

BaRe-ESA: A Riemannian Framework for Unregistered Human Body Shapes

Emmanuel Hartman¹, Emery Pierson^{2,4}, Martin Bauer^{1,3}, Nicolas Charon³, Mohamed Daoudi^{4,5}
Florida State University, Tallahassee, Florida, USA ¹,

University of Vienna, Vienna, Austria²

University of Houston, Houston, Texas, USA³

Univ. Lille, CNRS, Centrale Lille, Institut Mines-Télécom, UMR 9189 CRISAL, Lille, F-59000, France⁴,
IMT Nord Europe, Institut Mines-Télécom, Univ. Lille, Centre for Digital Systems, Lille, F-59000, France⁵

Abstract

We present **Basis Restricted Elastic Shape Analysis** (*BaRe-ESA*), a novel Riemannian framework for human body scan representation, interpolation and extrapolation. *BaRe-ESA* operates directly on unregistered meshes, i.e., without the need to establish prior point to point correspondences or to assume a consistent mesh structure. Our method relies on a latent space representation, which is equipped with a Riemannian (non-Euclidean) metric associated to an invariant higher-order metric on the space of surfaces. Experimental results on the FAUST and DFAUST datasets show that *BaRe-ESA* brings significant improvements with respect to previous solutions in terms of shape registration, interpolation and extrapolation. The efficiency and strength of our model is further demonstrated in applications such as motion transfer and random generation of body shape and pose.

1. Introduction

Over the past decade there has been an increased interest in analyzing the morphological variability of the human anatomy. In particular, the issue of modeling and retrieving changes in full body shape and pose has a wide range of graphics applications spanning from 3D human modeling, to augmented and virtual reality for animated films and computer games. In these applications one is interested in a framework that allows for operations such as shape interpolation [16], – the task of finding a plausible deformation between two given body scans – shape extrapolation, – the task of finding a plausible deformation given a body scan and a corresponding initial deformation (movement) – and random shape generation. Furthermore, one is interested in

obtaining a natural disentanglement of changes in the pose and in the shape of the human body [15], which in turn will allow for operations such as motion transfer [5].

Motivation: A major challenge in the context of human body shape analysis is the registration problem, i.e. raw scans of human bodies are usually not equipped with any point correspondences and will, in general, not even admit a consistent mesh structure. Traditionally this issue has been tackled by finding new representations of the meshes with a consistent mesh structure and estimated point correspondences using an external framework such as functional maps [38, 35]. This data preprocessing step is then followed with an independent framework for the analysis of parametrized shapes, such as the As-Rigid-As-Possible (ARAP) energy [45] or the as isometric as possible energy [27]. In recent years, in the context of general functional data analysis, it has been shown that such a separation into registration and subsequent analysis can introduce a significant bias into the resulting statistical analysis, and it has been acknowledged that a significant increase in performance can be obtained by using a unified approach [46, 25].

Contribution: Towards this aim we introduce a new pipeline that quantifies geometric differences between unregistered human body scans, i.e., that does not require prior point correspondences or consistent parametrization across the dataset. Furthermore, we are not only interested in a pure metric comparison of two individuals, but also in estimating plausible deformation processes from one human body to the other. To that end, we introduce a transformation model that allows to disentangle changes in the pose and in the shape of the human body so as to obtain realistic ways to interpolate from one scan to another, extrapolate a given motion or transpose it to a new individual and even to generate visually plausible random pose and shapes; see Figure 1 for a schematic visualization of our framework. It is important to highlight that, through the use of the varifold representation and kernel metrics, our approach does not re-

This work was supported by ANR projects 16-IDEX-0004 (ULNE) and by ANR-19-CE23-0020 (Human4D); by NSF grants DMS-1912037, DMS-1953244, DMS-1945224 and DMS-1953267, and by FWF grant FWF-P 35813-N. Corresponding Author: M. Bauer (bauer@math.fsu.edu)

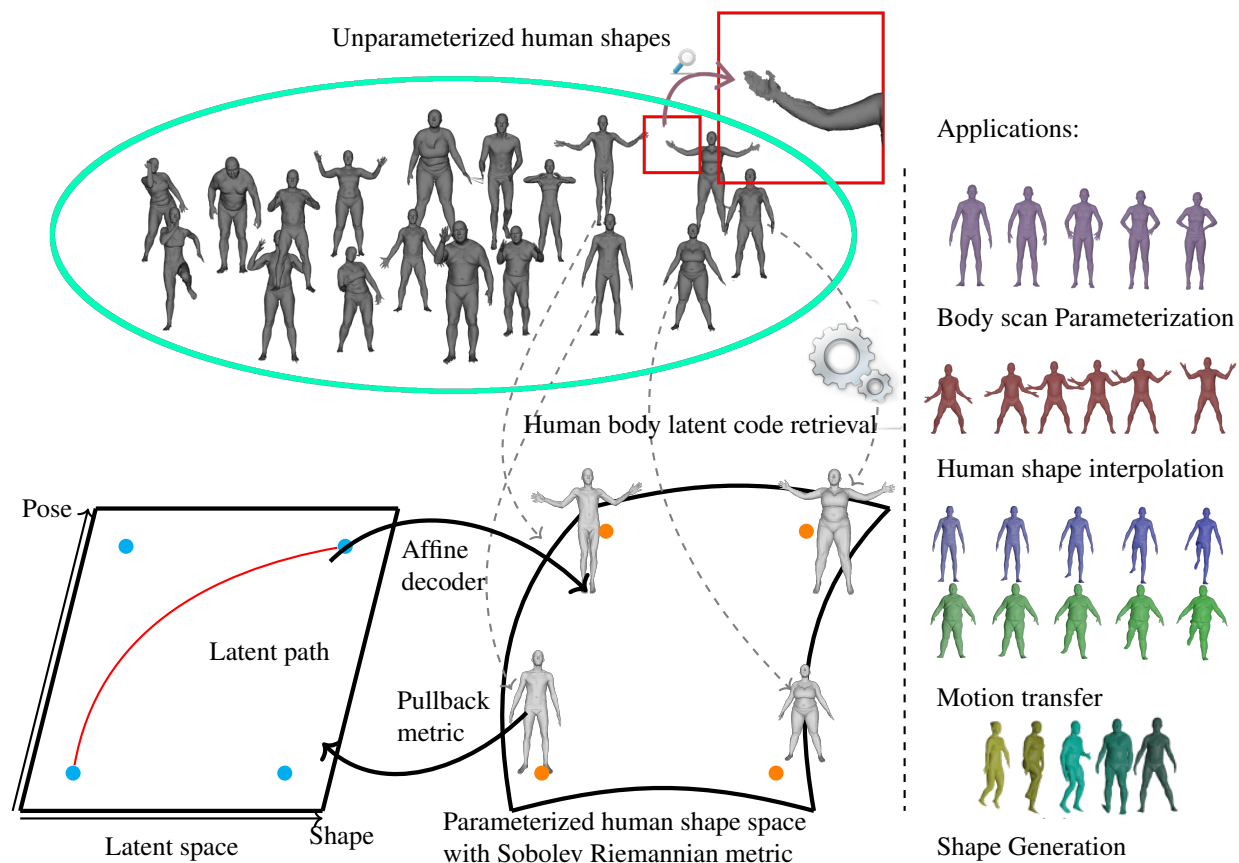


Figure 1: Overview of our method. We seek to represent unparameterized human bodies, with different mesh connectivity, and possible noise or topological changes in a disentangled latent space. We define our latent space as the sum of pose and shape spaces. The paths in the latent space are not linear but curved, corresponding to geodesics in the parameterized human body space. After retrieving the latent codes of the human bodies, we can use the space along with its Riemannian metric to solve several problems in human body deformation: inter-extrapolation, motion transfer, shape generation.

quire having a consistent mesh structure across the dataset and performs well on human body scans with different numbers of vertices and even under the presence of topological noise (e.g. holes in the scans). Furthermore, unlike current deep learning frameworks for human body analysis, the encoding component in our approach relies on pre-trained deformation bases for the shape and pose changes but is coupled with a non-Euclidean metric in the latent space. Thus the training of our model is notably simple and does not require large amount of training data. Moreover, as our results suggest, it leads to better properties when it comes to the interpretability of human body paths and the generalization to unseen data.

2. Related Work

Human body statistical spaces: Since Blanz *et al.*'s 3D Morphable Model for faces [8], statistical parametric mod-

els have been a preferred way for modeling 3D shape variations of human shapes (body, face and so on). In the particular domain of human bodies [1, 23, 41], the approaches consist of blending a skeleton transformation using several possible transformation models, along with a learned space of identity deformation. Recently, the SMPL model [34] is widely used for a variety of applications as it gives a valuable method of compressing the human body characteristics. The pipeline to retrieve the parameters of the SMPL model remains expensive and often requires manual intervention. As datasets of human body shapes grow larger, manual intervention becomes unfeasible and modern machine learning methods have been proposed to automate the processing of raw shape data.

Riemannian deformations of surfaces: In order to deform human shapes a common approach is to optimize a deformation relative to a given energy. Multiple deformation en-

ergies have been used for defining the space of plausible deformations for human (or humanoid) shapes. Some project the points into Riemannian representation space [33, 4], or define energies from Riemannian metrics, such as the As-Rigid-As-Possible (ARAP) [45], the as isometric as possible [27] or the framework of Pierson *et al.* [40], who minimize a first order Riemannian energy with respect to deformations learned from training data of human motions. All these methods are, however, parameterization dependent and thus point correspondences (and a consistent mesh structure) must be precomputed.

The field of *elastic shape analysis* (ESA) [46], which employs reparametrization invariant Riemannian metrics for defining deformation energies, has been proposed to overcome this drawback by solving for the registration (point correspondences) and deformation path in a unified fashion. This approach defines a shape as a point in the space of embeddings (or immersions) from a parameter space M into \mathbb{R}^3 , that can be quotiented by the action of reparameterizations of M and the group of Euclidean transformations. The induced manifold is equipped with the inherited Riemannian metric from the space of parametrized surfaces and is then used to measure the distance between two given shapes as well as interpolation by computing a geodesic that connects them [29, 25, 47, 50, 31]. These approaches, however, assume that the surfaces are given by analytic (spherical) representations. In the context of real data (raw body scans) this thus requires one to find such an analytic representation of the data in a pre-processing step [42, 30]. In the presence of imaging noise in particular, this is a highly non-trivial task, that can introduce significant bias and error in subsequent statistical analyses.

The varifold representation of surfaces [14, 26], can be used to define discrepancy loss functions between shapes that are not only invariant to parameterization, but also robust to scan inconsistencies such as potential noise or holes. They indeed have been applied successfully in the context of human shapes [39], although it was for sequence comparison and not deformation of surfaces. The approach of this paper follows some of the recent advances [6, 22] that combine Riemannian elastic metrics with varifolds discrepancy terms and thereby overcome the aforementioned difficulties.

Data-driven human body latent spaces: The deep learning era has introduced multiple generative models such as the popular Generative Adversarial Networks (GAN) [19] and Variational AutoEncoders (VAE) [28] that are able to model faithfully sample variations from a training set of images. Geometric deep learning [13], and in particular 3D autoencoders methods seek to extend those methods to 3D data. Those methods [12, 32, 24, 20] proved to be successful in order to find a low dimensional representation of the space compared to classical algorithms such as PCA.

In recent years, several autoencoder frameworks have been proposed to represent human bodies in a low-dimensional latent vector space while being independent of parameterization, thanks to different architecture such as PointNet [43, 3, 15, 36, 21] or more recently transformers, but their training cost remain high [49]. By using several deformation energy losses in the training set such as geodesic distances [15], or ARAP [24, 36], these methods train non-linear encoder and decoder networks to map the low dimensional latent vector space to the space of human bodies and vice versa. This allows one to map linear paths in the latent space to paths of plausible human body motions. The main drawback of these methods comes from their reliance on large datasets of parameterized surfaces and their failure to sufficiently learn the non-linear map from the flat space to the space of human bodies thereby lacking in generalizability when confronted to unseen data.

In contrast, our proposed pipeline, as illustrated by Figure 1, imposes an affine map, called the affine decoder, from a given low dimensional latent space to a corresponding space of human body transformations, which is based on the use of some pre-estimated bases to represent infinitesimal body shape and body pose deformations. Recent approaches has demonstrated that regularizing latent space [48, 18, 2] can improve the computations of paths for 3D shapes. To do so, unlike aforementioned deep learning architectures, we do not rely on the standard Euclidean metric in the latent space but instead on the pullback of a second order parametrization-invariant Sobolev metric in the space of deformation field. Thus, interpolating paths in latent space are no more straight lines, but are associated to (sub-Riemannian) geodesics in the surface space for this metric.

3. Riemannian Latent Space Method

The space of human shapes: In this article we model the space of all human shapes as a subset of the infinite dimensional space of template-based unparametrized surfaces, i.e., we view a human shape as an element of the quotient space $\mathcal{S} = \text{Imm} / \text{Diff}$, where $\text{Imm} = \{q \in C^\infty(\mathcal{T}, \mathbb{R}^3) : Tq \text{ is inj.}\}$ denotes the space of smooth maps (immersions) and where Diff is the group of all smooth, bijective maps (diffeomorphisms) on a human body template \mathcal{T} .

A Riemannian approach: To define our framework for the analysis of human shapes we will resort to the field of Riemannian geometry, more precisely we will view the space \mathcal{S} as an infinite dimensional manifold and equip it with a Riemannian metric. This will allow us to reduce tasks such as shape interpolation and extrapolation to geometric operations – the geodesic initial and boundary value problem. In order to define a Riemannian metric on \mathcal{S} we will follow the elastic shape analysis (ESA) paradigm of defining first a

reparametrization *invariant* Riemannian metric on the space of parametrized surfaces Imm , that will then descend to a Riemannian metric on the quotient space \mathcal{S} . Recall that a Riemannian metric G is a family of inner products

$$G_q : C^\infty(\mathcal{T}, \mathbb{R}^3) \times C^\infty(\mathcal{T}, \mathbb{R}^3) \rightarrow \mathbb{R}_+$$

that depends smoothly on the foot point $q \in \text{Imm}$. Here we have identified the tangent space $T_q \text{Imm}$ (the space of deformation vectors) with the space of smooth vector fields over \mathcal{T} .

Latent Space Model: For applications it is natural to restrict this model to only consider a certain set of admissible deformations applied to a template shape. In the context of shapes of human bodies, such a restriction ensures that the model only contains deformations that are "natural" i.e. deformations stay in the space of expected human body shapes. Therefore in our model, we restrict ourselves to surfaces q that can be written as linear combinations of a basis of admissible body pose deformations (shape basis), $\{h_i\}_{i=1}^m$, and a basis of admissible body type deformations (motion basis), $\{k_i\}_{i=1}^n$. Thus, all of the shapes we consider are determined by the latent space $\mathcal{L} \subset \mathbb{R}^{n+m}$ of coefficients for the combined basis. This latent space model is thus related to the space of immersions by the mapping $F : \mathcal{L} \rightarrow \text{Imm}$ via

$$F(\alpha^j) = \bar{q} + \sum_{i=1}^m \alpha_i^j h_i + \sum_{i=m+1}^{m+n} \alpha_i^j k_{i-m}.$$

where \bar{q} is our parameterized template human body shape, h_i is a basis of (realistic) body type (identity) deformations and where k_i is a basis of (realistic) body pose deformations (movements). This choice of bases is a crucial ingredient to obtain a good performance of the proposed method. In the present work we will construct them in a data driven way, which will be described in more details in Section 3.1. We then equip our latent space with a non-euclidean metric by defining the pullback of G via F . In particular, for $\alpha \in \mathcal{L}$ and $\beta, \eta \in \mathcal{L}$, we define the metric on the latent space by

$$\bar{G}_\alpha(\beta, \eta) := G_{F(\alpha)}(F(\beta) - \bar{q}, F(\eta) - \bar{q}).$$

Choice of Metric: There is a variety of different Riemannian metrics G satisfying the required invariance properties that have been proposed in the literature. In the current article we choose a metric from the family of second order Sobolev metrics with six parameters proposed in [22]. We will not present the exact formula of the Riemannian metric here, but refer the interested reader to supplementary material. For the purpose of the present article we only mention the following fundamental properties of our choice of Riemannian metric, which are central for the proposed applications: 1) the second order term in the metric pro-

vides enough regularity to prevent the occurrence of singularities along the solution of the interpolation and extrapolation problem described above, 2) the parameters in the metric provide enough flexibility to enforce certain behaviors of the resulting optimal deformations, eg. as isometric as possible deformations and 3) they naturally can be extended to the space of triangulated meshes using the methods of discrete differential geometry.

Interpolation as a geodesic boundary value problem: The shape interpolation problem between two surfaces (human shapes) is the task of finding an optimal deformation (path of immersions) between the two given surfaces. In our Riemannian setup this reduces to minimizing the energy

$$E(\alpha) = \int_0^1 \bar{G}_\alpha(\partial_t \alpha, \partial_t \alpha) dt$$

over all paths $\alpha : [0, 1] \rightarrow \mathcal{L}$ and diffeomorphisms $\varphi \in \text{Diff}$ such that $F(\alpha)(0) = q_0$ and $F(\alpha)(1) = q_1 \circ \varphi$, where q_0 and q_1 are arbitrarily chosen parametrizations of the unparametrized shapes. In order to apply this model to real data the main difficulty is the action of the diffeomorphism group on the endpoint constraint and the fact that raw body scans are, in general, not equipped with a consistent mesh structure, i.e., different scans can have different resolution, different mesh structures or even involve imaging errors resulting in e.g. holes or missing parts.

To circumvent these difficulties we will instead consider a relaxed formulation of the BVP given by the energy

$$\tilde{E}(\alpha) = \int_0^1 \bar{G}_\alpha(\partial_t \alpha, \partial_t \alpha) dt + \lambda \Gamma(F(\alpha)(0), q_0) + \lambda \Gamma(F(\alpha)(1), q_1), \quad (1)$$

where Γ is a reparametrization blind similarity measure, i.e., Γ satisfies the fundamental property $\Gamma(q_0, q_1) = 0$ if and only if q_0 and q_1 represent the same shape, i.e., there exists a $\varphi \in \text{Diff}$ such that $q_0 = q_1 \circ \varphi$. We have thus reduced the geodesic boundary value problem to an unconstrained minimization problem over all paths $\alpha : [0, 1] \rightarrow \mathcal{L}$. There are various different possibilities for this similarity term, such as the Hausdorff or Chamfer distance. We will rely on a similarity term derived from geometric measure theory, namely kernel metrics on the space of varifolds previously used in e.g. [26]. An important advantage of this framework is the fact that the resulting fidelity loss function remains differentiable with respect to the positions of the shapes' vertices. Additionally, after discretization, the varifold distance does not require the given meshes to have a consistent mesh structure which allows us to compare human body scans with different numbers of vertices and even different topologies, e.g. in the case of holes in the scans.

We then discretize the path α in time and compute the path energy using finite differences. Further, we may discretize the varifold data similarity term as in [26]. Details

of such a discretization are included in the supplementary materials. Thus, we have reduced the infinite dimensional minimization problem to a finite dimensional one, with the variables being the (discrete) path of latent variables. We tackle the resulting unconstrained non-convex minimization problem using the L-BFGS algorithm.

Extrapolation as a geodesic initial value problem: The shape extrapolation problem consists in predicting the future evolution of a surface (human body) given an initial deformation direction. In our Riemannian framework this reduces to solving the geodesic equation with given initial condition $q(0) = q_0$ (the initial pose) and $\partial_t q(0) = h$ (the direction of deformation). The geodesic equation is the first order optimality condition of the energy functional; it is non-linear PDE, that is second order in time t and forth order in space. For the exact formula of this equation, which is rather lengthy and not particularly insightful, we refer the interested reader to the literature, see eg. [7]. To solve such initial value problems in our latent space, we modify methods of discrete geodesic calculus [44] for our setting. We approximate the geodesic starting at α^0 in the direction of β with a PL path with $N + 1$ evenly spaced breakpoints. At the first step, we set $\alpha^1 = \alpha^0 + \frac{1}{N}\beta$ and find α^2 such that $F(\alpha^1)$ is the geodesic midpoint of $F(\alpha^0)$ and $F(\alpha^2)$, i.e., we solve for α^2 such that

$$\alpha^1 = \underset{\tilde{\alpha}}{\operatorname{argmin}} [\overline{G}_{\alpha^0}(\beta_0, \beta_0) + \overline{G}_{\tilde{\alpha}}(\tilde{\beta}, \tilde{\beta})]$$

where $\beta_0 = \tilde{\alpha} - \alpha^0$ and $\tilde{\beta} = \alpha^2 - \tilde{\alpha}$. Differentiating with respect to $\tilde{\alpha}$ and evaluating the resulting expression at α^1 , we obtain the system of equations

$$\begin{aligned} 2\overline{G}_{\alpha^0}(\beta_0, h_i) - 2\overline{G}_{\alpha^1}(\tilde{\beta}, h_i) + D_{\alpha^1}\overline{G}_{\cdot}(\tilde{\beta}, \tilde{\beta})_i &= 0, \\ 2\overline{G}_{\alpha^0}(\beta_0, k_i) - 2\overline{G}_{\alpha^1}(\tilde{\beta}, k_i) + D_{\alpha^1}\overline{G}_{\cdot}(\tilde{\beta}, \tilde{\beta})_{i+m} &= 0 \end{aligned} \quad (2)$$

where $\{h_i, k_i\}$ is our basis of deformations as introduced above. We denote the system of equations in (2) by $\Phi(\alpha^2; \alpha^1, \alpha^0) = 0$, where we stress again that α^0 and α^1 are here fixed and known. We solve this system of equations for α^2 using a nonlinear least squares approach,

$$\alpha^2 = \underset{\tilde{\alpha}}{\operatorname{argmin}} \|\Phi(\tilde{\alpha}; \alpha^1, \alpha^0)\|_2^2.$$

We repeat this process $N - 1$ times, thereby constructing the discrete solution up to time $t = 1$.

3.1. A data-driven basis of deformations

To construct the bases of movements and body type deformations we interpret registered mesh sequences of motions as paths in shape space whose tangent vectors are implicitly restricted to the space of valid motions of a human body. The tangent vectors of those sequences are used as training data, on which we perform principle component

analysis to obtain a tractable yet expressive basis for the valid pose deformations of a human body. We then collect meshes of the same pose from each identity (generally available in T or A-pose) and compute the (unrestricted) pairwise geodesics between these meshes with respect to our second-order Sobolev metric, where we use the Pytorch implementation of [22]. Note that these meshes show only moderate deformations and thus there are no difficulties with applying the unrestricted matching algorithm. We then collect the tangent vectors to these paths and perform again PCA to define our basis of shape deformations. For the analysis of pre-registered human body motions a similar procedure has been used in [40]. We discuss the limitations and benefits of this process in Section 8 of the supplementary material.

Parameter selection: The resulting basis size for our model is as follows (the number of elements was chosen experimentally): the motion basis has $n = 130$ elements, whereas the basis for the body type variation has only $m = 40$ elements. The coefficients for the H^2 -metric were chose to enforce close to isometric deformations that allow for some stretching and shearing to allow change in body type. Further, we added a small coefficient to the second-order term to further regularize the deformations. The final six parameters for the H^2 -metric are set to $(1, 1000, 100, 1, 1, 1)$. We perform sequential minimizations where the parameter σ of the varifold term is decreased from .4 to .025 and the balancing term λ is increased from 10^2 to 10^8 . A pseudo code of our main algorithm is available in the supplementary material.

4. Results

In this section we demonstrate the accuracy of our framework in five different experiments: the registration of human body scans, the interpolation and extrapolation of human body motions, random generation and motion transfer.

4.1. Datasets

Training dataset: To construct our basis we will make use of the publicly available Dynamic FAUST (D-FAUST) [11] dataset. This dataset contains high quality scans, along with registered meshes that will be used as training data. More specifically D-FAUST [11] is a 4D scan dataset of human motions. 10 individuals performed at most 14 in-place motions, recorded using multi view cameras and high quality 4D scanner, sampled at 60 Hz. Due to the high speed of the recording, D-FAUST scans contain several singularities in the reconstructed surface, such as holes or even artificial objects (eg. parts of walls). The SMPL mesh is registered using a registration pipeline based on image texture information, and a novel body motion model, allowing to have the registrations corresponding to each scan. A set of 7 long range sequence, divided in 10 representative mini-

sequences, are left for testing (see supp. material for the video of these sequences).

Test dataset: In addition to these 7 sequences from D-FAUST we tested our algorithm on the static FAUST dataset [10]. This is a 3D static scan dataset designed for human mesh registration, that contains scans of minimally clothed humans, similar to the ones in the D-FAUST dataset. Ten individuals (different from the D-FAUST training set) performed 30 different poses recorded using a high quality 3D scanner. We selected 9 unregistered poses of the training set that show no rotations, and use them as a testing set. We chose this dataset rather than other possible ones [1, 35] because the geometric closeness of minimally clothed humans allows us to evaluate the geometric quality of registrations (as opposed to the case of clothed humans) and it allows to test robustness against topological noise.

4.2. Evaluation and comparison

Before presenting our results, we discuss our procedure for evaluating and comparing different methods. We perform a thorough comparison of our approach with other state-of-the-art approaches for human body analysis that rely on latent space learning for registration, interpolation, and extrapolation tasks. To keep the section more focused, we deliberately restrict to those, and do not consider other methods that can potentially tackle the same tasks but without a low dimensional latent space [16], or that are specifically designed for other tasks [36]. We compare our approach to LIMP [15], which models shape deformations using a variational auto-encoder with geodesic constraints; ARAPReg [24], which models deformations using an auto-decoder with regularization through the as rigid as possible energy; and 3D-Coded [20], which is similar to LIMP but with lighter training and without geometric loss regularization. LIMP and 3D-Coded both utilize a PointNet architecture as an encoder, which enables invariance to parameterization. On the other hand, ARAPReg recovers latent vectors within a registered setting utilizing the L^2 metric, which assumes that the target meshes possess identical mesh structure as the model’s output. To make this framework viable for our application we replace the L^2 -metric by the varifold distance thereby extending ARAPReg to unregistered point clouds. We trained all three networks on the D-FAUST dataset using reported training details from the respective papers. We used 3D-Coded only for comparison in the latent code retrieval experiments as its interpolation and extrapolation results were notably bad.

In our experiments we evaluate the quality of the results using different similarity measures (distances) between the outputs of the different methods and the original scan. The “shape” matching is evaluated by comparing each method against the original scans using three different remeshing invariant similarity measures. First, we evaluate the methods

	MSE	Hausdorff	Chamfer	Varifold
LIMP	NA	0.23	0.098	0.073
ARAPReg	0.035	0.11	0.117	0.021
3D-Coded	0.053	0.07	0.020	0.023
BaRe-ESA	0.015	0.08	0.019	0.014

Table 1: Shape registration and reconstruction results. Each method is trained on D-FAUST and tested on Faust. Where applicable, we compute the mean squared error (MSE) between each method’s outputs and the ground truth registration of FAUST. We compute the reconstruction errors between the outputs of the methods and the original scans.

using the varifold metric introduced above. As our method minimizes this distance during the registration process, we include two additional metrics to avoid bias: the widely used Hausdorff distance, which provides a good insight for the quality of a mesh reconstruction, but can be sensitive to single outliers present in low quality scans and the Chamfer distance [17, 21], which is more robust to such outliers.

In our first experiment – latent code retrieval – we will in addition evaluate the quality of the obtained point correspondences – in this section we use data with given ground truth point correspondences. Therefore we will compute the mean squared error of the each method to the ground truth registrations of the testing set. Unfortunately, one method (LIMP) does not return the same mesh structure as the ground truth registrations and thus we could not compare it this way. For a detailed description of all these evaluation metrics we refer to the supplementary material.

4.3. Mesh Invariant Latent Code Retrieval

Given a scan, q_1 with arbitrary mesh structure, we retrieve the latent code of the shape class of q_1 by performing a relaxed geodesic matching problem between \bar{q} and q_1 with T time steps. This produces a geodesic path from the template to the shape class of the target mesh, thus the endpoint α^T is the latent code of the shape class of q_1 .

To demonstrate the effectiveness of our mesh latent code retrieval algorithm we test on the 90 meshes from the unregistered FAUST dataset (recall that this dataset was not contained in the training process). In this experiment, we construct latent code representations with BaRe-ESA, LIMP, 3D-Coded, and ARAPReg and measure the distance from the reconstructed meshes to the original scans using the evaluation methods outlined in Section 4.2. In Figure 2 we present a qualitative comparison of the obtained results. A quantitative comparison of the performance of the different methods is presented in Table 1 where we demonstrate that BaRe-ESA significantly outperforms the mesh autoencoder methods with respect to the registration and reconstruction evaluation metrics. In the supplementary material we also discuss the computational cost of our method.

	Interpolation									Extrapolation								
	Hausdorff			Chamfer			Varifold			Hausdorff			Chamfer			Varifold		
	LIMP	ARAPReg	BaRe-ESA	LIMP	ARAPReg	BaRe-ESA	LIMP	ARAPReg	BaRe-ESA	LIMP	ARAPReg	BaRe-ESA	LIMP	ARAPReg	BaRe-ESA	LIMP	ARAPReg	BaRe-ESA
punching	0.13	0.12	0.081	0.029	0.020	0.04	0.060	0.034	0.022	0.140	0.368	0.102	0.030	0.112	0.085	0.066	0.097	0.025
running on spot	0.28	0.14	0.076	0.112	0.080	0.069	0.072	0.052	0.025	0.287	0.309	0.152	0.11	0.177	0.125	0.071	0.079	0.027
running on spot b	0.20	0.13	0.082	0.125	0.101	0.045	0.068	0.040	0.025	0.264	0.415	0.222	0.138	0.179	0.116	0.071	0.083	0.027
shake arms	0.34	0.33	0.078	0.063	0.049	0.076	0.061	0.031	0.025	0.410	0.832	0.273	0.080	0.283	0.171	0.066	0.083	0.027
chicken wings	0.18	0.15	0.093	0.101	0.083	0.04	0.062	0.029	0.016	0.182	0.395	0.092	0.11	0.189	0.092	0.072	0.081	0.018
knees	0.060	0.35	0.051	0.182	0.266	0.036	0.097	0.066	0.016	1.13	0.282	0.104	0.23	0.205	0.081	0.321	0.072	0.016
knees b	0.52	0.054	0.084	0.168	0.107	0.029	0.067	0.019	0.016	0.489	0.319	0.244	0.17	0.215	0.065	0.066	0.071	0.017
jumping jacks	0.12	0.32	0.046	0.074	0.054	0.014	0.070	0.044	0.015	0.195	0.645	0.091	0.076	0.302	0.072	0.068	0.086	0.027
jumping jacks b	0.40	0.11	0.062	0.111	0.086	0.009	0.076	0.030	0.025	0.492	0.570	0.122	0.12	0.324	0.061	0.10	0.083	0.029
one leg jump	0.16	0.062	0.07	0.138	0.109	0.067	0.068	0.025	0.024	0.175	0.363	0.167	0.143	0.249	0.136	0.069	0.078	0.025
mean	0.30	0.19	0.072	0.106	0.097	0.042	0.070	0.039	0.020	0.391	0.452	0.157	0.118	0.214	0.100	0.103	0.082	0.023

Table 2: Full interpolation and extrapolation comparison on 10 D-FAUST sequences. The Hausdorff, Chamfer and varifold distance are computed against ground truth sequences.

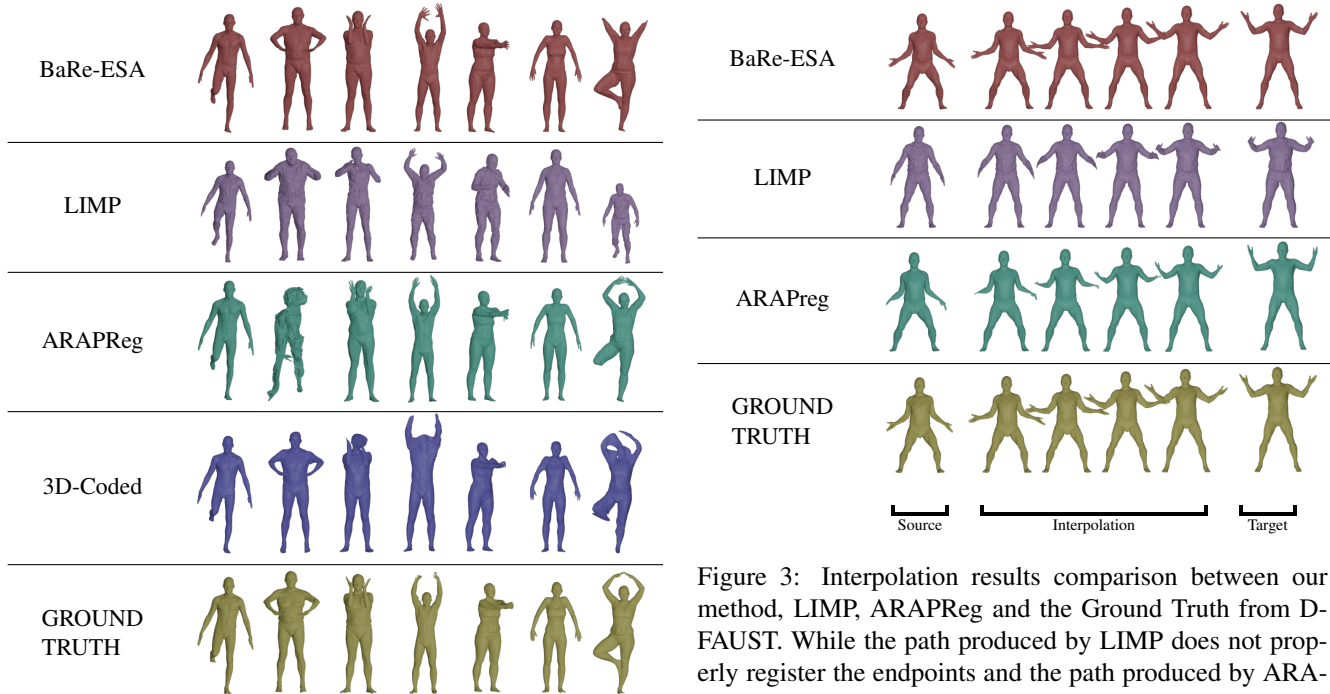


Figure 2: Registration of five elements of FAUST using three methods trained on D-FAUST. Both LIMP and ARAPReg fail to generalize to the new data and do not represent the ground truth in certain examples. BaRe-ESA consistently produces a decent representation in all examples, with a failure case displayed in the last column.

4.4. Interpolation

Next we turn our attention to the interpolation problem, i.e., the task of constructing a deformation between two different human body poses, that follows a “realistic” motion pattern. In our Riemannian setup, as described previously, this corresponds to solving the geodesic boundary value problem, i.e., we need to minimize the discrete energy given in (1). We use the start and end point of our 10 test mini-sequences as the input for our interpolation problem. This allows us to compare the obtained results to the full

Figure 3: Interpolation results comparison between our method, LIMP, ARAPReg and the Ground Truth from D-FAUST. While the path produced by LIMP does not properly register the endpoints and the path produced by ARAPReg does not stay in the space of human bodies, BaRe-ESA successfully produces a path of human shapes whose endpoints match the source and target shapes.

mini-sequences, seen as a ground truth motion (see the supplementary material for their corresponding animations). In Figure 3, we compare the results of our method, ARAPReg, and LIMP with the ground truth for one mini-sequence. Our method is successful at recovering the latent codes that represent the endpoints and producing interpolations that remain in the space of human shapes. We further perform a quantitative comparison of the methods by measuring the distance to the ground-truth sequences at each break point with respect to the evaluation metrics given in Section 4.2; these results are displayed in Table 2. One can clearly observe that our method again outperforms the other methods both qualitatively and quantitatively.




Figure 4: Extrapolation results comparison between our method, LIMP, ARAPreg and D-FAUST Ground Truth. While all methods capture the primary motion of lifting a leg, the extrapolations of LIMP and ARAPreg include extraneous motions of arms and slight changes in body type.

4.5. Extrapolation

In the following we consider the shape extrapolation problem, i.e., the task of predicting the future movement given a body shape and an initial movement (deformation). In our Riemannian setting this corresponds to solving the geodesic initial value problem using the method outlined in Section 3. From each of the 10 mini-sequences in our testing set, we recover the latent codes of the first two meshes in the sequence and then use the first latent code and the difference of the codes as input to our method. In Figure 4, we display the results of our extrapolation method, the extrapolations computed using LIMP and ARAPreg, and the original sequence (see the supplementary material for their corresponding animations). Our method is successful at producing extrapolations that capture the correct motion of the mesh without any extraneous motions that stay in space of human bodies. As with the interpolation comparison, we measure the distance to the ground-truth sequences at each break point and display the results of the quantitative comparison in Table 2. Similar as in the previous experiments, our method significantly outperforms the other methods.

4.6. Pose and Shape Disentanglement

By the construction of our basis, the first 130 dimensions of our latent space represent the body pose and the remaining 40 dimensions represent the body shape. Thus, our latent codes can be decomposed into coefficients that repre-




Figure 5: Motion Transfer: We display the original motion in the top row and the transfer of the motion to the target shapes in the second and third row.




Figure 6: 13 random shapes generated using a Gaussian mixture model on the space of initial velocities.

sent a change in body type and coefficients that represent a change in body pose. As a result, our framework is able to perform instant *motion transfer*: once a motion is represented as a sequence of latent codes we simply replace the shape coefficients of each element of the sequence with the shape coefficients of the target shape. An example of this method in action is displayed in Figure 5.

4.7. Random Shape Generation

Another possible application of our framework is random shape generation. The idea is to use a data driven distribution on the human shape tangent space. Therefore we first perform latent code retrieval on a subset of D-FAUST. We then compute, the initial tangent vector of each of these paths in the latent space, separated in pose and shape components. For each of these collections of tangent vectors, we fit a Gaussian mixture model, which is popular to generate human shapes [9, 37]. We used 10 and 6 components respectively, which proved to be sufficient to get visually satisfying random shapes. The generation process consists in sampling a pose and shape vector in the tangent space and to solve the corresponding geodesic initial value problem from the template in the direction of the generated vector. We display a selection of 13 generated shapes in Figure 6.

5. Conclusions, Limitations and Future work

In this paper, we proposed a novel framework for 3D human shape interpolation, extrapolation and generation, that relies on learning deformation bases for body type and body

pose changes combined with the use of a particular Riemannian structure on the latent space. Importantly, our method does not require surfaces to have consistent meshes and vertex correspondences and performs well even in the presence of imaging noise such as meshes with holes and/or missing parts. We showcased the different advantages of the proposed framework, in particular over recent deep learning approaches, for recovering and generating meaningful body shape trajectories. We point out, however, that this comes at the price of solving optimization problems to estimate interpolated or extrapolated geodesic paths for our metric. One possible way around this would be to train neural networks in a supervised setting to approximate solutions of those problems, which we leave as a subject of future investigation. Another potential limitation of the proposed framework is the need for sufficiently rich training data: in the current implementation we used the D-FAUST dataset for the generation of our motion and pose bases. In this dataset there is, however, only very limited movement of the fingers/hands present and consequently our motion basis is not able to faithfully represent such a movement. In future work we plan to use additional datasets for the generation of the bases, which in turn will allow us to get a more complete representation of all possible movements.

Finally we mention a simple and yet potentially relevant extension of our model, namely to introduce two distinct Sobolev Riemannian metrics on the shape change and the motion change deformation fields. This comes with the idea of adapting the metric to the different nature of those deformations, and thus even better disentangling these quantities.

References

- [1] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM SIGGRAPH 2005 Papers*, pages 408–416. 2005. [2](#), [6](#)
- [2] Matan Atzmon, David Novotny, Andrea Vedaldi, and Yaron Lipman. Augmenting implicit neural shape representations with explicit deformation fields. *arXiv preprint arXiv:2108.08931*, 2021. [3](#)
- [3] Tristan Aumentado-Armstrong, Stavros Tsogkas, Allan Jepson, and Sven Dickinson. Geometric disentanglement for generative latent shape models. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8180–8189, 2019. [3](#)
- [4] Sumukh Bansal and Aditya Tatu. Lie bodies based 3d shape morphing and interpolation. In *Proceedings of the 15th ACM SIGGRAPH European Conference on Visual Media Production*, pages 1–10, 2018. [3](#)
- [5] Jean Basset, Adnane Boukhayma, Stefanie Wuhrer, Franck Multon, and Edmond Boyer. Neural human deformation transfer. In *International Conference on 3D Vision, 3DV 2021, London, United Kingdom, December 1-3, 2021*, pages 545–554. IEEE, 2021. [1](#)
- [6] Martin Bauer, Nicolas Charon, Philipp Harms, and Hsi-Wei Hsieh. A numerical framework for elastic surface matching, comparison, and interpolation. *Int. J. Comput. Vis.*, 129(8):2425–2444, 2021. [3](#)
- [7] Martin Bauer, Philipp Harms, and Peter W Michor. Sobolev metrics on shape space of surfaces. *Journal of Geometric Mechanics*, 3(4):389, 2011. [5](#)
- [8] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 187–194, 1999. [2](#)
- [9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European conference on computer vision*, pages 561–578. Springer, 2016. [8](#)
- [10] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Piscataway, NJ, USA, June 2014. IEEE. [6](#)
- [11] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017. [5](#)
- [12] Giorgos Bouritsas, Sergiy Bokhnyak, Stylianos Ploumpis, Michael Bronstein, and Stefanos Zafeiriou. Neural 3d morphable models: Spiral convolutional networks for 3d shape representation learning and generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7213–7222, 2019. [3](#)
- [13] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017. [3](#)
- [14] Nicolas Charon and Alain Trouvé. The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM journal on Imaging Sciences*, 6(4):2547–2580, 2013. [3](#)
- [15] Luca Cosmo, Antonio Norelli, Oshri Halimi, Ron Kimmel, and Emanuele Rodolà. Limp: Learning latent shape representations with metric preservation priors. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 19–35. Springer, 2020. [1](#), [3](#), [6](#)
- [16] M. Eisenberger, D. Novotny, G. Kerchenbaum, P. Labatut, N. Neverova, D. Cremers, and A. Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [6](#)
- [17] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. [6](#)
- [18] Oren Freifeld and Michael J Black. Lie bodies: A manifold representation of 3d human shape. *ECCV (1)*, 7572:1–14, 2012. [3](#)

- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 3
- [20] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. 3d-coded : 3d correspondences by deep deformation. In *ECCV*, 2018. 3, 6
- [21] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3d-coded: 3d correspondences by deep deformation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 230–246, 2018. 3, 6
- [22] Emmanuel Hartman, Yashil Sukurdeep, Eric Klassen, Nicolas Charon, and Martin Bauer. Elastic shape analysis of surfaces with second-order sobolev metrics: a comprehensive numerical framework. *arXiv preprint arXiv:2204.04238*, 2022. 3, 4, 5
- [23] Nils Hasler, Carsten Stoll, Martin Sunkel, Bodo Rosenhahn, and H-P Seidel. A statistical model of human pose and body shape. In *Computer graphics forum*, volume 28, pages 337–346. Wiley Online Library, 2009. 2
- [24] Qixing Huang, Xiangru Huang, Bo Sun, Zaiwei Zhang, Junfeng Jiang, and Chandrajit Bajaj. Arapreg: An as-rigid-as possible regularization loss for learning deformable shape generators. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5815–5825, 2021. 3, 6
- [25] Ian H Jermyn, Sebastian Kurtek, Hamid Laga, and Anuj Srivastava. Elastic shape analysis of three-dimensional objects. *Synthesis Lectures on Computer Vision*, 12(1):1–185, 2017. 1, 3
- [26] Irène Kaltenmark, Benjamin Charlier, and Nicolas Charon. A general framework for curve and surface comparison and registration with oriented varifolds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3346–3355, 2017. 3, 4
- [27] Martin Kilian, Niloy J Mitra, and Helmut Pottmann. Geometric modeling in shape space. In *ACM SIGGRAPH 2007 papers*, pages 64–es. 2007. 1, 3
- [28] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. 3
- [29] Sebastian Kurtek, Eric Klassen, John C. Gore, Zhaohua Ding, and Anuj Srivastava. Elastic geodesic paths in shape space of parameterized surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(9):1717–1730, 2012. 3
- [30] Sebastian Kurtek, Anuj Srivastava, Eric Klassen, and Hamid Laga. Landmark-guided elastic shape analysis of spherically-parameterized surfaces. In *Computer graphics forum*, volume 32, pages 429–438. Wiley Online Library, 2013. 3
- [31] Hamid Laga, Marcel Padilla, Ian H Jermyn, Sebastian Kurtek, Mohammed Bennamoun, and Anuj Srivastava. 4d atlas: Statistical analysis of the spatio-temporal variability in longitudinal 3d shape data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [32] Clément Lemeunier, Florence Denis, Guillaume Lavoué, and Florent Dupont. Representation learning of 3d meshes using an autoencoder in the spectral domain. *Computers & Graphics*, 107:131–143, 2022. 3
- [33] Yaron Lipman, Olga Sorkine, David Levin, and Daniel Cohen-Or. Linear rotation-invariant coordinates for meshes. *ACM Transactions on Graphics (ToG)*, 24(3):479–487, 2005. 3
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 2
- [35] R. Marin, S. Melzi, E. Rodolà, and U. Castellani. Farm: Functional automatic registration method for 3d human bodies. *Computer Graphics Forum*, 39(1):160–173, 2020. 1, 6
- [36] Sanjeev Muralikrishnan, Siddhartha Chaudhuri, Noam Aigerman, Vladimir G. Kim, Matthew Fisher, and Niloy J. Mitra. GLASS: geometric latent augmentation for shape spaces. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2022. 3, 6
- [37] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *2018 international conference on 3D vision (3DV)*, pages 484–494. IEEE, 2018. 8
- [38] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):1–11, 2012. 1
- [39] Emery Pierson, Mohamed Daoudi, and Sylvain Arguillere. 3d shape sequence of human comparison and classification using current and varifolds. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 523–539, Cham, 2022. Springer Nature Switzerland. 3
- [40] Emery Pierson, Mohamed Daoudi, and Alice-Barbara Tumpach. A riemannian framework for analysis of human body surface. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2991–3000, January 2022. 3, 5
- [41] Leonid Pishchulin, Stefanie Wuhrer, Thomas Helten, Christian Theobalt, and Bernt Schiele. Building statistical shape spaces for 3d human modeling. *Pattern Recognition*, 67:276–286, 2017. 2
- [42] Emil Praun and Hugues Hoppe. Spherical parametrization and remeshing. *ACM transactions on graphics (TOG)*, 22(3):340–349, 2003. 3
- [43] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 77–85, 2017. 3

- [44] Martin Rumpf and Benedikt Wirth. Discrete geodesic calculus in shape space and applications in the space of viscous fluidic objects. *SIAM Journal on Imaging Sciences*, 6(4):2581–2602, 2013. [5](#)
- [45] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pages 109–116, 2007. [1](#), [3](#)
- [46] Anuj Srivastava and Eric P Klassen. *Functional and shape data analysis*, volume 1. Springer, 2016. [1](#), [3](#)
- [47] Zhe Su, Martin Bauer, Eric Klassen, and Kyle Gallivan. Simplifying transformations for a family of elastic metrics on the space of surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 848–849, 2020. [3](#)
- [48] Garvita Tiwari, Dimitrije Antić, Jan Eric Lenssen, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Pose-ndf: Modeling human pose manifolds with neural distance fields. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 572–589, Cham, 2022. Springer Nature Switzerland. [3](#)
- [49] Giovanni Trappolini, Luca Cosmo, Luca Moschella, Riccardo Marin, Simone Melzi, and Emanuele Rodolà. Shape registration in the time of transformers. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 5731–5744, 2021. [3](#)
- [50] Alice Barbara Tumpach, Hassen Drira, Mohamed Daoudi, and Anuj Srivastava. Gauge Invariant Framework for Shape Analysis of Surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):46–59, Jan. 2016. arXiv: 1506.03065. [3](#)