Structures for deep learning and topology optimization of functions on 3D shapes

Adrien Poulenard

Supervisor: Maks Ovsjanikov

LIX, Ecole Polytechnique

18 March 2020

Shape Analysis

- Design algorithms to analyse shapes : classification, detection, denoising, reconstruction, synthesis ...
- shapes comes from different acquisition devices & modeling tools.
- Different formats: voxels, polygon meshes, point clouds.

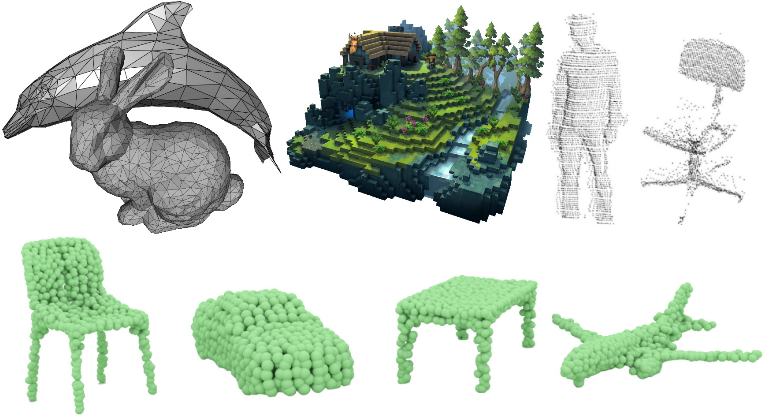

New challenges, geometric deep learning

- Applying deep learning techniques to shape analysis is challenging.
- Unlike images, shapes comes in various formats and don't have regular (fixed) structure.
- Learn variable domains instead of variable signals on a fixed domain.

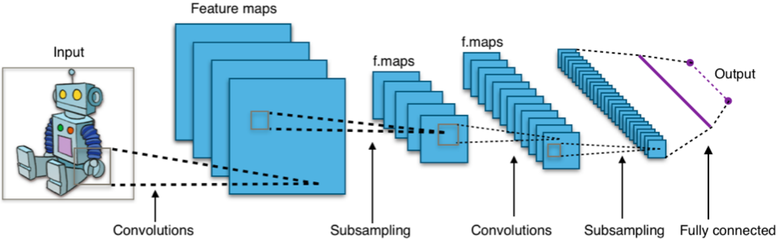

Plan

1. **Topological Function Optimization for Continuous Shape Matching.**
   Adrien Poulenard, Primoz Skraba, Maks Ovsjanikov.
   Symposium on Geometry Processing (SGP) 2018.
2. **Multi Directional Geodesic Neural Networks via Equivariant Convolution.**
   Adrien Poulenard, Maks Ovsjanikov.
   SIGGRAPH Asia 2018.
3. **Effective Rotation-Invariant Point CNN with Spherical Harmonics kernels.**
   Adrien Poulenard, Marie-Julie Rakotosaona, Yann Ponty, Maks Ovsjanikov.
   International Conference on 3D Vision (3DV) 2019.

Part I

Topological Function Optimization for Continuous Shape Matching

The functional map setting

A pointwise map between two shapes $T : \mathcal{N} \to \mathcal{M}$

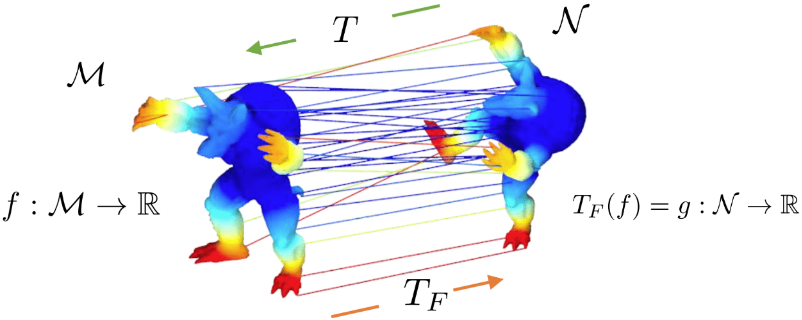

Induces a linear functional correspondence by pull back[1]
$T_F(f) := g$ where $g := f \circ T$.

[1] Ovsjanikov et al. Functional maps: a flexible representation of maps between shapes, TOG, 2012.

Problem: Functional maps tend to produce discontinuous correspondences

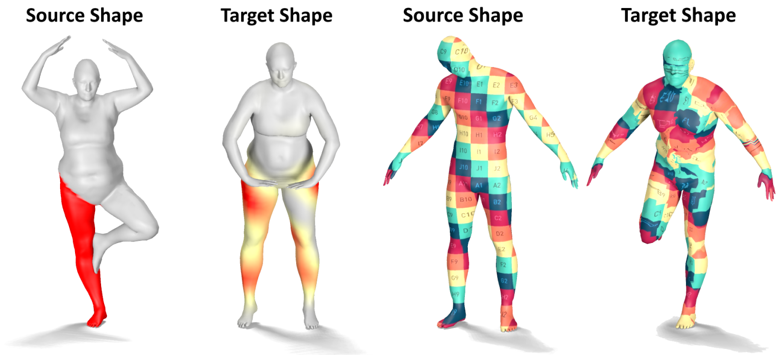

Transfer of indicator function via functional map (left). Texture transfer via point-to-point map obtained from functional map (right).

Intuition

- We want the image $g := \Phi_{\mathcal{N}} \mathbf{C} \Phi_{\mathcal{M}}^{+} f$ of any unimodal function $f : \mathcal{M} \to \mathbb{R}$ to be unimodal.
- We optimize $\mathbf{C}$ to remove least prominent maxima from g.

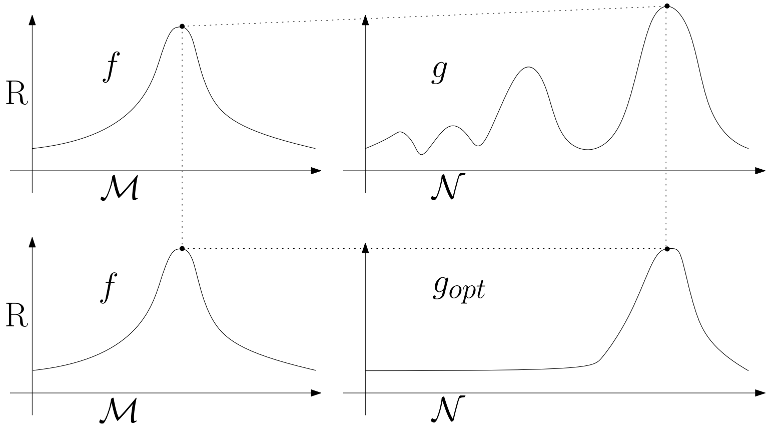

Where g_{opt} is the resulting function after optimizing $\mathbf{C}$.

Removing least prominent modes

$\mathrm{Pers}(g)$ being differentiable w.r.t. g it can be optimized by continuous optimization techniques. We can remove least prominent modes of a function g by optimising g to minimise $\mathrm{Pers}(g)$.

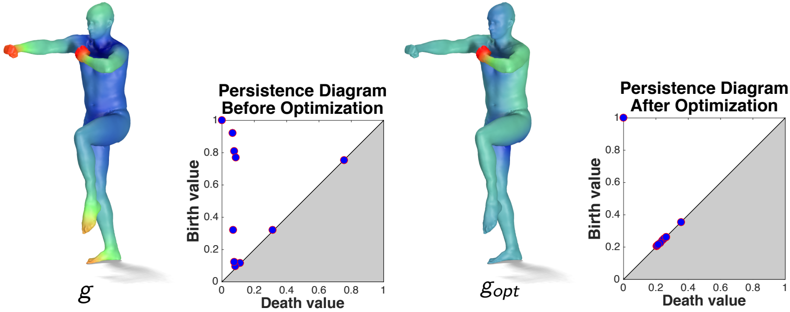

Improving functional maps

(a) Indicator function of a region

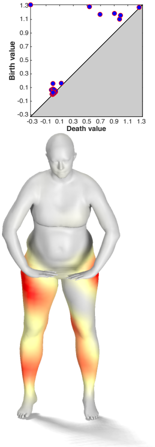

(b) Image before optimization

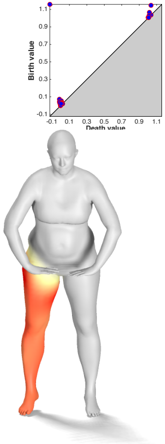

(c) Image after optimization

Part II

Multi Directional Geodesic Neural Networks via Equivariant Convolution

Problem setting

We would like to solve problems like classification, segmentation and matching, on shapes given as triangle meshes by using learning methods.

Shape matching

(a) GCNN (b) Ours (c) NN in descriptor space

Shape segmentation

(a) GCNN 3D (b) Ours 3D (c) Ground truth

CNNs on surfaces

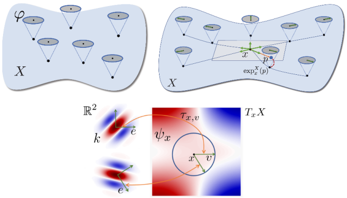

Classification

Fig. 12. Examples of CIFAR-10 images mapped to a sphere using elliptical mapping

Problem setting

The most successful learning techniques such as CNNs are not adapted to 3D shapes because they don't have global canonical coordinate systems.

Overview of Geodesic Convolution

The GCNN, (Boscaini et al., 2015) approach uses local polar coordinates induced by exponential maps to compare a signal on the surface to a kernel by mapping them on the tangent plane at every point.

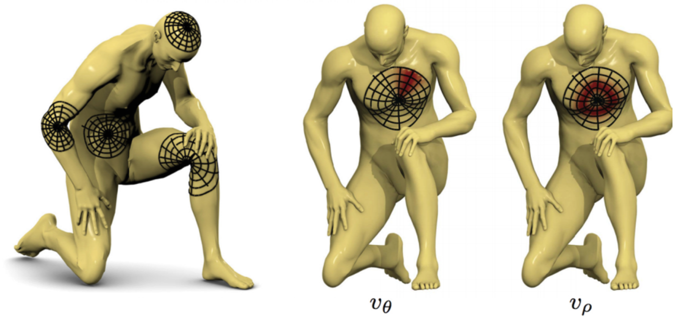

GCNN pipeline

Local ambiguity

Let $f : X \to \mathbb{R}$, Problem: Defining convolution by $\langle f \circ \exp_x, k_{x,u} \rangle_{L^2}$ is ambiguous as it depends on arbitrary $u \in T_x X$.

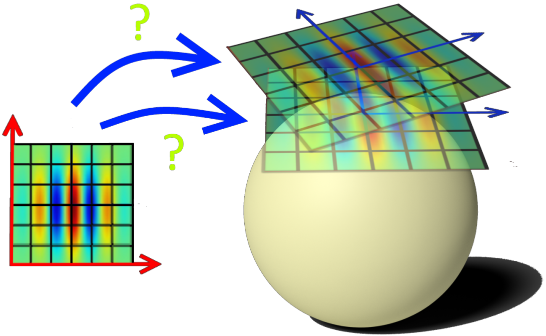

How to resolve ambiguity ?

$$f \odot k(x) := \max_u \sigma(\langle f \circ \exp_x, k_{x,u} \rangle_{L^2})$$

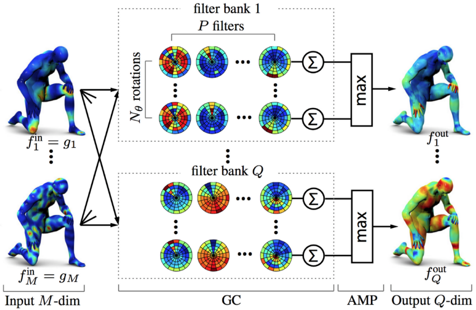

Geodesic Convolutional Neural Network (GCNN) architecture using angular max pooling. Problem: We lose directional information by taking the maximum response over direction u.

GCNN pipeline

Our approach

- Keep all responses $\langle f \circ \exp_x, k_{x,u} \rangle_{L^2}$ for all $u \in T_x X$.
- Work with directional functions $\varphi(x, u)$ depending on point x and direction $u \in T_x X$

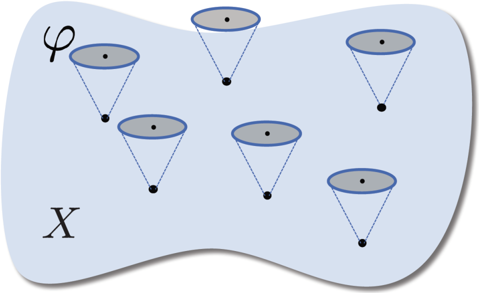

- Define a new "(multi)-directional convolution" operator taking a directional function φ and the same kernel k producing a new directional function $\varphi \star k$ over X.

Shape segmentation

Segmentation on the human body dataset introduced in Convolutional Neural Networks on Surfaces via Seamless Toric Covers (Maron et al., 2017). The dataset consists of 370 train shapes and 18 test shapes. **Goal**: Predict part label of vertices.

(a) GCNN (b) MDGCNN (c) Ground truth

Human shapes segmentation using 3D coordinates as input GCNN (Boscaini et al., 2015), MDGCNN (ours).

Part III

Effective Rotation-invariant Point CNN with Spherical Harmonics kernels

Problem setting

- Many datasets are **aligned** (canonical pose).
- Networks trained on aligned data **cannot generalise** to arbitrary poses.
- Require **data augmentation** by **random rotations** during training, **generalisation gap**.
- Instead we propose a **rotation invariant** design.

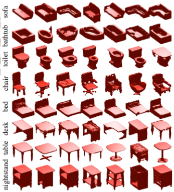

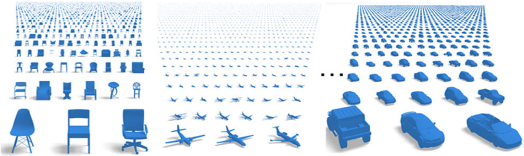

Our approach

Rotation invariant convolution kernels basis based on spherical harmonics $Y_{\ell m}$:

$$\kappa_{r\ell m}(x) := \underbrace{\exp\left(-\frac{\left|\|x\|_2 - \rho\frac{r}{n_R-1}\right|^2}{2\sigma^2}\right)}_{\text{radial comp}} \underbrace{Y_{\ell m}\left(\frac{x}{\|x\|_2}\right)}_{\text{directional comp}}$$

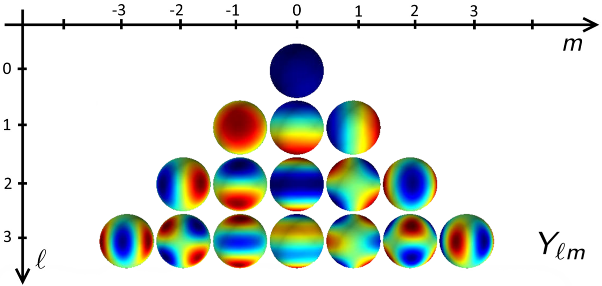

Segmentation

Example of RNA molecule segmentation:

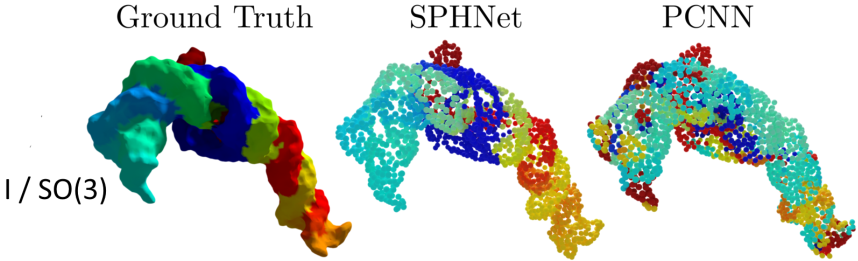

Thank you !