

Dynamic Neural Surfaces for Elastic 4D Shape Representation and Analysis

Awais Nizamani^{1*} Hamid Laga¹ Guanjin Wang¹ Farid Boussaid²
Mohammed Bennamoun² Anuj Srivastava³

¹Murdoch University ²The University of Western Australia ³Florida State University

Abstract

We propose a novel framework for the statistical analysis of genus-zero 4D surfaces, i.e., 3D surfaces that deform and evolve over time. This problem is particularly challenging due to the arbitrary parameterizations of these surfaces and their varying deformation speeds, necessitating effective spatiotemporal registration. Traditionally, 4D surfaces are discretized, in space and time, before computing their spatiotemporal registrations, geodesics, and statistics. However, this approach may result in suboptimal solutions and, as we demonstrate in this paper, is not necessary. In contrast, we treat 4D surfaces as continuous functions in both space and time. We introduce Dynamic Spherical Neural Surfaces (D-SNS), an efficient smooth and continuous spatiotemporal representation for genus-0 4D surfaces. We then demonstrate how to perform core 4D shape analysis tasks such as spatiotemporal registration, geodesics computation, and mean 4D shape estimation, directly on these continuous representations without upfront discretization and meshing. By integrating neural representations with classical Riemannian geometry and statistical shape analysis techniques, we provide the building blocks for enabling full functional shape analysis. We demonstrate the efficiency of the framework on 4D human and face datasets. The source code and additional results are available at <https://4d-dsns.github.io/DSNS/>.

1. Introduction

Statistical 3D shape analysis focuses on quantifying shape similarities and differences between 3D objects. It also aims at modelling, using probability distributions, the shape variability within and across object classes. 4D shape statistics, on the other hand, adds the temporal dimension. It aims to discover typical shape deformation patterns in a class of objects, and statistically model the variability of these deformations within and across classes. This is a very challenging problem since 3D shapes come with arbitrary discretization and thus need to be spatially registered onto each other. 4D shapes add another level of complexity since ob-

jects deform and grow at different rates. Thus, they need to be temporally aligned onto each other. Traditionally, 3D and 4D shapes are represented as discrete point clouds or triangular meshes. The spatiotemporal registration problem is then reduced to that of matching landmarks across shapes. Shape statistics such as means and modes of variation are then computed directly on these registered landmarks. This, however, leads to solutions that depend on the quality and resolution of the discretization.

In this paper, we treat 4D shapes as continuous functions in both space and time and develop a novel statistical analysis framework that operates directly on these continuous representations, without upfront discretization. This includes performing spatiotemporal registration, computing geodesics, and estimating the mean 4D shape of a set of 4D surfaces. We focus on genus-0 surfaces, which are abundant in nature, e.g., human bodies and body parts. We propose a novel surface-based neural representation, hereinafter referred to as Dynamic Spherical Neural Surface (D-SNS), that treats a 4D surface as a continuous mapping from a spherical domain and time to 3D. This mapping is then parameterized using Multi-Layer Perceptrons (MLPs). This will allow us to formulate the spatial registration problem as that of finding the optimal spatial reparameterizations of the neural functions. Similarly, temporal registration can be formulated as the problem of optimal temporal reparameterization of these neural functions. This can be efficiently implemented using an MLP and optimized using an elastic Riemannian metric that quantifies physical deformations of surfaces, i.e., bending and stretching. Additionally, we propose a framework that co-registers a set of 4D surfaces, represented with their D-SNSs, and simultaneously computes their 4D mean as a D-SNS. The main contributions of this paper can be summarized as follows;

- We introduce a new continuous representation for genus-0 4D surfaces, parameterized using Multi-Layer Perceptrons (MLPs). This framework treats both 3D and 4D shapes as functions, enabling functional shape analysis. We also demonstrate how this representation can be efficiently learned from discrete input 4D surfaces, with performance validated on datasets such as 4D facial

shapes [9, 29] and 4D human bodies [5, 20].

- We reformulate the spatiotemporal registration problem as one of optimally reparameterizing, in space and time, the neural representation. We measure the optimality using a physically-motivated Riemannian elastic metric.
- We show that by mapping our representation to the space of Square Root Normal Fields (SRNF) [13, 16] and Square Root Velocity Fields (SRVF) [31, 32], the complex Riemannian elastic metric simplifies to an \mathbb{L}^2 metric, which facilitates various downstream applications.
- We develop a comprehensive framework for performing spatiotemporal registration and computing geodesics and 4D means directly from the neural surface representation without discretization.

The remainder of the paper is organized as follows; Section 2 discusses the related work. Section 3 describes the proposed approach. Section 4 presents the results of the proposed approach, analyzes its performance, and compares it to the state-of-the-art. Section 5 concludes the paper.

2. Related Works

2.1. Statistical 3D and 4D shape analysis

Early statistical 3D shape analysis methods, based on morphable models [2, 4, 8], treat 3D shapes as a set of discrete landmarks in correspondence and use the \mathbb{L}^2 metric for computing geodesics and Principal Component Analysis (PCA) for their statistical analysis. These methods, however, assume that the 3D shapes are elements of an Euclidean shape space. As a result, they cannot handle 3D shapes that undergo large elastic deformations such as bending and stretching. SMPL model and its variants [18, 19, 27, 38] address this issue by explicitly modelling articulated motion using skeleton joints. SMPL-based methods, however, are class-specific and thus, they do not allow cross-category analysis.

A common property of these methods is that they assume that correspondences are given, *i.e.*, either they are manually specified or computed in a processing step using other techniques, *e.g.*, [6, 24, 25]. Jermyn *et al.* [13] and later Laga *et al.* [16], treat 3D shapes as elements of a Riemannian shape space equipped with an elastic metric that measures bending and stretching. This way, registration, geodesics, and summary statistics computation can be jointly formulated as an optimization problem under the same elastic metric. More importantly, they showed that by further mapping the shapes to the space of Square Root Normal Fields (SRNFs), the complex elastic metric becomes an \mathbb{L}^2 metric, thus, significantly simplifying downstream applications [15, 16]. The SRNF representation has been used for the analysis of genus-0 surfaces that undergo nonrigid motion. They have also been extended to the analysis of 4D surfaces [17] by treating them as time-parameterized trajec-

tories in the SRNF space, which has an \mathbb{L}^2 structure. Other methods represent 4D surfaces as the flow of deformation of the 3D volume and compute geodesics on a Riemannian manifold [3, 7, 11]. Unfortunately, computing deformations on 3D volumes is computationally expensive.

All these methods treat 4D shapes as discrete, in space and time, signals. In this paper, we consider, for the first time, 4D shapes as continuous functions and derive a comprehensive framework that operates directly on these functions. Discretization is only required for visualization.

2.2. Neural representations

Implicit neural representations such as DeepSDF [26], Neural Radiance Fields (NeRF) [21], and Neural Implicit Surfaces (NeuS) [33], leverage neural networks to represent the geometry and appearance of 3D shapes as continuous functions. They have been extensively used for 3D reconstruction but also have the potential to be used for 3D shape analysis tasks. In particular, some papers used neural networks to overfit individual SDFs [10]. This leads to a continuous and differentiable representation of individual shapes. While SDFs represent 3D shapes of arbitrary topologies, they are volumetric representations and thus expensive to evaluate since, in general, the surface of interest occupies just a tiny proportion of the 3D space. As such, several papers have explored neural representations of explicit surfaces. In particular, AtlasNet [12] models surfaces as a collection of parametric patches. Each patch is treated as a continuous mapping from a 2D domain to 3D and thus can be parameterized using an MLP. Morreale *et al.* [22] introduced Neural Surface Maps, which represent a surface as a mapping of a unit disk to 3D and then overfits to it an MLP. This representation, which enables surface-to-surface mapping, has been extended to computing semantic maps between genus-0 surfaces [23] and for 3D reconstruction [36]. Williamson *et al.* [34] showed how to compute differential properties of surfaces and perform geometry processing tasks on spherical neural representations without upfront discretization.

Our approach follows [1, 22]. However, our **first** key novelty is the generalization of Spherical Neural Surfaces to dynamic surfaces, *i.e.*, 3D surfaces that deform over time, and show how to learn it from discrete 4D surfaces. We refer to this representation as Dynamic Spherical Neural Surfaces (D-SNS). Our **second** key novelty is that we use this representation to develop a comprehensive framework for the statistical analysis of 4D surfaces represented with their D-SNS. Traditionally, statistical shape analysis is performed on discretized surfaces. We show that this is not necessary and propose a framework that enables for the first time to perform functional shape analysis.

3. Method

We introduce a novel neural representation, termed Dynamic Spherical Neural Surfaces, to represent 4D (*i.e.*, 3D + time) surfaces (Section 3.1). The representation is continuous in space and time and thus enables us to treat 4D surfaces as continuous functions. To enable the statistical analysis of such 4D surfaces, we introduce a Riemannian shape space, equipped with a Riemannian elastic metric that measures bending and stretching. 3D surfaces can be seen as points in the Riemannian shape space while 4D neural surfaces can be seen as trajectories in that space. The shape space, however, is of infinite dimension and has a nonlinear structure, which makes downstream analysis tasks complex and computationally expensive. Thus, we further map 3D surfaces to the space of Square-Root Normal Fields (SRNFs) [16], which has an \mathbb{L}^2 structure. More importantly, the \mathbb{L}^2 metric in the space of SRNFs is equivalent to the partial elastic metric in the original space. This way, 4D neural surfaces become trajectories in the SRNF space, which is Euclidean. Thus, spatiotemporal registration (Section 3.2), geodesics between neural 4D surfaces (Section 3.3), and summary shape statistics of neural 4D surfaces (Section 3.4) can be efficiently computed in the SRNF space and mapped back to the original space for visualization. Importantly, we show that these quantities can be directly computed on the neural representation without upfront discretization.

3.1. Dynamic Spherical Neural Surfaces (D-SNS)

We treat a manifold surface, after normalization for translation and scale, as a function $f : \Omega \rightarrow [-1, 1]^3$, where $\Omega \subset \mathbb{R}^3$ is a parameterization domain. Let \mathcal{C}_f be the pre-shape space of such functions. When dealing with closed genus-0 surfaces, Ω can be defined as a unit sphere \mathbb{S}^2 . Dynamic surfaces, *i.e.*, surfaces that deform over time can be represented in the same way by adding the temporal dimension, *i.e.*, $F : \Omega \times [0, 1] \rightarrow [-1, 1]^3$ where $[0, 1]$ is the normalized time domain. Traditionally, f and F are discretized both in space, and space and time, respectively [16, 17]. Thus, the performance of the subsequent analysis tasks depends on the quality of the discretization. In this paper, we exploit the power of neural networks as universal approximators to represent F as a continuous function in space and time. We refer to this novel neural representation as *Dynamic Spherical Neural Surfaces (D-SNS)*.

3.1.1 Representation

Given a discrete 4D surface where each time instance is a genus-0 triangular mesh, we first spherically parameterize each individual mesh using the approach of Praun and Hoppe [14, 28] and put them in correspondence using the approach of Laga *et al.* [16]. To learn a continuous representation, we overfit a Multi-Layer Perceptron (MLP) F_Θ

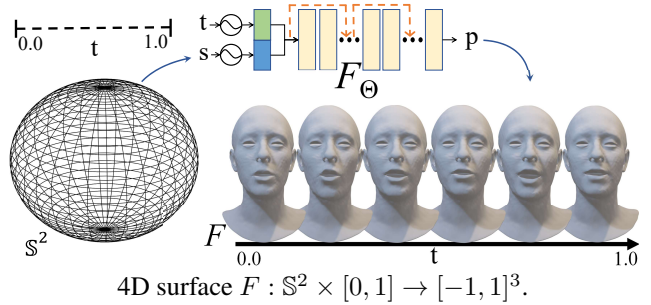


Figure 1. **Dynamic Spherical Neural Surface.** Given a discrete 4D surface of meshes parameterized by a unit sphere \mathbb{S}^2 and time t - we overfit an MLP F_Θ to create a continuous 4D surface F by minimizing the Mean Square Error between the ground truth and predicted surface points.

to this representation. The MLP maps a $s \in \mathbb{S}^2$ and a time $t \in [0, 1]$ to a 3D surface point $p \in [-1, 1]^3$:

$$F_\Theta : \mathbb{S}^2 \times [0, 1] \rightarrow [-1, 1]^3; \quad s, t \mapsto F_\Theta(s, t) = p. \quad (1)$$

Here, $s = (u, v)$ are the spherical coordinates and Θ the network parameters. We overfit the MLP to a discrete 4D surface by minimizing the Mean Square Error between the ground truth point p_i and the predicted point p_i^* :

$$L_N = \frac{1}{K} \frac{1}{N} \sum_{j=1}^K \sum_{i=1}^N \| F_\Theta(s_i, t_j) - p_{ij} \|^2. \quad (2)$$

Here, N is the number of sample points per surface and K is the number of temporal samples used for training.

One advantage of the proposed neural representation is that differential properties of 3D and 4D surfaces can be computed directly in the continuous domain, without discretization. In particular, the normal field of a 3D surface $F_\Theta(\cdot, t)$, which is essential for the spatiotemporal registration as will be seen in Section 3.2, is given by:

$$\nabla F_\Theta(\cdot, t) = \left(\frac{\partial F_\Theta(\cdot, t)}{\partial u}, \frac{1}{\sin u} \frac{\partial F_\Theta(\cdot, t)}{\partial v} \right)^\top, \quad (3)$$

where $u \in [0, \pi]$ and $v \in [0, 2\pi)$. Similarly, the tangent field to the 4D surface can be computed by taking the derivative with respect to time, *i.e.*, $\partial F_\Theta / \partial t$. These differential quantities can be automatically computed using the automatic differentiation functionality built into modern machine learning setups.

3.1.2 Network architecture

We use an MLP consisting of six residual blocks. Each block has two layers, with 1024 nodes each. The use of residual connections significantly enhances the network's




Figure 2. Illustration of the spatiotemporal registration framework proposed in this paper. "Enc" refers to positional encoding.

ability to represent coarse and fine geometric details. It also significantly improves convergence time and reduces the risk of converging to local minima. Inspired by Space-time NeRF [35], we apply positional encoding to the spatial domain \mathbb{S}^2 and to the temporal domain $[0, 1]$. To avoid sharp edges typically associated with the ReLU activation function, we use the SoftPlus activation function [37], which is smooth and continuous. This way, the neural network represents 4D surfaces with very high accuracy.

3.2. Spatiotemporal registration using D-SNS

In this section, we show how to spatiotemporally register two 4D surfaces F_1 and F_2 directly in the continuous domain using their D-SNS representation. We formulate the spatiotemporal registration problem as that of finding a rotation $O \in SO(3)$, a spatial diffeomorphism $\gamma : \mathbb{S}^2 \rightarrow \mathbb{S}^2$, and a temporal diffeomorphism $\zeta : [0, 1] \rightarrow [0, 1]$ that bring the two 4D surfaces as close as possible to each other. This is an optimization problem of the form:

$$(O^*, \gamma^*, \zeta^*) = \operatorname{argmin}_{O, \gamma, \zeta} \int_0^1 \mathcal{L}(F_1(t), O(F_2 \circ \zeta)(t) \circ \gamma) dt. \quad (4)$$

Here, \mathcal{L} is a measure of closeness. In this paper, we treat γ and ζ as continuous functions; see Figure 2. During training, we freeze the D-SNS networks that represent F_1 and F_2 , and only optimize the parameters of γ and ζ , using the metric \mathcal{L} as the loss function. This is, however, not straightforward and can be very time-consuming since the optimization needs to be performed over the entire temporal sequence and using a metric \mathcal{L} that is not linear. Thus, we first proceed with spatial registration (Section 3.2.1) followed by temporal registration (Section 3.2.2).

3.2.1 Spatial registration

For simplicity of notation, let $f_1 = F_1(t)$ and $f_2 = F_2(t)$. To spatially register f_2 onto f_1 , we need to find a spatial diffeomorphism $\gamma^* : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ and a rotation $O^* \in SO(3)$ that bring f_2 as close as possible to f_1 , *i.e.*,

$$(O^*, \gamma^*) = \operatorname{argmin}_{O, \gamma} d_C(f_1, O(f_2 \circ \gamma)). \quad (5)$$




Figure 3. Illustration of the spatial registration of Dynamic Spherical Neural Surfaces (D-SNSs). First, input D-SNSs are mapped to their SRNF representation. We then use the \mathbb{L}^2 metric in the SRNF space to elastically register the two surfaces, *i.e.*, finding the optimal rotation O^* and diffeomorphism γ^* that minimize the \mathbb{L}^2 metric in the SRNF space.




Figure 4. A spatial diffeomorphism can be formulated as an MLP with two layers: the input layer evaluates the harmonic basis at the query point $s \in \mathbb{S}^2$. The output layer then computes their weighted sum. The learnable parameters are the weights of the output layer. To perform spatial registration, we freeze the weights of the D-SNS networks and only optimize the weights of the diffeomorphism network using the loss of Eqn. (7).

The metric d_C should measure the amount of bending and stretching one needs to apply to f_2 in order to align it onto f_1 . Such nonlinear metric, however, is complex to work with since it leads to a nonlinear optimization. Jermyn *et al.* [13] showed that by mapping the surfaces to the space of Square Root Normal Fields (SRNFs), the elastic metric becomes an \mathbb{L}^2 metric. Formally, the SRNF map H of a surface f is given by:

$$H(f) = h, \text{ such that } h(s) = \frac{\mathbf{n}(s)}{\sqrt{\|\mathbf{n}(s)\|_2}}, \quad (6)$$

where \mathbf{n} is the normal field of the surface f and is computed directly from the SNS representation using Eqn. (3).

With this representation, the spatial registration in Eqn. (5) can be reformulated under the SRNF as an optimization problem of the form (see Figure 3):

$$(O^*, \gamma^*) = \operatorname{argmin}_{O, \gamma} \|h_1 - O(h_2 \circ \gamma)\|^2. \quad (7)$$

Here, $h_1 = H(f_1)$ and $h_2 = H(f_2)$. To solve this optimization problem, we represent the diffeomorphism γ , which is a function on \mathbb{S}^2 , using the weighted sum of spherical harmonic basis: $\gamma = \sum_{i=1}^n a_i B_i$ and estimate the optimal weights $\{a_i\}$ using gradient descent; see Figure 4. During training, we freeze the weights and jointly optimize for rotation and diffeomorphism in an iterative manner until convergence: we first fix γ and solve for the rotation O using Singular Value Decomposition. We then fix O and solve for γ by optimizing the loss of Eqn. (7) using gradient descent. We repeat this process until convergence.

3.2.2 Temporal registration

Given two spatially registered 4D surfaces and both represented using their D-SNS F_1 and F_2 , their temporal registration can be formulated as that of finding a time-warping function $\zeta : [0, 1] \rightarrow [0, 1]$ that brings F_2 as close as possible to F_1 . This is an optimization problem of the form:

$$\zeta^* = \underset{\zeta}{\operatorname{argmin}} d(F_1, F_2 \circ \zeta). \quad (8)$$

where d is a measure of closeness. We treat time warping as curve stretching and thus, d should measure elasticity. However, working directly with the 4D surfaces F_1 and F_2 is computationally expensive since they are of infinite dimension (given the continuous nature of the proposed representation). Also, the metric d should measure elasticity and thus is nonlinear, which makes the optimization problem of Eqn. (8) complex to solve.

To address the first problem, we map the 4D surfaces to a low dimensional space of curves using Principle Component Analysis. This way, a 4D surface F becomes a curve α in the low-dimensional PCA space. Now, to address the second problem, we further map the curve to their Square Root Velocity Fields (SRVFs). Srivastava *et al.* [30] showed that the partial elastic metric that measures bending and stretching of curves becomes an \mathbb{L}^2 metric in the SRVF space. Formally, the SRVF of a curve α is defined as its tangent field scaled by the square root of the tangent field, *i.e.*, :

$$q = Q(\alpha) = \frac{\dot{\alpha}}{\sqrt{\|\dot{\alpha}\|}}. \quad (9)$$

Here, $\dot{\alpha} = \frac{\partial \alpha}{\partial t}$ is the tangent vector field on α . Given two trajectories α_1 and α_2 , we represent them in their respective SRVFs q_1 and q_2 . Then, Eqn. (8) is reformulated as:

$$\zeta^* = \underset{\zeta}{\operatorname{argmin}} \|q_1 - q_2 \circ \zeta\|^2. \quad (10)$$

We treat ζ as a continuous function and implement it using an MLP that takes time $t \in \mathbb{R}$ and returns $\zeta(t)$, which we enforce to in $[0, 1]$ by applying to it a Sigmoid activation function. During training, we freeze the parameters of the




Figure 5. Illustration of the temporal registration process of two 4D surfaces F_1 and F_2 using the proposed Dynamic Spherical Neural Surfaces representation. The 4D surfaces are first mapped to the SRVF space, which has an Euclidean structure. Thus, we formulate the temporal registration of 4D surfaces as that of elastic registration of curves in the SRVF space.

networks of q_1 and q_2 and only update the parameters of ζ by minimizing the loss of Eqn. (10). However, to ensure that the MLP learns a diffeomorphism in the temporal domain $[0, 1]$, *i.e.*, $\zeta_v : [0, 1] \rightarrow [0, 1]$ such that $0 < \zeta(t) < 1$, we regularize the time-warping network by enforcing the first derivative of the network with respect to time t to be non-negative. This is done by adding the following regularization term to the loss function:

$$L_M = \int_0^1 \max\left(0, -\frac{\partial \zeta(t)}{\partial t}\right). \quad (11)$$

Hence, the total loss for temporal registration is:

$$L = \|q_1 - q_2 \circ \zeta\|^2 + \lambda L_M. \quad (12)$$

Figure 5 summarizes the temporal registration process. We use an MLP composed of two residual blocks. Each block has two layers of 32 neurons each. We train the network using mini-batches of discretized time intervals and change the input samples after every 200 epochs to learn a continuous temporal registration.

3.3. Geodesics

Let F_1 and F_2 be two 4D surfaces after their spatiotemporal registration, using the approach described in Section 3.2. The geodesic between the two 4D surfaces can now be defined as the shortest path, with respect to the elastic metric of Eqn. (4). However, instead of working with this complex metric, we first map the 4D surfaces to the SRVF space, following the approach described in Section 3.2.2. Let q_1 and q_2 be the SRVFs of F_1 and F_2 , respectively. Since the SRVF space is Euclidean, then the geodesic path Λ_q between q_1 and q_2 is a straight line, *i.e.*,

$$\Lambda_q(\tau) = (1 - \tau)q_1 + \tau q_2, \quad \tau \in [0, 1], \quad (13)$$

which is straightforward to compute. To visualize the geodesic paths Λ_q , we first map the obtained geodesics in the SRVF space back to the space of 4D surfaces through SRVF inversion, which has an analytical form [30], followed by PCA inversion. Figure 9 shows an example of a geodesic between two 4D surfaces corresponding to 4D human performing jumping actions.

3.4. Co-registration and Karcher mean

The mean 4D surface \bar{F} , also called Karcher mean, of a set 4D surfaces $\{F_1, \dots, F_n\}$ is the 4D surface that is the closest, with respect to the metric of Eqn. (4), to all the input 4D surfaces after their spatiotemporal co-registration. Thus, the joint problem of computing the mean 4D surface and co-registration is an optimization problem of the form:

$$\bar{F}, \{O_i^*, \gamma_i^*, \zeta_i^*\}_{i=1}^n = \operatorname{argmin}_{F, \{O_i, \gamma_i, \zeta_i\}_{i=1}^n} \sum_{i=1}^n \int_0^1 \mathcal{L}(F(t), O(F_i \circ \zeta)(t) \circ \gamma) dt. \quad (14)$$

Similar to the pairwise temporal registration, we solve this optimization in two stages. First, we separately perform the spatial registration of the input 4D surfaces by computing (O_i^*, γ_i^*) following the approach described in Section 3.2.1. Then, we map the spatially registered 4D surfaces to their SRVF space and compute the mean as well as the temporal warpings $\{\zeta_i^*\}_{i=1}^n$ using the \mathbb{L}^2 metric, which is equivalent to the elastic metric \mathcal{L} in the original space, *i.e.*, :

$$\bar{q}, \{\zeta_i^*\}_{i=1}^n = \operatorname{argmin}_{q, \{\zeta_i\}_{i=1}^n} \sum_{i=1}^n \|q - q_i \circ \zeta_i\|^2. \quad (15)$$

Once solved, the mean 4D surface \bar{F} can be computed analytically by applying the inverse SRVF [30] to \bar{q} .

4. Results

Implementation details. We trained our neural networks using NVIDIA GeForce RTX 4090 GPU with 2.4 GHz Intel Core i9. All networks are optimized using RMSProp with a learning rate of 10^{-4} and a momentum of 0.9. We trained D-SNS for 50K epochs, which takes between two to three hours depending on the number of vertices in the mesh, the spatial diffeomorphism for 500 iterations, which takes around two minutes, and the time warping network for up to 3,000 epochs, which takes between 10 to 30 minutes. The D-SNS and the time-warping network use SoftPlus activation. We initialize the weights of the time-warping network by overfitting it to the identity diffeomorphism.

Datasets. We have evaluated the proposed framework on the 4D face and 4D human body surfaces from the MPI DFAUST [5], VOCA [9], COMA [29], and CAPE [20]; see

Datasets	Mean ($\times 10^{-6}$)	Std ($\times 10^{-6}$)	Median ($\times 10^{-6}$)
DFAUST [5]	1.60	0.52	1.52
CAPE [20]	1.51	0.68	1.27
COMA [29]	2.29	1.68	1.66
VOCA [9]	0.89	0.51	0.79

Table 1. Evaluation of the performance of the proposed neural representation on four datasets. We use the average \mathbb{L}^2 distance between the ground truth discrete and the neural surface.

Datasets	Original temporal samples			30 temporal samples		
	Mean ($\times 10^{-5}$)	Std ($\times 10^{-5}$)	Median ($\times 10^{-5}$)	Mean ($\times 10^{-5}$)	Std ($\times 10^{-5}$)	Median ($\times 10^{-5}$)
DFAUST [5]	0.17	0.08	0.16	1.00	0.55	1.20
CAPE [20]	0.14	0.02	0.12	1.34	0.72	1.25
COMA [29]	0.23	0.10	0.20	0.14	0.02	0.14
VOCA [9]	0.10	0.04	0.10	0.11	0.06	0.08

Table 2. D-SNS interpolation results on original and 30 temporal samples, computed using the average \mathbb{L}^2 distance between the ground truth and the neural surfaces. Results are shown for five random 4D surfaces across four datasets. Note that, due to small errors, all values are shown as multipliers of 10^{-5} .

the Supplementary Material. The datasets come with registered triangulation meshes. We spherically parameterize these datasets using the implementation [14] of Praun and Hoppe’s approach [28]. We then generate random diffeomorphisms to simulate non-registered surfaces.

4.1. Dynamic Spherical Neural Surfaces results

In this experiment, we take the ground truth discrete 4D surfaces of DFAUST, COMA, CAPE, and VOCA datasets and train one D-SNS per 4D surface. To quantitatively evaluate the quality of the neural representation, we measure the average point-wise \mathbb{L}^2 distance between the ground truth discrete 4D surface and the D-SNS representation. Table 1 summarizes the mean, median, and standard deviation of the representation error over each of the four datasets. As one can see, the errors are in the order of 10^{-6} , which demonstrates that the novel neural representation can faithfully represent genus-0 4D surfaces. (Note that all the surfaces are normalized for translation and scale and thus they fit within a bounding sphere of radius 1.)

Figure 6 shows some results demonstrating the representation quality of the proposed D-SNS. In this figure, we picked up some surfaces and showed, in the form of a heat map, the error between the original surface and the neural surface. As one can see, the errors are less than 0.01. This shows that the proposed D-SNS can accurately represent complex 4D surfaces. Table 2, on the other hand, demonstrates its interpolation ability. In this experiment, the neural representation was trained only on 30 temporal samples of the entire sequence, yet, the method is able to interpolate those keyframes and generate plausibly smooth 4D surfaces. This demonstrates the D-SNS ability to suc-




Figure 6. We measure the pointwise error between the proposed D-SNS and the groundtruth discrete 4D surfaces. Here, we show some time frames with the error plotted as a heatmap. Observe that the proposed neural representation can accurately represent 4D surfaces, with a pointwise error that is less than 0.01. Please refer to the Supplementary Material for additional results.

cessfully interpolate temporally-sparse sequences.

4.2. Spatiotemporal registration and geodesics

Figure 7 shows an example of the spatiotemporal registration of pairs of 4D humans deforming at different rates. We show the 4D surfaces before and after their temporal registration. As we can see, our approach is capable of bringing the sequences in full alignment. Note that the spatiotemporal registration of human bodies is more complex than faces due to greater motion and body articulation, which leads to significant bending and stretching. Please refer to the Supplementary Material for additional visual results.

We quantitatively evaluate the proposed temporal registration against the method of Laga *et al.* [17] using the evaluation framework proposed in 4D Atlas [17], since, to the best of our knowledge, is the only method paper that dealt with the analysis of 4D surfaces. We take a 4D neural surface F_i and apply to it 10 random temporal diffeomorphisms ζ_{ij}^* resulting in 10 surfaces $\{F_j\}_{j=1}^{10}$ that differ in their execution speeds. Then, for each pair (F_i, F_j) , we compute, using the proposed approach, the optimal time warping ζ_{ij}^* that aligns F_j onto F_i . Ideally, ζ_{ij}^* should be the inverse of ζ_{ij} . Thus, we use the distance between ζ_{ij}^* and the inverse of ζ_{ij} as a measure of error, which we aggregate over the 10-random diffeomorphisms.

We perform this experiment on the four datasets, using our method and 4D Atlas [17]. We use two different temporal samplings (50 and 25 temporal samples) to analyze the robustness of both methods. Figure 8 presents the results in the form of box plots. This experiment shows that when the temporal sampling is dense, the two methods achieve similar performance, with a slight advantage to ours. However, when we reduce the temporal sampling (25 samples), we observe that the alignment error of 4D Atlas [17] increases while ours remains more or less stable. This is due to the fact that we treat 4D surfaces as continuous functions and

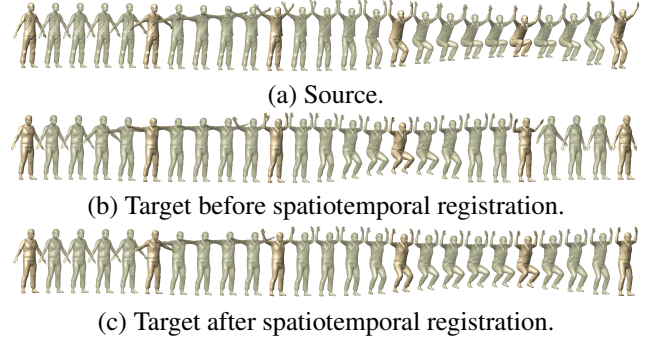


Figure 7. Example of the spatiotemporal registration of two 4D humans from the CAPE dataset. The proposed framework aligns the target 4D surface (b) onto the source (a). The target neural surface after registration (c) is fully aligned with the source. Please refer to the Supplementary Material for additional results.

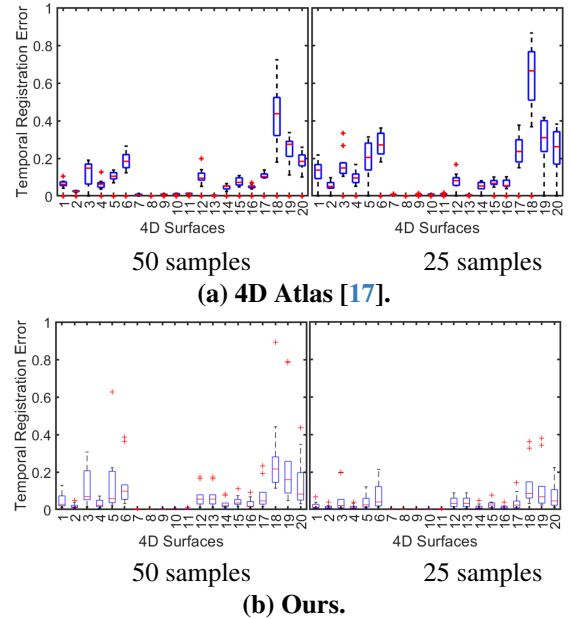


Figure 8. Boxplots providing the statistics on the accuracy of the temporal registration when using two different temporal samplings (50 and 25 temporal samples). We run the experiment on 20 pairs, each pair is randomly sampled with 10 different random temporal diffeomorphisms and report the statistics on the temporal registration error aggregated over the 10 runs. The red lines denote the median error, while the boxes represent its spread.

thus they are less sensitive to the sampling density.

We also evaluate the registration performance on different 4D surfaces performing the same action; see Table 3. In this experiment, we take 5 pairs from each data set, perform their temporal registration using 4D Atlas [17] and our method. We then measure the geodesic distance between the registered 4D surfaces. As each pair of 4D surfaces perform the same action, the smaller the geodesic distance, the

	CAPE [20]			DFAUST [5]			COMA [29]			VOCA [9]		
	Mean	Std	Median	Mean	Std	Median	Mean	Std	Median	Mean	Std	Median
4D Atlas	1.2938	0.4764	1.404	2.3785	1.1237	1.8665	0.0597	0.0165	0.0617	0.5026	0.1057	0.4060
Ours	0.3574	0.1697	0.3199	0.7726	0.2037	0.7210	0.0276	0.0191	0.0173	0.0994	0.0253	0.1030

Table 3. Comparison of the proposed spatiotemporal registration with 4D Atlas [17]. The evaluation is performed on five pairs performing the same actions. We measure the geodesic distance after the spatiotemporal registration using 4D Atlas [17] and our approach. The smaller this distance is, the better the alignment. The best results are shown in **bold**; see the Supplementary Material for additional results.



Figure 9. Example of a geodesic after the spatiotemporal registration between two 4D surfaces performing a jumping action. The first row shows the source 4D surface, the last row shows the target 4D surface, and the intermediate rows show intermediate 4D surfaces sampled at equidistance along the geodesic path between the source and target. The highlighted row corresponds to the mean 4D surface. The Supplementary Material provides the original 4D surfaces and geodesic before spatiotemporal registration. The Supplementary Material provides more results.

better the alignment. As we can see, our continuous framework is able to outperform 4D Atlas [17] on all datasets. We performed this experiment using five pairs of 4D surfaces, each pair performing the same action at varying speeds. The Supplementary Material provides more evaluation details.

4D geodesics. Figure 9 shows an example of a geodesic between 4D neural surfaces performing a jumping action. We show the geodesic after spatiotemporal registration of the source (top row) to the target surface (bottom row). The Supplementary Material shows the geodesic before spatiotemporal registration. Observe that before registration (see the Supplementary Material), the jumping actions are misaligned, and the geodesics are not as closely aligned compared to those after registration. We refer the reader to the Supplementary Material for more examples and results.

4.3. Co-registration and mean 4D surfaces

Figure 10 shows the 4D mean surface of a set of 4D neural surfaces, from the CAPE dataset, after their co-registration. These surfaces perform various actions at different execution speeds. The 4D surfaces in the first four rows perform a

squat action, and the following two rows perform a bending action. Note that the squat action is repetitive and the number of cycles differs from one 4D surface to another. Despite this complexity, the proposed approach is able to co-register the 4D surfaces and compute a plausible 4D mean that is as close as possible to all the other 4D surfaces; see the Supplementary Material for additional examples.

5. Conclusion

We proposed Dynamic Spherical Neural Surfaces as a novel continuous representation of genus-0 4D surfaces that is continuous in both space and time. We then formulated the 4D shape analysis problem as one of analyzing functions. This allows all analysis tasks to be performed directly on these continuous functions without unnecessary discretization. By linking neural networks with the classical Riemannian geometry for 4D shape analysis, this work establishes the first foundational approach that enables comprehensive functional analysis of 4D shapes.



Figure 10. Example of a mean 4D surface (highlighted in red) computed on six 4D surfaces from the CAPE dataset after their spatiotemporal registration. The input 4D surfaces perform different actions: the 4D surfaces in the first four rows perform a squat action while the last two perform a back and forward bending action; see the Supplementary Material for additional results.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In *International conference on machine learning*, pages 40–49. PMLR, 2018. 2
- [2] Brett Allen, Brian Curless, and Zoran Popović. The space of human body shapes: reconstruction and parameterization from range scans. *ACM TOG*, 22(3):587–594, 2003. 2
- [3] M Faisal Beg, Michael I Miller, Alain Trounevé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61:139–157, 2005. 2
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 2
- [5] 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)*, 2017. 2, 6, 8, 3
- [6] Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006. 2
- [7] Alexandre Bône, Olivier Colliot, and Stanley Durrleman. Learning the spatiotemporal variability in longitudinal shape data sets. *International Journal of Computer Vision*, 128, 2020. 2
- [8] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active shape models-their training and application. *CVIU*, 61(1):38–59, 1995. 2
- [9] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael Black. Capture, learning, and synthesis of 3D speaking styles. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10101–10111, 2019. 2, 6, 8, 3
- [10] Thomas Davies, Derek Nowrouzezahrai, and Alec Jacobson. On the effectiveness of weight-encoded neural implicit 3d shapes. In *International Conference on Machine Learning (ICML)*. PMLR, 2021. 2
- [11] Vianney Debavelaere, Alexandre Bône, Stanley Durrleman, and Stéphanie Allasonnière. Clustering of longitudinal shape data sets using mixture of separate or branching trajectories. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV*, page 66–74, Berlin, Heidelberg, 2019. Springer-Verlag. 2
- [12] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3D surface generation. In *IEEE CVPR*, pages 216–224, 2018. 2
- [13] 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. 2, 4
- [14] Sebastian Kurtek, Anuj Srivastava, Eric Klassen, and Hamid Laga. Landmark-guided elastic shape analysis of spherically-parameterized surfaces. In *Computer graphics forum*, pages 429–438. Wiley Online Library, 2013. 3, 6, 1
- [15] Hamid Laga. A survey on nonrigid 3D shape analysis. *Academic Press Library in Signal Processing, Volume 6*, pages 261–304, 2018. 2
- [16] Hamid Laga, Qian Xie, Ian H Jermyn, and Anuj Srivastava. Numerical inversion of srnf maps for elastic shape analysis of genus-zero surfaces. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2451–2464, 2017. 2, 3
- [17] Hamid Laga, Marcel Padilla, Ian H. Jermyn, Sebastian Kurtek, Mohammed Bennamoun, and Anuj Srivastava. 4d atlas: Statistical analysis of the spatiotemporal variability in longitudinal 3d shape data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1335–1352, 2023. 2, 3, 7, 8, 5, 6
- [18] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM TOG*, 36(6):194–1, 2017. 2
- [19] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A Skinned Multi-Person Linear Model. *ACM TOG*, 34(6):1–16, 2015. 2
- [20] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to Dress 3D People in Generative Clothing. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6, 8, 3
- [21] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*, pages 405–421. Springer, 2020. 2
- [22] Luca Morreale, Noam Aigerman, Vladimir G Kim, and Niloy J Mitra. Neural Surface Maps. In *IEEE CVPR*, pages 4639–4648, 2021. 2
- [23] Luca Morreale, Noam Aigerman, Vladimir G Kim, and Niloy J Mitra. Neural semantic surface maps. In *Computer Graphics Forum*, page e15005. Wiley Online Library, 2024. 2
- [24] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4), 2012. 2
- [25] Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas, Frédéric Chazal, and Alex Bronstein. Computing and processing correspondences with functional maps. In *SIG-GRAPH ASIA*, pages 1–60, 2016. 2
- [26] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *IEEE CVPR*, pages 165–174, 2019. 2
- [27] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive Body Capture: 3D Hands, Face, and Body From a Single Image. In *IEEE CVPR*, 2019. 2

- [28] Emil Praun and Hugues Hoppe. Spherical parametrization and remeshing. *ACM transactions on graphics (TOG)*, 22(3):340–349, 2003. [3](#), [6](#)
- [29] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. Generating 3D faces using convolutional mesh autoencoders. In *European Conference on Computer Vision (ECCV)*, pages 725–741, 2018. [2](#), [6](#), [8](#), [3](#)
- [30] Anuj Srivastava, Eric Klassen, Shantanu H Joshi, and Ian H Jermyn. Shape analysis of elastic curves in euclidean spaces. *IEEE transactions on pattern analysis and machine intelligence*, 33(7):1415–1428, 2010. [5](#), [6](#)
- [31] Jingyong Su, Sebastian Kurtek, Eric Klassen, and Anuj Srivastava. Statistical analysis of trajectories on riemannian manifolds: Bird migration, hurricane tracking and video surveillance. *The Annals of Applied Statistics*, 8(1), 2014. [2](#)
- [32] Zhe Su, Eric Klassen, and Martin Bauer. The square root velocity framework for curves in a homogeneous space, 2017. [2](#)
- [33] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. [2](#)
- [34] Romy Williamson and Niloy J. Mitra. Neural Geometry Processing via Spherical Neural Surfaces. In *arXiv:2407.07755*, 2024. [2](#)
- [35] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *IEEE CVPR*, pages 9421–9431, 2021. [4](#)
- [36] Jason Zhang, Gengshan Yang, Shubham Tulsiani, and Deva Ramanan. NeRS: Neural Reflectance Surfaces for Sparse-view 3D Reconstruction in the Wild. *Advances in Neural Information Processing Systems*, 34:29835–29847, 2021. [2](#)
- [37] Hao Zheng, Zhanlei Yang, Wenju Liu, Jizhong Liang, and Yanpeng Li. Improving deep neural networks using softplus units. In *2015 International joint conference on neural networks (IJCNN)*, pages 1–4. IEEE, 2015. [4](#)
- [38] Silvia Zuffi, Angjoo Kanazawa, Tanya Berger-Wolf, and Michael J Black. Three-d safari: Learning to estimate zebra pose, shape, and texture from images” in the wild”. In *IEEE CVPR*, pages 5359–5368, 2019. [2](#)