

Surface Networks

Ilya Kostrikov¹, Zhongshi Jiang¹, Daniele Panozzo^{*1}, Denis Zorin^{†1}, and Joan Bruna^{‡1,2}

¹Courant Institute of Mathematical Sciences, New York University

²Center for Data Science, New York University

Abstract

We study data-driven representations for three-dimensional triangle meshes, which are one of the prevalent objects used to represent 3D geometry. Recent works have developed models that exploit the intrinsic geometry of manifolds and graphs, namely the Graph Neural Networks (GNNs) and its spectral variants, which learn from the local metric tensor via the Laplacian operator.

Despite offering excellent sample complexity and built-in invariances, intrinsic geometry alone is invariant to isometric deformations, making it unsuitable for many applications. To overcome this limitation, we propose several upgrades to GNNs to leverage extrinsic differential geometry properties of three-dimensional surfaces, increasing its modeling power. In particular, we propose to exploit the Dirac operator, whose spectrum detects principal curvature directions — this is in stark contrast with the classical Laplace operator, which directly measures mean curvature. We coin the resulting models Surface Networks (SN).

We prove that these models define shape representations that are stable to deformation and to discretization, and we demonstrate the efficiency and versatility of SNs on two challenging tasks: temporal prediction of mesh deformations under non-linear dynamics and generative models using a variational autoencoder framework with encoders/decoders given by SNs.

1. Introduction

3D geometry analysis, manipulation and synthesis plays an important role in a variety of applications from engineering to computer animation to medical imaging. Despite the

vast amount of high-quality 3D geometric data available, data-driven approaches to problems involving complex geometry have yet to become mainstream, in part due to the lack of data representation regularity which is required for traditional convolutional neural network approaches. While in computer vision problems inputs are typically sampled on regular two or three-dimensional grids, surface geometry is represented in a more complex form and, in general, cannot be converted to an image-like format by parametrizing the shape using a single planar chart. Most commonly an irregular triangle mesh is used to represent shapes, capturing its main topological and geometrical properties.

Similarly to the regular grid case (used for images or videos), we are interested in data-driven representations that strike the right balance between expressive power and sample complexity. In the case of CNNs, this is achieved by exploiting the inductive bias that most computer vision tasks are locally stable to deformations, leading to localized, multiscale, stationary features. In the case of surfaces, we face a fundamental modeling choice between *extrinsic* versus *intrinsic* representations. Extrinsic representations rely on the specific embedding of surfaces within a three-dimensional ambient space, whereas intrinsic representations only capture geometric properties specific to the surface, irrespective of its parametrization. Whereas the former offer arbitrary representation power, they are unable to easily exploit inductive priors such as stability to local deformations and invariance to global transformations.

A particularly simple and popular extrinsic method [30, 31] represents shapes as point clouds in \mathbb{R}^3 of variable size, and leverages recent deep learning models that operate on input sets [38, 37]. Despite its advantages in terms of ease of data acquisition (they no longer require a mesh triangulation) and good empirical performance on shape classification and segmentation tasks, one may wonder whether this simplification comes at a loss of precision as one considers more challenging prediction tasks.

In this paper, we develop an alternative pipeline that applies neural networks directly on triangle meshes, building

^{*}DP was supported in part by the NSF CAREER award IIS-1652515, a gift from Adobe, and a gift from nTopology.

[†]DZ was supported in part by the NSF awards DMS-1436591 and IIS-1320635.

[‡]JB was partially supported by Samsung Electronics (Improving Deep Learning using Latent Structure) and DOA W911NF-17-1-0438. Corresponding author: bruna@cims.nyu.edu

on *geometric deep learning*. These models provide data-driven intrinsic graph and manifold representations with inductive biases analogous to CNNs on natural images. Models based on Graph Neural Networks [34] and their spectral variants [6, 9, 22] have been successfully applied to geometry processing tasks such as shape correspondence [27]. In their basic form, these models learn a deep representation over the discretized surface by combining a latent representation at a given node with a local linear combination of its neighbors’ latent representations, and a point-wise nonlinearity. Different models vary in their choice of linear operator and point-wise nonlinearity, which notably includes the graph Laplacian, leading to spectral interpretations of those models.

Our contributions are three-fold. First, we extend the model to support extrinsic features. More specifically, we exploit the fact that surfaces in \mathbb{R}^3 admit a first-order differential operator, the *Dirac* operator, that is stable to discretization, provides a direct generalization of Laplacian-based propagation models, and is able to detect principal curvature directions [8, 16]. Next, we prove that the models resulting from either Laplace or Dirac operators are stable to deformations and to discretization, two major sources of variability in practical applications. Last, we introduce a generative model for surfaces based on the variational autoencoder framework [21, 33], that is able to exploit non-Euclidean geometric regularity.

By combining the Dirac operator with input coordinates, we obtain a fully differentiable, end-to-end feature representation that we apply to several challenging tasks. The resulting Surface Networks – using either the Dirac or the Laplacian, inherit the stability and invariance properties of these operators, thus providing data-driven representations with certified stability to deformations. We demonstrate the model efficiency on a temporal prediction task of complex dynamics, based on a physical simulation of elastic shells, which confirms that whenever geometric information (in the form of a mesh) is available, it can be leveraged to significantly outperform point-cloud based models.

Our main contributions are summarized as follows:

- We demonstrate that Surface Networks provide accurate temporal prediction of surfaces under complex non-linear dynamics, motivating the use of geometric shape information.
- We prove that Surface Networks define shape representations that are stable to deformation and to discretization.
- We introduce a generative model for 3D surfaces based on the variational autoencoder.

A reference implementation of our algorithm is available at <https://github.com/jiangzhongshi/SurfaceNetworks>.

2. Related Work

Learning end-to-end representations on irregular and non-Euclidean domains is an active and ongoing area of research. [34] introduced graph neural networks as recursive neural networks on graphs, whose stationary distributions could be trained by backpropagation. Subsequent works [23, 37] have relaxed the model by untying the recurrent layer weights and proposed several nonlinear updates through gating mechanisms. Graph neural networks are in fact natural generalizations of convolutional networks to non-Euclidean graphs. [6, 15] proposed to learn smooth spectral multipliers of the graph Laplacian, albeit with high computational cost, and [9, 22] resolved the computational bottleneck by learning polynomials of the graph Laplacian, thus avoiding the computation of eigenvectors and completing the connection with GNNs. We refer the reader to [5] for an exhaustive literature review on the topic. GNNs are finding application in many different domains. [2, 7] develop graph interaction networks that learn pairwise particle interactions and apply them to discrete particle physical dynamics. [10, 19] study molecular fingerprints using variants of the GNN architecture, and [12] further develop the model by combining it with set representations [38], showing state-of-the-art results on molecular prediction. The resulting models, so-called Message-Passing Neural Networks, also learn the diffusion operator, which can be seen as generalizations of the Dirac model on general graphs.

In the context of computer graphics, [25] developed the first CNN model on meshed surfaces using intrinsic patch representations, and further generalized in [4] and [27]. This last work allows for flexible representations via the so-called pseudo-coordinates and obtains state-of-the-art results on 3D shape correspondence, although it does not easily encode first-order differential information. These intrinsic models contrast with Euclidean models such as [42, 40], that have higher sample complexity, since they need to learn the underlying invariance of the surface embedding. Point-cloud based models are increasingly popular to model 3d objects due to their simplicity and versatility. [30, 31] use set-invariant representations from [38, 37] to solve shape segmentation and classification tasks. More recently, [24] proposes to learn surface convolutional network from a canonical representation of planar flat-torus, with excellent performance on shape segmentation and classification, although such canonical representations may introduce exponential scale changes that can introduce instabilities. Finally, [11] proposes a point-cloud generative model for 3D shapes, that incorporates invariance to point permutations, but does not encode geometrical information as our shape generative model. Learning variational deformations is an important problem for graphics applications, since it enables negligible and fixed per-frame cost [29], but it is currently limited to 2D deformations using point handles.

In contrast, our method easily generalizes to 3D and learns dynamic behaviours.

3. Surface Networks

This section presents our surface neural network model and its basic properties. We start by introducing the problem setup and notations using the Laplacian formalism (Section 3.1), and then introduce our model based on the Dirac operator (Section 3.2).

3.1. Laplacian Surface Networks

Our first goal is to define a trainable representation of discrete surfaces. Let $\mathcal{M} = \{V, E, F\}$ be a triangular mesh, where $V = (v_i \in \mathbb{R}^3)_{i \leq N}$ contains the node coordinates, $E = (e_{i,j})$ corresponds to edges, and F is the set of triangular faces. We denote as Δ the discrete Laplace-Beltrami operator (we use the popular cotangent weights formulation, see [5] for details).

This operator can be interpreted as a local, linear high-pass filter in \mathcal{M} that acts on signals $x \in \mathbb{R}^{d \times |V|}$ defined on the vertices as a simple matrix multiplication $\tilde{x} = \Delta x$. By combining Δ with an *all-pass* filter and learning generic linear combinations followed by a point-wise nonlinearity, we obtain a simple generalization of localized convolutional operators in \mathcal{M} that update a feature map from layer k to layer $k + 1$ using trainable parameters A_k and B_k :

$$x^{k+1} = \rho(A_k \Delta x^k + B_k x^k), \quad A_k, B_k \in \mathbb{R}^{d_{k+1} \times d_k}. \quad (1)$$

By observing that the Laplacian itself can be written in terms of the graph weight similarity by diagonal renormalization, this model is a specific instance of the graph neural network [34, 5, 22] and a generalization of the spectrum-free Laplacian networks from [9]. As shown in these previous works, convolutional-like layers (1) can be combined with graph coarsening or pooling layers.

In contrast to general graphs, meshes contain a low-dimensional Euclidean embedding that contains potentially useful information in many practical tasks, despite being extrinsic and thus not invariant to the global position of the surface. A simple strategy to strike a good balance between expressivity and invariance is to include the node canonical coordinates as input channels to the network: $x^1 := V \in \mathbb{R}^{|V| \times 3}$. The mean curvature can be computed by applying the Laplace operator to the coordinates of the vertices:

$$\Delta x^1 = -2H\mathbf{n}, \quad (2)$$

where H is the mean curvature function and $\mathbf{n}(u)$ is the normal vector of the surface at point u . As a result, the Laplacian neural model (1) has access to mean curvature and normal information. Feeding Euclidean embedding coordinates into graph neural network models is related to the use of generalized coordinates from [27]. By cascading K

layers of the form (1) we obtain a representation $\Phi_\Delta(\mathcal{M})$ that contains generic features at each node location. When the number of layers K is of the order of $\text{diam}(\mathcal{M})$, the diameter of the graph determined by \mathcal{M} , then the network is able to propagate and aggregate information across the whole surface.

Equation (2) illustrates that a Laplacian layer is only able to extract isotropic high-frequency information, corresponding to the mean variations across all directions. Although in general graphs there is no well-defined procedure to recover anisotropic local variations, in the case of surfaces some authors ([4, 1, 27]) have considered anisotropic extensions. We describe next a particularly simple procedure to increase the expressive power of the network using a related operator from quantum mechanics: the Dirac operator, that has been previously used successfully in the context of surface deformation [8] and shape analysis [16].

3.2. Dirac Surface Networks

The Laplace-Beltrami operator Δ is a second-order differential operator, constructed as $\Delta = -\text{div}\nabla$ by combining the gradient (a first-order differential operator) with its adjoint, the divergence operator. In an Euclidean space, one has access to these first-order differential operators separately, enabling oriented high-pass filters.

For convenience, we embed \mathbb{R}^3 to the imaginary quaternion space $\text{Im}(\mathbb{H})$ (see Appendix A in the Suppl. Material for details). The Dirac operator is then defined as a matrix $D \in \mathbb{H}^{|F| \times |V|}$ that maps (quaternion) signals on the nodes to signals on the faces. In coordinates,

$$D_{f,j} = \frac{-1}{2|\mathcal{A}_f|} e_j, \quad f \in F, j \in V,$$

where e_j is the opposing edge vector of node j in the face f , and \mathcal{A}_f is the area (see Appendix A) using counter-clockwise orientations on all faces.

To apply the Dirac operator defined in quaternions to signals in vertices and faces defined in real numbers, we write the feature vectors as quaternions by splitting them into chunks of 4 real numbers representing the real and imaginary parts of a quaternion; see Appendix A. Thus, we always work with feature vectors with dimensionalities that are multiples of 4. The Dirac operator provides first-order differential information and is sensitive to local orientations. Moreover, one can verify [8] that

$$\text{Re } D^* D = \Delta,$$

where D^* is the adjoint operator of D in the quaternion space (see Appendix A). The adjoint matrix can be computed as $D^* = M_V^{-1} D^H M_F$ where D^H is a conjugate transpose of D and M_V, M_F are diagonal mass matrices with one third of areas of triangles incident to a vertex and face areas respectively.

The Dirac operator can be used to define a new neural surface representation that alternates layers with signals defined over nodes with layers defined over faces. Given a d -dimensional feature representation over the nodes $x^k \in \mathbb{R}^{d \times |V|}$, and the faces of the mesh, $y^k \in \mathbb{R}^{d \times |F|}$, we define a d' -dimensional mapping to a face representation as

$$y^{k+1} = \rho(C_k D x^k + E_k y^k), C_k, E_k \in \mathbb{R}^{d_{k+1} \times d_k}, \quad (3)$$

where C_k, E_k are trainable parameters. Similarly, we define the adjoint layer that maps back to a d -dimensional signal over nodes as

$$x^{k+1} = \rho(A_k D^* y^{k+1} + B_k x^k), A_k, B_k \in \mathbb{R}^{d_{k+1} \times d_k}, \quad (4)$$

where A_k, B_k are trainable parameters. A surface neural network layer is thus determined by parameters $\{A, B, C, E\}$ using equations (3) and (4) to define $x^{k+1} \in \mathbb{R}^{d_{k+1} \times |V|}$. We denote by $\Phi_D(\mathcal{M})$ the mesh representation resulting from applying K such layers (that we assume fixed for the purpose of exposition).

The Dirac-based surface network is related to edge feature transforms proposed on general graphs in [12], although these edge measurements cannot be associated with derivatives due to lack of proper orientation. In general graphs, there is no notion of square root of Δ that recovers oriented first-order derivatives.

4. Stability of Surface Networks

Here we describe how Surface Networks are geometrically stable, because surface deformations become additive noise under the model. Given a continuous surface $S \subset \mathbb{R}^3$ or a discrete mesh \mathcal{M} , and a smooth deformation field $\tau : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, we are particularly interested in two forms of stability:

- Given a discrete mesh \mathcal{M} and a certain non-rigid deformation τ acting on \mathcal{M} , we want to certify that $\|\Phi(\mathcal{M}) - \Phi(\tau(\mathcal{M}))\|$ is small if $\|\nabla\tau(\nabla\tau)^* - \mathbf{I}\|$ is small, i.e. when the deformation is nearly rigid; see Theorem 4.1.
- Given two discretizations \mathcal{M}_1 and \mathcal{M}_2 of the same underlying surface S , we would like to control $\|\Phi(\mathcal{M}_1) - \Phi(\mathcal{M}_2)\|$ in terms of the resolution of the meshes; see Theorem 4.2.

These stability properties are important in applications, since most tasks we are interested in are stable to deformation and to discretization. We shall see that the first property is a simple consequence of the fact that the mesh Laplacian and Dirac operators are themselves stable to deformations. The second property will require us to specify under which conditions the discrete mesh Laplacian $\Delta_{\mathcal{M}}$ converges to the Laplace-Beltrami operator Δ_S on S . Unless it is clear

from the context, in the following Δ will denote the discrete Laplacian.

Theorem 4.1 *Let \mathcal{M} be a N -node mesh and $x, x' \in \mathbb{R}^{|V| \times d}$ be input signals defined on the nodes. Assume the nonlinearity $\rho(\cdot)$ is non-expansive ($|\rho(z) - \rho(z')| \leq |z - z'|$). Then*

- $\|\Phi_{\Delta}(\mathcal{M}; x) - \Phi_{\Delta}(\mathcal{M}; x')\| \leq \alpha_{\Delta} \|x - x'\|$, where α_{Δ} depends only on the trained weights and the mesh.
- $\|\Phi_D(\mathcal{M}; x) - \Phi_D(\mathcal{M}; x')\| \leq \alpha_D \|x - x'\|$, where α_D depends only on the trained weights and the mesh.
- Let $|\nabla\tau|_{\infty} := \sup_u \|\nabla\tau(u)(\nabla\tau(u))^* - \mathbf{I}\|$, where $\nabla\tau(u)$ is the Jacobian matrix of $u \mapsto \tau(u)$. Then $\|\Phi_{\Delta}(\mathcal{M}; x) - \Phi_{\Delta}(\tau(\mathcal{M}); x)\| \leq \beta_{\Delta} |\nabla\tau|_{\infty} \|x\|$, where β_{Δ} is independent of τ and x .
- Denote by $|\widetilde{\nabla\tau}|_{\infty} := \sup_u \|\nabla\tau(u) - \mathbf{I}\|$. Then $\|\Phi_D(\mathcal{M}; x) - \Phi_D(\tau(\mathcal{M}); x)\| \leq \beta_D |\widetilde{\nabla\tau}|_{\infty} \|x\|$, where β_D is independent of τ and x .

Properties (a) and (b) are not specific to surface representations, and are a simple consequence of the non-expansive property of our chosen nonlinearities. The constant α is controlled by the product of ℓ_2 norms of the network weights at each layer and the norm of the discrete Laplacian operator. Properties (c) and (d) are based on the fact that the Laplacian and Dirac operators are themselves stable to deformations, a property that depends on two key aspects: first, the Laplacian/Dirac is localized in space, and next, that it is a high-pass filter and therefore only depends on relative changes in position.

One caveat of Theorem 4.1 is that the constants appearing in the bounds depend upon a bandwidth parameter given by the reciprocal of triangle areas, which increases as the size of the mesh increases. This corresponds to the fact that the spectral radius of $\Delta_{\mathcal{M}}$ diverges as the mesh size N increases.

In order to overcome this problematic asymptotic behavior, it is necessary to exploit the smoothness of the signals incoming to the surface network. This can be measured with Sobolev norms defined using the spectrum of the Laplacian operator. Given a mesh \mathcal{M} of N nodes approximating an underlying surface S , and its associated cotangent Laplacian $\Delta_{\mathcal{M}}$, consider the spectral decomposition of $\Delta_{\mathcal{M}}$ (a symmetric, positive definite operator):

$$\Delta_{\mathcal{M}} = \sum_{k \leq N} \lambda_k e_k e_k^T, e_k \in \mathbb{R}^N, 0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N.$$

Under normal uniform convergence¹ [39], the spectrum of $\Delta_{\mathcal{M}}$ converges to the spectrum of the Laplace-Beltrami operator Δ_S of S . If S is bounded, it is known from the Weyl

¹which controls how the normals of the mesh align with the surface normals; see [39].

law [41] that there exists $\gamma > 0$ such that $k^{-\gamma(S)} \lesssim \lambda_k^{-1}$, so the eigenvalues λ_k do not grow too fast. The smoothness of a signal $x \in \mathbb{R}^{|V| \times d}$ defined in \mathcal{M} is captured by how fast its spectral decomposition $\hat{x}(k) = e_k^T x \in \mathbb{R}^d$ decays [36]. We define $\|x\|_{\mathcal{H}}^2 := \sum_k \lambda(k)^2 \|\hat{x}(k)\|^2$ is Sobolev norm, and $\beta(x, S) > 1$ as the largest rate such that its spectral decomposition coefficients satisfy

$$\|\hat{x}(k)\| \lesssim k^{-\beta}, (k \rightarrow \infty). \quad (5)$$

If $x \in \mathbb{R}^{|V| \times d}$ is the input to the Laplace Surface Network of R layers, we denote by $(\beta_0, \beta_1, \dots, \beta_{R-1})$ the smoothness rates of the feature maps $x^{(r)}$ defined at each layer $r \leq R$.

Theorem 4.2 Consider a surface S and a finite-mesh approximation \mathcal{M}_N of N points, and Φ_Δ a Laplace Surface Network with parameters $\{(A_r, B_r)\}_{r \leq R}$. Denote by $d(S, \mathcal{M}_N)$ the uniform normal distance, and let x_1, x_2 be piece-wise polyhedral approximations of $\bar{x}(t)$, $t \in S$ in \mathcal{M}_N , with $\|\bar{x}\|_{\mathcal{H}(S)} < \infty$. Assume $\|\bar{x}^{(r)}\|_{\mathcal{H}(S)} < \infty$ for $r \leq R$.

- (a) If x_1, x_2 are two functions such that the R feature maps $x_i^{(r)}$ have rates $(\beta_0, \beta_1, \dots, \beta_{R-1})$, then

$$\|\Phi_\Delta(x_1; \mathcal{M}_N) - \Phi_\Delta(x_2; \mathcal{M}_N)\|^2 \leq C(\beta) \|x_1 - x_2\|^{h(\beta)}, \quad (6)$$

with $h(\beta) = \prod_{r=1}^R \frac{\beta_r - 1}{\beta_r - 1/2}$, and where $C(\beta)$ does not depend upon N .

- (b) If τ is a smooth deformation field, then $\|\Phi_\Delta(x; \mathcal{M}_N) - \Phi_\Delta(x; \tau(\mathcal{M}_N))\| \leq C \|\nabla \tau\|^{h(\beta)}$, where C does not depend upon N .
- (c) Let \mathcal{M} and \mathcal{M}' be N -point discretizations of S , If $\max(d(\mathcal{M}, S), d(\mathcal{M}', S)) \leq \epsilon$, then $\|\Phi_\Delta(\mathcal{M}; x) - \Phi_\Delta(\mathcal{M}', x')\| \leq C \epsilon^{h(\beta)}$, where C is independent of N .

This result ensures that if we use as generator of the SN an operator that is consistent as the mesh resolution increases, the resulting surface representation is also consistent. Although our present result only concerns the Laplacian, the Dirac operator also has a well-defined continuous counterpart [8] that generalizes the gradient operator in quaternion space. Also, our current bounds depend explicitly upon the smoothness of feature maps across different layers, which may be controlled in terms of the original signal if one considers nonlinearities that demodulate the signal, such as $\rho(x) = |x|$ or $\rho(x) = \text{ReLU}(x)$. These extensions are left for future work. Finally, a specific setup that we use in experiments is to use as input signal the canonical coordinates of the mesh \mathcal{M} . In that case, an immediate application of the previous theorem yields



Figure 1. Height-Field Representation of surfaces. A 3D mesh $\mathcal{M} \subset \mathbb{R}^3$ (right) is expressed in terms of a “sampling” 2D irregular mesh $\tilde{\mathcal{M}} \subset \mathbb{R}^2$ (left) and a depth scalar field $f : \tilde{\mathcal{M}} \rightarrow \mathbb{R}$ over $\tilde{\mathcal{M}}$ (center).

Corollary 4.3 Denote $\Phi(\mathcal{M}) := \Phi_{\mathcal{M}}(V)$, where V are the node coordinates of \mathcal{M} . Then, if $A_1 = 0$,

$$\|\Phi(\mathcal{M}) - \Phi(\tau(\mathcal{M}))\| \leq \kappa \max(\|\nabla \tau\|_\infty, \|\nabla^2 \tau\|)^{h(\beta)}. \quad (7)$$

5. Generative Surface Models

State-of-the-art generative models for images, such as generative adversarial networks [32], pixel autoregressive networks [28], or variational autoencoders [21], exploit the locality and stationarity of natural images in their probabilistic models, in the sense that the model satisfies $p_\theta(x) \approx p_\theta(x_\tau)$ by construction, where x_τ is a small deformation of a given input x . This property is obtained via encoders and decoders with a deep convolutional structure. We intend to exploit similar geometric stability priors with SNs, owing to their stability properties described in Section 4. A mesh generative model contains two distinct sources of randomness: on the one hand, the randomness associated with the underlying continuous surface, which corresponds to shape variability; on the other hand, the randomness of the discretization of the surface. Whereas the former contains the essential semantic information, the latter is not informative, and to some extent independent of the shape identity. We focus initially on meshes that can be represented as a depth map over an (irregular) 2D mesh, referred as *height-field* meshes in the literature. That is, a mesh $\mathcal{M} = (V, E, F)$ is expressed as $(\tilde{\mathcal{M}}, f(\tilde{\mathcal{M}}))$, where $\tilde{\mathcal{M}} = (\tilde{V}, \tilde{E}, \tilde{F})$ is now a 2D mesh and $f : \tilde{V} \rightarrow \mathbb{R}$ is a *depth-map* encoding the original node locations V , as shown in Figure 1.

In this work, we consider the variational autoencoder framework [21, 33]. It considers a mixture model of the form $p(\mathcal{M}) = \int p_\theta(\mathcal{M} | h) p_0(h) dh$, where $h \in \mathbb{R}^{|S|}$ is a vector of latent variables. We train this model by optimizing the variational lower bound of the data log-likelihood:

$$\min_{\theta, \psi} \frac{1}{L} \sum_{l \leq L} -\mathbb{E}_{h \sim q_\psi(h | \mathcal{M}_l)} \log p_\theta(\mathcal{M}_l | h) + D_{KL}(q_\psi(h | \mathcal{M}_l) || p_0(h)). \quad (8)$$

We thus need to specify a conditional generative model $p_\theta(\mathcal{M} | h)$, a prior distribution $p_0(h)$ and a variational approximation to the posterior $q_\psi(h | \mathcal{M})$, where θ and ψ denote respectively generative and variational trainable parameters. Based on the height-field representation,

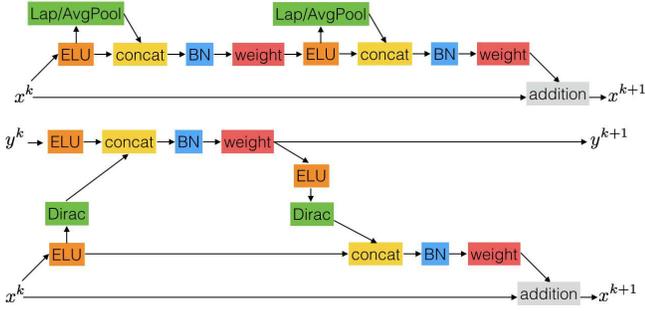


Figure 2. A single ResNet-v2 block used for Laplace, Average Pooling (top) and Dirac models (bottom). The green boxes correspond to the linear operators replacing convolutions in regular domains. We consider Exponential Linear Units (ELU) activations (orange), Batch Normalization (blue) and ‘ 1×1 ’ convolutions (red) containing the trainable parameters; see Eqs (1, 3 and 4). We slightly abuse language and denote by x^{k+1} the output of this 2-layer block.

we choose for simplicity a separable model of the form $p_\theta(\mathcal{M} | h) = p_\theta(f | h, \tilde{\mathcal{M}}) \cdot p(\tilde{\mathcal{M}})$, where $\tilde{\mathcal{M}} \sim p(\tilde{\mathcal{M}})$ is a homogeneous Poisson point process, and $f \sim p_\theta(f | h, \tilde{\mathcal{M}})$ is a normal distribution with mean and isotropic covariance parameters given by a SN:

$$p_\theta(f | h, \tilde{\mathcal{M}}) = \mathcal{N}(\mu(h, \tilde{\mathcal{M}}), \sigma^2(h, \tilde{\mathcal{M}})\mathbf{1}),$$

with $[\mu(h, \tilde{\mathcal{M}}), \sigma^2(h, \tilde{\mathcal{M}})] = \Phi_D(\tilde{\mathcal{M}}; h)$. The generation step thus proceeds as follows. We first sample a 2D mesh $\tilde{\mathcal{M}}$ independent of the latent variable h , and then sample a depth field over $\tilde{\mathcal{M}}$ conditioned on h from the output of a decoder network $\Phi_D(\tilde{\mathcal{M}}; h)$. Finally, the variational family q_ψ is also a Normal distribution whose parameters are obtained from an encoder Surface Neural Network whose last layer is a global pooling that removes the spatial localization: $q_\psi(h | \mathcal{M}) = \mathcal{N}(\bar{\mu}, \bar{\sigma}^2\mathbf{1})$, with $[\bar{\mu}, \bar{\sigma}] = \bar{\Phi}_D(\mathcal{M})$.

6. Experiments

For experimental evaluation, we compare models built using ResNet-v2 blocks [14], where convolutions are replaced with the appropriate operators (see Fig. 2): (i) a point cloud based model from [37] that aggregates global information by averaging features in the intermediate layers and distributing them to all nodes; (ii) a Laplacian Surface network with input canonical coordinates; (iii) a Dirac Surface Network model. We report experiments on generative models using an unstructured variant of MNIST digits (Section 6.1), and on temporal prediction under non-rigid deformation models (Section 6.2).

6.1. MeshMNIST

For this task, we construct a MeshMNIST database with only *height-field* meshes (Sec. 5). First, we sample points on a 2D plane ($[0, 27] \times [0, 27]$) with Poisson disk sampling with $r = 1.0$, which roughly generates 500 points,

and apply a Delaunay triangulation to these points. We then overlay the triangulation with the original MNIST images and assign to each point a z coordinate bilinearly interpolating the grey-scale value. Thus, the procedure allows us to define a sampling process over 3D height-field meshes.

We used VAE models with decoders and encoders built using 10 ResNet-v2 blocks with 128 features. The encoder converts a mesh into a latent vector by averaging output of the last ResNet-v2 block and applying linear transformations to obtain mean and variance, while the decoder takes a latent vector and a 2D mesh as input (corresponding to a specific 3D mesh) and predicts offsets for the corresponding locations. We keep variance of the decoder as a trainable parameter that does not depend on input data. We trained the model for 75 epochs using Adam optimizer [20] with learning rate 10^{-3} , weight decay 10^{-5} and batch size 32. Figures 3,4 illustrate samples from the model. The geometric encoder is able to leverage the local translation invariance of the data despite the irregular sampling, whereas the geometric decoder automatically adapts to the specific sampled grid, as opposed to set-based generative models.

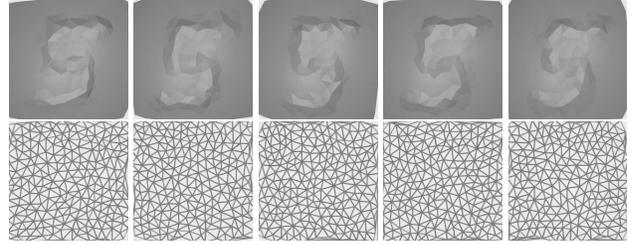


Figure 3. Samples generated for the same latent variable and different triangulations. The learned representation is independent of discretization/triangulation (Poisson disk sampling with $p=1.5$).

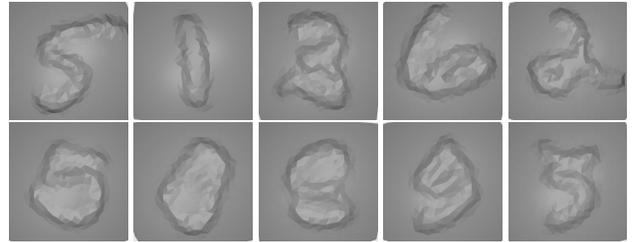


Figure 4. Meshes from the dataset (first five). And meshes generated by our model (last five).

6.2. Spatio-Temporal Predictions

One specific task we consider is temporal predictions of non-linear dynamics. Given a sequence of frames $X = X^1, X^2, \dots, X^n$, the task is to predict the following frames $Y = Y^1 = X^{n+1}, Y^2, \dots, Y^m = X^{n+m}$. As in [26], we use a simple non-recurrent model that takes a concatenation of input frames X and predicts a concatenation of frames Y . We condition on $n = 2$ frames and predict the next

Model	Receptive field	Number of parameters	Smooth L1-loss (mean per sequence (std))
MLP	1	519672	64.56 (0.62)
PointCloud	-	1018872	23.64 (0.21)
Laplace	16	1018872	17.34 (0.52)
Dirac	8	1018872	16.84 (0.16)

Table 1. Evaluation of different models on the temporal task

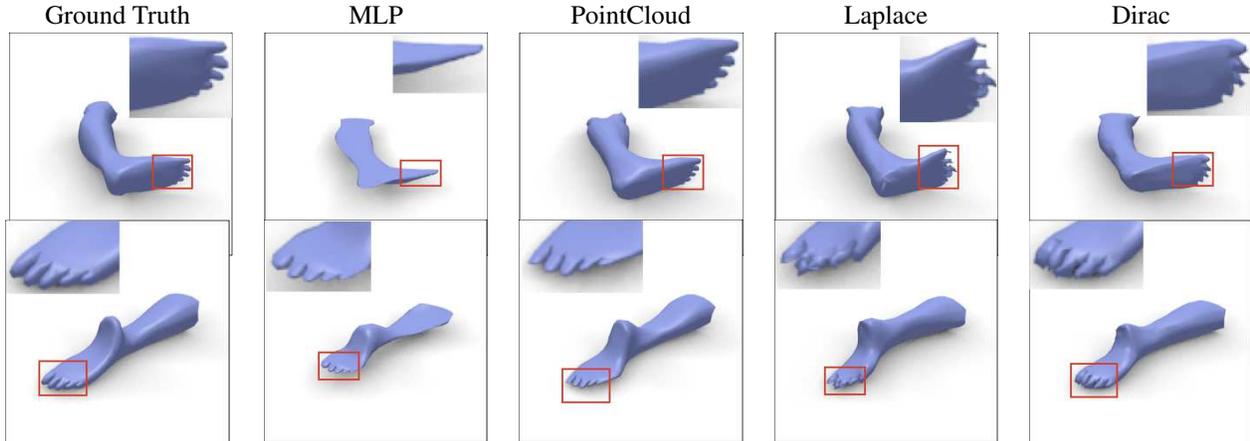


Figure 5. Qualitative comparison of different models. We plot 30th predicted frames correspondingly for two sequences in the test set. Boxes indicate distinctive features. For larger crops, see Figure 6

$m = 40$ frames. In order to generate data, we first extracted 10k patches from the MPI-Faust dataset[3], by selecting a random point and growing a topological sphere of radius 15 edges (i.e. the 15-ring of the point). For each patch, we generate a sequence of 50 frames by randomly rotating it and letting it fall to the ground. We consider the mesh a thin elastic shell, and we simulate it using the As-Rigid-As-Possible technique [35], with additional gravitational forces [17]. Libigl [18] has been used for the mesh processing tasks. Sequences with patches from the first 80 subjects were used in training, while the 20 last subjects were used for testing. The dataset and the code are available on request. We restrict our experiments to temporal prediction tasks that are deterministic when conditioned on several initial frames. Thus, we can train models by minimizing smooth-L1 loss [13] between target frames and output of our models.

We used models with 15 ResNet-v2 blocks with 128 output features each. In order to cover larger context for Dirac and Laplace based models, we alternate these blocks with Average Pooling blocks. We predict offsets to the last conditioned frame and use the corresponding Laplace and Dirac operators. Thus, the models take 6-dimensional inputs and produce 120-dimensional outputs. We trained all models using the Adam optimizer [20] with learning rate 10^{-3} , weight decay 10^{-5} , and batch size 32. After 60k steps we decreased the learning rate by a factor of 2 every 10k steps. The models were trained for 110k steps in overall.

Table 1 reports quantitative prediction performance of different models, and Figure 5 displays samples from the

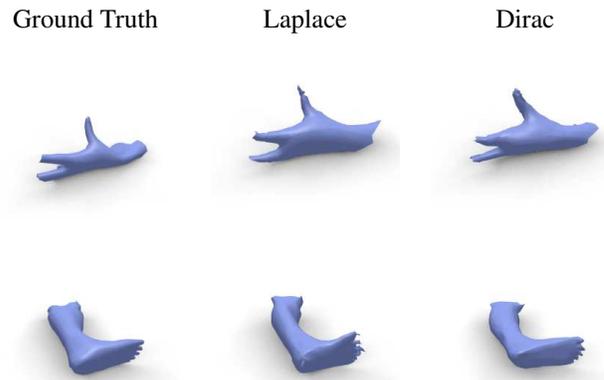


Figure 6. Dirac-based model visually outperforms Laplace-based models in the regions of high mean curvature.

prediction models at specific frames. The set-to-set model [38, 37], corresponding to a point-cloud representation used also in [30], already performs reasonably well on the task, even if the visual difference is noticeable. Nevertheless, the gap between this model and Laplace-/Dirac-based models is significant, both visually and quantitatively. Dirac-based model outperforms Laplace-based model despite the smaller receptive field. Videos comparing the performance of different models are available in the additional material.

Figure 6 illustrates the effect of replacing Laplace by Dirac in the formulation of the SN. Laplacian-based models, since they propagate information using an isotropic operator, have more difficulties at resolving corners and pointy

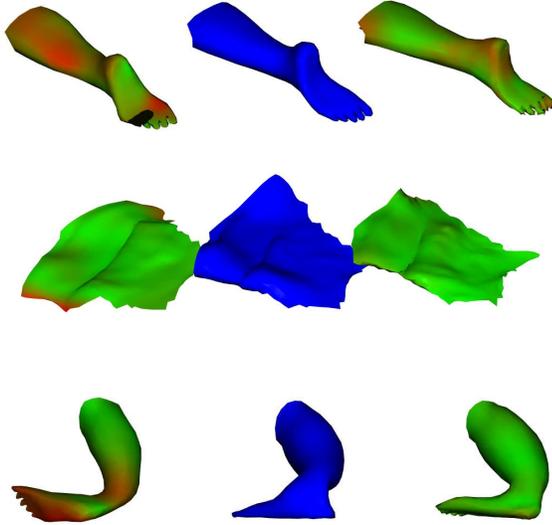


Figure 7. From left to right: PointCloud (set2set), ground truth and Dirac based model. Color corresponds to mean squared error between ground truth and prediction: green - smaller error, red - larger error.

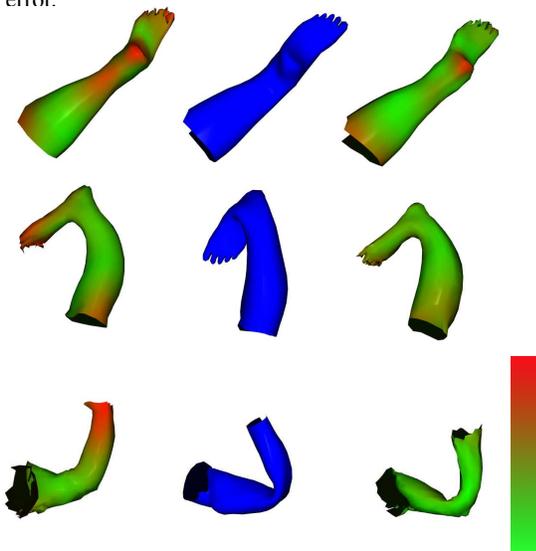


Figure 8. From left to right: Laplace, ground truth and Dirac based model. Color corresponds to mean squared error between ground truth and prediction: green - smaller error, red - larger error.

structures than the Dirac operator, that is sensitive to principal curvature directions. However, the capacity of Laplace models to exploit the extrinsic information only via the input coordinates is remarkable and more computationally efficient than the Dirac counterpart. Figures 7 and 8 overlay the prediction error and compare Laplace against Dirac and PointCloud against Dirac respectively. They confirm first that SNs outperform the point-cloud based model, which often produce excessive flattening and large deformations, and next that first-order Dirac operators help resolve areas with high directional curvature. We refer to the supplementary material for additional qualitative results.

7. Conclusions

We have introduced Surface Networks, a deep neural network that is designed to naturally exploit the non-Euclidean geometry of surfaces. We have shown how a first-order differential operator (the Dirac operator) can detect and adapt to geometric features beyond the local mean curvature, the limit of what Laplacian-based methods can exploit. This distinction is important in practice, since areas with high directional curvature are perceptually important, as shown in the experiments. That said, the Dirac operator comes at increased computational cost due to the quaternion calculus, and it would be interesting to instead *learn* the operator, akin to recent Message-Passing NNs [12] and explore whether Dirac is recovered.

Whenever the data contains good-quality meshes, our experiments demonstrate that using intrinsic geometry offers vastly superior performance to point-cloud based models. While there are not many such datasets currently available, we expect them to become common in the next years, as scanning and reconstruction technology advances and 3D sensors are integrated in consumer devices. SNs provide efficient inference, with predictable runtime, which makes them appealing across many areas of computer graphics, where a fixed, per-frame cost is required to ensure a stable framerate, especially in VR applications. Our future plans include applying Surface Networks precisely to having automated, data-driven mesh processing, and generalizing the generative model to arbitrary meshes, which will require an appropriate multi-resolution pipeline.

References

- [1] M. Andreux, E. Rodolà, M. Aubry, and D. Cremers. Anisotropic Laplace-Beltrami operators for shape analysis. In *Proc. NORDIA*, 2014. 3
- [2] P. Battaglia, R. Pascanu, M. Lai, D. J. Rezende, et al. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*, pages 4502–4510, 2016. 2
- [3] F. Bogo, J. Romero, M. Loper, and M. J. Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014. 7
- [4] D. Boscaini, J. Masci, E. Rodolà, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 3189–3197, 2016. 2, 3
- [5] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *arXiv preprint arXiv:1611.08097*, 2016. 2, 3
- [6] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. *Proc. ICLR*, 2013. 2

- [7] M. B. Chang, T. Ullman, A. Torralba, and J. B. Tenenbaum. A compositional object-based approach to learning physical dynamics. *ICLR*, 2016. 2
- [8] K. Crane, U. Pinkall, and P. Schröder. Spin transformations of discrete surfaces. In *ACM Transactions on Graphics (TOG)*. ACM, 2011. 2, 3, 5
- [9] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3837–3845, 2016. 2, 3
- [10] D. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gómez-Bombarelli, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Neural Information Processing Systems*, 2015. 2
- [11] H. Fan, H. Su, and L. Guibas. A point set generation network for 3d object reconstruction from a single image. *arXiv preprint arXiv:1612.00603*, 2016. 2
- [12] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017. 2, 4, 8
- [13] R. Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1440–1448, 2015. 7
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016. 6
- [15] M. Henaff, J. Bruna, and Y. LeCun. Deep convolutional networks on graph-structured data. *arXiv:1506.05163*, 2015. 2
- [16] K. C. Hsueh-Ti Derek Liu, Alec Jacobson. A dirac operator for extrinsic shape analysis. *Computer Graphics Forum*, 2017. 2, 3
- [17] A. Jacobson. *Algorithms and Interfaces for Real-Time Deformation of 2D and 3D Shapes*. PhD thesis, ETH, Zürich, 2013. 7
- [18] A. Jacobson, D. Panozzo, et al. libigl: A simple C++ geometry processing library, 2016. <http://libigl.github.io/libigl/>. 7
- [19] S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 2016. 2
- [20] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representation*, 2015. 6, 7
- [21] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2, 5
- [22] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 2, 3
- [23] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015. 2
- [24] H. Maron, M. Galun, N. Aigerman, M. Trope, N. Dym, E. Yumer, V. Kim, and Y. Lipman. Convolutional neural networks on surfaces via seamless toric covers. In *SIGGRAPH*, 2017. 2
- [25] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45, 2015. 2
- [26] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. 6
- [27] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. *arXiv preprint arXiv:1611.08402*, 2016. 2, 3
- [28] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016. 5
- [29] R. Poranne and Y. Lipman. Simple approximations of planar deformation operators. Technical report, ETHZ, 2015. 2
- [30] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *arXiv preprint arXiv:1612.00593*, 2016. 1, 2, 7
- [31] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. 1, 2
- [32] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 5
- [33] D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015. 2, 5
- [34] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009. 2, 3
- [35] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, 2007. 7
- [36] D. A. Spielman. Spectral graph theory and its applications. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 29–38. IEEE, 2007. 5
- [37] S. Sukhbaatar, R. Fergus, et al. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, pages 2244–2252, 2016. 1, 2, 6, 7
- [38] O. Vinyals, S. Bengio, and M. Kudlur. Order matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391*, 2015. 1, 2, 7
- [39] M. Wardetzky. Convergence of the cotangent formula: An overview. In *Discrete Differential Geometry*, pages 275–286. 2008. 4
- [40] L. Wei, Q. Huang, D. Ceylan, E. Vouga, and H. Li. Dense human body correspondences using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1544–1553, 2016. 2
- [41] H. Weyl. Über die asymptotische verteilung der eigenwerte. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1911:110–117, 1911. 5
- [42] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1912–1920, 2015. 2

Supplementary Material for Surface Networks

Ilya Kostrikov, Zhongshi Jiang, Daniele Panozzo, Denis Zorin, Joan Bruna
New York University

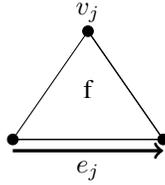
March 29, 2018

Abstract

This note contains the appendices of the paper Surface Networks.

A The Dirac Operator

The quaternions \mathbb{H} is an extension of complex numbers. A quaternion $q \in \mathbb{H}$ can be represented in a form $q = a + bi + cj + dk$ where a, b, c, d are real numbers and i, j, k are quaternion units that satisfy the relationship $i^2 = j^2 = k^2 = ijk = -1$.



As mentioned in Section 3.1, the Dirac operator used in the model can be conveniently represented as a quaternion matrix:

$$D_{f,j} = \frac{-1}{2|\mathcal{A}_f|} e_j, \quad f \in F, j \in V,$$

where e_j is the opposing edge vector of node j in the face f , and \mathcal{A}_f is the area, as illustrated in Fig. A, using counter-clockwise orientations on all faces.

The Deep Learning library PyTorch that we used to implement the models does not support quaternions. Nevertheless, quaternion-valued matrix multiplication can be replaced with real-valued matrix multiplication where each entry $q = a + bi + cj + dk$ is represented as a 4×4 block

$$\begin{bmatrix} a & -b & -c & -d \\ b & a & -d & c \\ c & d & a & -b \\ d & -c & b & a \end{bmatrix}$$

and the conjugate $q^* = a - bi - cj - dk$ is a transpose of this real-valued matrix:

$$\begin{bmatrix} a & b & c & d \\ -b & a & d & -c \\ -c & -d & a & b \\ -d & c & -b & a \end{bmatrix}.$$

B Theorem 4.1

B.1 Proof of (a)

We first show the result for the mapping $x \mapsto \rho(Ax + B\Delta x)$, corresponding to one layer of Φ_Δ . By definition, the Laplacian Δ of \mathcal{M} is

$$\Delta = \text{diag}(\bar{\mathcal{A}})^{-1}(U - W),$$

where $\bar{\mathcal{A}}_j$ is one third of the total area of triangles incident to node j , and $W = (w_{i,j})$ contains the cotangent weights [9], and $U = \text{diag}(W\mathbf{1})$ contains the node aggregated weights in its diagonal.

From [4] we verify that

$$\begin{aligned} \|U - W\| &\leq \sqrt{2} \max_i \left\{ \sqrt{U_i^2 + U_i \sum_{i \sim j} U_j w_{i,j}} \right\} \\ &\leq 2\sqrt{2} \sup_{i,j} w_{i,j} \sup_j d_j \\ &\leq 2\sqrt{2} \cot(\alpha_{\min}) d_{\max}, \end{aligned} \quad (1)$$

where d_j denotes the degree (number of neighbors) of node j , α_{\min} is the smallest angle in the triangulation of \mathcal{M} and S_{\max} the largest number of incident triangles. It results that

$$\|\Delta\| \leq C \frac{\cot(\alpha_{\min}) S_{\max}}{\inf_j \bar{\mathcal{A}}_j} := L_{\mathcal{M}},$$

which depends uniquely on the mesh \mathcal{M} and is finite for non-degenerate meshes. Moreover, since $\rho(\cdot)$ is non-expansive, we have

$$\begin{aligned} \|\rho(Ax + B\Delta x) - \rho(Ax' + B\Delta x')\| &\leq \|A(x - x') + B\Delta(x - x')\| \\ &\leq (\|A\| + \|B\|L_{\mathcal{M}}) \|x - x'\|. \end{aligned} \quad (2)$$

By cascading (2) across the K layers of the network, we obtain

$$\|\Phi(\mathcal{M}; x) - \Phi(\mathcal{M}; x')\| \leq \left(\prod_{k \leq K} (\|A_k\| + \|B_k\|L_{\mathcal{M}}) \right) \|x - x'\|,$$

which proves (a). \square

B.2 Proof of (b)

The proof is analogous, by observing that $\|D\| = \sqrt{\|\Delta\|}$ and therefore

$$\|D\| \leq \sqrt{L_{\mathcal{M}}}. \quad \square$$

B.3 Proof of (c)

To establish (c) we first observe that given three points $p, q, r \in \mathbb{R}^3$ forming any of the triangles of \mathcal{M} ,

$$\begin{aligned} \|p - q\|^2 (1 - |\nabla\tau|_\infty)^2 &\leq \|\tau(p) - \tau(q)\|^2 \leq \|p - q\|^2 (1 + |\nabla\tau|_\infty)^2 \\ \mathcal{A}(p, q, r)^2 (1 - |\nabla\tau|_\infty C \alpha_{\min}^{-2} - o(|\nabla\tau|_\infty^2)) &\leq \mathcal{A}(\tau(p), \tau(q), \tau(r))^2 \leq \mathcal{A}(p, q, r)^2 (1 + |\nabla\tau|_\infty C \alpha_{\min}^{-2} + o(|\nabla\tau|_\infty^2)) \end{aligned} \quad (3)$$

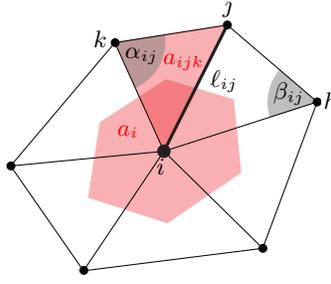


Figure 1: Triangular mesh and Cotangent Laplacian (figure reproduced from [2])

Indeed, (3) is a direct consequence of the lower and upper Lipschitz constants of $\tau(u)$, which are bounded respectively by $1 - |\nabla\tau|_\infty$ and $1 + |\nabla\tau|_\infty$. As for (4), we use the Heron formula

$$\mathcal{A}(p, q, r)^2 = s(s - \|p - q\|)(s - \|p - r\|)(s - \|r - q\|),$$

with $s = \frac{1}{2}(\|p - q\| + \|p - r\| + \|r - q\|)$ being the half-perimeter. By denoting s_τ the corresponding half-perimeter determined by the deformed points $\tau(p), \tau(q), \tau(r)$, we have that

$$s_\tau - \|\tau(p) - \tau(q)\| \leq s(1 + |\nabla\tau|_\infty) - \|p - q\|(1 - |\nabla\tau|_\infty) = s - \|p - q\| + |\nabla\tau|_\infty(s + \|p - q\|) \text{ and}$$

$$s_\tau - \|\tau(p) - \tau(q)\| \geq s(1 - |\nabla\tau|_\infty) - \|p - q\|(1 + |\nabla\tau|_\infty) = s - \|p - q\| - |\nabla\tau|_\infty(s + \|p - q\|),$$

and similarly for the $\|r - q\|$ and $\|r - p\|$ terms. It results in

$$\begin{aligned} \mathcal{A}(\tau(p), \tau(q), \tau(r))^2 &\geq \mathcal{A}(p, q, r)^2 \left[1 - |\nabla\tau|_\infty \left(1 + \frac{s + \|p - q\|}{s - \|p - q\|} + \frac{s + \|p - r\|}{s - \|p - r\|} + \frac{s + \|r - q\|}{s - \|r - q\|} \right) - o(|\nabla\tau|_\infty^2) \right] \\ &\geq \mathcal{A}(p, q, r)^2 \left[1 - C|\nabla\tau|_\infty \alpha_{\min}^{-2} - o(|\nabla\tau|_\infty^2) \right], \end{aligned}$$

and similarly

$$\mathcal{A}(\tau(p), \tau(q), \tau(r))^2 \leq \mathcal{A}(p, q, r)^2 \left[1 + C|\nabla\tau|_\infty \alpha_{\min}^{-2} - o(|\nabla\tau|_\infty^2) \right].$$

By noting that the cotangent Laplacian weights can be written (see Fig. 1) as

$$w_{i,j} = \frac{-\ell_{ij}^2 + \ell_{jk}^2 + \ell_{ik}^2}{\mathcal{A}(i, j, k)} + \frac{-\ell_{ij}^2 + \ell_{jh}^2 + \ell_{ih}^2}{\mathcal{A}(i, j, h)},$$

we have from the previous Bilipschitz bounds that

$$\tau(w_{i,j}) \leq w_{i,j} [1 - C|\nabla\tau|_\infty \alpha_{\min}^{-2}]^{-1} + 2|\nabla\tau|_\infty [1 - C|\nabla\tau|_\infty \alpha_{\min}^{-2}]^{-1} \left(\frac{\ell_{ij}^2 + \ell_{jk}^2 + \ell_{ik}^2}{\mathcal{A}(i, j, k)} + \frac{\ell_{ij}^2 + \ell_{jh}^2 + \ell_{ih}^2}{\mathcal{A}(i, j, h)} \right),$$

$$\tau(w_{i,j}) \geq w_{i,j} [1 + C|\nabla\tau|_\infty \alpha_{\min}^{-2}]^{-1} - 2|\nabla\tau|_\infty [1 + C|\nabla\tau|_\infty \alpha_{\min}^{-2}]^{-1} \left(\frac{\ell_{ij}^2 + \ell_{jk}^2 + \ell_{ik}^2}{\mathcal{A}(i, j, k)} + \frac{\ell_{ij}^2 + \ell_{jh}^2 + \ell_{ih}^2}{\mathcal{A}(i, j, h)} \right),$$

which proves that, up to second order terms, the cotangent weights are Lipschitz continuous to deformations.

Finally, since the mesh Laplacian operator is constructed as $\text{diag}(\bar{\mathcal{A}})^{-1}(U - W)$, with $\bar{\mathcal{A}}_{i,i} = \frac{1}{3} \sum_{j,k:(i,j,k) \in F} \mathcal{A}(i,j,k)$, and $U = \text{diag}(W\mathbf{1})$, let us show how to bound $\|\Delta - \tau(\Delta)\|$ from

$$\bar{\mathcal{A}}_{i,i}(1 - \alpha_{\mathcal{M}}|\nabla\tau|_{\infty} - o(|\nabla\tau|_{\infty}^2)) \leq \tau(\bar{\mathcal{A}}_{i,i}) \leq \bar{\mathcal{A}}_{i,i}(1 + \alpha_{\mathcal{M}}|\nabla\tau|_{\infty} + o(|\nabla\tau|_{\infty}^2)) \quad (5)$$

and

$$w_{i,j}(1 - \beta_{\mathcal{M}}|\nabla\tau|_{\infty} - o(|\nabla\tau|_{\infty}^2)) \leq \tau(w_{i,j}) \leq w_{i,j}(1 + \beta_{\mathcal{M}}|\nabla\tau|_{\infty} + o(|\nabla\tau|_{\infty}^2)). \quad (6)$$

Using the fact that $\bar{\mathcal{A}}, \tau(\bar{\mathcal{A}})$ are diagonal, and using the spectral bound for $k \times m$ sparse matrices from [3], Lemma 5.12,

$$\|Y\|^2 \leq \max_i \sum_{j: Y_{i,j} \neq 0} |Y_{i,j}| \left(\sum_{r=1}^l |Y_{r,j}| \right),$$

the bounds (5) and (6) yield respectively

$$\begin{aligned} \tau(\bar{\mathcal{A}}) &= \bar{\mathcal{A}}(\mathbf{1} + \epsilon_{\tau}), \text{ with } \|\epsilon_{\tau}\| = o(|\nabla\tau|_{\infty}), \text{ and} \\ \tau(U - W) &= U - W + \eta_{\tau}, \text{ with } \|\eta_{\tau}\| = o(|\nabla\tau|_{\infty}). \end{aligned}$$

It results that, up to second order terms,

$$\begin{aligned} \|\Delta - \tau(\Delta)\| &= \|\tau(\bar{\mathcal{A}})^{-1}(\tau(U) - \tau(W)) - \bar{\mathcal{A}}^{-1}(U - W)\| \\ &= \left\| (\bar{\mathcal{A}}[\mathbf{1} + \epsilon_{\tau}])^{-1} [U - W + \eta_{\tau}] - \bar{\mathcal{A}}^{-1}(U - W) \right\| \\ &= \left\| (\mathbf{1} - \epsilon_{\tau} + o(|\nabla\tau|_{\infty}^2)) \bar{\mathcal{A}}^{-1}(U - W + \eta_{\tau}) - \bar{\mathcal{A}}^{-1}(U - W) \right\| \\ &= \|\epsilon_{\tau}\Delta + \bar{\mathcal{A}}^{-1}\eta_{\tau}\| + o(|\nabla\tau|_{\infty}^2) \\ &= o(|\tau|_{\infty}), \end{aligned}$$

which shows that the Laplacian is stable to deformations in operator norm. Finally, by denoting \tilde{x}_{τ} a layer of the deformed Laplacian network

$$\tilde{x}_{\tau} = \rho(Ax + B\tau(\Delta)x),$$

it follows that

$$\|\tilde{x} - \tilde{x}_{\tau}\| \leq \|B(\Delta - \tau(\Delta))x\| \quad (7)$$

$$\leq C\|B\|\|\nabla\tau\|_{\infty}\|x\|. \quad (8)$$

Also,

$$\begin{aligned} \|\tilde{x} - \tilde{y}_{\tau}\| &\leq \|A(x - y) + B(\Delta x - \tau(\Delta)y)\| \\ &\leq (\|A\| + \|B\|\|\Delta\|)\|x - y\| + \|\Delta - \tau(\Delta)\|\|x\| \\ &\leq \underbrace{(\|A\| + \|B\|\|\Delta\|)}_{\delta_1} \|x - y\| + \underbrace{C\|\nabla\tau\|_{\infty}}_{\delta_2} \|x\|, \end{aligned} \quad (9)$$

and therefore, by plugging (9) with $y = \tilde{x}_{\tau}$, K layers of the Laplacian network satisfy

$$\begin{aligned} \|\Phi(x; \Delta) - \Phi(x; \tau(\Delta))\| &\leq \left(\prod_{j \leq K-1} \delta_1(j) \right) \|\tilde{x} - \tilde{x}_{\tau}\| + \left(\sum_{j < K-1} \prod_{j' \leq j} \delta_1(j') \delta_2(j) \right) \|\nabla\tau\|_{\infty} \|x\| \\ &\leq \left[C \left(\prod_{j \leq K-1} \delta_1(j) \right) \|B\| + \left(\sum_{j < K-1} \prod_{j' \leq j} \delta_1(j') \delta_2(j) \right) \right] \|\nabla\tau\|_{\infty} \|x\|. \quad \square. \end{aligned}$$

B.4 Proof of (d)

The proof is also analogous to the proof of (c), with the difference that now the Dirac operator is no longer invariant to orthogonal transformations, only to translations. Given two points p, q , we verify that

$$\|p - q - \tau(p) - \tau(q)\| \leq \widetilde{|\tau|}_\infty \|p - q\| ,$$

which, following the previous argument, leads to

$$\|D - \tau(D)\| = o(\widetilde{|\tau|}_\infty) . \quad (10)$$

C Theorem 4.2

C.1 Proof of part (a)

The proof is based on the following lemma:

Lemma C.1 *Let $x_N, y_N \in \mathcal{H}(\mathcal{M}_N)$ such that $\forall N, \|x_N\|_{\mathcal{H}} \leq c, \|y_N\|_{\mathcal{H}} \leq c$. Let $\hat{x}_N = \mathcal{E}_N(x_N)$, where \mathcal{E}_N is the eigendecomposition of the Laplacian operator Δ_N on \mathcal{M}_N , with associated eigenvalues $\lambda_1 \dots \lambda_N$ in increasing order. Let $\gamma > 0$ and β be defined as in (??) for x_N and y_N . If $\beta > 1$ and $\|x_N - y_N\| \leq \epsilon$ for all N ,*

$$\|\Delta_N(x_N - y_N)\|^2 \leq C\epsilon^{2 - \frac{1}{\beta-1/2}} , \quad (11)$$

where C is a constant independent of ϵ and N .

One layer of the network will transform the difference $x_1 - x_2$ into $\rho(Ax_1 + B\Delta x_1) - \rho(Ax_2 + B\Delta x_2)$. We verify that

$$\|\rho(Ax_1 + B\Delta x_1) - \rho(Ax_2 + B\Delta x_2)\| \leq \|A\| \|x_1 - x_2\| + \|B\| \|\Delta(x_1 - x_2)\| .$$

We now apply Lemma C.1 to obtain

$$\begin{aligned} \|\rho(Ax_1 + B\Delta x_1) - \rho(Ax_2 + B\Delta x_2)\| &\leq \|A\| \|x_1 - x_2\| + C \|B\| \|x_1 - x_2\|^{\frac{\beta-1}{\beta-1/2}} \\ &\leq \|x_1 - x_2\|^{\frac{\beta-1}{\beta-1/2}} \left(\|A\| \|x_1 - x_2\|^{(2\beta-1)^{-1}} + C \|B\| \right) \\ &\leq C(\|A\| + \|B\|) \|x_1 - x_2\|^{\frac{\beta-1}{\beta-1/2}} , \end{aligned}$$

where we redefine C to account for the fact that $\|x_1 - x_2\|^{(2\beta-1)^{-1}}$ is bounded. We have just showed that

$$\|x_1^{(r+1)} - x_2^{(r+1)}\| \leq f_r \|x_1^{(r)} - x_2^{(r)}\|^{g_r} \quad (12)$$

with $f_r = C(\|A_r\| + \|B_r\|)$ and $g_r = \frac{\beta_r-1}{\beta_r-1/2}$. By cascading (12) for each of the R layers we thus obtain

$$\|\Phi_\Delta(x_1) - \Phi_\Delta(x_2)\| \leq \left[\prod_{r=1}^R f_r^{\prod_{r' > r} g_{r'}} \right] \|x_1 - x_2\|^{\prod_{r=1}^R g_r} , \quad (13)$$

which proves (??) \square .

Proof of (11): Let $\{e_1, \dots, e_N\}$ be the eigendecomposition of Δ_N . For simplicity, we drop the subindex N in the signals from now on. Let $\hat{x}(k) = \langle x, e_k \rangle$ and $\tilde{x}(k) = \lambda_k \hat{x}(k)$; and analogously for y . From the Parseval identity we have that $\|x\|^2 = \|\hat{x}\|^2$. We express $\|\Delta(x - y)\|$ as

$$\|\Delta(x - y)\|^2 = \sum_{k \leq N} \lambda_k^2 (\hat{x}(k) - \hat{y}(k))^2. \quad (14)$$

The basic principle of the proof is to cut the spectral sum (14) in two parts, chosen to exploit the decay of $\tilde{x}(k)$. Let

$$F(x)(k) = \frac{\sum_{k' \geq k} \tilde{x}(k')^2}{\|x\|_{\mathcal{H}}^2} = \frac{\sum_{k' \geq k} \tilde{x}(k')^2}{\sum_{k'} \tilde{x}(k')^2} = \frac{\sum_{k' \geq k} \lambda_{k'}^2 \hat{x}(k')^2}{\sum_{k'} \lambda_{k'}^2 \hat{x}(k')^2} \leq 1,$$

and analogously for y . For any cutoff $k_* \leq N$ we have

$$\begin{aligned} \|\Delta(x - y)\|^2 &= \sum_{k \leq k_*} \lambda_k^2 (\hat{x}(k) - \hat{y}(k))^2 + \sum_{k > k_*} \lambda_k^2 (\hat{x}(k) - \hat{y}(k))^2 \\ &\leq \lambda_{k_*}^2 \epsilon^2 + 2(F(x)(k_*) \|x\|_{\mathcal{H}}^2 + F(y)(k_*) \|y\|_{\mathcal{H}}^2) \\ &\leq \lambda_{k_*}^2 \epsilon^2 + 2F(k_*) (\|x\|_{\mathcal{H}}^2 + \|y\|_{\mathcal{H}}^2) \\ &\leq \lambda_{k_*}^2 \epsilon^2 + 4F(k_*) D^2, \end{aligned} \quad (15)$$

where we denote for simplicity $F(k_*) = \max(F(x)(k_*), F(y)(k_*))$. By assumption, we have $\lambda_k^2 \lesssim k^{2\gamma}$ and

$$F(k) \lesssim \sum_{k' \geq k} k'^{2(\gamma-\beta)} \simeq k^{1+2(\gamma-\beta)}.$$

By denoting $\tilde{\beta} = \beta - \gamma - 1/2$, it follows that

$$\|\Delta(x - y)\|^2 \lesssim \epsilon^2 k_*^{2\gamma} + 4D^2 k_*^{-2\tilde{\beta}} \quad (16)$$

Optimizing for k_* yields

$$\epsilon^2 2\gamma k_*^{2\gamma-1} - 2\tilde{\beta} 4D^2 k_*^{-2\tilde{\beta}-1} = 0, \text{ thus}$$

$$k_* = \left[\frac{4\beta D^2}{\gamma \epsilon^2} \right]^{\frac{1}{2\gamma+2\tilde{\beta}}}. \quad (17)$$

By plugging (17) back into (16) and dropping all constants independent of N and ϵ , this leads to

$$\|\Delta(x - y)\|^2 \lesssim \epsilon^{2-\frac{1}{\gamma+\tilde{\beta}}} = \epsilon^{2-\frac{1}{\beta-1/2}},$$

which proves part (a) \square .

C.2 Proof of part (b)

We will use the following lemma:

Lemma C.2 *Let $\mathcal{M} = (V, E, F)$ is a non-degenerate mesh, and define*

$$\eta_1(\mathcal{M}) = \sup_{(i,j) \in E} \frac{\bar{\mathcal{A}}_i}{\bar{\mathcal{A}}_j}, \quad \eta_2(\mathcal{M}) = \sup_{(i,j,k) \in F} \frac{\ell_{ij}^2 + \ell_{jk}^2 + \ell_{ik}^2}{\mathcal{A}(i,j,k)}, \quad \eta_3(\mathcal{M}) = \alpha_{\min}. \quad (18)$$

Then, given a smooth deformation τ and x defined in \mathcal{M} , we have

$$\|(\Delta - \tau(\Delta))x\| \leq C|\nabla\tau|_\infty\|\Delta x\|, \quad (19)$$

where C depends only upon η_1, η_2 and η_3 .

In that case, we need to control the difference $\rho(Ax + B\Delta x) - \rho(Ax + B\tau(\Delta)x)$. We verify that

$$\|\rho(Ax + B\Delta x) - \rho(Ax + B\tau(\Delta)x)\| \leq \|B\|\|(\Delta - \tau(\Delta))x\|.$$

By Lemma C.2 it follows that $\|(\Delta - \tau(\Delta))x\| \leq C|\nabla\tau|_\infty\|\Delta x\|$ and therefore, by denoting $x_1^{(1)} = \rho(Ax + B\Delta x)$ and $x_2^{(1)} = \rho(Ax + B\tau(\Delta)x)$, we have

$$\|x_1^{(1)} - x_2^{(1)}\| \leq C|\nabla\tau|_\infty\|\Delta x\| = C|\nabla\tau|_\infty\|x\|_{\mathcal{H}}. \quad (20)$$

By applying again Lemma C.1, we also have that

$$\begin{aligned} \|\Delta x_1^{(1)} - \tau(\Delta)x_2^{(1)}\| &= \|\Delta x_1^{(1)} - (\Delta + \tau(\Delta) - \Delta)x_2^{(1)}\| \\ &= \|\Delta(x_1^{(1)} - x_2^{(1)}) + (\tau(\Delta) - \Delta)x_2^{(1)}\| \\ &\leq C\|x_1^{(1)} - x_2^{(1)}\|^{\frac{\beta_1-1}{\beta_1-1/2}} + |\nabla\tau|_\infty\|x_2^{(1)}\|_{\mathcal{H}} \\ &\lesssim C|\nabla\tau|_\infty^{\frac{\beta_1-1}{\beta_1-1/2}}, \end{aligned}$$

which, by combining it with (20) and repeating through the R layers yields

$$\|\Phi_\Delta(x, \mathcal{M}) - \Phi_\Delta(x, \tau(\mathcal{M}))\| \leq C|\nabla\tau|_\infty^{\prod_{r=1}^R \frac{\beta_r-1}{\beta_r-1/2}}, \quad (21)$$

which concludes the proof \square .

Proof of (19): The proof follows closely the proof of Theorem ??, part (c). From (5) and (6) we have that

$$\begin{aligned} \tau(\bar{\mathcal{A}}) &= \bar{\mathcal{A}}(\mathbf{I} + G_\tau), \text{ with } |G_\tau|_\infty \leq C(\eta_2, \eta_3)|\nabla\tau|_\infty, \text{ and} \\ \tau(U - W) &= (\mathbf{I} + H_\tau)(U - W), \text{ with } |H_\tau|_\infty \leq C(\eta_2, \eta_3)|\nabla\tau|_\infty. \end{aligned}$$

It follows that, up to second order $o(|\nabla\tau|_\infty^2)$ terms,

$$\begin{aligned} \tau(\Delta) - \Delta &= \tau(\bar{\mathcal{A}})^{-1}(\tau(U) - \tau(W)) - \bar{\mathcal{A}}^{-1}(U - W) \\ &= (\bar{\mathcal{A}}[\mathbf{1} + G_\tau])^{-1}[(\mathbf{I} + H_\tau)(U - W)] - \bar{\mathcal{A}}^{-1}(U - W) \\ &\simeq \bar{\mathcal{A}}^{-1}H_\tau(U - W) + G_\tau\Delta. \end{aligned} \quad (22)$$

By writing $\bar{\mathcal{A}}^{-1}H_\tau = \widetilde{H}_\tau\bar{\mathcal{A}}^{-1}$, and since $\bar{\mathcal{A}}$ is diagonal, we verify that

$$(\widetilde{H}_\tau)_{i,j} = (H_\tau)_{i,j} \frac{\mathcal{A}_{i,i}}{\mathcal{A}_{j,j}}, \text{ with}$$

$\frac{\mathcal{A}_{i,i}}{\mathcal{A}_{j,j}} \leq \eta_1$, and hence that

$$\bar{\mathcal{A}}^{-1}H_\tau(U - W) = \widetilde{H}_\tau\Delta, \text{ with } |\widetilde{H}_\tau|_\infty \leq C(\eta_1, \eta_2, \eta_3)|\nabla\tau|_\infty. \quad (23)$$

We conclude by combining (22) and (23) into

$$\begin{aligned} \|(\Delta - \tau(\Delta))x\| &= \|(G_\tau + \widetilde{H}_\tau)\Delta x\| \\ &\leq C'(\eta_1, \eta_2, \eta_3)|\nabla\tau|_\infty\|\Delta x\|, \end{aligned}$$

which proves (19) \square

C.3 Proof of part (c)

This result is a consequence of the consistency of the cotangent Laplacian to the Laplace-Beltrami operator on S [9]:

Theorem C.3 ([9], Thm 3.4) *Let \mathcal{M} be a compact polyhedral surface which is a normal graph over a smooth surface S with distortion tensor \mathcal{T} , and let $\bar{\mathcal{T}} = (\det \mathcal{T})^{1/2} \mathcal{T}^{-1}$. If the normal field uniform distance $d(\mathcal{T}, \mathbf{1}) = \|\bar{\mathcal{T}} - \mathbf{1}\|_\infty$ satisfies $d(\mathcal{T}, \mathbf{1}) \leq \epsilon$, then*

$$\|\Delta_{\mathcal{M}} - \Delta_S\| \leq \epsilon. \quad (24)$$

If $\Delta_{\mathcal{M}}$ converges uniformly to Δ_S , in particular we verify that

$$\|x\|_{\mathcal{H}(\mathcal{M})} \rightarrow \|x\|_{\mathcal{H}(S)}.$$

Thus, given two meshes $\mathcal{M}, \mathcal{M}'$ approximating a smooth surface S in terms of uniform normal distance, and the corresponding irregular sampling x and x' of an underlying function $\bar{x} : S \rightarrow \mathbb{R}$, we have

$$\|\rho(Ax + B\Delta_{\mathcal{M}}x) - \rho(Ax' + B\Delta_{\mathcal{M}'}x')\| \leq \|A\|\|x - x'\| + \|B\|\|\Delta_{\mathcal{M}}x - \Delta_{\mathcal{M}'}x'\|. \quad (25)$$

Since \mathcal{M} and \mathcal{M}' both converge uniformly normally to S and \bar{x} is Lipschitz on S , it results that

$$\|x - \bar{x}\| \leq L\epsilon, \text{ and } \|x' - \bar{x}\| \leq L\epsilon,$$

thus $\|x - x'\| \leq 2L\epsilon$. Also, thanks to the uniform normal convergence, we also have convergence in the Sobolev sense:

$$\|x - \bar{x}\|_{\mathcal{H}} \lesssim \epsilon, \quad \|x' - \bar{x}\|_{\mathcal{H}} \lesssim \epsilon,$$

which implies in particular that

$$\|x - x'\|_{\mathcal{H}} \lesssim \epsilon. \quad (26)$$

From (25) and (26) it follows that

$$\begin{aligned} \|\rho(Ax + B\Delta_{\mathcal{M}}x) - \rho(Ax' + B\Delta_{\mathcal{M}'}x')\| &\leq 2\|A\|L\epsilon + \\ &\quad + \|B\|\|\Delta_{\mathcal{M}}x - \Delta_S\bar{x} + \Delta_S\bar{x} - \Delta_{\mathcal{M}'}x'\| \\ &\leq 2\epsilon(\|A\|L + \|B\|). \end{aligned} \quad (27)$$

By applying again Lemma C.1 to $\tilde{x} = \rho(Ax + B\Delta_{\mathcal{M}}x)$, $\tilde{x}' = \rho(Ax' + B\Delta_{\mathcal{M}'}x')$, we have

$$\|\tilde{x} - \tilde{x}'\|_{\mathcal{H}} \leq C\|\tilde{x} - \tilde{x}'\|_{\beta_1^{-1/2}}^{\beta_1-1} \lesssim \epsilon^{\frac{\beta_1-1}{\beta_1-1/2}}.$$

We conclude by retracing the same argument as before, reapplying Lemma C.1 at each layer to obtain

$$\|\Phi_{\mathcal{M}}(x) - \Phi_{\mathcal{M}'}(x')\| \leq C\epsilon^{\prod_{r=1}^R \frac{\beta_r-1}{\beta_r-1/2}}. \quad \square.$$

D Proof of Corollary 4.3

We verify that

$$\begin{aligned}\|\rho(B\Delta x) - \rho(B\tau(\Delta)\tau(x))\| &\leq \|B\|\|\Delta x - \tau(\Delta)\tau(x)\| \\ &\leq \|B\|\|\Delta(x - \tau(x)) + (\Delta - \tau(\Delta))(\tau(x))\| \\ &\leq \|B\|(\|\Delta(x - \tau(x))\| + \|(\Delta - \tau(\Delta))(\tau(x))\|).\end{aligned}$$

The second term is $o(\|\nabla\tau\|_\infty)$ from Lemma C.2. The first term is

$$\|x - \tau(x)\|_{\mathcal{H}} \leq \|\Delta(\mathbf{I} - \tau)\| \|x\| \leq \|\nabla^2\tau\| \|x\|,$$

where $\|\nabla^2\tau\|$ is the uniform Hessian norm of τ . The result follows from applying the cascading argument from last section. \square

E Preliminary Study: Metric Learning for Dense Correspondence

As an interesting extension, we apply the architecture we built in Experiments 6.2 directly to a dense shape correspondence problem.

Similarly as the graph correspondence model from [8], we consider a Siamese Surface Network, consisting of two identical models with the same architecture and sharing parameters. For a pair of input surfaces $\mathcal{M}_1, \mathcal{M}_2$ of N_1, N_2 points respectively, the network produces embeddings $E_1 \in \mathbb{R}^{N_1 \times d}$ and $E_2 \in \mathbb{R}^{N_2 \times d}$. These embeddings define a trainable similarity between points given by

$$s_{i,j} = \frac{e^{\langle E_{1,i}, E_{2,j} \rangle}}{\sum_{j'} e^{\langle E_{1,i}, E_{2,j'} \rangle}}, \quad (28)$$

which can be trained by minimizing the cross-entropy relative to ground truth pairs. A diagram of the architecture is provided in Figure 2.

In general, dense shape correspondence is a task that requires a blend of intrinsic and extrinsic information, motivating the use of data-driven models that can obtain such tradeoffs automatically. Following the setup in Experiment 6.2, we use models with 15 ResNet-v2 blocks with 128 output features each, and alternate Laplace and Dirac based models with Average Pooling blocks to cover a larger context: The input to our network consists of vertex positions only.

We tested our architecture on a reconstructed (i.e. changing the mesh connectivity) version of the real scan of FAUST dataset [1]. The FAUST dataset contains 100 real scans and their corresponding ground truth registrations. The ground truth is based on a deformable template mesh with the same ordering and connectivity, which is fitted to the scans. In order to eliminate the bias of using the same template connectivity, as well as the need of a single connected component, the scans are reconstructed again with [5]. To foster replicability, we release the processed dataset in the additional material. In our experiment, we use 80 models for training and 20 models for testing.

Since the ground truth correspondence is implied only through the common template mesh, we compute the correspondence between our meshes with a nearest neighbor search between the point cloud and the reconstructed mesh. Consequently, due to the drastic change in vertex replacement after the remeshing, only 60-70 percent of labeled matches are used. Although making it more challenging, we believe this setup is close to a real case scenario, where acquisition noise and occlusions are unavoidable.

Our preliminary results are reported in Figure 3. For simplicity, we generate predicted correspondences by simply taking the mode of the softmax distribution for each reference node i : $\hat{j}(i) = \arg \max_j s_{i,j}$, thus avoiding a refinement step that is standard in other shape correspondence pipelines. The MLP model uses no context whatsoever and provides a baseline that captures the prior information from input coordinates alone. Using contextual information (even extrinsically as in point-cloud model) brings significant improvements, but these results may be substantially improved by encoding further prior knowledge. An example of the current failure of our model is depicted in Figure 5, illustrating that our current architecture does not have sufficiently large spatial context to disambiguate between locally similar (but globally inconsistent) parts.

We postulate that the FAUST dataset [1] is not an ideal fit for our contribution for two reasons: (1) it is small (100 models), and (2) it contains only near-isometric deformations, which do not require the generality offered by our network. As demonstrated in [7], the correspondence performances can be dramatically improved by constructing basis that are invariant to the deformations. We look forward to the emergence of new geometric datasets, and we are currently developing a capture setup that will allow us to acquire a more challenging dataset for this task.

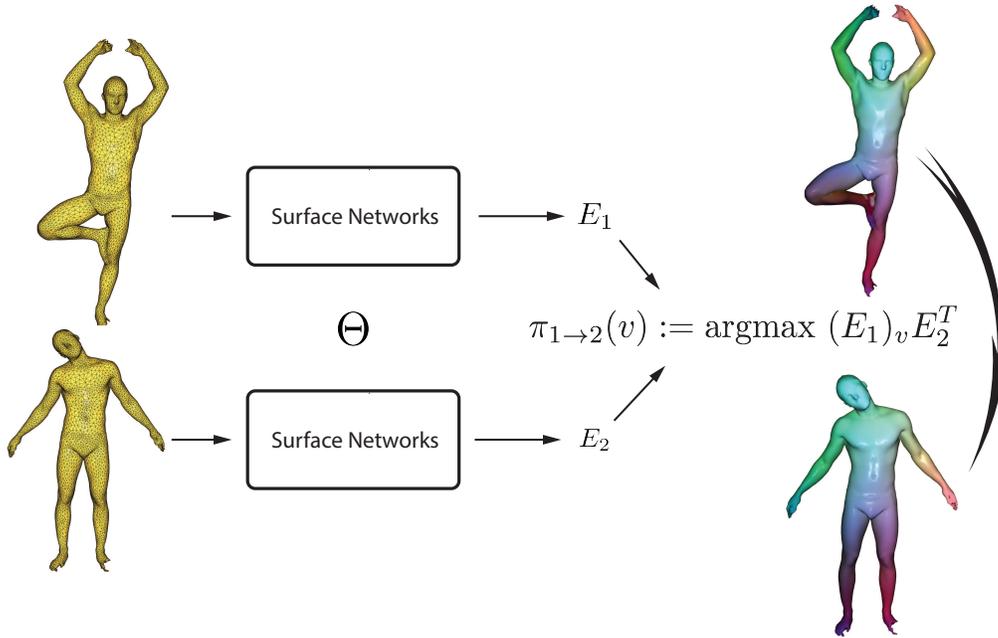


Figure 2: Siamese network pipeline: the two networks take vertex coordinates of the input models and generate a high dimensional feature vector, which are then used to define a map from \mathcal{M}_1 to \mathcal{M}_2 . Here, the map is visualized by taking a color map on \mathcal{M}_2 , and transferring it on \mathcal{M}_1

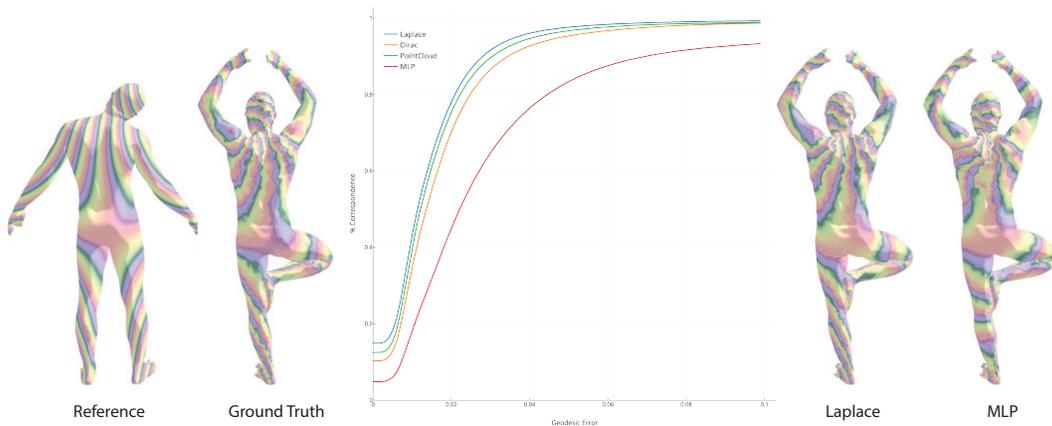


Figure 3: Additional results from our setup. Plot in the middle shows rate of correct correspondence with respect to geodesic error [6]. We observe that Laplace is performing similarly to Dirac in this scenario. We believe that the reason is that the FAUST dataset contains only isometric deformations, and thus the two operators have access to the same information. We also provide visual comparison, with the transfer of a higher frequency colormap from the reference shape to another pose.

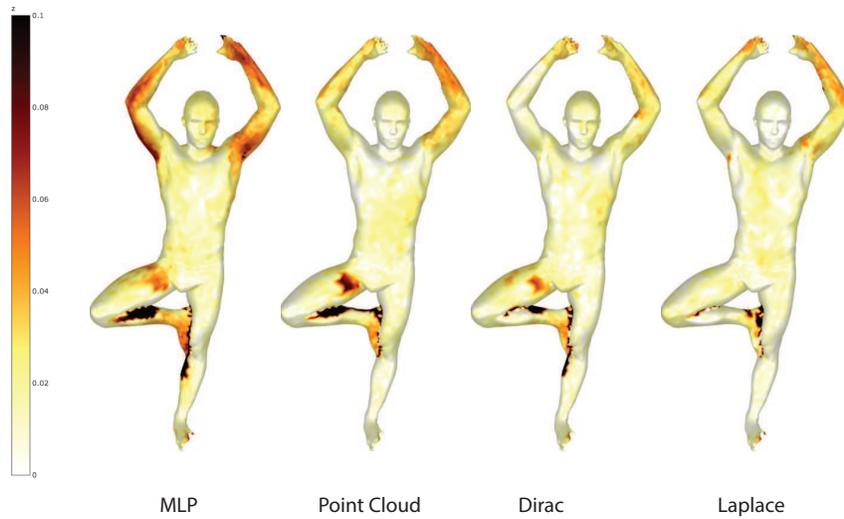


Figure 4: Heat map illustrating the point-wise geodesic difference between predicted correspondence point and the ground truth. The unit is proportional to the geodesic diameter, and saturated at 10%.

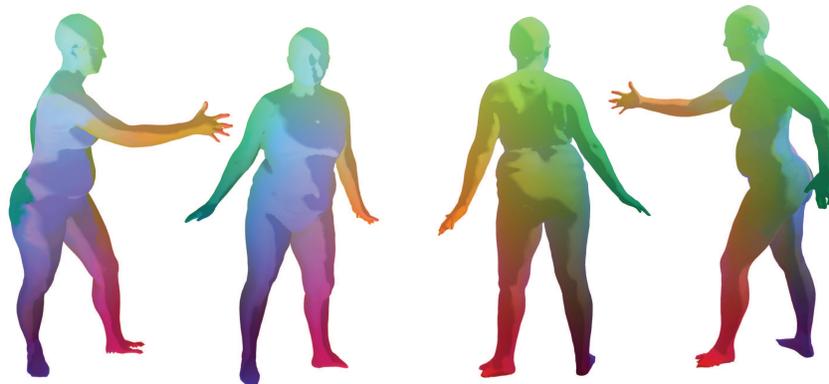


Figure 5: A failure case of applying the Laplace network to a new pose in the FAUST benchmark dataset. The network confuses between left and right arms. We show the correspondence visualization for front and back of this pair.

F Further Numerical Experiments



Figure 6: Qualitative comparison of different models. We plot 1th, 10th, 20th, 30th and 40th predicted frame correspondingly.

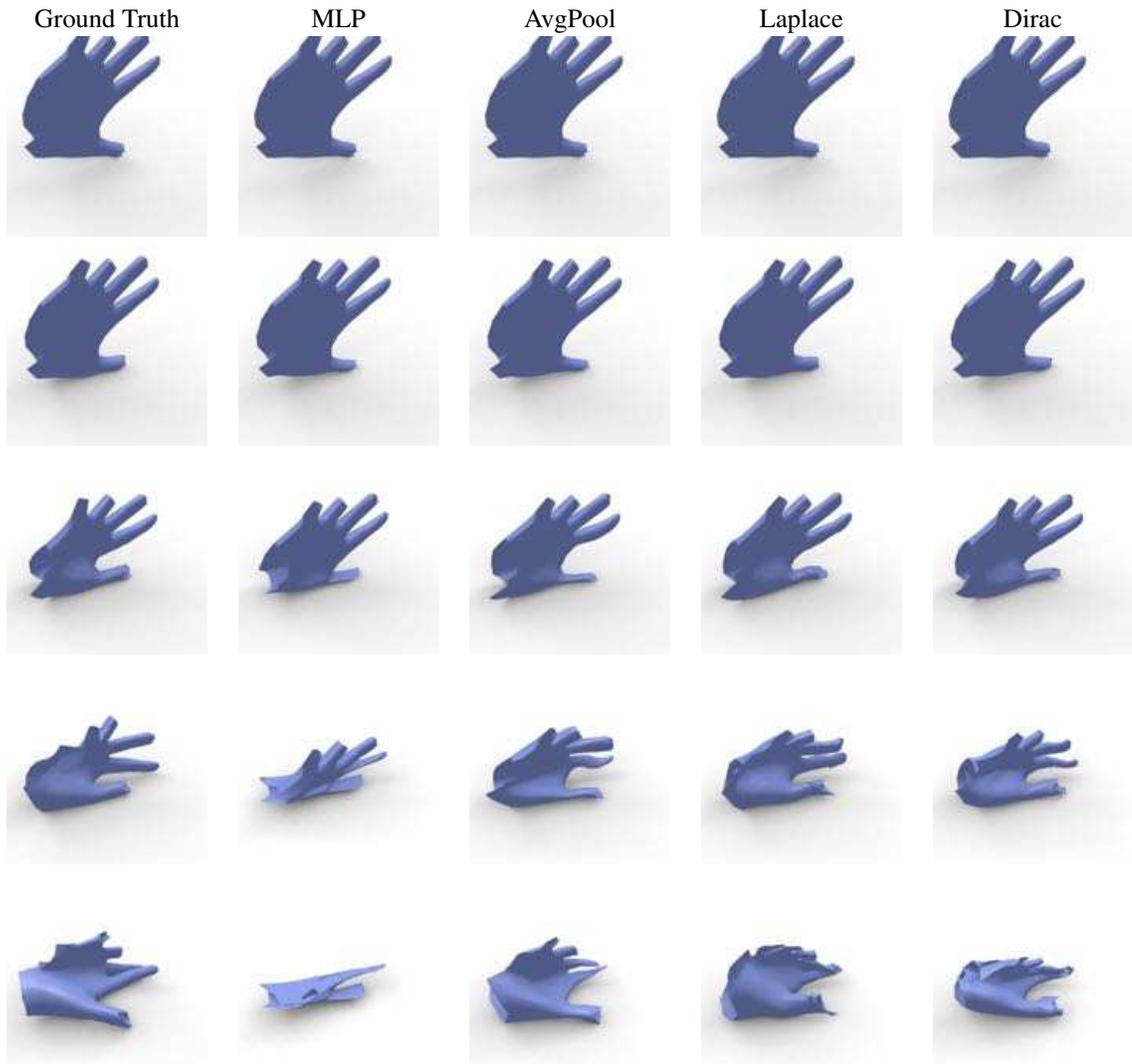


Figure 7: Qualitative comparison of different models. We plot 1th, 10th, 20th, 30th and 40th predicted frame correspondingly.

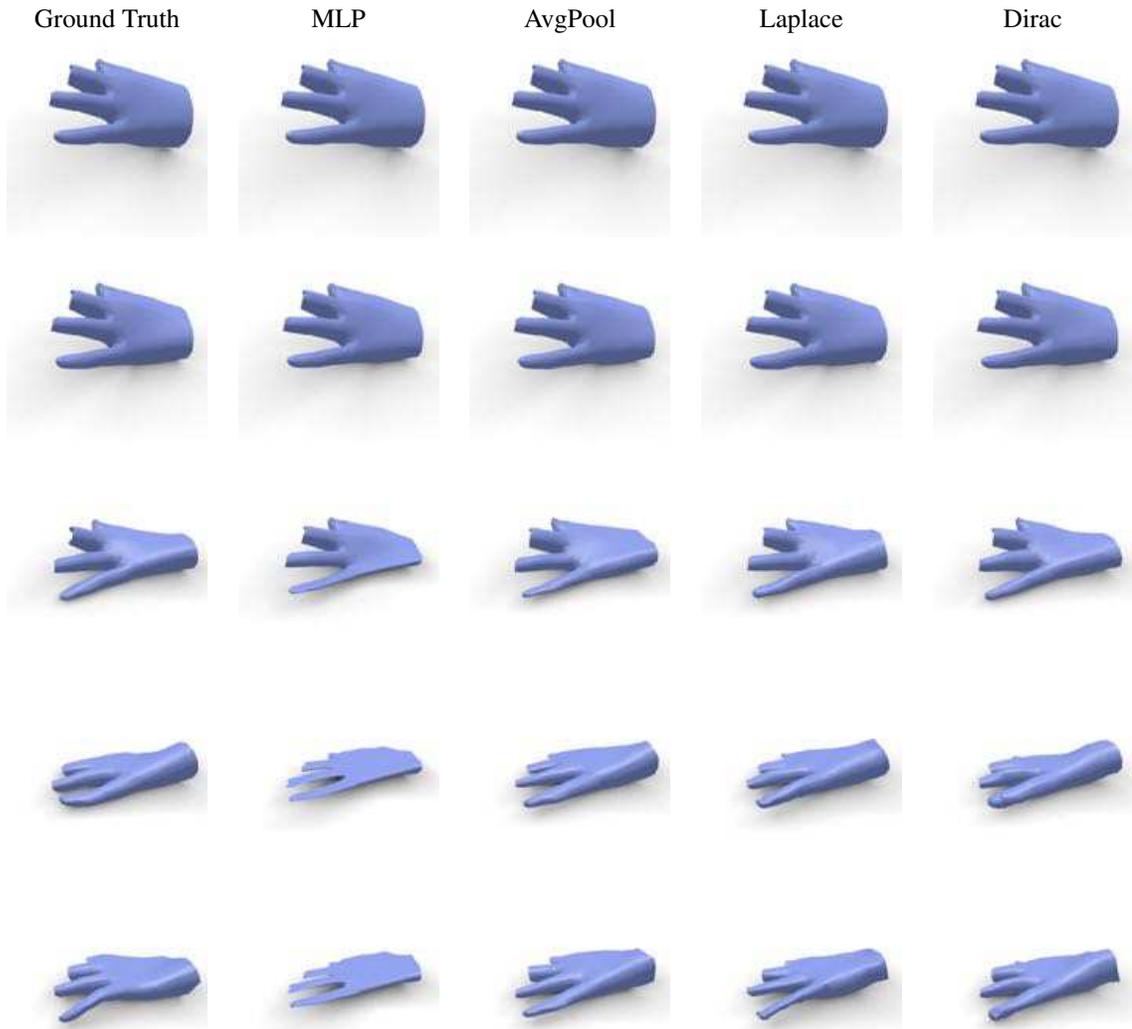


Figure 8: Qualitative comparison of different models. We plot 1th, 10th, 20th, 30th and 40th predicted frame correspondingly.

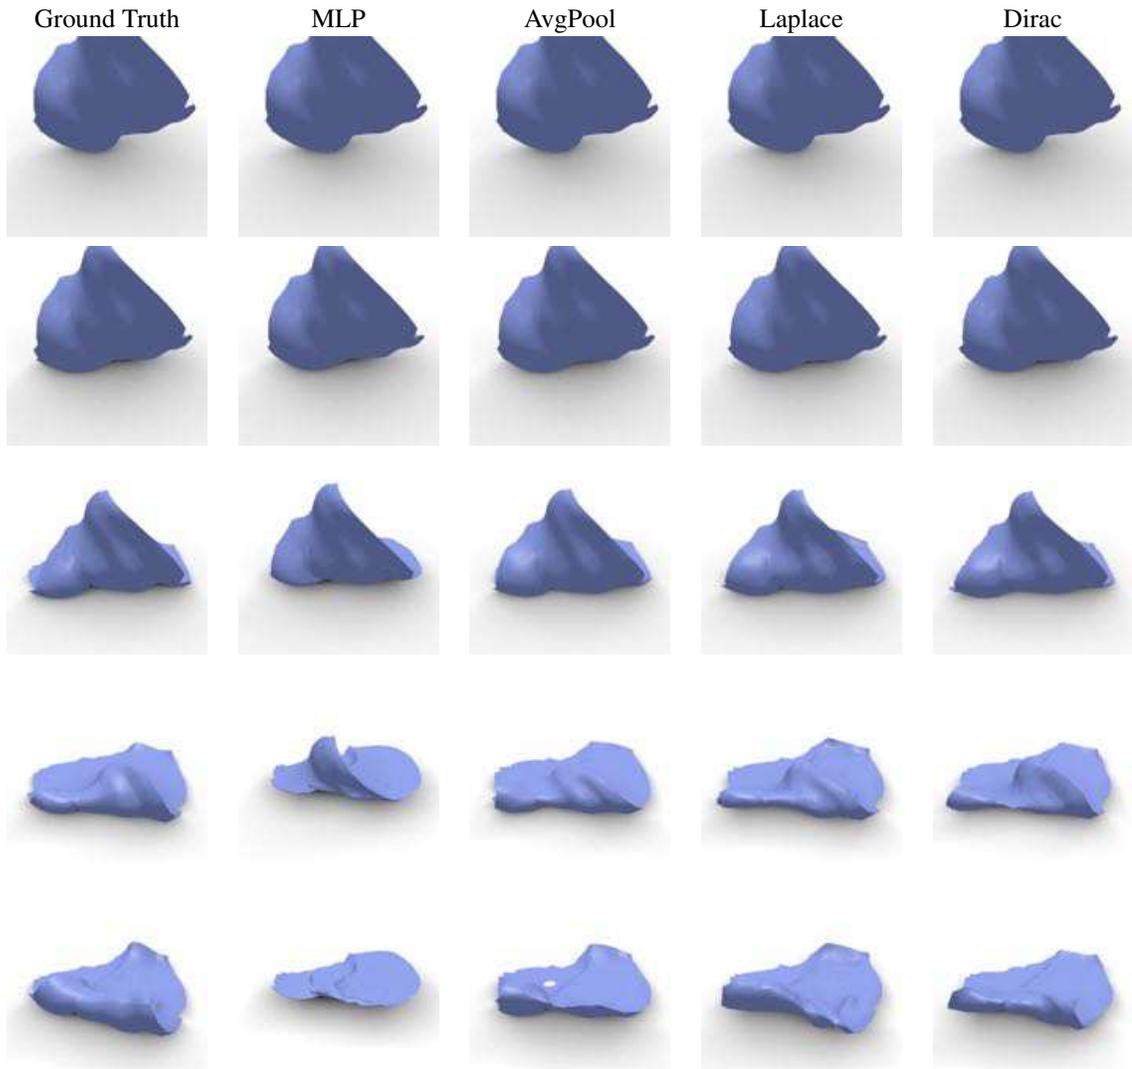


Figure 9: Qualitative comparison of different models. We plot 1th, 10th, 20th, 30th and 40th predicted frame correspondingly.

Ground Truth

Laplace

Dirac



Figure 10: Dirac-based model visually outperforms Laplace-based models in the regions of high mean curvature.

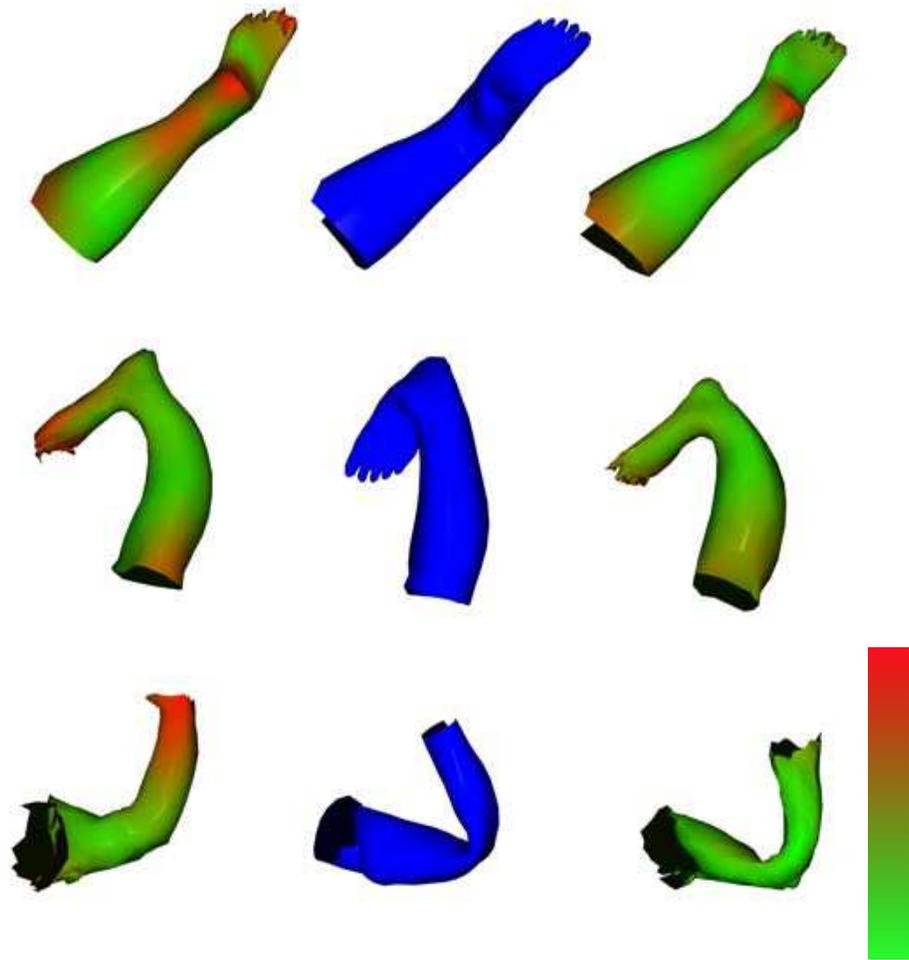


Figure 11: From left to right: Laplace, ground truth and Dirac based model. Color corresponds to mean squared error between ground truth and prediction: green - smaller error, red - larger error.

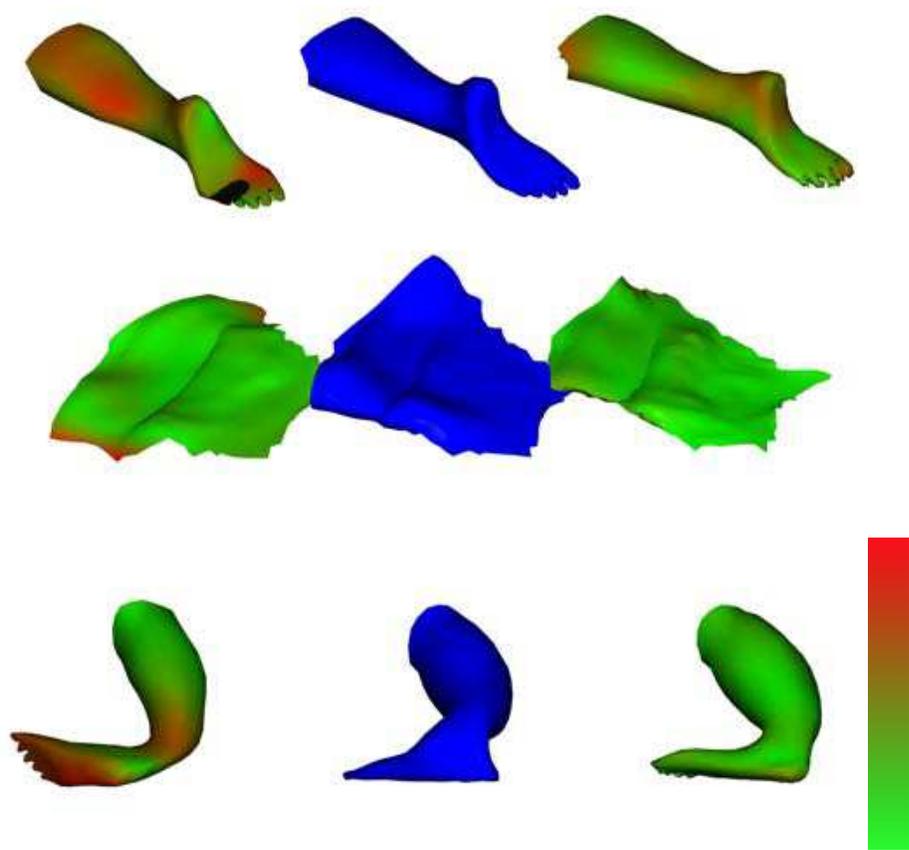


Figure 12: From left to right: set-to-set, ground truth and Dirac based model. Color corresponds to mean squared error between ground truth and prediction: green - smaller error, red - larger error.

References

- [1] F. Bogo, J. Romero, M. Loper, and M. J. Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014. 10
- [2] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *arXiv preprint arXiv:1611.08097*, 2016. 3
- [3] D. Chen and J. R. Gilbert. Obtaining bounds on the two norm of a matrix from the splitting lemma. *Electronic Transactions on Numerical Analysis*, 21:28–46, 2005. 4
- [4] K. C. Das. Extremal graph characterization from the upper bound of the laplacian spectral radius of weighted graphs. *Linear Algebra and Its Applications*, 427(1):55–69, 2007. 2
- [5] Y. Hu, Q. Zhou, X. Gao, A. Jacobson, D. Zorin, and D. Panozzo. Tetrahedral meshing in the wild. *Submitted to ACM Transaction on Graphics*, 2018. 10
- [6] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)*, volume 30, page 79. ACM, 2011. 11
- [7] O. Litany, T. Remez, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5660–5668, 2017. 10
- [8] A. Nowak, S. Villar, A. S. Bandeira, and J. Bruna. A note on learning algorithms for quadratic assignment with graph neural networks. *arXiv preprint arXiv:1706.07450*, 2017. 10
- [9] M. Wardetzky. Convergence of the cotangent formula: An overview. In *Discrete Differential Geometry*, pages 275–286. 2008. 2, 8