Dirichlet Energy for Analysis and Synthesis of Soft Maps

Justin Solomon* Leonidas Guibas* Adrian Butscher*†

*Geometric Computing Group at Stanford University
†Max Planck Center for Visual Computing and Communication

Symposium on Geometry Processing — July 2013

— Introduction —

Mappings Between Shapes

Let M_0 and M be smooth surfaces discretized as triangle meshes. We consider discrete representations of smooth maps $\phi : M_0 \to M$.

The of maps of interest should satisfy certain properties:

- Geometric
  - $\rightarrow$ Bijective
  - $\rightarrow$ Continuous
  - $\rightarrow$ Preserves fine details
- Semantic
  - $\rightarrow$ Meaningful
  - $\rightarrow$ Preserves features
  - $\rightarrow$ Satisfies user constraints

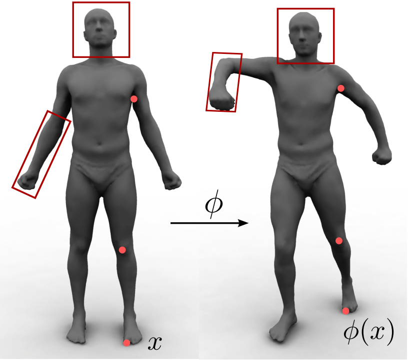

Difficulties with Point-to-Point Representations

An obvious discrete representation for a map is a vertex-to-vertex correspondence. This is inherently combinatorial and has drawbacks.

- The vast majority of vx-to-vx maps are in no way desirable.
- Continuity cannot be properly defined and quantified.
- The mesh itself interferes at the smallest scale!

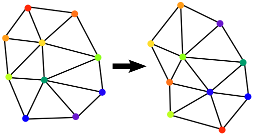

Thus: Vx-to-vx maps involve

- Subsampling.
- Measuring pairwise distances and adjacency relationships.
- This leads to problems!

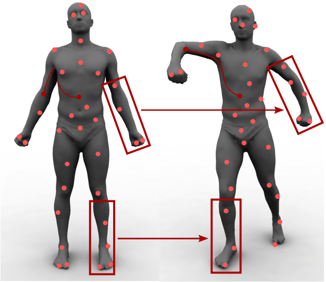

Continuity

In principle: These problems should be detectable via a failure of continuity somewhere. Continuity should have a regularizing effect.

- Why? Think of a result like the Intermediate Value Theorem.

The problem: Vertex-to-vertex representations are not adequate for quantifying continuity at this infinitesimal scale.

Possible resolution: An alternate representation for smooth maps.

- It should make sense for smooth surfaces yet be easily discretized, and should be convergent under mesh refinement.
- Continuity should make sense both discretely and in the smooth limit, and should be quantifiable.
- We should still be able to incorporate desirable map properties.

Soft Maps

We propose a representation that takes a probabilistic appoach.

Definition: A soft map from M_0 to M is a map $\mu : M_0 \to Prob(M)$.

I.e. every point on the source surface M_0 maps to a probability distribution of potential matches on the target surface M.

- Interpretation:

$$\mu_x = \begin{bmatrix} \text{Probability that } y \in M \\ \text{corresponds to } x \in M_0 \end{bmatrix}$$

- Recall SGP 2012.

  (Then: approximation by histograms.
  Now: the limit as the bin size $\to 0$.)

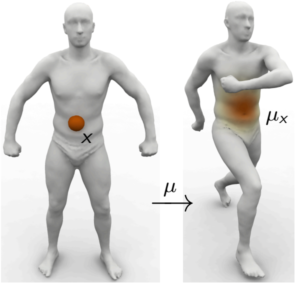

Advantages of Soft Maps

- They can be defined via scalar functions on $M_0 \times M$.

  $\to$ Each μ_x has a positive density that integrates to one.

- They generalize point-to-point maps $\phi : M_0 \to M$.

  $\to$ The associated density is sharply peaked at $\phi(x)$.

- They permit blurring and superposition.

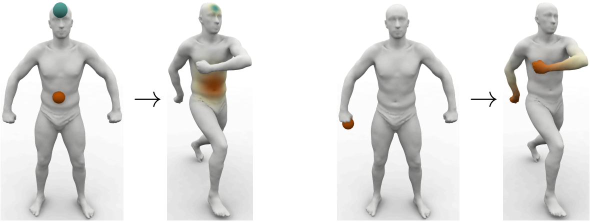

The "ideal" soft map is a convex combination of a small number (associated with symmetries) of blurred point-to-point maps.

— Soft Map Energies —

Quantifying Continuity

Recall: Dirichlet energies quantify the "degree of continuity" of mappings between domains in many different contexts.

Examples:

- For $f : M_0 \to \mathbb{R}$

$$\mathcal{E}_D(f) := \int_{M_0} \|\nabla_0 f(x)\|^2 dx$$

- For $\phi : M_0 \to M$

$$\mathcal{E}_D(\phi) := \int_{M_0} \|\nabla_0 \phi(x)\|^2 dx$$

Quantifying Continuity

Recall: Dirichlet energies quantify the "degree of continuity" of mappings between domains in many different contexts.

Examples:

- For $f : M_0 \to \mathbb{R}$

$$\mathcal{E}_D(f) := \int_{M_0} \|\nabla_0 f(x)\|^2 dx$$

- For $\phi : M_0 \to M$

$$\mathcal{E}_D(\phi) := \int_{M_0} \|\nabla_0 \phi(x)\|^2 dx$$

A generalization: These are all instances of a general framework for maps $\phi : (M_0, dist_0) \to (M, dist)$ between any metric spaces:

$$\mathcal{E}_D(\phi) := \int_{M_0} \left(\lim_{\varepsilon \to 0} \frac{1}{Area(B_\varepsilon(x))} \int_{B_\varepsilon(x)} \frac{dist^2(\phi(x), \phi(x'))}{dist_0{}^2(x, x')} dx' \right) dx$$

The Wasserstein Metric on *Prob*(M)

Key idea: We can view *Prob*(M) as a metric space.

- The theory of Optimal Transportation gives us a metric on *Prob*(M) called the **Wasserstein metric** with **quadratic cost**.

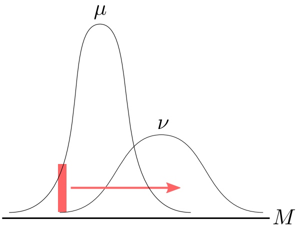

Let $\mu, \nu \in Prob(M)$.

The distance $W_2(\mu, \nu)$ is the cost of the optimal way of transporting mass from the distribution μ to the distribution ν.

The transportation cost for an individual particle from y to y' is $dist^2(y, y')$.

- This is also called the Earth Mover's Distance (with quadratic cost).

Consequence: We can define a Dirichlet energy for soft maps.

The Dirichlet Energy of a Soft Map

Definition:

Let $\mu : M_0 \to Prob(M)$ be a soft map.

The Dirichlet energy of μ is the quantity

$$\mathcal{E}_D(\mu) := \int_{M_0} \left(\lim_{\varepsilon \to 0} \frac{1}{Area(B_\varepsilon(x))} \int_{B_\varepsilon(x)} \frac{W_2^2(\mu_x, \mu_{x'})}{dist_0^2(x, x')} dx' \right) dx$$

Key properties:

- The Dirichlet energy is convex in μ.
- It generalizes the Dirichlet energy for maps. So if ϕ is a map and μ_ϕ is the associated soft map then $\mathcal{E}_D(\mu_\phi) = \mathcal{E}_D(\phi)$.
- The Dirichlet energy of any constant soft map is zero.

Simplification of the Dirichlet Energy

Problem: This form of the Dirichlet energy is difficult to work with.

Theorem: The following simplification holds.

Consider a soft map with smooth positive density $\rho(x, y)$. Then the Dirichlet energy of ρ satisfies

$$\mathcal{E}_D(\rho) = \iint_{M_0 \times M} \rho(x, y) \|\nabla Q(x, y)\|^2 dy \, dx \, .$$

The quantity Q is vectorial and lives on $M_0 \times M$.

It is defined as follows. For each x and direction V, then $Q(x, y)$ satisfies the linear PDE in the y-variables given by

$$\nabla \cdot \big(\rho \nabla \langle Q, V \rangle\big) = -\langle \nabla_0 \rho, V \rangle$$

Interpretation of Q

We call Q the **transportation potential** of the soft map. We can interpret it in terms of conservative mass flow.

I.e. we view each $\rho(x, \cdot)$ is a swarm of particles. Now:

- Displace x in the direction V to an infinitesimally near x'.
- The mass distribution $\rho(x, \cdot)$ changes into $\rho(x', \cdot)$.

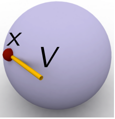

Source

$\rightarrow$

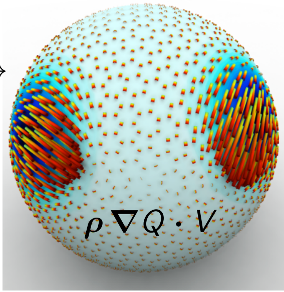

Target

- Assume it's by optimal transport.
- The particle at y has velocity equal to

$$\nabla Q(x, y) \cdot V$$

- The Wasserstein distance relates to the kinetic energy.

$$\frac{W_2^2(\mu_x, \mu_{x'})}{dist_0^2(x, x')} \approx \int_M \rho(x, y) \|\nabla Q(x, y) \cdot V\|^2 dy$$

Soft Map Bijectivity

An issue: All constant soft maps all have the same minimal Dirichlet energy equal to zero. Can we tell them apart?

Idea: Measure the equidistribution of probabilistic mass pushed forward from M_0 to M. Quantify as follows.

- We can interpret the integral $b(y) := \int_{M_0} \rho(x, y) dx$ as the probability that y receives mass from *somewhere* in M_0.
- So if the square integral

$$\mathcal{E}_b(\rho) := \int_M \left[\int_{M_0} \rho(x, y)\, dx \right]^2 dy$$

  is small, then each $\rho(x, \cdot)$ is as spread out as possible and each point of M receives an equal amount of mass from M_0.
- We call $\mathcal{E}_b(\rho)$ the bijectivity energy of ρ.

— Soft Map Analysis and Synthesis —

Map and Soft Map Analysis

The two energies and their densities that we have introduced can be used for soft map analysis.

We can study:

- The soft map of a pt-to-pt map.
- Or a soft map coming from shape descriptor differences, of the form

$$\rho(x, y) \propto e^{-(d_1(x)-d_2(y))^2/\sigma^2}$$

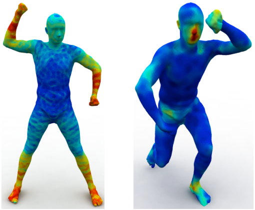

Dirichlet and bijectivity energy densities on M_0 and M, resp.

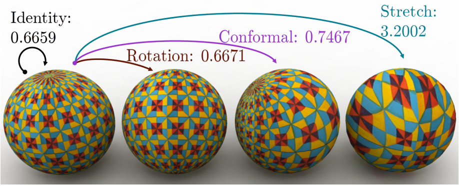

Energies of various self-maps of the sphere.

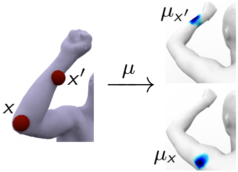

Unfavourable stretching in WKS revealed by the Dirichlet energy

Local Correspondence Extraction

Recall: Choose x and a direction V. Let the mass distribution $\rho(x, \cdot)$ change optimally with x.

Then the particle at y moves with velocity $\nabla Q(x, y) \cdot V$.

So what: We get a method for extracting point correspondences.

- Choose a path $x(\varepsilon)$ s.t. $x(0) = x_{init}$ and decide on a point $y_{init} \in M$ that should correspond to x_{init}.
- Integrate the velocity field

$$\dot{y} = \nabla Q(x, y) \cdot \dot{x}$$

- Get a path $y(\varepsilon)$ in M with initial data $y(0) = y_{init}$.
- The paths $x(\varepsilon)$ and $y(\varepsilon)$ are now in correspondence.

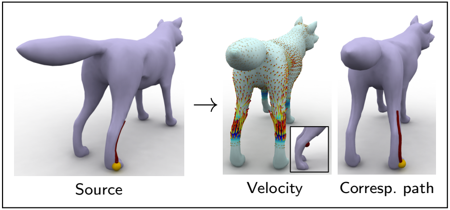

Generating Soft Maps

Goal: Generate soft maps by solving a constrained optimization problem in the space of soft maps. It's convex!

$$\text{minimize} \quad \mathcal{E}(\mu) := \mathcal{E}_D(\mu) + \lambda \mathcal{E}_b(\mu)$$

And: To avoid the constant soft map, we must impose constraints.

- E.g. a few points or subsets of M_0 and M must correspond.
- This is similar to the harmonic maps problem.

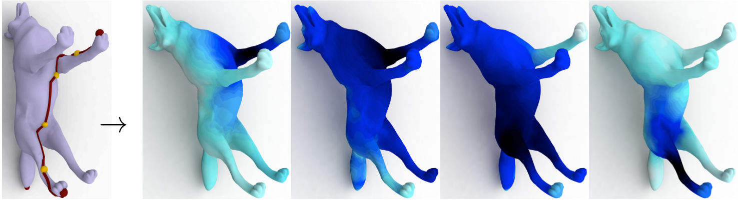

Source, red constraints

Optimal soft map distributions associated to the yellow points.

Conclusion and Future Work

What we have done:

- Introduced a representation for maps that supports a Dirichlet energy for measuring continuity.
- Used this representation for map analysis and synthesis.

What we would like to do next:

- More efficient computation of Q.
- Decomposition of ρ into a convex combination of soft maps associated to maps.
- Map extraction at multiple scales.

Thank you!

Acknowledgements: NSF grants IIS 914833, CCF 1161480, DMS 1228304, NSF/NIH grant 900700, AFOSR grant 1156110-1-TACAD, a Google Research Award, the National Defense Science and Engineering Graduate Fellowship, the Hertz Foundation Fellowship, the NSF GRF program, a Stanford Bio-X travel grant, and the Max Planck Center for Visual Computing and Communications.