

Principal Geodesic Analysis in the Space of Discrete Shells

B. Heeren^{1*}, C. Zhang^{2*}, M. Rumpf¹, and W. Smith²

¹University of Bonn, Institute for Numerical Simulation, Germany
²University of York, Department of Computer Science, United Kingdom




Figure 1: We show how to learn nonlinear, physically plausible modes of shape variation (bottom right) from a set of highly varying training shapes, which can be used for projection onto a low dimensional submanifold and thus sparse representation by a small set of weights. This model can be used to solve problems such as reconstruction of dense body shapes from motion capture markers (top) providing compressed animations for the reconstructed shapes via time sequences of the weights (bottom left) even when the captured data is sparse, noisy and comes from a different body shape. (Training data: 50 meshes of Dyna [PMRMB15]; Key parameters: $K = 4$, $J = 10$).

Abstract

Important sources of shape variability, such as articulated motion of body models or soft tissue dynamics, are highly nonlinear and are usually superposed on top of rigid body motion which must be factored out. We propose a novel, nonlinear, rigid body motion invariant Principal Geodesic Analysis (PGA) that allows us to analyse this variability, compress large variations based on statistical shape analysis and fit a model to measurements. For given input shape data sets we show how to compute a low dimensional approximating submanifold on the space of discrete shells, making our approach a hybrid between a physical and statistical model. General discrete shells can be projected onto the submanifold and sparsely represented by a small set of coefficients. We demonstrate two specific applications: model-constrained mesh editing and reconstruction of a dense animated mesh from sparse motion capture markers using the statistical knowledge as a prior.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling

1. Introduction

Compact models of the shape variability of a class of 3D objects are useful in a wide range of analysis and synthesis applications across graphics and vision. Such statistical models learnt from data

provide constraint for analysis problems, compress high dimensional data to a low dimensional space and ensure plausibility of

* B. Heeren and C. Zhang contributed equally to this work.

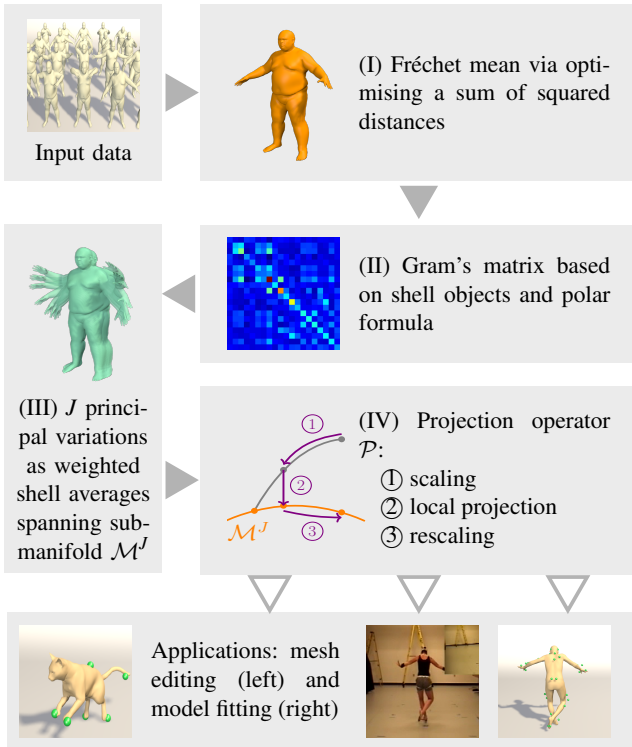
synthesised results. Specifically, they can be used for non-rigid registration, reconstruction from incomplete, noisy or 2D data, mesh editing, performance-driven animation and deformation transfer. To meet these applications, we address in this paper a number of important challenges:

▷ First, many important sources of shape variability are highly nonlinear. For example, nonrigid deformation (such as articulation, bending, and stretching) and nonlinear shape changes (such as weight variation or shape differences between individuals).

▷ Second, models should be physically plausible so that unrealistic shapes are avoided and to enable meaningful interpolation between and extrapolation beyond the training samples.

▷ Third, deformations must be modelled independently of rigid body motion. Methods that rely on factoring out rigid body motion by alignment require a choice of alignment metric, the choice of which influences the final model. Moreover, for nonrigid deformations a meaningful rigid alignment may not exist.

The natural concept to deal with these requirements is a Riemannian shape manifold. The key ingredient for our model is a discrete geodesic (i.e. a geodesic path discretised in time) in the space of discrete shells (a triangle mesh-based approximation of the thin shell physical model). From this starting point, we propose time-discrete statistics on manifolds and make the following key contributions (summarised in the flowchart):



(I) Starting from a set of input shells, we define in Section 4 a discrete geodesic average (i.e. Fréchet mean) as the minimiser of the sum of squared discrete geodesic distances to the input shapes. (II) Then, we use the polar formula for scalar products to introduce an approximate Gram matrix defined directly on discrete shapes and not as usual on infinitesimal shape variations. (III) Given the eigenvectors, principal variations are defined as weighted shell averages




Figure 2: Key properties of the discrete shell model: equivalence classes of discrete shells $[s]$ incorporating rigid body motion invariance (a), with a physically sound bending (left) and membrane (right) energy density (b).

on the manifold. They are the nonlinear counterpart of (infinitesimal) principal components and span a finite dimensional submanifold (cf. Section 5). (IV) Arbitrary shells can be projected onto this submanifold to provide low dimensional representations. This projection can be used to constrain the admissible set of shapes in different shape optimisation applications. In Sections 6 and 7 we exemplarily use the model for mesh editing and dense reconstruction from motion capture data (cf. Fig. 1). The model ensures that the results exhibit physically realistic deformations while remaining statistically plausible.

We work directly with meshes and do not require problem-specific articulated skeletons yet our approach is able to handle many different kinds of nonlinear deformation. The discrete shell model (see Fig. 2b) provides highly plausible interpolations and extrapolations within the nonlinear shape manifold, meaning that we can build rich models from very sparse training samples. The shell space in which we work is a space of equivalence classes of shapes that differ by rigid body motions (see Fig. 2a) and we take special care to transfer this invariance to our time-discrete statistics. Therefore, our whole framework is rigid body motion invariant and does not require a choice of alignment metric or a preprocessing alignment step.

2. Related work

Elastic shape modelling. Physically-based elastic energy models have been widely used for simulation, interpolation, mesh editing and, more recently, statistical modelling. The classical model for elastically deformable surfaces is the shell model, originally introduced in a graphics context by Terzopoulos et al. [TPBF87], for thin, flexible materials. Grinspun et al. [GHDS03] introduced the discrete shell model in which a triangle mesh is a spatially-discrete representation of the mid-surface of a shell. The model was used for simulation of deformable materials under physical forces. In the direction of improving efficiency, the as-rigid-as-possible (ARAP) framework [SA07] is based on alternating minimisation over vertex positions and local rotations of an energy that measures deviation from rigidity. Von Radziewsky et al. [vRESH16] recently showed how model reduction can be used to efficiently evaluate elastic deformation models, including the discrete shell energy. This enables elastic models to be used in realtime applications (see also [vTSSH15]). Like our model, Zhang et al. [ZHR15] sought to construct a statistical shape model in shell space. The same nonlinear elastic deformation energy is used, however the model is built in

a linear space of vertex displacements and so is not rigid body motion invariant and does not have an underlying Riemannian model.

Articulated models. The natural representation for deformations due to articulation is a skeleton model comprised of joint locations and relative orientations. Heap and Hogg [HH96] extended classical 2D landmark-based statistical modelling into the articulated domain by building linear models over joint angles rather than vertex positions. For 3D shapes, skeletons are used to deform dense surface models (usually meshes) via a process known as skinning [LCF00]. In their Shape Completion and Animation of People (SCAPE) framework, Anguelov et al. [ASK*05] learn a combined pose and deformation model and a model of the variability in body shape. A skeleton is used to drive mesh deformation using a method based on deformation transfer [SP04] and variations in body shape are learnt using a linear model of bodies in a standard pose. The Dyna model [PMRMB15] is built on top of SCAPE and adds a linear dynamics model whose coefficients depend upon the skeleton pose. The SMPL model [LMR*15] shows that pose-dependent blend shapes can depend linearly on the rotation matrices of the skeleton joints yet still achieve high realism of pose dependent shape and dynamics. A drawback of all of these approaches is that articulated models must be handcrafted for a specific object class and cannot capture general deformations.

Triangle deformation models. A popular approach is to build models based on the statistics of triangle deformations [ACPH06, HLRB12, CLZ13]. Instead of being trained to reproduce the input meshes directly, they are trained to reproduce the local deformations that produced those meshes. Unlike elastic models, these are not physically-motivated. Sumner and Popović [SP04] express deformation in terms of affine transformation and a displacement - the same as the deformation model used in SCAPE. Sumner et al. [SZGP05] used deformation gradients for mesh-based inverse kinematics. Hasler et al. [HSS*09] use a nonlinear representation of triangle deformations with 15 DoF which captures the relationship between pose and shape. Freifeld and Black [FB12] derive a 6D Lie group representation of triangle deformations with no redundant degrees of freedom. None of these approaches are rigid body motion invariant. Fröhlich and Botsch [FB11] additionally introduce a bending term, expressing deformations in terms of changes to geometric quantities (triangle edge lengths and the dihedral angle between adjacent triangles). Gao et al. [GLL*16] introduce a rotation-invariant mesh difference representation in which plausible deformations often form a near linear subspace. The deformations produced by all of these approaches will not in general be realisable by a connected triangle mesh. Hence, these models require a further step to solve for the mesh that best fits the desired deformations, which might be unsatisfactory from a theoretical standpoint.

Riemannian shape modeling. There have been numerous attempts to cast shape modelling or statistical shape analysis in a Riemannian setting; e.g. [FLPJ04, Pen06, KMP07]. Kilian et al. [KMP07] showed how to compute geodesic paths between triangle meshes using a metric that measures changes in triangle edge lengths. Frequently, the underlying metric is based on measuring the lack of isometry, e.g. via a (linearised) elastic energy acting on the Cauchy-Green strain tensor of an associated infinitesimal

mesh deformation [SP04, ASK*05, ACPH06, HSS*09, FB12, HLRB12, CLZ13, PMRMB15]. To avoid irregular, isometric shape deformations an additional regularisation is required. Heeren et al. [HRWW12] take a similar approach but use the discrete shell model which includes a bending term and leads to time-discrete geodesic paths with physical meaning (they minimise the dissipation of thin shell elastic energy). The resulting shell space was subsequently further explored [HRS*14] by introducing time-discrete versions of Riemannian concepts such as the exponential and logarithmic maps and parallel transport.

Shape collection analysis. The classical *statistical shape modelling* approach deals with objects represented by a configuration of landmark points. A point in Kendall’s shape space [Ken84] corresponds to a configuration of landmarks in which rigid body motion has been “factored out”. Linear Principal Components Analysis (PCA) in this space is used to extract the important modes of shape variation. PCA has been extended to manifold valued data in the form of Principal Geodesic Analysis (PGA) [FLPJ04]. PGA models are now widely used. Fletcher et al. [FLPJ04] originally proposed the approach for modelling medially-defined anatomical objects. Freifeld and Black [FB12] used PGA to build statistical models on their Lie group representation of triangle deformations. Tournier et al. [TWC*09] used PGA to build a statistical skeleton model. Tycowicz et al. [vTAMZ18] presented a non-Euclidean statistical analysis of triangle meshes (represented by means of deformation gradients, cf. [SP04]) and consider medical applications such as shape-based classification of morphological disorders. Besides modelling shape variation, PGA has also been applied to distributions of probability measures using an optimal transport metric [SC15]. Whereas all these approaches model variability using only distances between objects, Rustamov et al. [ROA*13] began a line of work which uses shape maps to more richly characterise differences between shapes and differences between differences. Boscaini et al. [BEKB15] showed how to reconstruct shapes from these shape difference operators, enabling shape analogy synthesis and style transfer. While the original shape differences operator captures only intrinsic distortion, Cormen et al. [CSBC*17] use offset surfaces to capture extrinsic distortion.

3. Preliminaries

One can consider the space of shapes, e.g. triangle meshes, as a *Riemannian manifold* \mathcal{M} with a metric g . Then, for a path $\mathfrak{s} : [0, 1] \rightarrow \mathcal{M}$ the *path energy* is given by

$$\mathcal{E}[\mathfrak{s}] = \int_0^1 g_{\mathfrak{s}(t)}(\dot{\mathfrak{s}}(t), \dot{\mathfrak{s}}(t)) dt, \quad (1)$$

where the velocity $\dot{\mathfrak{s}}(t)$ at time t is an infinitesimal variation of $\mathfrak{s}(t)$. Given two points $s_A, s_B \in \mathcal{M}$ a path \mathfrak{s} minimising (1) among all paths with $\mathfrak{s}(0) = s_A$ and $\mathfrak{s}(1) = s_B$ is called a (shortest) *geodesic* connecting s_A and s_B and we have

$$\text{dist}^2(s_A, s_B) = \min_{\mathfrak{s}(0)=s_A, \mathfrak{s}(1)=s_B} \mathcal{E}[\mathfrak{s}]. \quad (2)$$

Note that minimizers of (1) also minimize the *length functional* $\mathcal{L}[\mathfrak{s}] = \int_0^1 \sqrt{g_{\mathfrak{s}}(\dot{\mathfrak{s}}, \dot{\mathfrak{s}})} dt$. In contrast to \mathcal{L} , the path energy is *not* independent of reparameterization and minimizers $t \mapsto \mathfrak{s}(t)$ have con-

stant absolute velocity, i.e. for all $t \in [0, 1]$ we have

$$g_{\mathfrak{s}(t)}(\dot{\mathfrak{s}}(t), \dot{\mathfrak{s}}(t)) = g_{\mathfrak{s}(0)}(\dot{\mathfrak{s}}(0), \dot{\mathfrak{s}}(0)) = \text{dist}^2(\mathfrak{s}(0), \mathfrak{s}(1)). \quad (3)$$

The geometric logarithm $\log_{s_A} s_B$ is defined as the initial velocity $v = \dot{\mathfrak{s}}(0)$ and for fixed $s_A \in \mathcal{M}$ there is a 1-to-1 correspondence between v and s_B (for s_B close to s_A). The corresponding inverse mapping is denoted as exponential map, i.e. $\exp_{s_A}(v) = s_B$.

Principal geodesic analysis. Let us briefly recall classical Principal Components Analysis (PCA) on \mathbb{R}^N before we consider Riemannian manifolds. For data points $s^1, \dots, s^n \in \mathbb{R}^N$ the arithmetic average is given by

$$\bar{s} = \arg \min_{s \in \mathbb{R}^N} \sum_{i=1}^n \|s - s^i\|_{\mathbb{R}^N}^2 = \frac{1}{n} \sum_{i=1, \dots, n} s^i. \quad (4)$$

Then *Gram's matrix* is defined by $G = \frac{1}{n} DD^T \in \mathbb{R}^{n,n}$, where $D \in \mathbb{R}^{n,N}$ represents the *data matrix* whose i th row is given by $(s^i - \bar{s})^T \in \mathbb{R}^{1,N}$. In particular, the entries of G depend on the underlying (Euclidean) scalar product as $G_{ij} = \frac{1}{n} \langle s^i - \bar{s}, s^j - \bar{s} \rangle_{\mathbb{R}^N}$. Since G is a symmetric and positive semi-definitive matrix we obtain non-negative eigenvalues $\{\lambda_j\}_j$ and corresponding orthonormal eigenvectors $\{w_j\}_j$, i.e. $Gw_j = \lambda_j w_j$ for $j = 1, \dots, n$. Finally, the principal modes of variation of the data set $\{s^1 - \bar{s}, \dots, s^n - \bar{s}\}$ are obtained via $v_j = \lambda_j^{-1/2} D^T w_j \in \mathbb{R}^N$ for $j = 1, \dots, n$.

PCA in Euclidean space easily translates to Riemannian manifolds [FLPJ04]. To this end, one considers data points s^1, \dots, s^n on the manifold \mathcal{M} and performs a classical PCA for the logarithms $\log_{\bar{s}} s^j$ of the input shapes s^j with respect to their Fréchet average \bar{s} – the Riemannian counterpart of the arithmetic average. Thereby, the tangent vector $u^j = \log_{\bar{s}} s^j$ represents the geometric variation of s^j relative to the average \bar{s} in an infinitesimal sense. Here, the metric $g_{\bar{s}}$ is taken into account as the scalar product on these infinitesimal shape variations. Thus, Gram's matrix is defined by $G_{ij} = \frac{1}{n} g_{\bar{s}}(u^i, u^j)$ and—as before—its spectral decomposition leads to the pairing $(v_j, \lambda_j)_{j=1, \dots, n}$ which is called Principal Geodesic Analysis (PGA).

Discrete Riemannian calculus on the space of shells. Rumpf and Wirth [RW15] introduced a discrete Riemannian calculus on Hilbert manifolds. Using (2) on consecutive pairs of interpolated shapes $s_k = \mathfrak{s}(k/K)$ for $k = 0, \dots, K$ and the Cauchy-Schwarz inequality one obtains

$$\mathcal{E}[\mathfrak{s}] \geq K \sum_{k=1}^K \text{dist}^2(s_{k-1}, s_k). \quad (5)$$

Note that (5) becomes an equality iff \mathfrak{s} is already a geodesic path. Now, the key ingredient of the discrete calculus is a functional $\mathcal{W} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ which *locally* approximates the squared Riemannian distance, i.e.

$$\mathcal{W}[s, \bar{s}] = \text{dist}^2(s, \bar{s}) + \mathcal{O}\left(\text{dist}^3(s, \bar{s})\right) \quad (6)$$

and replacing dist^2 by \mathcal{W} in (5) leads to the definition of a *discrete path energy*

$$E[\mathfrak{s}] := K \sum_{k=1}^K \mathcal{W}[s_{k-1}, s_k], \quad (7)$$

where \mathfrak{s} denotes a polygonal path with vertices $s_k = \mathfrak{s}(k/K)$ for $k = 0, \dots, K$. A minimiser of (7) for fixed endpoints is referred to as *discrete K -geodesic*, where the minimisation of (7) is with respect to the $K-1$ inner vertices $\{s_1, \dots, s_{K-1}\} \subset \mathcal{M}$. It is shown in [RW15] that under suitable assumptions discrete K -geodesics converge to continuous geodesics for $K \rightarrow \infty$.

Here, we pick up the discrete calculus on the space of discrete shells [HRWW12, HRS*14]. For a fixed mesh topology a discrete shell can be identified with the vector of vertex positions in \mathbb{R}^{3M} , where M is the number of vertices. The space of discrete shells $\mathcal{M} \subset \mathbb{R}^{3M}$ is then equipped with a metric which measures the energy dissipation caused by infinitesimal membrane distortion and normal bending (cf. Fig. 2b). The definition of the metric is based on an elastic deformation energy $\mathcal{W}[s, \bar{s}]$ for thin shells needed to deform $s \in \mathcal{M}$ into $\bar{s} \in \mathcal{M}$. To account for the physical properties of thin elastic shells, \mathcal{W} splits into a membrane and a bending distortion energy (cf. Fig. 2b), i.e.

$$\mathcal{W}[s, \bar{s}] = \mathcal{W}_{\text{mem}}[s, \bar{s}] + \mathcal{W}_{\text{bend}}[s, \bar{s}]. \quad (8)$$

Thereby, the bending energy is taken from [GHDS03]. The discrete shell model is physically valid for thin shell materials. More generally, it proves useful for modelling a much wider class of objects by capturing two important modes of deformation: bending and stretching. Concretely, the membrane and bending energies are defined as follows (Here, quantities with a tilde always refer to the deformed configuration):

$$\mathcal{W}_{\text{mem}}[s, \bar{s}] = \delta \sum_{t \in \mathcal{T}(s)} a_t W(\mathcal{G}_t), \quad \mathcal{W}_{\text{bend}}[s, \bar{s}] = \delta^3 \sum_{e \in \mathcal{E}(s)} \frac{(\theta_e - \tilde{\theta}_e)^2}{a_e} l_e^2,$$

where $\mathcal{T}(s)$ and $\mathcal{E}(s)$ denote the set of triangles resp. edges of s , $\delta > 0$ is the physical thickness and $W : \mathbb{R}^{2,2} \rightarrow \mathbb{R}$ is the hyperelastic energy density given by Eq. (8) in [HRWW12]. Furthermore, $\mathcal{G}_t \in \mathbb{R}^{2,2}$ is a two-dimensional representation of the Cauchy Green strain tensor of the deformation of the triangle t , a_t is the triangle volume of t , l_e is the edge length of e , θ_e is the dihedral angle at e and $a_e = \frac{1}{3}(a_t + a_{t'})$ is an area weight associated with $e = t \cap t'$. In detail, if $e_0, e_1, e_2 \in \mathbb{R}^3$ are edges of a triangle t , a discrete first fundamental form on t is given by $g_t = [e_2 | -e_1]^T [e_2 | -e_1] \in \mathbb{R}^{2,2}$, which yields the representation $\mathcal{G}_t = g_t^{-1} \tilde{g}_t$.

In order to retrieve the underlying metric one can apply Rayleigh's paradigm by replacing strains by strain rates for a second order approximation of this energy. Indeed, due to [HRS*14, Thm 1] the Hessian of (8) actually induces a Riemannian metric on the space of discrete shells modulo rigid body motions. In particular, the deformation energy \mathcal{W} represents a consistent approximation of the induced (squared) Riemannian distance as in (6).

4. Discrete principal geodesic analysis

Based on these preliminaries we now derive a principal geodesic analysis on the space of discrete shells. The central building blocks are a discrete geodesic average, an approximation of Gram's matrix, and the computation of principal modes of variation.

A critical observation of the discrete shell space introduced by Heeren et al. [HRWW12] is its rigid body motion invariance incorporated in (8), i.e.

$$\mathcal{W}[s, \bar{s}] = \mathcal{W}[s, R\bar{s} + b] \quad (9)$$

for $R \in SO(3)$ (the space of rotation matrices in \mathbb{R}^3) and $b \in \mathbb{R}^3$. Indeed, a discrete shell is no longer a single triangular mesh s but an equivalence class of shells $[s] = SO(3)s + \mathbb{R}^3$, cf. the sketch in Fig. 2a. As a consequence the shape manifold \mathcal{M} is a space of such equivalence classes. For simplicity we stick to the notation s instead of $[s]$. Then tangent vectors – as they appear in the classical principal geodesic analysis – are equivalence classes as well, where the associated Lie algebra $so(3)$ has to be taken into account. This renders the computational treatment of the shell manifold’s tangent bundle very cumbersome. In what follows, we will derive a rigid body motion invariant, discrete principal geodesic calculus based on elastic energy \mathcal{W} . Thus, in all components of our algorithm we will solely treat discrete shells and avoid any direct tangent vector computation.

Discrete geodesic average. Let s^1, \dots, s^n be discrete input shells in \mathcal{M} . The Riemannian average on the manifold \mathcal{M} – called Fréchet mean – is obtained by using in (4) the Riemannian distance (2) in place of the Euclidean distance. Further replacing \mathcal{E} by the discrete path energy (7) in (2) yields the definition of a *discrete geodesic average*

$$\bar{s} = \arg \min_{s \in \mathcal{M}} \sum_{i=1}^n \min_{\substack{s^i(0)=s, \\ s^i(1)=s^i}} E[s^i], \quad (10)$$

where the interior minimisation is over a polygonal spider consisting of all polygonal paths s^i connecting the average $s^i(0) = \bar{s}$ and the input shapes $s^i(1) = s^i$ for $i = 1, \dots, n$ as shown in Fig. 4. Obviously, \bar{s} is invariant with respect to rigid body motions due to (9).

Approximation of Gram’s matrix. Next, we substitute metric evaluations on tangent vectors in the definition of Gram’s matrix by evaluations of the squared distance directly on discrete shells, and then in a second step by the corresponding local approximation (6) as follows. Due to property (3) of geodesic paths we obtain

$$g(u^j, u^j) = \text{dist}^2(\bar{s}, s^j) = \sigma^{-2} \text{dist}^2(\bar{s}, s^j(\sigma)) \approx \sigma^{-2} \mathcal{W}[\bar{s}, s^j(\sigma)]$$

for the tangent vectors $u^j = \log_{\bar{s}}(s^j)$ in the standard PGA. Here, $\sigma > 0$ is some generic scaling factor and $s^j : [0, 1] \rightarrow \mathcal{M}$ is the geodesic connecting \bar{s} and s^j . Note that we have used the shortcut notation $g = g_{\bar{s}}$ (here and in the following). For the off-diagonal entries of Gram’s matrix we take into account the polar formula

$$g(u^j, u^i) = \frac{1}{2} \left(g(u^j, u^j) + g(u^i, u^i) - g(u^j - u^i, u^j - u^i) \right)$$

and an analogous, now also in the first replacement approximate, identity

$$g(u^j - u^i, u^j - u^i) \approx \sigma^{-2} \text{dist}^2(s^j(\sigma), s^i(\sigma)) \approx \sigma^{-2} \mathcal{W}[s^j(\sigma), s^i(\sigma)],$$

to replace evaluations of the metric with (approximative) squared distances on \mathcal{M} . Finally, we replace $s^j(\sigma)$ by $s_1^j = \mathbf{I}(\bar{s}, s^j, \sigma)$, where \mathbf{I} denotes the discrete geodesic interpolation operator, as described in Section 8, and $\sigma = \sigma(K) = \frac{1}{K}$ is a suitable choice, which retrieves the first node along the discrete K -geodesic from \bar{s} to s^j . Altogether we define the entries of an approximative Gram’s matrix G (which




Figure 3: Convergence of the discrete Gram matrix and its eigenvectors and eigenvalues as $K \rightarrow \infty$ for the SCAPE dataset shown in Fig. 5. We show RMS relative error, using $K_{max} = 16$ as pseudo ground truth. Second order convergence is illustrated by the green triangle.

actually depends on K) as

$$G_{ij} = \frac{\mathcal{W}[\bar{s}, s_1^i] + \mathcal{W}[\bar{s}, s_1^j] - \frac{1}{2} \left(\mathcal{W}[s_1^i, s_1^j] + \mathcal{W}[s_1^j, s_1^i] \right)}{2n\sigma^2} \quad (11)$$

for $i, j = 1, \dots, n$. The additional symmetrisation in the last terms ensures symmetry of G . Again, due to the rigid body motion invariance (9) the resulting G does not depend on the chosen representation of the equivalence classes of discrete shells. As before we obtain approximate eigenvalues $\{\lambda_j\}_j$ and corresponding (orthonormal) eigenvectors $\{w_j\}_j \subset \mathbb{R}^n$ with $Gw_j = \lambda_j w_j$ for $j = 1, \dots, n$. Applying the convergence theory for the discrete calculus developed in [RW15] one obtains that \bar{s} converges to the Fréchet mean and G converges to the original Riemannian Gram matrix $\left(\frac{1}{n} g(u^j, u^i) \right)_{ij}$ for $K \rightarrow \infty$. We demonstrate this convergence empirically in Fig. 3.

Principal variations instead of principal components. Next, we replace the principal component (eigenmode) v_j in the tangent space at the Fréchet mean by a (nonlinear) discrete principal variation on the shape manifold \mathcal{M} . Let us start with a straightforward observation. For some $\alpha \in \mathbb{R}^n$ with $\sum_{i=1, \dots, n} \alpha_i = 1$, we consider the linear combination $u[\alpha] = \sum_{i=1, \dots, n} \alpha_i u^i$ of tangent vectors $u^i = \log_{\bar{s}} s^i$. Then $u[\alpha]$ can be characterized as the minimizer of the quadratic functional $u \mapsto \sum_{i=1, \dots, n} \alpha_i g(u - u^i, u - u^i)$. Using Taylor expansion in σ for a given $\alpha \in \mathbb{R}^n$ this implies that for

$$p^\sigma[\alpha] := \arg \min_{p \in \mathcal{M}} \sum_{i=1, \dots, n} \alpha_i \text{dist}^2(s^i(\sigma), p), \quad (12)$$

the rescaled logarithm $\frac{1}{\sigma} \log_{\bar{s}} p^\sigma[\alpha]$ converges to $u[\alpha]$ for $\sigma \rightarrow 0$. Hence, $p^\sigma[\alpha] \in \mathcal{M}$ can be considered as a *nonlinear* variation of the Fréchet mean corresponding to the *linear* infinitesimal variation $u[\alpha]$ in the tangent space at the Fréchet mean.

Again, we replace dist^2 in (12) by its local approximation \mathcal{W} as well as $s^i(\sigma)$ by the discrete geodesic interpolation $s_1^i = \mathbf{I}(\bar{s}, s^i, \sigma)$ with $\sigma = 1/K$ and obtain

$$p[\alpha] := \arg \min_{p \in \mathcal{M}} \sum_{i=1, \dots, n} \alpha_i \mathcal{W}[s_1^i, p] \quad (13)$$

for some coefficient vector $\alpha \in \mathbb{R}^n$. However, we have to proceed




Figure 4: Submanifold \mathcal{M}^J (yellow) and polyhedron $\mathcal{C}^J \subset \mathcal{M}^J$ (with red boundary) spanned by nonlinear combinations of principal variations $\{p_j\}_j$. Note that the input shapes $\{s_k\}_k \subset \mathcal{M}$ do not lie on \mathcal{M}^J in general. The polygonal spider connecting input shapes and Fréchet mean is drawn in grey.




Figure 5: Time-discrete PGA models built on TOSCA cats [BBK08] and SCAPE [ASK*05]. Input training shapes (yellow), mean shape (orange) and first five principal variations (green).

with special care in particular for entries of α that might be negative. If $\alpha_i < 0$ for some i we replace α_i by $|\alpha_i|$ and $s_1^i = \mathbf{I}(\bar{s}, s^i, \sigma)$ by its discrete geometric reflection at \bar{s} involving extrapolation via a discrete exponential map (see Section 8), i.e. $\mathbf{I}(\bar{s}, s^i, -\sigma)$. This is necessary because \mathcal{W} is no longer quadratic and there is no a priori control of the growth of \mathcal{W} for general coefficients $\alpha_i \in \mathbb{R}$. Thus, without this modification existence of minimisers in (13) are not guaranteed. Finally, we define *discrete principal variations* by choosing α to be the eigenvectors $w_j = (w_{j,i})_{i=1,\dots,n}$ of the approximate Gram’s matrix G , i.e.

$$p_j := \arg \min_{p \in \mathcal{M}} \sum_{i=1,\dots,n} |w_{j,i}| \mathcal{W} \left[\mathbf{I}(\bar{s}, s^i, \text{sgn}(w_{j,i})/K), p \right] \quad (14)$$

for $j = 1, \dots, J$, where we have rescaled $w_j \in \mathbb{R}^n$ such that its entries sum to 1 (which does not affect the minimiser).

Due to the convergence of the discrete Fréchet mean and the discrete Gram matrix we expect that for an eigenvalue of multiplicity 1 for $K \rightarrow \infty$ the eigenvalues λ_j converge to their continuous counterparts and $K \log_{\bar{s}} p_j$ converges (up to scaling) to a representative of the corresponding principal component v_j .

Evaluation. In Fig. 5 we show two time-discrete PGA models (average and first five principal variations for $K = 4$). We visualise




Figure 6: Model compactness with respect to \mathcal{W} for models shown in Fig. 5 (left and centre) and Fig. 1 (right). Number of model dimensions on x axis, proportion of variance captured on y axis.

the j th principal variation by using the geodesic interpolation operator $t \mapsto \mathbf{I}(\bar{s}, p_j, \pm t)$ to sample along the one dimensional principal geodesic and overlay the resulting shapes. Note that they clearly correspond to nonlinear motions present in the training data. In Fig. 6 we show model compactness as a function of the number of retained modes for these two models and the one used in Fig. 8 and Fig. 12. Note that, in all three cases, we are able to compress a significant proportion of the variance into a small number of modes.

5. Submanifold projection

In this section we define a local submanifold “spanned” by the principal variations defined in (14) as illustrated in Fig. 4. This is the nonlinear counterpart of the linear subspace spanned by the principal components in classical PCA or standard PGA. The projection onto this submanifold returns a discrete shell which is uniquely determined by a small set of weights and approximates the input shape on the basis provided by our Riemannian statistical analysis.

Defining the submanifold. We consider (14) for the J dominant principal variations and also their associated reflections $p_{-j} = \mathbf{I}(\bar{s}, p_j, -1)$ (the sign of a principal component is arbitrary so our submanifold includes variations in both directions). At first we define the convex Riemannian polyhedron induced by the vertices $\{p_j | j = -J, \dots, -1, 1, \dots, J\}$. Discrete shells on the polyhedron are obtained by computing “variational Riemannian” combinations of the p_j for weights $\alpha = (\alpha_{-J}, \dots, \alpha_{-1}, \alpha_1, \dots, \alpha_J) \in \mathbb{R}^{2J}$ subject to $\sum_{j=-J,\dots,J} \alpha_j = 1$ and $\alpha_j \geq 0$, i.e.

$$\mathcal{C}^J = \left\{ \arg \min_{p \in \mathcal{M}} \sum_{j=-J}^J \alpha_j \mathcal{W}[p_j, p] \mid \sum_{j=-J}^J \alpha_j = 1, \alpha_j \geq 0 \right\} \quad (15)$$

with the notational convention $\alpha_0 = 0$ and p_0 staying undefined. Note that in particular $\bar{s} \in \mathcal{C}^J$, e.g. for $\alpha_j = \alpha_{-j} = \frac{1}{2}$ and $\alpha_i = \alpha_{-i} = 0$ if $i \neq j$ for an arbitrary $j \in \{1, \dots, J\}$ as an example that different choices of α might represent the same shell on \mathcal{C}^J .

Now, the actual submanifold \mathcal{M}^J is defined via discrete geodesic extrapolation of the convex polyhedron \mathcal{C}^J using the interpolation \mathbf{I} for a discrete shell $p \in \mathcal{C}^J$ and times $t > 0$ (cf. Fig. 4):

$$\mathcal{M}^J := \left\{ \mathbf{I}(\bar{s}, p, t) \mid p \in \mathcal{C}^J, t > 0 \right\}. \quad (16)$$

We might allow for non vanishing α_0 and $p_0 := \bar{s}$, which does not alter the definition of \mathcal{C}^J . Note that (15) can also be constructed directly from the input data without computing principal variations first. This might be reasonable if the number of input data is small.




Figure 7: Right: Projection of an unseen shape s onto the model space \mathcal{M}^J : scale s to s_{loc} , project s_{loc} locally to $\mathcal{P}_{loc}[s_{loc}] \in \mathcal{C}^J$, and finally rescale to get $\mathcal{P}[s] \in \mathcal{M}^J$. Left: Model fitting of s^* driven by sparse landmarks $X \in \mathbb{R}^{3L}$ depending on fitting parameter $\gamma > 0$.

The tangent space to \mathcal{M}^J at \bar{s} is spanned by the $\log_{\bar{s}} p_j$ (which converge to v_j for $K \rightarrow \infty$). Altogether, we get that

$$K \log_{\bar{s}}(\mathcal{M}^J) \rightarrow \text{span}(\{v_1, \dots, v_J\}) \quad \text{for } K \rightarrow \infty.$$

That both principal variations p_j and their reflections p_{-j} are indispensable to our submanifold construction reflects the fact that the infinitesimal counterpart, the principal components v_j , generate one dimensional geodesic subspaces and not just geodesic rays.

Defining a projection onto the submanifold. In what follows, we will derive a suitable projection of a given discrete shell $s \in \mathcal{M}$ on the (approximate) submanifold \mathcal{M}^J as defined in (16). The classical Riemannian projection or projection onto the submanifold defined as exponential map of the subspace of the tangent space spanned by the dominant J principal components v_1, \dots, v_J would work as follows: First compute an infinitesimal representation $v = \log_{\bar{s}} s$ of s in the tangent space at the Fréchet mean, then project v (locally) onto the subspace $\text{span}\{v_1, \dots, v_J\}$ via the formula $v^J = \sum_{j=1, \dots, J} g_{\bar{s}}(v, v_j) v_j$ and finally compute the actual projection $\mathcal{P}[s] = \exp_{\bar{s}} v^J$. Note that this closed-form projection identity for v^J only holds if $\{v_1, \dots, v_J\}$ is an orthonormal system.

Once more the incorporation of rigid body motion invariance is a very delicate undertaking. Just replacing the metric $g_{\bar{s}}(\cdot, \cdot)$ by the approximation used in the definition of the discrete Gram matrix in (11) does not lead to a satisfactory solution. Indeed, the expected orthogonality relation $g_{\bar{s}}(v_i, v_j) = \delta_{ij}$ holds only approximately and that deteriorates the Gram-Schmidt orthogonalisation procedure to compute the *linear* projection v^J (see paragraph above). Instead, we propose to perform a nonlinear projection on the approximating manifold \mathcal{M}^J consisting of three elementary steps: *scaling*, *local projection*, and *rescaling*. These steps are illustrated in Fig. 7 and defined in detail as follows.

[Scaling] Firstly, we scale the given shape s in order to make sure that it can be locally projected onto the polyhedron \mathcal{C}^J (i.e. we ensure $\alpha_j \geq 0$). This is done by means of the discrete geodesic interpolation (see Section 8), i.e. we define $s_{loc} = \mathbf{I}(\bar{s}, s, \rho)$ where

$$\rho := \kappa \frac{\min_j \text{dist}(\bar{s}, p_j)}{\text{dist}(\bar{s}, s)} \quad (17)$$

for sufficiently small $\kappa > 0$. The resulting scaling factor ρ is in general not a multiple of $\frac{1}{\kappa}$. Hence, a discrete geodesic interpolation $\mathbf{I}(\cdot, \cdot, t)$ for general $t \in \mathbb{R}$ is needed (see also Section 8).

[Local projection] Secondly, we aim at computing a local projection as the best approximation of s_{loc} on \mathcal{C}^J . Let us first review the projection onto a convex set $\mathcal{C} = \{\sum_j \alpha_j q_j \mid \sum_j \alpha_j = 1, \alpha_j \geq 0\}$ in Euclidean space for a given set of points $q_1, \dots, q_J \in \mathbb{R}^N$. For some arbitrary point $p \in \mathbb{R}^N$ the projection can be written as

$$\mathcal{P}_{Euc1}[p] = \arg \min_{q \in \mathcal{C}} \text{dist}^2(p, q),$$

where $\text{dist}^2(\cdot, \cdot)$ is the squared Euclidean distance. Note that the projection coincides with the usual *orthogonal* projection onto the linear space $\text{span}(q_1, \dots, q_J) \supset \mathcal{C}$ if $\mathcal{P}_{Euc1}[p]$ is an interior point in \mathcal{C} (in the relative topology of \mathcal{C}). This formulation translates one-to-one to the local projection of a shell $s_{loc} \in \mathcal{M}$ onto $\mathcal{C}^J \subset \mathcal{M}^J$ for small κ , again by replacing dist^2 by the local approximation \mathcal{W} . We define

$$\mathcal{P}_{loc}[s_{loc}] = \arg \min_{q \in \mathcal{C}^J} \mathcal{W}[s_{loc}, q], \quad (18)$$

where the constraint $q \in \mathcal{C}^J$ is equivalent to

$$q \in \left\{ \arg \min_{p \in \mathcal{M}^J} \sum_{j=-J}^J \alpha_j \mathcal{W}[p_j, p] \mid \sum_{j=-J}^J \alpha_j = 1, \alpha_j \geq 0 \right\}. \quad (19)$$

In our applications $\kappa = \frac{1}{2}$ in (17) already implies that $\mathcal{P}_{loc}[s_{loc}]$ is an interior point in \mathcal{C}^J .

[Rescaling] Finally, we rescale the local projection to define the desired projection

$$\mathcal{P}[s] = \mathbf{I}(\bar{s}, \mathcal{P}_{loc}[s_{loc}], 1/\rho). \quad (20)$$

By means of this nonlinear projection method we are able to represent an arbitrary shape s in terms of $2J + 1$ scalar variables, i.e. $\alpha \in [0, 1]^{2J}$ to represent $\mathcal{P}_{loc}[s_{loc}] \in \mathcal{C}^J$ and $\rho > 0$ as in (17), which allows for a substantial compression rate. For example, we visualise α for $J = 5$ in Fig. 1 (bottom, left).

Let us emphasise that the constrained optimisation problem incorporated in the projection \mathcal{P}_{loc} does not require any treatment of tangent vectors and is built on the rigid body motion invariant energy functional \mathcal{W} .

Evaluation. We show a qualitative example of submanifold projection in Fig. 8. The input shape (gray) is projected onto the submanifold obtained by building a discrete PGA model (with $K = 4$) using the Dyna dataset (model shown in Fig. 1). We vary the model dimensionality over $J = 5, 11, 17$ and show the approximated shape in yellow. The subtleties of the shape are correctly reconstructed as J increases, yielding a smooth residual energy. We evaluate the generalisation ability of our model in Fig. 9. We compare against [FB12] with 60 dimensions retained, the data-driven approach of [GLL*16] using all training shapes and the Shell PCA model [ZHR15]. Using only 20 dimensions, our model generalises almost as well as [GLL*16] and outperforms the other two models substantially.



Figure 8: Qualitative visualisation of input shape (gray) projected onto model in Fig. 1 with (cols 2-4) $J = 5, 11, 17$ dimensions. Col 5 shows residual energy of projection with $J = 17$.




Figure 9: Leave-one-out evaluation of generalisation error on the SCAPE data set compared to [GLL*16] (using all shapes), [FB12] (60 dimensions) and [ZHR15]

6. Mesh editing via hard constraints

Our method can be used for model-based mesh editing. Assume we are given a discrete PGA model and a set of *handle vertex positions*. Now, one positions (a subset of) the handle vertices manually and asks for a shell obeying the new handle positions while being a physically plausible deformation of a shell lying on the statistical submanifold. Using the submanifold projection introduced in Section 5 we define this shell as the minimiser s of the energy

$$\mathcal{W}[s, \mathcal{P}[s]] \quad (21)$$

subject to the constraint positions of the deformed handle vertices. Thus, we ask for the “closest” (in terms of the elastic energy functional \mathcal{W}) discrete shell s to the nonlinear submanifold associated with the dominant J principal variations of our training data. Note that (21) is again an approximation to the actual (squared) distance.

Depending on the application one can either regard s or $\mathcal{P}[s]$ as a solution. Indeed, s exactly obeys the prescribed handle vertex positions but $s \notin \mathcal{M}^J$ in general, whereas $\mathcal{P}[s] \in \mathcal{M}^J$ and can be represented by the $2J$ weights α_j but the constraint of the prescribed handle vertex positions is usually fulfilled only approximately. Note that this mesh editing tool comes with a selection of a particular representative s from its equivalence class $[s]$, which is determined by the handle vertex positions (as long as there are at least 3 handle vertices not lying on a line).

Fig. 10 shows mesh editing results for five comparison methods and our proposed approach. [SA07] and [SSP07] are classical mesh editing approaches that use only a single reference mesh. The




Figure 10: Comparison of mesh editing results. (a) initial pose, (b) [SA07], (c) [SSP07], (d) [SZGP05], (e) [FB11], (f-g) [GLL*16], (h-i) Ours (with $K = 4, J = 20$).




Figure 11: Mesh editing with five (a-b) vs. six (c-d) handle positions to be fitted, where the handle at the tail is shifted.

challenging configuration of handles causes these methods to fail dramatically. [SZGP05], [FB11] and [GLL*16] are data-driven and use the same set of training shapes as we use to build our model. These provide more natural results but [SZGP05] and [FB11] produce significant distortions and self-intersections while even the state of the art [GLL*16] loses details, causes the arms to thin and the back to curve and deforms the head. Our result preserves details and retains plausible arms and head and a straight back. Note though that the thickening of the left foot is an artefact. This is a result of the training data not including examples with such severe bending at the hip. To fit the handle on top of the foot, the solution deforms the foot rather than further bending the upper leg.

To obtain the desired result of the edit, it might be necessary to take into account sufficiently many handles as indicated in Fig. 11. Here, we consider the cat model (cf. Fig. 5) first with five handles and fit to modified handle positions in which the tail tip is moved. To minimise in particular the bending energy our method significantly bends the whole object. This can easily be prevented by adding a sixth handle on the back of the cat (cf. Fig. 11, c and d).

7. Model fitting via soft constraints

In this section we relax the hard constraint for the handle vertices in the mesh editing application by means of a soft penalty approach. In particular, this allows us to reconstruct a discrete shell from (potentially noisy) input data from a motion capture device. In this case, the input data is given as a vector of L sparse marker positions, i.e. $x = (x_\ell)_{\ell=1, \dots, L}$, corresponding to vertex positions $\mathcal{X}_\ell(s)$




Figure 12: *Qualitative results of fitting to motion capture data. Frames from original sequence (top) shown with corresponding reconstruction (bottom, using the same model as Fig. 8 with $K = 4$, $J = 10$).*

on the mesh s . Knowing these correspondences, we measure the mismatch of some discrete shell $s \in \mathcal{M}$ and the given landmarks by $\mathcal{F}_x[s] = \sum_{\ell=1}^L \|\mathcal{X}_\ell(s) - x_\ell\|_{\mathbb{R}^3}^2$. Now, we consider $\mathcal{W}[s, \mathcal{P}[s]]$ as a prior for the identification of a reconstructed discrete shell. Hence, we seek a minimizer s of the model fitting energy given by

$$\mathcal{F}_x[s] + \gamma \mathcal{W}[s, \mathcal{P}[s]] \quad (22)$$

for some weight $\gamma > 0$, which controls the proximity of s with respect to our submanifold \mathcal{M}^J for given training data, cf. Fig. 7 (left). Again this ansatz comes with a selection of a particular representative s from its equivalence class $[s]$, which is driven by the data term. For the numerical solution of this problem, we make use of the following alternating scheme (based on the initial guess $\mathcal{P}[s] = \bar{s}$): First, we minimize (22) in s for fixed $\mathcal{P}[s]$. If necessary, we re-compute $\mathcal{P}[s]$ (see Sec. 5) and go back to the first step. In our application this scheme quickly converges and only very few iterations already give very satisfactory fitting results. In practice, we use two iterations for the results shown.

In Figures 12 and 13 we show qualitative results of fitting to 41 markers in sequences from the CMU mocap dataset and 89 markers from MPI MoSh dataset [LMB14] respectively. Fig. 12 shows a result in which the learnt body model has quite different geometry to that of the performer. Note that the video frames are just shown for comparison - we use only the 3D marker data as input. Our fitted model is still able to capture the dynamic poses of the performance.

In Fig. 13 we compare against [LMB14]. It should be noted that this method uses a model of substantially higher complexity than ours. It is trained on 3,803 body scans in neutral pose and 1,832 body scans in dynamic poses and uses a 19 parameter skeleton model and retains up to 300 dimensions of the statistical deformation model (10 used in Fig. 13). Our result is obtained using a model trained on 20 scans of a single person (chosen to match the body shape of the performer), is entirely mesh-based (we have no articulation model) and we also retain only 10 principal variations. Nevertheless, our results are qualitatively very similar.




Figure 13: *Comparison of reconstruction from motion capture data with the MoSh model [LMB14]. Although MoSh (top) is trained on more than 5,000 scans and uses an additional skeleton model, our method with $K = 4$ (bottom) obtains similar results using 10 principal variations only, trained on a subset of 20 shapes from Dyna.*


8. Computational tools

Here, we collect all algorithmic ingredients of the presented approach and discuss their computational complexity.

Discrete geodesic interpolation. A discrete K -geodesic is defined as the minimiser of the discrete path energy (7). Thus, the unknowns s_1, \dots, s_{K-1} determining the polygonal path \mathfrak{s} solve the system of Euler–Lagrange equations

$$\mathcal{W}_{,2}[s_{k-1}, s_k] + \mathcal{W}_{,1}[s_k, s_{k+1}] = 0 \quad (23)$$

for $k = 1, \dots, K-1$ with $s_0 = s_A \in \mathcal{M}$ and $s_K = s_B \in \mathcal{M}$ being fixed. Here, \mathcal{W}_i denotes the variation with respect to the i th argument. For $t = k/K$ for some $0 \leq k \leq K$ we set $\mathbf{I}(s_A, s_B, k/K) = s_k$. For $t = m/K$ with arbitrary $m \in \mathbb{Z}$ we define a discrete extrapolation by an iterative scheme based on the following induction: Assume $k \geq K$, such that s_{k-1} and s_k are already known, then we compute s_{k+1} to be the solution of (23). Likewise, for $k \leq 0$, such that s_k and s_{k+1} are already known, we define s_{k-1} to be the solution of (23). With these extrapolated discrete shells at hand we define $\mathbf{I}(s_A, s_B, t)$ for arbitrary multiples t of $\frac{1}{K}$. Finally, for general $t \in \mathbb{R}$ we denote $t(K) = tK - \lfloor tK \rfloor$ (where the floor function $\lfloor \cdot \rfloor$ returns the largest integer less than or equal to the



argument) and define $\mathbf{I}(s_A, s_B, t)$ as the midpoint s of a discrete 3-geodesic $(s_{\lfloor tK \rfloor}, s, s_{\lfloor tK \rfloor + 1})$ minimising

$$(1 - t(K))\mathcal{W}[s_{\lfloor tK \rfloor}, s] + t(K)\mathcal{W}[s, s_{\lfloor tK \rfloor + 1}], \quad (24)$$

where $s_m = \mathbf{I}(s_A, s_B, m/K)$ for $m \in \mathbb{Z}$, as described above. In particular, $\mathbf{I}(s_A, s_B, -1)$ defines a discrete Riemannian reflection of s_B about s_A . Computationally, we use Newton's method to solve the nonlinear system (23) with $(K-1)3M$ unknowns, which yields

the evaluation of $\mathbf{I}(\cdot, \cdot, t)$ for $t = k/K$ with $1 \leq k \leq K-1$. A single step of extrapolation requires to solve (23) which is obtained by solving a nonlinear system with $3M$ variables.

Discrete Fréchet mean. The most costly task is the computation of the discrete Fréchet mean \bar{s} defined in (10). The degrees of freedom (dofs) are the shells defining the n polygonal paths s^i (with $3n(K-1)M$ dofs) connecting the input shells s^i and \bar{s} (with its $3M$ dofs). Each arc of the polygonal spider has to solve the system of Euler–Lagrange equations for a single discrete K -geodesic (i.e. (23) for $0 < k < K$) and the coupling at the center is described by the Euler–Lagrange equation

$$0 = \sum_{i=1}^n \beta_i \mathcal{W}_{,2}[s^i(1/K), \bar{s}] \quad (25)$$

with $\beta_i = 1/n$ and $s^i(1/K)$ is the first discrete shell along the discrete path from \bar{s} to the i th input shape s^i . This coupled problem is again solved by Newton’s method.

Gram’s matrix and spectral analysis. The evaluation of the approximate Gram’s matrix (11) only consists of scaling based on the discrete geodesic interpolation and of evaluations of \mathcal{W} . The spectral decomposition of $G \in \mathbb{R}^{n,n}$ can easily be solved.

Principal variations. The computation of principal variations p_j via (14) again involves a discrete geodesic scaling as well as computing a local weighted average similar to (25) but with non-constant weights $\beta_i = |w_{j,i}|$.

Submanifold projection. We solve the constrained optimisation problem (18) using a Quasi Newton method. To this end, we define the objective functional $\mathcal{J}[\alpha] = \mathcal{W}[s_{loc}, q[\alpha]]$ for $\alpha \in \mathbb{R}^{2J}$ and $q = q[\alpha] \in \mathbb{R}^{3M}$ as the (locally unique) minimiser of

$$q \mapsto \mathcal{A}[\alpha, q] = \sum_{j=-J, \dots, J} \alpha_j \mathcal{W}[p_j, q] \quad (26)$$

for fixed $\alpha_0 = 0$. To apply a Quasi Newton scheme we have to evaluate the partial derivatives of \mathcal{J} with respect to α_i . For $\alpha \in \mathbb{R}^{2J}$ the constraint on $q = q[\alpha]$ is given by $\mathcal{G}[\alpha, q] := \partial_q \mathcal{A}[\alpha, q] = 0$. The solution of (18) is linked to a saddle point of $\mathcal{L}[q, \alpha; \mu] := \mathcal{W}[s_{loc}, q] + \mathcal{G}[\alpha, q] \cdot \mu$, where μ is the vector of Lagrange multipliers. This leads to the nonlinear system $0 = D_{(q, \alpha, \mu)} \mathcal{L}[q, \alpha; \mu]$, i.e.

$$0 = D_q \mathcal{L}[q, \alpha; \mu] = \mathcal{W}_{,2}[s_{loc}, q] + \sum_{j=-J, \dots, J} \alpha_j \mathcal{W}_{,22}[p_j, q] \cdot \mu, \quad (27)$$

$$0 = D_{\alpha_j} \mathcal{L}[q, \alpha; \mu] = \mathcal{W}_{,2}[p_j, q] \cdot \mu, \quad j = 1, \dots, J, \quad (28)$$

$$0 = D_{\mu} \mathcal{L}[q, \alpha; \mu] = \sum_{j=-J, \dots, J} \alpha_j \mathcal{W}_{,2}[p_j, q], \quad (29)$$

where $\mathcal{W}_{,22}$ denotes the Hessian with respect to the second argument. By a classical result of constrained optimisation the right hand side of (28) returns the derivatives of the cost functional \mathcal{J} with respect to α_j . Thus, to evaluate $\partial_{\alpha_j} \mathcal{J}$ we first solve the nonlinear equation (29) for q via Newton’s method, the linear equation (27) for μ via the conjugate gradient method and then apply (28) to obtain $\partial_{\alpha_j} \mathcal{J}[\alpha] = \mathcal{W}_{,2}[[p_j, q] \cdot \mu$.

Multilevel algorithms. Solving a nonlinear system in $O(3MnK)$ variables directly is inefficient at least for larger M , n , and K . For

this reason, we use a multi-resolution approach for all nonlinear optimisation problems above. First, we coarsen all of the input shapes simultaneously by applying an iterative edge collapse approach based on the minimisation of the quadric error metric [GH97] and computed groupwise, as in [MG03], to preserve the dense correspondence between input shapes. We then solve the nonlinear optimisation problem on resulting meshes with reduced resolution with < 1000 vertices. Afterwards, the coarse solution is then prolonged to the original resolution, using the prolongation scheme from [FB11]. Then a fine scale optimisation can optionally be performed using the prolonged result as initialisation. For a discussion of the accuracy of this approach we refer to [FB11, Table 2].

Furthermore, for the computation of the discrete Fréchet mean we make use of an alternating relaxation and a cascadic approach along the discrete curves of the “spider”. For the alternating scheme we first relax the average by solving (25). Secondly, we relax the n geodesic paths (while fixing the average) by solving (23) for $k = 1, \dots, K-1$. For the cascadic approach in time, we begin with $K = 1$ such that (25) is solved for $s^i(1/K) = s^i$. Then, at each refinement, we subdivide the geodesic paths such that $K \leftarrow 2K$. In detail, we set $s_{2k} = s_k$ for $k = K, \dots, 0$ and initialise the new intermediate shapes s_{2k+1} as the discrete geodesic average of s_{2k} and s_{2k+2} for $k < K$.

Timings. The components of our approach, projection onto the model or fitting the model to data could not be performed in real-time based on the current implementation. For proof of concept, the results in this paper were prepared using a prototype implementation in MATLAB. We make this implementation available as open source to aid reproducibility and to enable others to build and fit their own Shell PGA models (<https://github.com/cazhang/shellPGA>). In Table 1, timings of all experiments are shown to give some idea of computational cost using MATLAB. Model building is performed on a linux machine with 12 cores (Intel Xeon CPU E5-2680 2.4GHz) for parallel computing geodesic paths. All other results are computed on a single CPU (Intel Core i7-6700 3.4GHz). Timings of both offline model building (i.e. computing the discrete Fréchet mean and the principal variations) as well as the online model fitting or editing are shown. For model fitting as shown in Fig. 12 and Fig. 13, averages over all frames are reported. To give some idea of computational speed-up, we have re-computed some experiments in C++ (on a Dell Intel Core i7-2600 3.4GHz). For example, the results shown in Fig. 8 can be obtained in roughly 2 minutes offline and 30s online cost.

9. Conclusions

We have shown how to perform principal geodesic analysis in the space of discrete shells. In so doing, we derived an alternate formulation of PGA that avoids performing any operations in the tangent space and works directly with objects lying on the manifold. The whole approach is based on an elastic energy functional measuring membrane and bending distortion. The result is a physically-guided statistical shape model, that is able to generalise across datasets containing large nonlinear articulations and deformations. The central tool - the projection onto a submanifold of discrete shells - is well suited as the key ingredient in mesh editing or model fitting.

Once again, a metric derived from an elastic thin shell model pro-

	Dataset	n / J	Offline	Online
Fig. 5	SCAPE TOSCA_Cat	71 / - 10 / -	73 m 66 m	- -
Fig. 8	Dyna_50009	29 / 10	21 m	70 s
Fig. 9	SCAPE	70 / 20	72 m	232 s
Fig. 10	SCAPE	71 / 20	73 m	919 s
Fig. 11	TOSCA_Cat	10 / 5	66 m	154 s
Fig. 12	Dyna_50009	29 / 10	21 m	100 s
Fig. 13	Dyna_50021	20 / 10	99 m	321 s

Table 1: Timings obtained with our prototype MATLAB implementation for fixed $K = 4$, but different numbers of training shapes n and principal variations J .

vides a representation of volumetric objects and their deformations which retains physical plausibility. In particular, Fig. 13 shows that our results are comparable to MoSh [LMB14] which models bones and muscles explicitly. If the training data set contains large bending distortions at joint locations (see e.g. the armpits in Fig. 2b), this will be picked up by the first few principal variations since they account for a lot of the variance in the Gram matrix (see Fig. 1 and Fig. 5). For example, one can see in Fig. 10 that joints are fairly easy to bend while showing realistic muscle deformation.

In comparison to the original PGA model [FLPJ04], which deals with a low dimensional medial axis description, we consider high dimensional shape manifolds. Furthermore, we extend PGA to the time-discrete setting and introduce a rigid body motion invariant distance measure. This invariance is a substantial advantage over the Shell PCA model [ZHR15], which is based on vertex displacements and hence alignment-dependent. To this end, the Shell PCA model [ZHR15] only allows for small deformations, i.e. mesh editing and model fitting applications are out of reach of this purely elastic PCA approach.

There are many avenues for future work. It would be interesting to translate the concept of the Mahalanobis distance to our submanifold so that we have a notion of the likelihood of a reconstructed shape. Although we have used the space of discrete shells as our motivating example, our proposed time-discrete PGA may have other applications in machine learning with a modified energy functional \mathcal{W} approximating an alternate measure of squared distance with a potentially different invariance principle. In terms of efficiency, the model reduction technique proposed in [vRESH16] would be ideal for speeding up our method. Since our submanifold works with convex combinations of principal variation shapes, a subspace of deformations trained on samples from the submanifold would dramatically reduce the computational cost and probably allow for real-time performance.

Acknowledgments

The authors are grateful to Mirela Ben-Chen, Michael Black, Klaus Hildebrandt, Peter Schröder and Max Wardetzky for their detailed and valuable feedback at different stages of this project, as well as to Josua Sassen for proofreading and helping with coding accelerations. The data in Fig. 10 as well as all results shown for comparison were kindly provided by Jie Yang. B. Heeren and M. Rumpf

acknowledge support by the FWF in Austria under the grant S117 (NFN) and by the Hausdorff Center.

References

- [ACPH06] ALLEN B., CURLESS B., POPOVIĆ Z., HERTZMANN A.: Learning a correlated model of identity and pose-dependent body shape variation for real-time synthesis. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2006), pp. 147–156. 3
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: shape completion and animation of people. In *ACM Trans. Graph.* (2005), vol. 24, pp. 408–416. 3, 6
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 6
- [BEKB15] BOSCAINI D., EYNARD D., KOUROUNIS D., BRONSTEIN M. M.: Shape-from-operator: Recovering shapes from intrinsic operators. *Comput. Graph. Forum* 34, 2 (2015), 265–274. 3
- [CLZ13] CHEN Y., LIU Z., ZHANG Z.: Tensor-based human body modeling. In *Proc. CVPR* (2013), pp. 105–112. 3
- [CSBC*17] CORMAN E., SOLOMON J., BEN-CHEN M., GUIBAS L., OVSJANIKOV M.: Functional characterization of intrinsic and extrinsic geometry. *ACM Trans. Graph.* 36, 2 (2017), 14. 3
- [FB11] FRÖHLICH S., BOTSCH M.: Example-driven deformations based on discrete shells. In *Comput. Graph. Forum* (2011), vol. 30, pp. 2246–2257. 3, 8, 10
- [FB12] FREIFELD O., BLACK M.: Lie bodies: A manifold representation of 3d human shape. In *Proc. ECCV* (2012). 3, 7, 8
- [FLPJ04] FLETCHER P. T., LU C., PIZER S. M., JOSHI S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Med. Imaging* 23, 8 (2004), 995–1005. 3, 4, 11
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proc. SIGGRAPH* (1997), pp. 209–216. 10
- [GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2003), pp. 62–67. 2, 4
- [GLL*16] GAO L., LAI Y.-K., LIANG D., CHEN S.-Y., XIA S.: Efficient and flexible deformation representation for data-driven surface modeling. *ACM Trans. Graph.* 35, 5 (2016), 158. 3, 7, 8
- [HH96] HEAP T., HOGG D.: Extending the point distribution model using polar coordinates. *Image Vis. Comput.* 14, 8 (1996), 589–599. 3
- [HLRB12] HIRSHBERG D., LOPER M., RACHLIN E., BLACK M.: Coregistration: Simultaneous alignment and modeling of articulated 3d shape. In *Proc. ECCV* (2012), pp. 242–255. 3
- [HRS*14] HEEREN B., RUMPF M., SCHRÖDER P., WARDETZKY M., WIRTH B.: Exploring the geometry of the space of shells. In *Comput. Graph. Forum* (2014), vol. 33, pp. 247–256. 3, 4
- [HRWW12] HEEREN B., RUMPF M., WARDETZKY M., WIRTH B.: Time-discrete geodesics in the space of shells. In *Comput. Graph. Forum* (2012), vol. 31, pp. 1755–1764. 3, 4
- [HSS*09] HASLER N., STOLL C., SUNKEL M., ROSENHAHN B., SEIDEL H.-P.: A statistical model of human pose and body shape. In *Comput. Graph. Forum* (2009), vol. 28, pp. 337–346. 3
- [Ken84] KENDALL D. G.: Shape manifolds, Procrustean metrics, and complex projective spaces. *Bull. London Math. Soc.* 16, 2 (1984), 81–121. 3
- [KMP07] KILIAN M., MITRA N. J., POTTMANN H.: Geometric modeling in shape space. In *ACM Trans. Graph.* (2007), vol. 26, p. 64. 3
- [LCF00] LEWIS J. P., CORDNER M., FONG N.: Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *Proc. SIGGRAPH* (2000), pp. 165–172. 3

- [LMB14] LOPER M., MAHMOOD N., BLACK M. J.: MoSh: Motion and shape capture from sparse markers. *ACM Trans. Graph.* 33, 6 (2014), 220. 9, 11
- [LMR*15] LOPER M., MAHMOOD N., ROMERO J., PONS-MOLL G., BLACK M. J.: SMPL: A skinned multi-person linear model. *ACM Trans. Graph.* 34, 6 (2015), 248:1–248:16. 3
- [MG03] MOHR A., GLEICHER M.: *Deformation Sensitive Decimation*. Tech. rep., University of Wisconsin, 2003. 10
- [Pen06] PENNEC X.: Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *J. Math. Imaging Vis.* 25, 1 (2006), 127–154. 3
- [PMRMB15] PONS-MOLL G., ROMERO J., MAHMOOD N., BLACK M. J.: Dyna: A model of dynamic human shape in motion. *ACM Trans. Graph.* 34, 4 (2015), 120:1–120:14. 1, 3
- [ROA*13] RUSTAMOV R. M., OVSJANIKOV M., AZENCOT O., BENCHEN M., CHAZAL F., GUIBAS L.: Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph.* 32, 4 (2013), 72. 3
- [RW15] RUMPF M., WIRTH B.: Variational time discretization of geodesic calculus. *IMA J. Numer. Anal.* 35, 3 (2015), 1011–1046. 4, 5
- [SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Proc. Eurographics Symposium on Geometry Processing (2007)*, pp. 109–116. 2, 8
- [SC15] SEGUY V., CUTURI M.: Principal geodesic analysis for probability measures under the optimal transport metric. In *Advances in Neural Information Processing Systems (2015)*, pp. 3312–3320. 3
- [SP04] SUMNER R. W., POPOVIĆ J.: Deformation transfer for triangle meshes. *ACM Trans. Graph.* 23, 3 (2004), 399–405. 3
- [SSP07] SUMNER R. W., SCHMID J., PAULY M.: Embedded deformation for shape manipulation. *ACM Trans. Graph.* 26, 3 (2007), 80. 8
- [SZGP05] SUMNER R. W., ZWICKER M., GOTSMAN C., POPOVIĆ J.: Mesh-based inverse kinematics. *ACM Trans. Graph.* 24, 3 (2005), 488–495. 3, 8
- [TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. In *Proc. SIGGRAPH (1987)*, vol. 21, pp. 205–214. 2
- [TWC*09] TOURNIER M., WU X., COURTY N., ARNAUD E., REVERET L.: Motion compression using principal geodesics analysis. In *Comput. Graph. Forum (2009)*, vol. 28, pp. 355–364. 3
- [vRESH16] VON RADZIEWSKY P., EISEMANN E., SEIDEL H.-P., HILDEBRANDT K.: Optimized subspaces for deformation-based modeling and shape interpolation. *Computers & Graphics* 58 (2016), 128–138. 2, 11
- [vTAMZ18] VON TYCOWICZ C., AMBELLAN F., MUKHOPADHYAY A., ZACHOW S.: An efficient riemannian statistical shape model using differential coordinates. *Med. Image Anal.* 43 (2018). 3
- [vTSSH15] VON TYCOWICZ C., SCHULZ C., SEIDEL H.-P., HILDEBRANDT K.: Real-time nonlinear shape interpolation. *ACM Trans. Graph.* 34, 3 (2015), 34:1–34:10. 2
- [ZHR15] ZHANG C., HEEREN B., RUMPF M., SMITH W. A.: Shell PCA: statistical shape modelling in shell space. In *Proc. ICCV (2015)*, pp. 1671–1679. 2, 7, 8, 11