediction runtime for the computation of the functional map C for a new pair of shapes was approximately 0.12 seconds using CUDA 9.0 (CPU: Intel Core i7-4790 3.60 GHz; GPU: NVIDIA GeForce GTX 660 2 GB).

Creation of ground-truth correspondences. As any supervised learning method, in the training phase FMNet requires examples of pairs for which the ground-truth correspondence of their elements is known [Fig 3]. In the case of medical data, manual labeling for all corresponding points is infeasible and impractical, and thus we rely on other automatic methods for defining the correspondences. In order to create training examples of paired descriptors, 50 subjects within the training set of the OAI ZIB dataset are randomly selected, for which the standard ICP method (Rusinkiewicz and Levoy, 2001) is used to compute the mappings between the elements (vertices) of the paired meshes. A total of $50 \times 50 = 2500$ training examples were constructed for each knee

component, by taking every possible combination of subjects within the dataset.

Although the definition of correspondences using an automatic method, instead of manual labeling, reduces the quality of the ground truth (possibly affecting the accuracy of the network to learn anatomically correct matches), in our experiments, we have observed that the network generalizes well even with imprecise input examples. This occurs mainly due to the fact that the randomly selected training examples can be bidirectional, meaning that, for a pair of shapes $\langle \mathcal{X}, \mathcal{Y} \rangle$, the mappings $\tau_i : \mathcal{X} \rightarrow \mathcal{Y}$ and $\tau_j : \mathcal{Y} \rightarrow \mathcal{X}$ can be used in both directions during the training session. The mappings created from the ICP are not invertible. However, the network adjusts accordingly to compensate for the matching points that differ if the input mapping is inverted.

Perhaps a better way of creating ground-truth correspondences would be by exploiting volumetric anatomical information from image registration algorithms that are subject to regularization constraints. Such solutions take advantage of the rich local information content and therefore are expected to be anatomically much more valid than simple geometrical criteria of extracted surfaces. However, methods that address this problem are beyond the scope of this work which focuses on 3D geometrical objects in the form of point clouds or meshes and does not require the availability of volumetric image intensity information.

Construction of Statistical Shape Models. The definition of a common set of vertices across different shapes is a necessary step in the construction of SSMs. Therefore the order of vertices is fixed according to that of a reference subject within the OAI dataset (chosen to be close to the average subject) that does not exhibit a severe pathological condition.

Next, all knee structures for each subject are concatenated and aligned jointly as a single point cloud by performing Generalized Procrustes Analysis. The modes of shape variation are subsequently



Fig. 3. Correspondence results obtained by the FMNet model on a pair of femoral cartilages. Corresponding points are assigned the same color. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.) The resulting functional map C (right) is sparse and diagonally dominant, indicating that the two shapes are approximately isometric.




Fig. 4. Correlation coefficients resulting from the Canonical Correlation Analysis (CCA) of the five knee structures – Femur Bone (FB), Tibia Bone (TB), Femoral Cartilage (FC), Medial Tibial Cartilage (TCM) and Lateral Tibial Cartilage (TCL).

computed with the weighted variant of the classical PCA and RPCA, separately. Depending on the choice of inter-structure relation weights w_{ij} , two separate cases can be examined.

- If the matrix $\Omega = I$ (in Eq. (9)), a single multi-shape PDM model is constructed, containing all the knee components. The resulting modes describe the shape variation of all the structures combined.
- If the matrix Ω contains the correlation coefficients resulting from the CCA (Fig. 4), five individual multi-shape SSMs (i.e., as many as the number of anatomical structures) are constructed, where each SSM describes the variation of a particular structure and its relationship with the other structures.

The constructed models are fitted to a new shape complex. [Fig.5] To

evaluate the performance of the constructed SSMs, the fitted shapes are first transformed from 3D triangle meshes back to a multi-labeled volumetric image based on ray intersection methods (Patil and Ravi, 2005), and then they are compared with the ground-truth labelmap using appropriate measures. In many cases, there appear overlaps or intersections between the fitted shapes of the different components of the multi-shape model. This issue can be attributed to the representation of the multiple structures with a single point cloud in which some of the structures are in contact, and therefore strongly affected by small deformations produced during reconstruction. Such conflicts are resolved in the volumetric image representation of the fitted SSM by assigning the overlapping area to the smaller structures (i.e., the cartilages). Although this approach might seem naïve, in practice the conflict areas are small in volume (one voxel layer between two structures), and thus in most




Fig. 5. Visualization of the reconstructed structures (colored contours) overlaid with the ground-truth (colored regions) (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.) in the volumetric image representation.

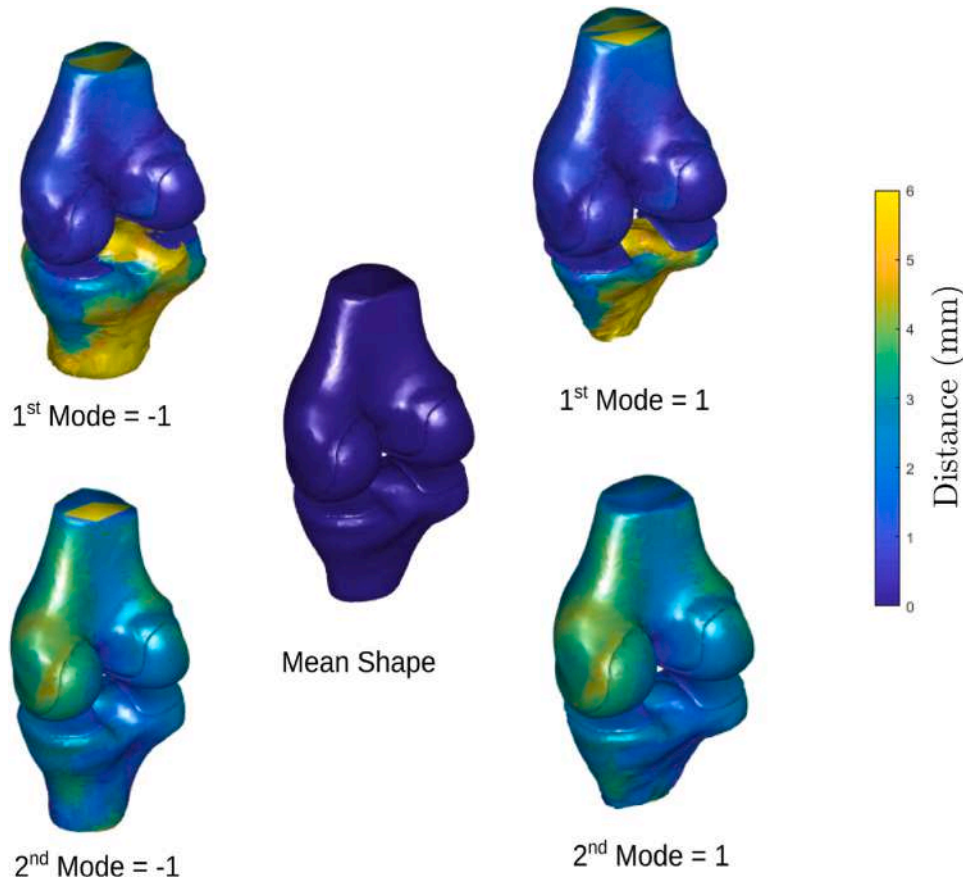


Fig. 6. Statistical shape model of the human knee joint consisting of five different structures. A new shape can be estimated by changing the variable b in $\hat{X} = \bar{X} + P_c b$ (shape differences are illustrated based on the distance, measured in mm , from the mean shape ($b = 0$) shown at the center). Top row shows how the first mode of variation ($b_1 = \pm 1$) affects the estimated shape. Bottom row shows how the second mode of variation ($b_2 = \pm 1$) affects the estimated shape.

cases there is no significant effect on the accuracy.

SSM-based regularization of segmentation results. One common application of SSMs is to regularize the results of automatic segmentation methods [Fig. 6]. Let, \mathcal{M}_S denote the initial segmentation mask depicting all the knee compartments of interest. After segmentation of the knee compartments in the image space, the previously constructed PCA-based and RPCA-based SSMs are fitted to meshes extracted from the segmentation masks, and then transformed back to a multi-labeled volumetric image \mathcal{M}_{SSM} .

We seek for a solution that refines the initial segmentation mask \mathcal{M}_S based on the incorporated shape model \mathcal{M}_{SSM} and some topological and regularization constraints in the image space. Specifically, we apply the following steps:

- Fill holes in \mathcal{M}_S that are located within the contour of the fitted model.
- Remove unconnected components of \mathcal{M}_S that are located outside of \mathcal{M}_{SSM} .
- Remove voxels that have distance larger than the 95% percentile of the calculated distances between \mathcal{M}_S and \mathcal{M}_{SSM} .

4. Results

This section is devoted to the evaluation of the individual components of our methods, and their application in two biomedical applications. Note that, even if we focus particularly on the knee complex, the methods are general and can easily be applied to other datasets as well.

4.1. Effect of registration technique on shape variation

The first of our series of experiments is dedicated in examining the effect of the different shape matching techniques on the ability of the constructed SSMs to capture the shape variation in a given dataset. We analyze three separate approaches based on the Functional Maps framework. Namely, we examine the following cases:

- Solving FM as the Orthogonal Procrustes problem (Eq. (5)).
- Solving FM as the Manifold Optimization problem (Eq. (6)).
- Solving FM using the FMNet model (Litany et al., 2017).

All results are compared with those calculated with the ICP method (Rusinkiewicz and Levoy, 2001), which is one of the most widely used methods for point-set registration. All the FM-based approaches are post-processed with the correspondence refinement technique described in Ovsjanikov et al. (2012).

Upon the calculation of shape correspondences, multi-shape SSMs are constructed with the 253 subjects of the OAI ZIB training set, without considering inter-structure correlation (i.e., $\Omega = I$), using the classical PCA and RPCA approaches. The produced models are applied to reconstruct the remaining 254 subjects in the testing set. The shape

matching performance is evaluated by the capability of the constructed models to represent new shapes. In Fig. 7 we present the results of the average DICE scores for the different shape correspondence techniques using the PCA-based and RPCA-based SSMs, respectively.

The first observation is that the FMNet model produces better results compared to any other method used in this experiment. Although all methods are based on the FM framework, we see that the ability of the neural network to extract features from the functions defined on the manifolds and use them to compute correspondences, increases significantly the generalization properties of the constructed SSMs at the expense of increased computational complexity (Table 2). Another interesting aspect of the FMNet is that even though it is trained on examples constructed using the ICP method, it surpasses the performance of the ICP in explaining shape variation. This observation is more prominent for the cartilage structures, which are rather thin-structured and small in volume, indicating that standard methods like the ICP are not the most appropriate for computing point correspondences on more subtle shape variation.

In addition, we compared the compactness of the PCA-based and RPCA-based SSMs using the four described correspondence methods. The results are presented in Fig. 8 and show that the ICP method presents higher compactness for both PCA and RPCA-based SSMs, whereas the FMNet achieves intermediate level of compactness for PCA-based models, and slightly lower level for the RPCA-based models. The difference in compactness is mainly attributed to the first component which explains significantly higher variance when ICP is used. However, this does not necessarily mean that the explained variance originates from correctly computed matches across the training shapes using the ICP.

Finally, we provide measurements on the specificity index for the constructed PCA-based and RPCA-based SSMs, respectively. First, we present the specificity index versus the increasing number of modes used to generate random shape instances from the constructed SSMs, illustrated in Fig. 9. A total of 500 random shape instances are generated for each case by selecting random values of b in the range $(-3\sqrt{\lambda_m}, +3\sqrt{\lambda_m})$, where λ_m is the m -th eigenvalue resulted from solving SVD as described in Eq. (9). In the case of PCA-based SSMs, it is observed that the specificity index is smaller only when a small number of modes is used. Whereas in the case of RPCA-based SSMs, the specificity index of the FMNet and ICP are comparable and approximately constant with increasing number of modes. The same applies for the FM Procrustes and FM Optimization methods, which show smaller specificity indexes compared to the FMNet and ICP. After computing the total number of modes required to explain the 98% variance, the final specificity index for each case is presented in Table 1. The results indicate that the specificity indexes are comparable across the various implemented techniques, with the ICP presenting slightly better specificity for the PCA-based SSMs and the FM-based correspondence methods being superior when RPCA is used for the construction of the SSMs.

In conclusion, the FMNet presents higher accuracy compared to the




Fig. 7. Evaluation of the effect of four different shape correspondence methods on the generalization properties of SSMs. The SSMs are built using classical PCA (left) and Robust PCA (right) considering the multi-shape PDM scheme for the five knee components.

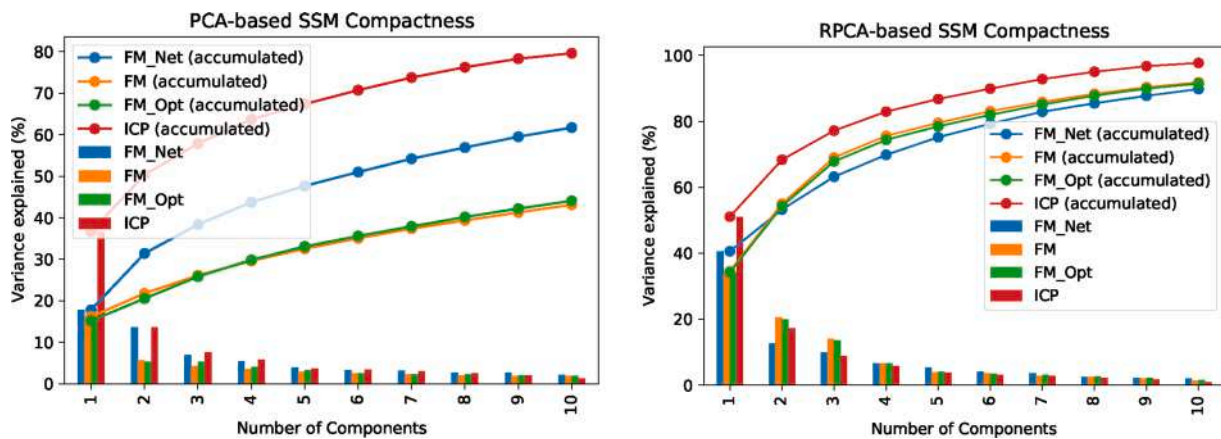


Fig. 8. Comparison of the compactness level of the constructed PCA-based (left) and RPCA-based (right) SSMs using the four shape correspondence methods. The explained variance and the accumulated variance are illustrated with increasing number of modes.

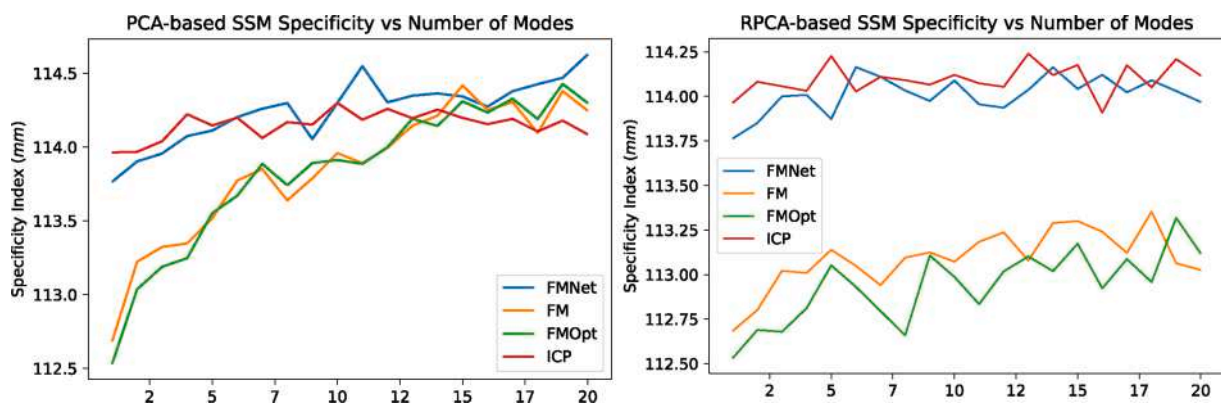


Fig. 9. Evaluation of the effect of the four shape correspondence methods on the specificity index of the PCA-based (left) and RPCA-based (right) SSMs, with increasing number of modes. A total of 500 random shape instances are generated with the use of the SSM in each case by selecting random values of b in the range $(-3\sqrt{\lambda_m}, +3\sqrt{\lambda_m})$, where λ_m is the m -th eigenvalue obtained by solving SVD in Eq. (9).

Table 1
Specificity index (measured in mm) of the PCA-based and RPCA-based SSMs, using the four shape correspondence methods.

	PCA	RPCA
FM Procrustes	115.52	113.12
FM Optimization	115.58	112.95
FMNet	114.41	114.04
ICP	114.35	114.06

Table 2
Average execution time (in seconds) of the individual correspondence methods for a pair of shapes of the knee complex.

ICP	FM procrustes	FM optimization	FMNet
0.36	1.75	13.11	2.74

other shape matching techniques, while maintaining comparable levels of specificity, though at the cost of being less compact. As a note, in many applications (including shape reconstruction and missing structure estimation presented in the following sections), the accuracy of the constructed SSMs is considered more important than the compactness and the specificity indexes.

Advantages and Limitations. The main advantage of the FMNet is that it does not necessarily require complete objects (different forms of partiality can be tackled if adequately represented in the training set

(Litany et al., 2017)). However, there have still been a few cases where the method failed to recover correspondences under extreme partiality. The inefficiency to resolve precisely shape correspondence given partial shapes is actually a particular limitation of the FM-based methods. This is due to the use of Laplacian eigenfunctions, which are inherently sensitive to topological changes, and the fact that we did not implement explicitly any extension of the framework to account for shape partiality. Thus complete shapes cannot be correctly matched to partial shapes, and vice-versa. For example, in Fig. 10 we show an instance where some elements of a complete shape have no corresponding points due to shape partiality. Thus the mapping assigns these elements to those of the other shape that express similar local shape characteristics, without actually representing anatomically correct correspondences.

However, since the solution under shape partiality is not invariant to the direction of the transformation, strong inconsistencies between the forward mapping (from the complete to the partial shape) and the inverse mapping (from the partial to the complete shape) could reveal regions with strong topological changes, such as missing parts. These outlier regions might even prove to be useful for disease localization, as performed in imaging-based approaches that detect pathology as deviation from the expected distribution (Erus et al., 2014; Zacharaki and Bezerianos, 2011).

4.2. Improving object segmentation

In this experiment we wanted to investigate whether our developed SSMs could be used to improve image segmentation results. For that

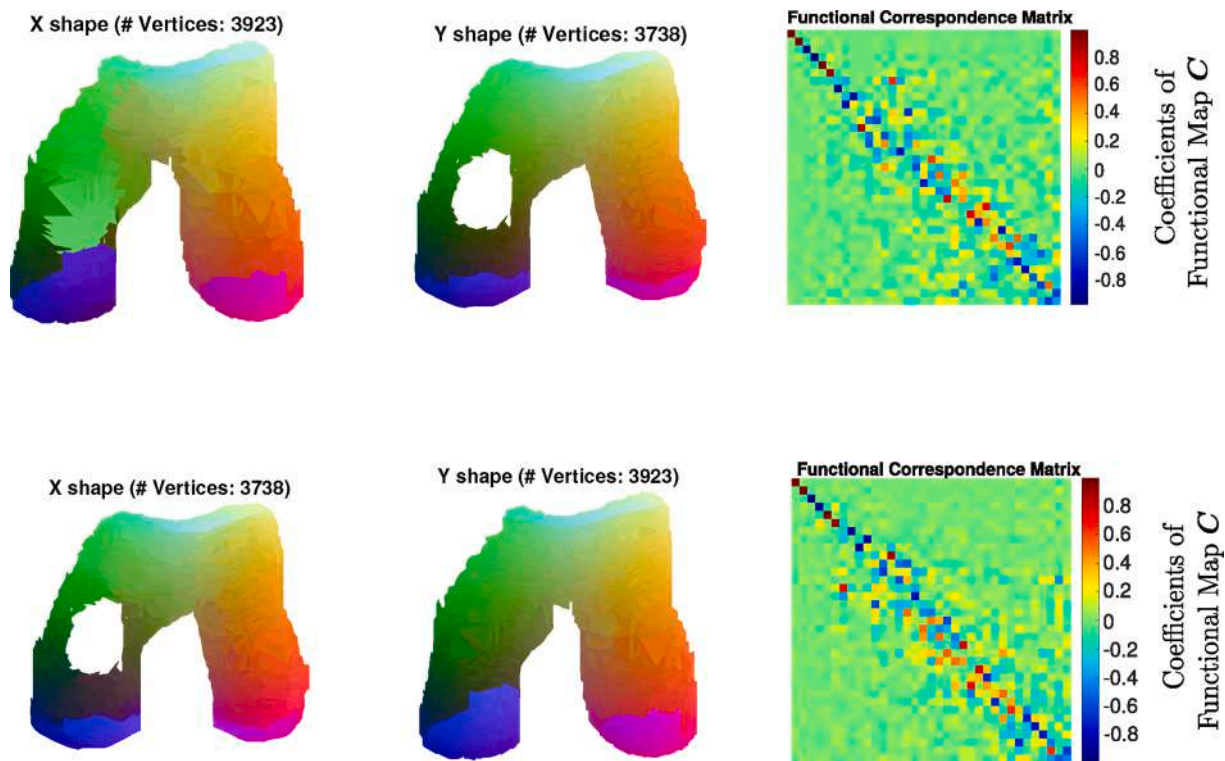


Fig. 10. Example of shape matching in case of extreme shape partiality due to osteoarthritis. In the top row, the FM framework fails to map elements of one shape (\mathcal{X}) to another (\mathcal{Y}), as the mapping occurs from the full to the partial shape. In this case, points with no counterparts are mapped to regions of the partial shape that present highest similarity. In the bottom row, the mapping occurs from the partial to the full shape, in which corresponding points exist.

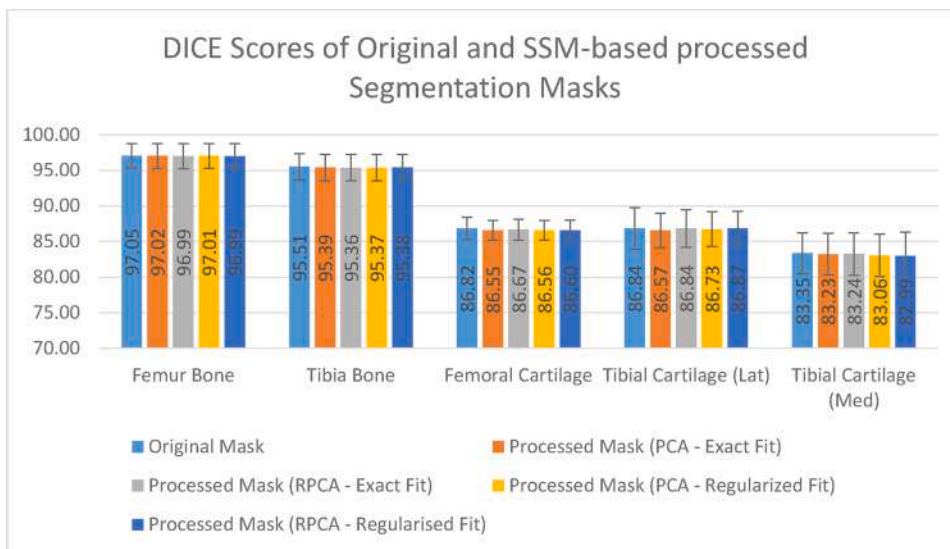


Fig. 11. Average DICE scores of the regularized segmentation masks \mathcal{M}_R . The maximum decrease in average DICE score (0.36%) was observed for the TCM using the RPCA-based SSM model with the regularized fitting scheme.

purpose, we produced image-based multi-label segmentation masks of the bones and cartilages from the OpenKnee dataset using a method (Nikolopoulos et al., 5678) (recently implemented in our group) based on a multi-atlas segmentation with label fusion and corrective learning approach (Wang and Yushkevich, 2013). We fitted the constructed PCA-based and RPCA-based SSMs (with $\Omega = I$) to meshes extracted from the segmentation masks and we applied the topological and regularization constraints in the image space, as described in Section 3. We distinguish the cases where the models are fitted using Eq. (11) (exact fitting process) and Eq. (13) (regularized fitting process).

After the SSM-based processing step, the DICE scores and Hausdorff distances of the initial masks \mathcal{M}_S and the regularized masks \mathcal{M}_R , compared to the ground truth masks outlined by human experts, are calculated and presented in Figs. 11 and 12, respectively. It can be observed that the differences in DICE scores are marginal, but the improvement in the Hausdorff distances is quite noticeable, resulting in more robust estimation of the ground-truth label maps. In more details, the change in DICE is small when the shape of a structure changes only locally, whereas the average Hausdorff distance decreases considerably, i.e., 2.22 – 2.27 mm (13.55% – 14.11%) for the FB, 3.04 – 3.16 mm

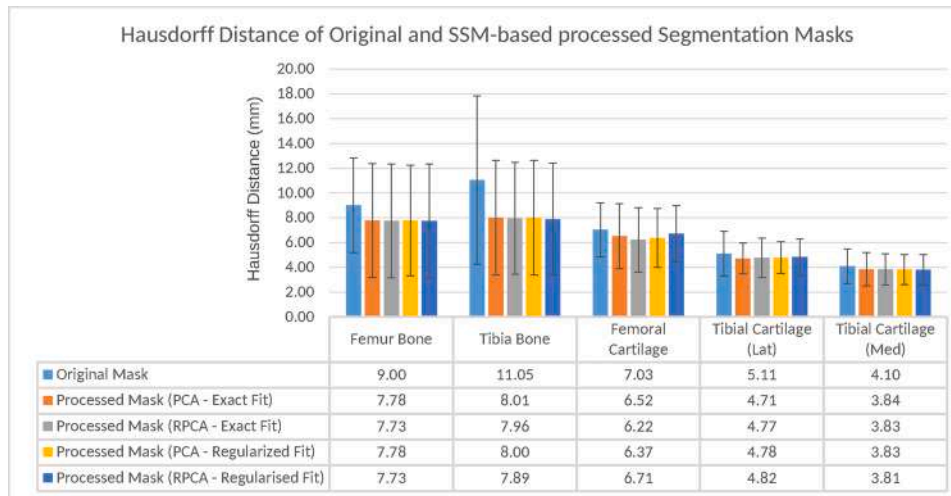


Fig. 12. Average Hausdorff distance of the regularized segmentation masks \mathcal{M}_R . The maximum reduction in Hausdorff distance was achieved with RPCA and regularized fit for 3 (out of 5) structures, i.e., the FB (11%), the TB (28.6%) and the TCM (7.83%).

Table 3

Average execution time (in seconds) of (exact and regularized) fitting process using the constructed PCA and RPCA-based SSMs.

PCA (Exact fit)	PCA (Regularized fit)	RPCA (Exact fit)	RPCA (Regularized fit)
15.5	672	5	15.8

(27.75% – 28.6%) for the TB, 0.32 – 0.81 mm (4.55% – 11.52%) for the FC, 0.29 – 0.4 mm (5.67% – 7.83%) for the TCL, and 0.26 – 0.29 mm (6.34% – 7.07%) for the TCM, indicating a better segmentation outcome. With the only exception of TCL, the RPCA-based SSMs produce better results compared to the classic PCA-based models. When combined with the regularized shape fitting process, the RPCA-based models achieve maximum decrease in Hausdorff distance for the FB, TB and TCM. As a note, one should also consider the additive effect of interpolation artefacts from the change of coordinate systems, i.e., from a regular volumetric grid used in image-based segmentation to an unstructured point cloud for shape reconstruction, and vice versa. The differences in the resolution level between the two coordinate systems might smooth out local extremities.

The execution time of the fitting process for the classical PCA-based and RPCA-based SSMs, using the exact or regularization fitting scheme are presented in Table 3. It is observed that the average processing time of the PCA-based model using the regularization fitting scheme is much longer than the other methods. This is mainly attributed to the larger number of variation modes required by the model to explain the predefined variation. Specifically, the PCA models are described using 135 modes, whereas the RPCA models are described using only 29 modes.

4.3. Prediction of missing structures

In the last experimental setup, we try to exploit the ability of SSMs to encode information about the shape variation of neighboring structures in order to infer the shape and pose of a missing structure in a new shape complex. In particular, using the individual SSMs that encode the interdependence between structures through CCA, we conduct shape inference about each missing structure modeled by the corresponding SSM. That is, knowing how the shape variation of the known structures affects the shape variation of the unknown one, we calculate an estimate of the missing structure that is close to the shape observed in the training set. For example, in the case of severe varus osteoarthritis, if the bone attrition in the medial tibia increases, the cartilage degeneration in the lateral femoral condyle and tibial plateau tends to increase too

(Nakagawa et al., 2015).

For a new shape complex we assume that only a single structure is missing, although the methods can be used for multiple missing structures as well. The procedure begins by first fitting the SSMs using the optimization scheme described in Eq. (13) where the target shape is described by the matrix $Y_i = (y_1; \dots; y_{i-1}; y_{i+1}; \dots; y_m)^T$ containing the coordinates of the four out of five structures in a new multi-structure complex. Since we are interested in estimating structures that are difficult to segment, we assume that the missing structure y_i is one of the cartilage tissues. The calculated parameters b can be used for the full SSM containing all five structures to produce an estimate \tilde{y}_i of the missing structure.

In order to assess the effect of regularization, we present in Fig. 13 the average DICE score and Hausdorff distance for each reconstructed missing structure, across 10 randomly selected test subjects, for regularization weight w ranging from 0 to 1. We selected a rather small number of subject to illustrate our concept because the computational time for multiple experiments (PCA, RPCA, 3 cartilage structures, multiple regularization weights) was very large. Nevertheless, the subjects were selected to include evenly distributed easy and difficult cases. This was performed by sorting the subjects according to the obtained fitting accuracy without regularization and then uniform sampling across the whole range. The graphs show the average (across subjects) performance (blue line) and the standard deviation (cyan region) for each regularization weight. It can be observed that small values of the regularization weight can improve the reconstruction outcome, while larger values tend to reduce accuracy. A single value, that exhibits optimal performance in terms of both DICE scores and Hausdorff distances (namely $1e^{-10}$), was selected and held constant throughout the following experiments.

We predict the shape and pose of the femoral and tibial cartilages for all the subjects within the testing set of the OAI ZIB dataset, using the PCA-based and RPCA-based models and we measure the average DICE scores of the predicted structures with respect to the ground-truth. The prediction accuracy is also compared with results derived from shape reconstruction by exact fitting presented in the first experiment (Section 4.1) where all shapes were known. The results for the PCA-based and RPCA-based models are shown in Fig. 14. The difference in accuracy between shape reconstruction and prediction was expected since for the latter no information was used for the pose, orientation or conformation of the unknown structure. In addition, the cartilages exhibit severe cases of osteoarthritis, which is difficult to predict based on the knowledge of the remaining structures.

In Fig. 15 we illustrate the spatial distribution of the distance error of

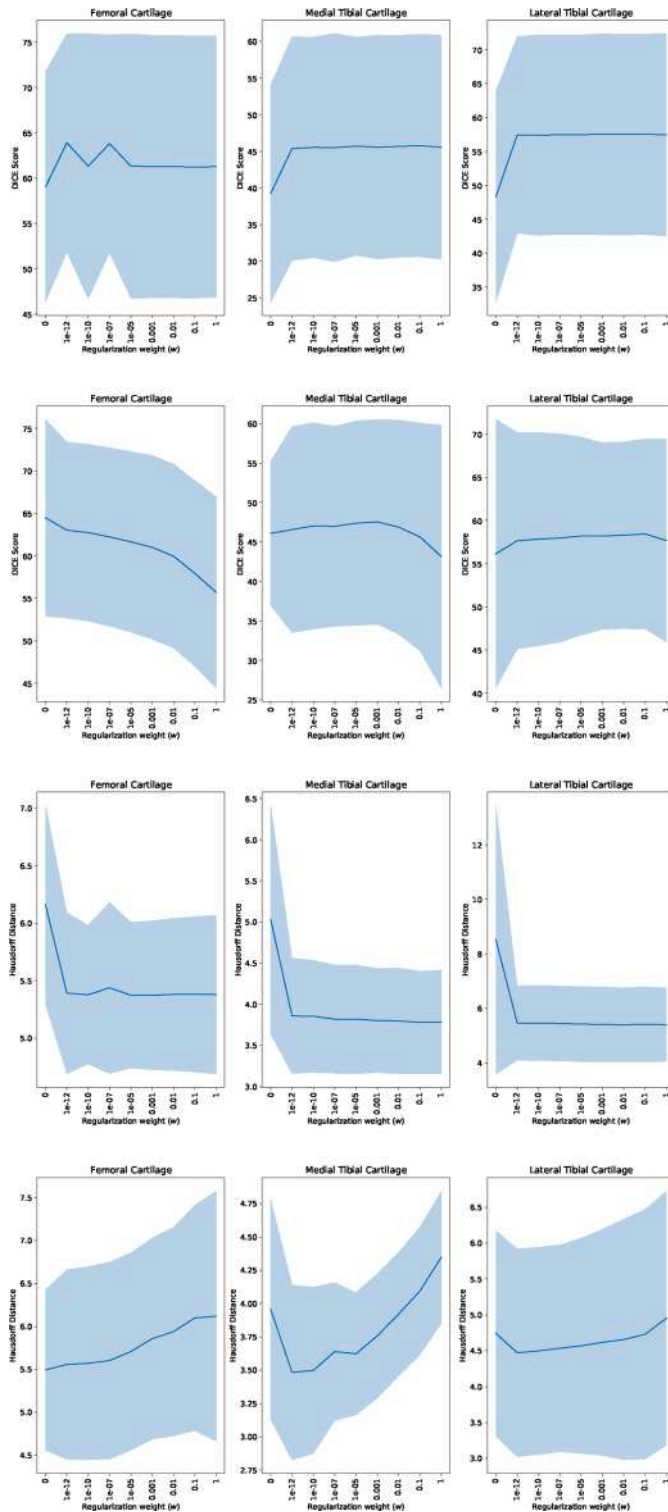


Fig. 13. Effect of the regularization weight w in Eq. (13) on the average DICE scores and Hausdorff distances after the reconstruction of missing structures, using the PCA and RPCA SSMs. Shaded regions represent the standard deviation from the mean value.

the predicted shapes from the ground-truth mapped in the common space of the mean shape of the PCA-based SSM and averaged over all test subjects. Largest error values tend to occur at the boundary regions of the structures with higher curvature, indicating that the model has higher uncertainty in the representation of finer details in the shape of the structures. Moreover, we observe that in contrast to the femoral

cartilage and lateral tibial cartilage, all regions of medial tibial cartilage are predicted with high consistency. This might be attributed to higher variability of the former two in comparison to the latter. In any case, the aim of this approach is not to accurately represent the missing shape, but rather to provide a first estimate of the unknown structure that can be used as an initialization process of other automatic segmentation algorithms. Lastly, recall that this procedure can be transferred across any imaging modality, so the shape inference can be conducted in cases where cartilages, or any other modeled structures, cannot be easily distinguished in the available images, such as, e.g., in computed tomography.

5. Discussion and conclusions

In this work we provided an overview of the principles of 3D shape modeling and presented a novel approach that combines statistical and deep learning techniques for reconstruction and estimation of anatomical compartments. The advantages of our approach were illustrated in the challenging task of modelling the human knee under various conditions (normal or pathological state and complete or partial shapes). We specifically focused on building multi-structure shape models, while taking into consideration the relationship between the individual anatomical structures through correlation analysis. This framework can be considered as a generalization of the methods used for statistical shape modeling, which efficiently characterize individual structure variability and relationship between different structures. The constructed models can also be parameterized differently to be equivalent to a wide variety of alternative SSM approaches, including the classic PDM.

The most fundamental requirement for the construction of statistical models is the identification of pairwise correspondences across the shapes. In particular, our main focus has been on the recently proposed framework of Functional Maps which generalizes the notion of shape matching. By creating different representations of the shapes using shape descriptors, and projecting them on an orthonormal basis in the function space of the shapes, the problem of shape correspondence becomes a system of linear algebraic equations. We investigated the performance of different approaches exploiting the Functional Maps, and we compared them with standard methods found in literature. In addition, we have demonstrated that by incorporating machine learning techniques, and specifically deep neural networks, it is possible to increase the captured shape variability and concomitantly the generalization properties of SSMs. To the best of our knowledge there are no other studies examining the potential of Deep Functional Maps to enhance the ability of SSMs to capture shape variation, avoiding the need of high correspondence quality in advance. Note that Functional Maps require the shapes to differ only by isometric deformations, an assumption that might not hold in the case of real anatomical structures that usually differ significantly among different subjects. However, since our main focus has been on the individual anatomic compartments of the knee complex, it is reasonable to assume that the variation of these smaller structures across individuals can be expressed by isometric deformations and that the intrinsic geometry of the shape is preserved. This assumption is mainly violated in case of a pathology like osteoarthritis.

An important application of SSMs is their adaptation to an initial shape. Model fitting allows to transfer any knowledge encoded in the statistical model directly to a new shape instance. In the case of defective shapes obtained from imprecise image segmentation, in order to avoid overfitting, we introduced an optimization scheme that constrains the allowed variation to plausible shapes. We applied this regularized approach on structures of the knee complex acquired from MRI data, and showed an improvement in accuracy over multi-atlas based image segmentation. In particular, the Hausdorff distance was reduced on average by 13.8% across all structures. From an image segmentation perspective, volumetric approaches, e.g., based on deep learning (Ambellan et al., 2019), are much more stable than point cloud-based approaches, as they

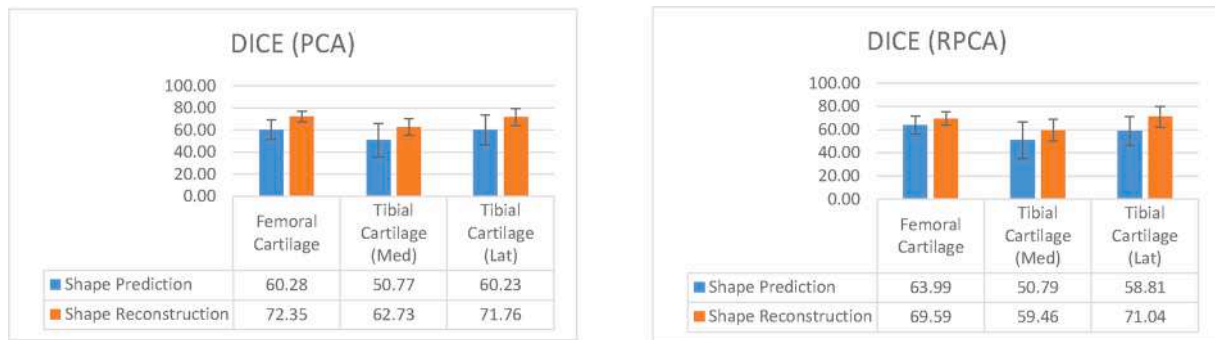


Fig. 14. Comparison of the DICE scores of the SSM models of femoral and tibial cartilages when the resulting fitted shapes are obtained using Shape Prediction (blue) (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.) and Shape Reconstruction (orange), respectively.



Fig. 15. Spatial distribution of the distance error of the predicted shapes (FC (left), TCM (middle) and TCL (right)) compared to the ground-truth geometries averaged for 15 random subjects in the common space of the mean shape of the PCA-based SSM.

have significantly more information content, such as the texture of surrounding structures. Nevertheless, medical image segmentation based on deep learning requires a large amount of data with delineated structures for training, whereas the construction of SSMs requires a significantly smaller amount of data, and the data do not even have to be volumetric (i.e., 3D images). Therefore, our aim was to illustrate the use of SSMs as a supplementary step that could help improve segmentation in the case of data-related limitations, as in the case of the OpenKnee dataset which contains only 7 subjects that can be used for training.

Finally, we took advantage of the ability of SSMs to encode the relationship between neighboring structures in order to estimate the shape and pose of a missing or unknown structure in a new data complex. This has the potential of conducting shape inference on image modalities that are not sensitive in visualizing certain biological tissues, or it can be used as an initialization process for other automatic segmentation methods.

One problem that arises in the context of multi-structure shape modeling is the joint alignment of the structures as a single point-cloud. Besides the shape differences, the alignment of a global point cloud also neglects subtle pose differences of one shape in respect to another. This affects significantly the ability of linear shape models to represent new shapes, especially smaller and flexible structures like the cartilages, which exhibit considerable amount of natural variability.

In future work, we intend to tackle the problem of joint modeling of multiple shapes, by modeling the rigid transformations and the shape variation of individual structures. Creating separate statistical models for the pose and shape parameters can potentially allow us to better characterize locality of each structure and thereby capture more accurately the shape variation. Also, we will further examine the use of regularization constraints and shape priors to better address the problem of shape partiality. Finally, we envision an approach for statistical shape analysis in the functional eigenspace, where shape variability is

described through changes in the harmonic content of shape descriptors.

Supplementary Material

Online information related to this publication, including source code, data, and results, can be found at the following link: https://gitlab.com/vvr/publications/regularized_knee_ssm

Authors' contribution

Filip Konstantinos: Methodology, Investigation, Software Implementation, Data curation, Writing, Original draft preparation.

Zacharaki I. Evangelia: Conceptualization, Investigation, Supervision, Reviewing and Editing.

Moustakas Konstantinos: Supervision, Reviewing and Editing.

Declaration of Competing Interest

None of the authors had any conflict of interest regarding this manuscript.

Acknowledgment

The authors would like to thank their colleagues F.P. Nikolopoulos for preparation of the segmented images and D. Stanev for valuable discussions on mesh processing. This research has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH - CREATE - INNOVATE (project code: T1EDK - 04234) and by the EC Horizon 2020 project OACTIVE Grant Agreement No.777159. There was no additional external funding received for this study. The funders had no role in the study design and analysis, decision

to publish, or preparation of the manuscript.

References

- Ambellan, F., Tack, A., Ehkle, M., Zachow, S., 2019. Automated segmentation of knee bone and cartilage combining statistical shape knowledge and convolutional neural networks: data from the osteoarthritis initiative. *Med. Image Anal.* 52, 109–118.
- Baldwin, M.A., Langenderfer, J.E., Rullkoetter, P.J., Laz, P.J., 2010. Development of subject-specific and statistical shape models of the knee using an efficient segmentation and mesh-morphing approach. *Comput. Methods. ProgramS Biomed.* 97 (3), 232–240.
- Biasotti, S., Cerri, A., Bronstein, A., Bronstein, M., 2016. Recent trends, applications, and perspectives in 3d shape similarity assessment. *Computer Graphics Forum*, Vol. 35. Wiley Online Library, pp. 87–119.
- Boscaini, D., Masci, J., Rodolà, E., Bronstein, M., 2016. Learning shape correspondence with anisotropic convolutional neural networks. *Advances in Neural Information Processing Systems*, pp. 3189–3197.
- Boumal, N., Mishra, B., Absil, P.-A., Sepulchre, R., 2014. Manopt, a matlab toolbox for optimization on manifolds. *J. Mach. Learn. Res.* 15 (1), 1455–1459.
- Bredbenner, T.L., Eliason, T.D., Potter, R.S., Mason, R.L., Havill, L.M., Nicoletta, D.P., 2010. Statistical shape modeling describes variation in tibia and femur surface geometry between control and incidence groups from the osteoarthritis initiative database. *J. biomech.* 43 (9), 1780–1786.
- Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P., 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Process. Mag.* 34 (4), 18–42.
- Candès, E.J., Li, X., Ma, Y., Wright, J., 2011. Robust principal component analysis? *JACM* 58 (3), 11.
- Cerrolaza, J.J., Reyes, M., Summers, R.M., González-Ballester, M.Á., Linguraru, M.G., 2015. Automatic multi-resolution shape modeling of multi-organ structures. *Med. Image Anal.* 25 (1), 11–21.
- Cerrolaza, J.J., Summers, R.M., Linguraru, M.G., 2016. Soft multi-organ shape models via generalized pca: a general framework. *International Conference on Medical Image Computing and Computer - Assisted Intervention* 219–228.
- Cerrolaza, J.J., Picazo, M.L., Humbert, L., Sato, Y., Rueckert, D., Ballester, M.Á.G., Linguraru, M.G., 2019. Computational anatomy for multi-organ analysis in medical imaging: a review. *Med. Image Anal.* 56, 44–67.
- Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1995. Active shape models-their training and application. *Comput. Vis. Image Underst.* 61 (1), 38–59.
- Cootes, T.F., Taylor, C.J., et al., 2004. *Statistical Models of Appearance for Computer Vision*.
- Dryden, I., Mardia, K., 1998. *Statistical Analysis of Shape*. Wiley.
- Erus, G., Zacharaki, E.I., Davatzikos, C., 2014. Individualized statistical learning from medical image databases: application to identification of brain lesions. *Med. Image Anal.* 18 (3), 542–554.
- Fitzpatrick, C.K., Baldwin, M.A., Laz, P.J., FitzPatrick, D.P., Lerner, A.L., Rullkoetter, P. J., 2011. Development of a statistical shape model of the patellofemoral joint for investigating relationships between shape and function. *J. Biomech.* 44 (13), 2446–2452.
- Goodall, C., 1991. Procrustes methods in the statistical analysis of shape. *J. R. Stat. Soc. Series B* 285–339.
- Gower, J.C., 1975. Generalized procrustes analysis. *Psychometrika* 40 (1), 33–51.
- M. J. Greenacre, *Theory and applications of correspondence analysis*.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 770–778.
- Heimann, T., Meinzer, H.-P., 2009. Statistical shape models for 3d medical image segmentation: a review. *Med. Image Anal.* 13 (4), 543–563.
- Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R., 2013. Coupled quasi-harmonic bases. *Computer Graphics Forum*, Vol. 32. Wiley Online Library, pp. 439–448.
- Kovnatsky, A., Bronstein, M.M., Bresson, X., Vandergheynst, P., 2015. Functional correspondence by matrix completion. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 905–914.
- Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M., Cremers, D., 2016. Non-rigid puzzles. *Computer Graphics Forum*, Vol. 35. Wiley Online Library, pp. 135–143.
- Litany, O., Remez, T., Rodolà, E., Bronstein, A., Bronstein, M., 2017a. Deep functional maps: structured prediction for dense shape correspondence. *Proceedings of the IEEE International Conference on Computer Vision* 5659–5667.
- Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M., 2017b. Fully spectral partial shape matching. *Computer Graphics Forum*, pp. 247–258. Vol. 36 Wiley Online Library.
- Lorensen, W.E., Cline, H.E., 1987. Marching cubes: a high resolution 3d surface construction algorithm. *ACM Siggraph Computer Graphics* Vol. 21, 163–169. ACM.
- Marquardt, D.W., 1963. An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Ind. Appl. Math.* 11 (2), 431–441.
- Masci, J., Boscaini, D., Bronstein, M., Vandergheynst, P., 2015. Geodesic convolutional neural networks on riemannian manifolds. *Proceedings of the IEEE International Conference on Computer Vision Workshops* 37–45.
- Mohamed, A., Kyriacou, S.K., Davatzikos, C., 2001. A statistical approach for estimating brain tumor induced deformation. *Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis MMBIA* 2001 52–59.
- Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M., 2017. Geometric deep learning on graphs and manifolds using mixture model cnns. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 5115–5124.
- Nakagawa, Y., Mukai, S., Yabumoto, H., Tarumi, E., Nakamura, T., 2015. Cartilage degeneration and alignment in severe varus knee osteoarthritis. *Cartilage* 6 (4), 208–215.
- Nikolopoulos, F.P., Zacharaki, E.I., Stanev, D., Moustakas, K., 2020. Personalized knee geometry modeling based on multi-atlas segmentation and mesh refinement. *IEEE Access* 8, 56766–56781.
- Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L., 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graphics (TOG)* 31 (4), 30.
- Ovsjanikov, M., Corman, E., Bronstein, M., Rodola, E., Ben-Chen, M., Guibas, L., Chazal, F., Bronstein, A., 2017. Computing and processing correspondences with functional maps. *ACM Siggraph 2017 Courses* 1–62.
- Patil, S., Ravi, B., 2005. Voxel-based representation, display and thickness analysis of intricate shapes. *Ninth International Conference on Computer Aided Design and Computer Graphics (CAD-CG'05)*. IEEE pp. 6–pp.
- Rao, A., Aljabar, P., Rueckert, D., 2008. Hierarchical statistical shape analysis and prediction of sub-cortical brain structures. *Med. Image Anal.* 12 (1), 55–68.
- Rao, C., Fitzpatrick, C.K., Rullkoetter, P.J., Maletsky, L.P., Kim, R.H., Laz, P.J., 2013. A statistical finite element model of the knee accounting for shape and alignment variability. *Med. Eng. Phys.* 35 (10), 1450–1456.
- Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M., 2018. Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graphics (TOG)* 37 (6), 1–16.
- Reuter, M., Biasotti, S., Giorgi, D., Patané, G., Spagnuolo, M., 2009. Discrete laplace-beltrami operators for shape analysis and segmentation. *Comput. Graph.* 33 (3), 381–390.
- Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D., 2017a. Partial functional correspondence. *Computer Graphics Forum*, Vol. 36. Wiley Online Library, pp. 222–236.
- Rodolà, E., Möller, M., Cremers, D., 2017b. Regularized pointwise map recovery from functional correspondence. *Computer Graphics Forum*, Vol. 36. Wiley Online Library, pp. 700–711.
- Roufousse, J.-M., Sharma, A., Ovsjanikov, M., 2019. Unsupervised deep learning for structured shape matching. *Proceedings of the IEEE International Conference on Computer Vision* 1617–1627.
- Rusinkiewicz, S., Levoy, M., 2001. Efficient variants of the icp algorithm. *3dim. IEEE*, p. 145.
- Sahillioglu, Y., 2020. Recent advances in shape correspondence. *Vis. Comput.* 36 (8), 1705–1721.
- Saito, A., Nakada, M., Oost, E., Shimizu, A., Watanabe, H., Nawano, S., 2013. A statistical shape model for multiple organs based on synthesized-based learning. *International MICCAI Workshop on Computational and Clinical Challenges in Abdominal Imaging* 280–289. Springer.
- Sun, J., Ovsjanikov, M., Guibas, L., 2009. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, Vol. 28. Wiley Online Library, pp. 1383–1392.
- Tombari, F., Salti, S., Di Stefano, L., 2010. Unique signatures of histograms for local surface description. *European Conference on Computer Vision*. Springer, pp. 356–369.
- Van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D., 2011. A survey on shape correspondence. *Computer Graphics Forum*, Vol. 30. Wiley Online Library, pp. 1681–1707.
- Wang, H., Yushkevich, P., 2013. Multi-atlas segmentation with joint label fusion and corrective learning - an open source implementation. *Front. Neuroinform.* 7, 27.
- Williams, T.G., Holmes, A.P., Waterton, J.C., Maciewicz, R.A., Hutchinson, C.E., Moots, R.J., Nash, A.F., Taylor, C.J., 2010. Anatomically corresponded regional analysis of cartilage in asymptomatic and osteoarthritic knees by statistical shape modelling of the bone. *IEEE Trans. Med. Imaging* 29 (8), 1541–1559.
- Zacharaki, E.I., Bezerianos, A., 2011. Abnormality segmentation in brain images via distributed estimation. *IEEE Trans. Inf. Technol. Biomed.* 16 (3), 330–338.
- Zimmer, W., Berquist, T., McLeod, R., Sim, F., Pritchard, D., Shives, T., Wold, L., May, G., 1985. Bone tumors: magnetic resonance imaging versus computed tomography. *Radiology* 155 (3), 709–718.

Konstantinos Filip received the Diploma degree in electrical and computer engineering and the M.Sc. degree in biomedical engineering in 2016 and 2019, respectively, both from the University of Patras. He is currently a research associate at the Visualization and Virtual Reality Group at the Department of Electrical and Computer Engineering, University of Patras, Greece. His main research interests include computational geometry, statistical modeling and machine learning, as well as computational modeling, simulation, and analysis of musculoskeletal systems.

Dr. Evangelia I. Zacharaki (female) (HdR'17, PhD'04, M.Eng'99) is a senior research scientist in the VVR Group of UOP since 2018. Prior to that she has worked as a research associate at the Section of Biomedical Image Analysis, UPenn, USA (2005 - 2009), the Medical Physics Department, UOP (2009 - 2012), and the Center for Visual Computing, CentraleSupélec/INRIA, France (2015-2017). In 2017 she obtained an HdR (highest accreditation to coordinate research) with specialization in informatics from Université Paris-Est. Her research interests focus on machine learning, computational and statistical modeling for representation, fusion and analysis of high-dimensional biomedical data. Dr. Zacharaki has received a Marie Curie IRG and has participated in more than 15 international (NIH/NIA), european (FP7, H2020) and national (GSRT, FRM) research projects. She serves as member of 4 editorial boards and is currently the Guest Editor of the Special

Issue "Multi-Sensor Fusion of Biomedical Data: Application to Diagnosis and Treatment" of the Sensors journal. She has co-authored in total 105 scientific publications in refereed international journals, book chapters and peer reviewed international conferences receiving more than 2320 citations (h-index=22, g-index=47, 09/2020).

Konstantinos Moustakas received the Diploma degree and the PhD in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2003 and 2007 respectively. During 2007-2011 he served as a post-doctoral research fellow in the Information Technologies Institute, Centre for Research and Technology Hellas. He is currently an Associate Professor in the Electrical and Computer Engineering Department of the University of Patras and Head of the Visualization and Virtual Reality Group. He is the director of the Wire Communications Laboratory and the Master's Program "Biomedical

Engineering" of the University of Patras. His main research interests include virtual, augmented and mixed reality, 3D geometry processing, virtual physiological human modeling, biomedical engineering, information visualization, physics-based simulations. During the latest years, he has been the (co)author of more than 210 papers in refereed journals, edited books, and international conferences. He serves as a regular reviewer for several technical journals and has participated to more than 20 research and development projects funded by the EC and the Greek Secretariat of Research and Technology. He was the coordinator of the GameCar Horizon2020 project, scientific coordinator of the NoTremor FP7 project, while he also chaired the scientific board of the myAirCoach and FrailSafe H2020 projects. He is a Senior Member of the IEEE and the IEEE Computer Society, IEEE Signal Processing Society and Member of Eurographics.