An Optimal Transportation Approach For Discrete Mappings Between Meshed Surfaces

Adrian Butscher

Stanford University and
Max Planck Institute for Informatics

August 2013

Outline

1. A brief sketch of the discrete shape matching problem.
2. A new representation for maps that is more convenient for the discretization of smooth maps between surfaces.
   → Soft maps
3. Dirichlet energy of a soft map.
4. A variational formulation of the discretized mapping problem.
5. Preliminary computational results and theoretical questions.

$*$ Joint work with Justin Solomon and Leonidas Guibas.

The Shape Matching Problem

Given surfaces M_0 and M represented discretely as triangle meshes, we seek a discrete representation of a smooth map $\phi : M_0 \to M$.

Criteria for a "good" map:

- Geometric
  - → Bijective
  - → Continuous
  - → Locally non-distorting
- Semantic
  - → Feature-preserving
  - → Meaningful

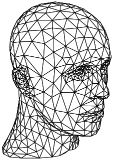

Triangle mesh: a standard discrete representation of embedded surfaces (Flexible, simple, processable)

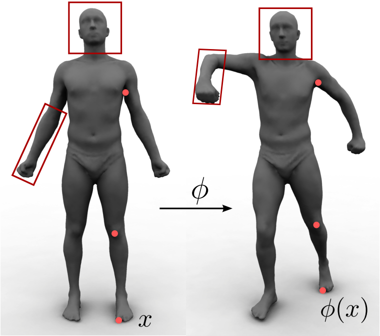

Example: feature- and region- preserving map between humans in different poses

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.

- Huge search space.
- 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!

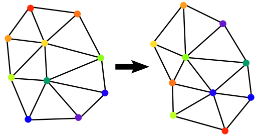

So: These issues are often tackled by

- Subsampling.
- Pairwise distances and adjacency.

But: Many problems, e.g. symmetry.

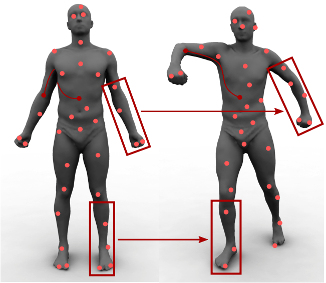

Continuity

In principle: These problems should be detectable (thus preventable!) at the infinitesimal level in some way.

- Loss of continuity (etc. — such as loss of local injectivity).

But: Vx-to-vx representations are not adequate at this 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 (etc.) 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 $\mu : M_0 \to Prob(M)$.

In this setting, every point of M_0 maps to a probability distribution of potential matches on M.

In other words:

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

for all subsets $U \subseteq M$.

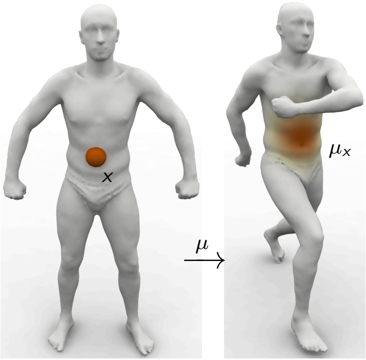

Advantages of Soft Maps

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

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

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

  $\rightarrow$ The associated soft map is $x \mapsto \delta_{\phi(x)}(y)\, dy$.

- They permit blurring and superposition.

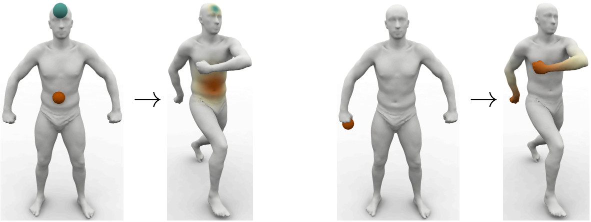

The "ideal" soft map is a convex combination of a small number (associated with symmetries) of blurred point-to-point maps.

Constraints on Soft Maps

Soft maps can handle the "traditional" constraints on pt-to-pt maps.

- **Descriptor matching.**

  Let $f_0 : M_0 \to \mathbb{R}$ and $f : M \to \mathbb{R}$ be descriptor functions that we expect should match. Then we can require

  $$f_0 \ \text{ must be close to } \ \mathbb{E}_\mu(f) := \int_M f(y) d\mu_x(y)$$

  which is the expected value of f at x under μ_x.

- **Region constraints.**

  Let $U_0 \subseteq M_0$ and $U \subseteq M$ be regions that we expect should match. Then we can require

  $$\left.\begin{array}{rr} & supp(\mu_x) \subseteq U \\ OR & \mu_x\big|_U = \text{given} \end{array}\right\} \ \forall x \in U_0$$

Quantifying Continuity for Soft Maps

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

- E.g. harmonic functions, geodesics, harmonic maps.

In general: These are instances of a universal framework for maps $\phi : (\mathcal{M}_0, dist_0) \to (\mathcal{M}, dist)$ between any metric spaces:

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

Our idea: We can apply this framework to soft maps if we take $\mathcal{M}_0 = M_0$ with geodesic distance and $\mathcal{M} = Prob(M)$ with the Wasserstein distance.

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

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

Key properties:

- Measures the "degree of continuity" of the map $x \mapsto \mu_x$.
- Convex in μ.
- Generalizes the Dirichlet energy for maps.

  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

But: This expression is cumbersome. Instead, we use a simpler one.

Let μ be a soft map with smooth density $\rho > 0$. Then

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

where Q is a section of $T^* M_0 \otimes C^\infty(M)$ and is defined by:

- For $(x, V) \in TM_0$ let q be the function $y \mapsto Q(x, y) \cdot V$.
- Then q satisfies the weak form of the equation

$$\left.\begin{aligned} \nabla \cdot (\rho(x, \cdot) \nabla q) &= -\langle \nabla_0 \rho(x, \cdot), V \rangle \\ \int_M q(y) \rho(x, y) dy &= 0 \end{aligned}\right\} \begin{array}{l} \text{One equation in } y \\ \text{for each } (x, V). \\ \text{Linear in } V. \end{array}$$

Formal Derivation

Preliminaries: Let ν and $\tilde{\nu}$ be two probability measures on M. The theory of optimal transportation gives us the following:

- A W_2-optimal map $T : M \to M$ with $T\#\nu = \tilde{\nu}$ of the form

$$T(y) := \exp_y(\nabla\phi(y)) \quad \text{for a cvx function } \phi : M \to \mathbb{R}$$

- The Wasserstein distance is $W_2^2(\nu, \tilde{\nu}) = \int_M \|\nabla\phi(y)\|^2 d\nu(y)$.

Next: Apply these results to a soft map μ.

- Choose nearby points x and $x' := \exp_x(\varepsilon V)$ for $V \in T_x M_0$.
- Take $\nu := \mu_x$ and $\tilde{\nu} := \mu_{x'}$ with optimap T_ε and potential ϕ_ε.
- Expand in ε.

Hope: With some work, this derivation can be made rigorous and extended to a much less regular class of measures.

Interpretation of Q

We have interpreted Q in terms of conservative mass flow.

- Each $\rho(x, \cdot)$ is a swarm of particles.
- Consider the path given by $x'(\varepsilon) := \exp_x(\varepsilon V)$.
- Consider the optimal maps of $\rho(x, \cdot)$ into $\rho(x'(\varepsilon), \cdot)$.
- The instantaneous velocity of a particle at y equals $\nabla Q(x, y) \cdot V$

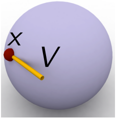

Source

$\rightarrow$

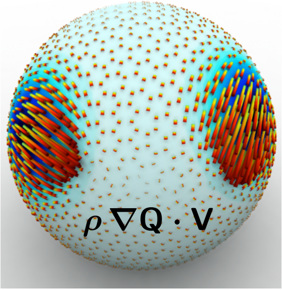

Target

- The Wasserstein distance is the instantaneous total kinetic energy.

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

Optimal Soft Maps

Goal: We would like to pose a constrained optimization problem in a space of soft maps.

- Inspiration: a harmonic map problem.

The energy: should promote "smoothness in the x-variable" via the Dirichlet energy.

But: The global minimum of $\mathcal{E}_D$ is $\mu = const$ with $\mathcal{E}_D(\mu) = 0$. How can we avoid the constant soft map?

- Add a descriptor matching term to the energy.
- Add region constraints.

Then: Develop a convergent discretization.

A Convex Optimization Problem

So: We would like to solve a discretization of the convex problem

Minimize

$$\mathcal{E}(\mu) := \mathcal{E}_D(\mu) + \sum_s \Big\| \underbrace{f_0^{(s)} - \mathbb{E}_\mu(f^{(s)})}_{\text{Descriptor functions}} \Big\|^2_{L^2(M_0)} + \cdots$$

subject to region constraints.

Some typical results:

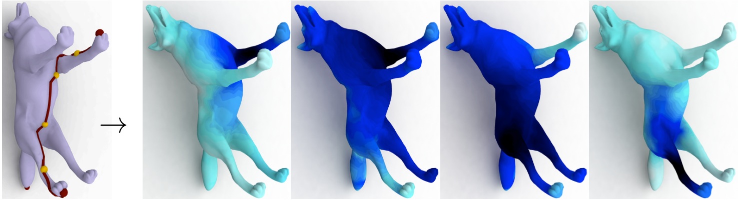

Source, red constraints

Optimal soft map distributions associated to the yellow points.

Discretization

All objects introduced so far can be represented via scalar functions so discretization can be done using a Finite Element Method.

- We introduce PL basis functions $\beta_{0i} : M_0 \to \mathbb{R}_+$ and $\beta_j : M \to \mathbb{R}_+$ where $\int_M \beta_j(y)dy = 1$ for all j.
- Then we work with soft maps of the form

$$d\mu_x(y) := \sum_{ij} C_{ij}\beta_{0i}(x)\beta_j(y)dy$$

$$\text{with} \quad C_{ij} \geq 0 \; \forall\, i,j \quad \text{and} \quad \sum_j C_{ij} = 1 \; \forall\, j$$

- Region constraints are linear in C.
- A similar discretization can be carried out for Q.
- Solving for Q and optimizing for ρ — linear algebra problems!

Theoretical Questions

Elementary questions:

- In what space can we solve this problem (both the continuous and discretized versions)?
- Characterization of the minimum (Euler-Lagrange equations)?
- Some exact solutions, or other intuition for the minimum?
- The qualitative behaviour of the PDE for Q? Especially at points where $\rho = 0$ or where $\rho =$ singular?
- Regularity of the solution?
- Convergence as the discretization is refined?
- Stability of the solution under perturbations of M_0 and M?

Theoretical Questions

Deeper questions:

- Are there conditions that guarantee solutions of the "ideal" form (convex combination of blurred maps)?
  - → How to quantify "blurriness" and avoid overly blurry solutions?
  - → Does "inconsistency" in the constraints correlate with "blurriness" of the solution in some way?
  - → How to extract the maps?
- The trivial solution, in the absence of soft/hard constraints, is

$$d\mu_x(y) = \rho(y)dy$$

The "constant" soft map with $\mathcal{E}_D(\mu) = 0$. A global minimum!

  - → How to avoid a trivial solution? How many constraints?
- How to correctly discretize this problem so that computational costs are reduced? Will involve a smart theoretical approach!

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.