

Soft Maps Between Surfaces

Justin Solomon

Andy Nguyen

Adrian Butscher

Mirela Ben-Chen

Leonidas Guibas

Geometric Computing Group, Stanford University




Figure 1: Soft maps from one model to another computed using our optimization technique. The colored patches on the leftmost model are mapped to the colored distributions over the models on the right. These soft maps acknowledge discrete left-right and front-back symmetries as well as localized ambiguities including slippage along the pig's back.

Abstract

The problem of mapping between two non-isometric surfaces admits ambiguities on both local and global scales. For instance, symmetries can make it possible for multiple maps to be equally acceptable, and stretching, slippage, and compression introduce difficulties deciding exactly where each point should go. Since most algorithms for point-to-point or even sparse mapping struggle to resolve these ambiguities, in this paper we introduce soft maps, a probabilistic relaxation of point-to-point correspondence that explicitly incorporates ambiguities in the mapping process. In addition to explaining a continuous theory of soft maps, we show how they can be represented using probability matrices and computed for given pairs of surfaces through a convex optimization explicitly trading off between continuity, conformity to geometric descriptors, and spread. Given that our correspondences are encoded in matrix form, we also illustrate how low-rank approximation and other linear algebraic tools can be used to analyze, simplify, and represent both individual and collections of soft maps.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Introduction

A natural problem in geometry processing is that of finding a smooth map between two surfaces. A reliable algorithm for finding such a map can be used in pipelines for texture or annotation transfer, segmentation, morphing, and surface editing, among other applications within graphics. Outside of graphics, ongoing research in vision and other fields makes use of shape maps to create links between new inputs and

previously-analyzed data; for instance, a robot navigating an unknown environment may try to map objects it encounters to ones in some given database of objects it can manipulate.

Unless shapes are rigid motions or isometric deformations of each other, it is difficult to define a single “best” map between most pairs of surfaces at point-to-point granularity. This difficulty arises because there are at least two geometric sources of ambiguity complicating the mapping problem:

Global ambiguity: Symmetric shapes may admit multiple geometrically equivalent maps; for instance, human models often have left-right symmetries that generate two equivalent maps in terms of the amount of geodesic distance distortion they induce. Note that the shapes might be symmetric under a rigid transformation of space or under an *intrinsic* symmetry (as in e.g. [OSG08,RBK10]).

Local ambiguity: Shapes that are not exact isometric deformations of one another may admit an informative map at some level of coarseness but not at the point-to-point level due to scaling, slippage, or the absence of identical details. For instance, generating maps between a horse and a dog model makes sense at the segment level because both animals have similar limb structures, even if the structures within those limbs are different.

Additional ambiguities can result from a lack of context. Without knowledge of the process used to obtain the target from the source, it is impossible to know which maps are semantically relevant, regardless of geometric cues.

Given these fundamental problems, a limitation of many mapping algorithms is that they attempt to find a point-to-point map with no more than a geometric prior, whether it be rigidity, conformality, isometry, elasticity, or otherwise. These methods are forced to make somewhat arbitrary decisions as to the user's desired map or, worse, unsuccessfully attempt to combine often disjoint acceptable maps. Thus, mapping algorithms that process a variety of shapes should incorporate uncertainty when the mapping problem is itself ambiguous. These ambiguities can be resolved with domain-specific knowledge, semantic information, or other cues, or they can be used to prompt the user for guidance.

In this paper, we propose *soft maps*, which generalize point-to-point maps by embracing uncertainty as a fundamental part of the mapping process. In this setup, maps are expressed as conditional distributions of a distribution on the product of the two surfaces, which we call a *soft correspondence*. In other words, we attach to each pair of regions on the two surfaces a probability indicating the likelihood that these regions should be mapped to one another.

A principled discretization of this formulation is possible by partitioning the surfaces into unions of patches and representing a soft correspondence as a probability distribution over the product of the patch index sets. This construction results in a discretized soft correspondence represented as a matrix whose rows and columns are the forward and backward soft maps after normalization. Point-to-point maps can be represented as soft maps using permutation matrices.

The soft map framework is capable of handling both local and global ambiguities, as illustrated in Figure 1. Here, given only geometric information, we can compute a soft mapping where the front hoof of a pig model is mapped $\sim 50\%$ to each of the front hooves on a different pig; the back of the source pig is mapped to a larger region on the target pig's back,

since the lack of distinguishing geometric features makes a more precise mapping impossible. Similarly, the map in Figure 1 between human models acknowledges their approximate left-right and front-back symmetries.

One important property of soft maps is *continuity*, which must be redefined probabilistically to ensure that nearby points on one surface yield nearby distributions on the other. We define infinitesimal and discretized notions of soft map continuity and show how Earth Mover's Distances can measure the discrete continuity of a soft map. With this metric and others describing a soft map's spread and alignment with geometric features, we provide a convex optimization approach for computing soft correspondences.

Once we have computed a single soft map between two shapes or multiple soft maps within a collection of shapes, the fact that soft maps are expressible as matrices allows for straightforward analysis using the standard linear-algebraic toolkit. We show how to make use of Principal Components Analysis (PCA) and related methods to reveal structure within collections of soft maps, obtaining low-dimensional representations of individual soft maps and pairwise maps between members of a collection of shapes.

1.1. Contributions

We introduce a probabilistic interpretation of maps between surfaces and show how they can be obtained, manipulated, and interpreted. In particular, we provide:

- A theoretical definition of soft maps, including a definition of continuity and a principled discretization (§3)
- A convex optimization method for finding soft maps that explicitly trades off between continuity, matching of geometric descriptors, and softness (§4)
- Model reduction methods for analyzing individual and collections of soft maps (§5)

2. Previous Work

The literature on mapping between surfaces is vast, and we refer the reader to [vKZHC010, CLM*11, BBK08] for general summaries of previous work. The idea of computing a mapping by minimizing descriptor distances and preserving continuity is common to many of the works surveyed here.

Recent work on mapping reveals several approaches incorporating geometric cues and matching strategies. [BBK06] embeds one surface into another using Generalized Multidimensional Scaling to minimize distortion. [LF09, KLCF10, KLF11] explore and combine maps from the group of Möbius transforms of a surface, which can be constructed efficiently and include isometries. [GBAL09, SOG09] introduces the *heat kernel signature* (HKS), assigning a pointwise signature based on heat flow, and [OMMG10] uses a related technique to find maps with guaranteed behavior for nearly-isometric surfaces. The HKS construction is applied to the

wave equation in [ASC11] for experimentally more informative signatures. The algorithms in these and other papers find a *single* sparse or full map, whereas our new *representation* of a map can encode multiple correspondences.

The idea of a “fuzzy map” in terms of probability matrices was introduced in [WL78] using simple heuristics to construct and update maps. More recently, the Möbius transformations sampled in [LF09] generate a “fuzzy correspondence matrix” guiding point-to-point matching. Fuzzy schemes are also used to relax point-to-point mappings as in [BBM05,RCB97]. A probabilistic approach to mapping is taken in [ASP*04, TBW*11], though here distributions are *over* non-soft point-to-point mappings and thus subject to the rigidity of point-to-point schemes. Our optimization for finding soft maps has commonalities with theirs since their energy can be separated into unitary and binary terms, although ours is convex and thus not prone to local minima.

Some existing approaches use convex optimizations that are related to ours. The continuous relaxation of the integer program in [WSSC11] could be viewed as a soft map, although the output is harder to interpret. The relaxation of the graph isomorphism problem in [SU97] provides some analogous constructions to the constraints on and desired properties of soft maps for graphs; a related construction on hypergraphs is provided in [ZS08]. Earth Mover’s Distances (EMD) also have been applied to optimizations for several related vision and geometry problems. For instance, [LD11] uses them to construct a distance metric between surfaces invariant to Möbius transformations, and [HZTA04] uses EMD to guide image registration.

The work closest related to soft mapping, however, is that on measure couplings and Gromov-Wasserstein distances [Mém07, Mém09, Mém11]. Here, correspondences are discretizations of measure couplings, or probability distributions over the product of two surfaces whose marginalizations to the surfaces yield areas. The method is only acceptable for nearly-isometric surfaces admitting area-preserving correspondences. Furthermore, the optimization problem for finding measure couplings is non-convex with multiple local optima when either surface is symmetric.

Applications of soft maps overlap significantly with those of point-to-point mapping. Some overlapping applications are better suited to the probabilistic context. For instance, annotation transfer and other tasks operating on shapes at a coarse level can use soft maps directly. Additionally, since soft maps can encode multiple point-to-point maps, methods like [NBCW*11] for finding consistent maps within a collection may have a higher chance of success.

3. Soft Maps

We introduce soft maps between embedded surfaces $M_1, M_2 \subset \mathbb{R}^3$ and show how their continuous definition leads to a discretization on triangle meshes. We then show how to

quantify key properties of soft maps that will be incorporated into our linear programming construction in Section 4

3.1. Continuous Definition

Let M_1 and M_2 be two surfaces embedded in \mathbb{R}^3 with area measures μ_1 and μ_2 , resp.; we will assume the surfaces are rescaled so that $\mu_1(M_1) = \mu_2(M_2) = 1$. We can view $M_1 \times M_2$ as a four-dimensional manifold with volume measure $\mu_1 \otimes \mu_2$ induced by μ_1 and μ_2 .

The basic object we consider is a probability measure $P_{12} \in \text{Prob}(M_1 \times M_2)$, which we call a *soft correspondence* between M_1 and M_2 (see Figure 2(a)). We view $P_{12}(U \times V)$ as the probability that a pair of points $p_1 \in U \subseteq M_1$ and $p_2 \in V \subseteq M_2$ are related to each other. With this interpretation, the *uniform distribution* in which all pairs of points (p_1, p_2) are deemed equally likely to be related is given by the measure $\mu_1 \otimes \mu_2$, while the relationship $y = f(x)$ induced by the mapping $f : M_1 \rightarrow M_2$ is encoded by a δ -measure whose support is the surface $\{x, f(x) : x \in M_1\} \subseteq M_1 \times M_2$.




Figure 2: (a) A soft correspondence as a probability measure on the product of the surfaces M_1 and M_2 ; the soft maps (b) from M_1 to M_2 and (c) from M_2 to M_1 .

Given a soft correspondence P_{12} between M_1 and M_2 , we use conditional probabilities to define *soft maps*. In particular, for $U \subseteq M_1$ and $V \subseteq M_2$, we obtain the conditional distributions $P_{U2} \in \text{Prob}(M_2)$ and $P_{1V} \in \text{Prob}(M_1)$ given by

$$P_{U2}(B) \equiv \frac{P_{12}(U \times B)}{P_{12}(U \times M_2)} \text{ and } P_{1V}(A) \equiv \frac{P_{12}(A \times V)}{P_{12}(M_1 \times V)}$$

for any $A \subseteq M_1, B \subseteq M_2$. We define the distributions P_{U2} and P_{1V} to be the soft maps from U into M_2 and from V into M_1 , resp. The quantity $P_{U2}(A)$ gives the probability that a randomly chosen point $p_2 \in M_2$ will be in A given that it possesses a counterpart in U ; an equivalent interpretation holds for P_{1V} (see Figure 2(b,c)).

3.2. Discretization

We partition M_1 and M_2 into disjoint subsets $U_1, \dots, U_{N_1} \subseteq M_1$ and $V_1, \dots, V_{N_2} \subseteq M_2$ with $M_1 = \cup_{i=1}^{N_1} U_i$ and $M_2 = \cup_{i=1}^{N_2} V_i$. In practice, we model M_1 and M_2 with triangle meshes and divide them into patches using geodesic Voronoi cells about geodesic farthest-point samples. An approximate version of P_{12} with piecewise constant density on each $U_i \times V_j$ with respect to $\mu_1 \otimes \mu_2$ is:

$$\hat{P}_{12}(E) \equiv \sum_{ij} a_{ij} \cdot \frac{\mu_1 \otimes \mu_2((U_i \times V_j) \cap E)}{\mu_1 \otimes \mu_2(U_i \times V_j)} \quad \forall E \subseteq M_1 \times M_2$$

Substituting $E = U_i \times V_j$ for fixed i, j and enforcing $P_{12}(U_i \times V_j) = \hat{P}_{12}(U_i \times V_j)$ shows that $a_{ij} = P_{12}(U_i \times V_j)$. These satisfy $\sum_{ij} a_{ij} = 1$ since the partitions of M_1 and M_2 are disjoint, so the matrix $A = (a_{ij})$ defines a probability distribution over $\{U_1, \dots, U_{N_1}\} \times \{V_1, \dots, V_{N_2}\}$. In this construction, the conditional probabilities \hat{P}_{U_i} and \hat{P}_{V_j} are given by the normalized rows or columns of A . Note that point-to-point maps can be represented by taking A to be a permutation matrix and varying the sizes of the partitions.




Figure 3: (a) Two maps from S^1 to itself, (b) a soft correspondence superposing the two maps on the torus $S^1 \times S^1$, and (c) the corresponding matrix A .

It can be difficult to visualize distributions on the four-dimensional product $M_1 \times M_2$. If M_1 and M_2 are curves, however, the two-dimensional product $M_1 \times M_2$ can be visualized on the plane or using a toroidal topology. We thus show an example in Figure 3 of a soft correspondence between the circle S^1 and itself on the torus $S^1 \times S^1$. This distribution represents a convex combination of two point-to-point maps, illustrating the expressive power of soft maps.

3.3. Geometric Properties of Soft Maps

Descriptor Matching. Often with each region $U \subseteq M_s$ for each $s = 1, 2$ we can associate a descriptor $\phi_s(U) \in \mathbb{R}^k$ describing the geometry of U . We measure a soft map's non-conformity to these descriptors using the functional:

$$E_\phi(A) = \sum_{ij} a_{ij} \phi_{ij} \quad (1)$$

where $\phi_{ij} = \|\phi_1(U_i) - \phi_2(V_j)\|$. The quantity $E_\phi(A)$ is the expected mismatch ϕ with respect to the probability distribution determined by A . If $E_\phi(A)$ is small, the correspondences suggested by A align with the regions matched by ϕ . We take ϕ_{ij} to be the mean L^1 difference between wave kernel signatures [ASC11] of vertices in U_i and V_j .

Optimizing E_ϕ would place a single 1 in the entry of A corresponding to the smallest value of ϕ_{ij} . Restricting A to have equal row and column sums is not a strong enough regularizer, as this (when $N_1 = N_2$) yields a linear assignment problem [BDM09], whose solution is a permutation matrix. These techniques do not incorporate higher-order relationships between pairs of points. Thus, we must introduce additional energy terms to find a reasonable soft map.

Continuity. A desirable soft map should have some degree of continuity, which we define using the generalized

Earth Mover's Distance (GEMD) as in [PW09] for the distance between rows or columns of A . That is, suppose $c = (c_1, \dots, c_N)$ and $c' = (c'_1, \dots, c'_N)$ with $c_i, c'_i \geq 0 \forall i$ and the cost for transport between buckets i and j is $d_{ij} > 0$. Furthermore, since we won't assume $\sum_i c_i = 1$ or $\sum_j c'_j = 1$, we take a "garbage bin" cost $d \geq d_{ij} \forall i, j$. Then we define

$$E_{GEMD}(c, c') = \min_{f_{ij}, g_i, g'_j} \left[d \left(\sum_i g_i + \sum_j g'_j \right) + \sum_{ij} d_{ij} f_{ij} \right]$$

$$\text{subject to } g_i + \sum_j f_{ij} = c_i \forall i \quad g'_j + \sum_i f_{ij} = c'_j \forall j$$

$$f_{ij}, g_i, g'_j \geq 0 \forall i, j \quad (2)$$

E_{GEMD} measures the minimum flow $\sum_{ij} f_{ij}$ from one distribution to another with ground distance d_{ij} with the option of disposing mass g at cost d . It is easy to show that E_{GEMD} is a distance metric between distributions $c, c' \in (\mathbb{R}^+)^N$.

Continuity for soft maps means that nearby regions on one surface should map to nearby distributions with respect to GEMD on the other. Thus we define both $E_1 = \{(i, j) : U_i \text{ and } U_j \text{ are neighbors on } M_1\}$ and E_2 analogously for M_2 , and we measure the continuity of a map using the energy

$$E_{cont}(A) = \max \left\{ \max_{(i,j) \in E_1} E_{GEMD}(A_{row i}, A_{row j}), \max_{(i,j) \in E_2} E_{GEMD}(A_{col i}, A_{col j}) \right\}.$$

The costs are given by geodesic distances between Voronoi centroids. While E_{cont} is nonlinear in A , it can be computed using a linear program in A and the flow values f_{ij}, g_i, g'_j .

The above energy is well-founded in the continuous version of the soft maps formalism. To see how, we first define a map $\Psi : M_1 \rightarrow \text{Meas}(M_2)$ from points in M_1 to measures on M_2 as follows. If $P_{12} \in \text{Prob}(M_1 \times M_2)$ satisfies $dP_{12} = \rho d\mu_1 \otimes d\mu_2$ for some density $\rho : M_1 \times M_2 \rightarrow \mathbb{R}$, then we let $\Psi(x) \equiv \rho(x, \cdot) d\mu_2$. The distance between such measures can be defined using a variant of the L^1 -Wasserstein distance between probability measures and computed by solving a Monge-Kantorovich optimal transportation problem [Vil03]. This is the continuous version of the GEMD. Note that this formulation makes available other optimal transportation distances such as the L^p -Wasserstein distances. Now if M_1 carries the topology induced by geodesic distance and $\text{Meas}(M_2)$ the topology induced by continuous GEMD, then the local Lipschitz constant of Ψ at $x \in M_1$ is

$$\text{Lip}_x(\Psi) = \lim_{r \rightarrow 0} \sup_{x' \in B_r(x) \setminus \{x\}} \frac{E_{GEMD}(\Psi(x), \Psi(x'))}{\text{dist}(x, x')}$$

and the Lipschitz constant of Ψ is $\sup_{x \in M_1} \text{Lip}_x(\Psi)$. Therefore E_{cont} can be understood, up to a distance factor on the order of the distance between Voronoi centers, as a measure of the Lipschitz constant of Ψ and of the map $\Psi' : M_2 \rightarrow \text{Meas}(M_1)$ constructed as above with the roles of M_1 and M_2 reversed. From this point of view, our approach to

soft mappings can be seen as a generalization of Lipschitz-minimizing maps (e.g. [MST06]) to a probabilistic setting.

Sharpness. Suppose M_1 and its accompanying patches admit an exact self-symmetry $\tau : M_1 \rightarrow M_1$, as in the left-right symmetry of an idealized human model. Then, modifying $P_{12} \in \text{Prob}(M_1 \times M_2)$ via $\tilde{P}_{12}(U \times V) \equiv P_{12}(\tau(U) \times V)$ will not affect a soft map’s descriptor matching or continuity. In this case, it is unclear whether an optimization trading off between continuity and descriptor matching should return P_{12} , \tilde{P}_{12} , or a convex combination thereof.

Thus, one final useful quantity for understanding soft maps characterize how spread out the probabilities are in the mapping matrix. We measure this *sharpness* via

$$E_s(A) = \sum_{ij} a_{ij}^2 = \|A\|_{Fro}^2. \quad (3)$$

This quantity is convex in the entries of A and achieves its maximum value at a permutation matrix and its minimal value at the uniform distribution. It allows us to distinguish between different soft maps that equally obey continuity and descriptor matching constraints. Since we do not want to add artificial information that is not signalled by the geometry, we seek the maximally spread map of all possible ones with similar descriptor matching and soft continuity.

4. Creating Soft Maps

4.1. Finding Soft Correspondences via Optimization

We propose optimizing the following energy functional to find soft correspondences between M_1 and M_2 :

$$E_{map}(A) = E_\phi(A) + \lambda E_{cont}(A) + \beta E_s(A) \quad (4)$$

where a “softness factor” $\lambda \in [0, \infty)$ and a “sharpness factor” $\beta \in [0, \infty)$ control the tradeoff between descriptor matching, continuity, and sharpness.

While optimizing E_ϕ alone yields a trivial matrix A matching one pair of patches, adding E_{cont} allows the optimal correspondence to be *soft*. It encourages a tradeoff between the quality of the descriptor match and continuity needed to resolve local and global ambiguities in the matching problem. The parameter λ controls this tradeoff; as λ becomes sufficiently large, A becomes the uniform $1/N_1N_2$ matrix satisfying $E_{cont}(A) = 0$. Large β also encourages the uniform matrix, but in our case as explained in Section 4.2 we actually choose relatively small β so that the E_s term acts as a tie-breaker when the other terms are ambivalent about multiple maps. This is analogous to the role of Tikhonov regularization in least squares problems [TA77].

As we showed in Section 3.3, each term in (4) is computable using a convex program in A and some auxiliary variables. To reduce the number of variables, we use the formulation of EMD stated in [Tak10], approximating geodesic distances on the surface with distances along the patch connectivity graph with edge lengths using geodesic distances

between patch centers; we find this approximation yields much faster output with little to no difference in map quality.

4.2. Finding λ and β

For fixed λ, β , the optimization problem (4) provides a convex formulation for obtaining a soft map. We choose a small value of β and propose an automatic choice of λ as follows.




Figure 4: The two terms in (4) as functions of λ for a typical pair of meshes. Note the sharp transition to the uniform distribution when λ becomes large.

As illustrated in Figure 4, for fixed β , the transition of A to a uniform matrix as λ increases is not smooth but rather has a phase transition. We find that the largest λ before the map becomes uniform yields the best soft maps without instabilities due to values ϕ_{ij} being close but not identical.

When $\beta = 0$ we can examine the dual of the resulting linear program to reveal a modification that computes the exact λ value we desire and its accompanying soft map. This surprising result, outlined in Appendix A, avoids binary search on λ and instead allows our computation to occur in a single step. We also show how to modify the resulting objective to deal with the $\beta > 0$ case for small β .

Figure 5 shows examples of soft maps computed using the technique above. Our method reliably and robustly generates continuous maps that respect geometric criteria. Maps between near-isometric shapes tend to produce the sharpest soft maps, while those between non-isometric but related shapes (e.g. ants to quadrupeds) contain larger spread.

4.3. Implementation Details

Our method for computing soft maps was implemented in C++ using IBM CPLEX as a convex solver [IBM]. Running on a four-core 2.4 GHz i7 notebook processor with 8 GB of memory, the optimization from Section 4 took two to five minutes to generate each of the 80-patch maps in Figure 5. Figures in this paper take the garbage cost d in (3.3) to be $10\times$ the maximal geodesic distance between adjacent Voronoi centers. The main part of our pipeline taking appreciable time is the convex optimization for computing maps.

5. Analysis and Decomposition of Soft Maps

Unlike point-to-point maps, the probability matrices we have computed are amenable to linear-algebraic analysis to understand the information they encode.




Figure 5: Examples of soft maps (source models on the left and targets on the right). To save space, we superimpose individual soft maps like those in Figure 1 in a single target image. Multiple copies of a target are shown when the map is too spread-out for superposition. The first two examples here show this technique applied to the same initial examples from Figure 1.

5.1. Bases for Mapping

Suppose we are given a soft correspondence matrix $A = (a_{ij})$ computed using the method in Section 4. Since our correspondence is written in the identity basis, each row or column of A encodes a map from a single patch on M_1 or M_2 , resp. This representation is not necessarily the most compact. For instance, if a surface admits a symmetry that cannot be resolved by ϕ , then symmetric patches always should be coupled. Similarly, if there is not enough evidence to distinguish nearby patches, they also can be coupled.

Once we have computed A , however, we can seek bases on M_1 and M_2 that better respect such couplings *a posteriori*. In particular, projecting the uniform vector of ones out of the columns of A and performing a singular decomposition mimics the steps of Principal Components Analysis, yielding an orthogonal basis (including the uniform vector) for the column space of A . This basis provides a simple representation of the couplings exhibited in the column marginals of A , and the singular values provide an indication of the importance of each basis vector. Explicitly including the uniform vector allows us to guarantee that our basis can represent probability distributions. A similar process can be carried out on the rows of A for a basis on the target surface.

Figure 6 shows eight members of the basis M_2 resulting from SVD analysis of a map in Figure 5. The basis reveals patterns on M_2 that should be mapped together, respecting symmetries that are not disambiguated by ϕ ; the basis on M_1 is similar. Such bases indicate the mapping resolution and couplings that should be expected for continuous maps respecting a given ϕ . Figure 7




Figure 7: Singular values for the basis in Figure 6.

Figure 8 shows a low-rank map using the basis from Figure 6; it is nearly indistinguishable from the original map.

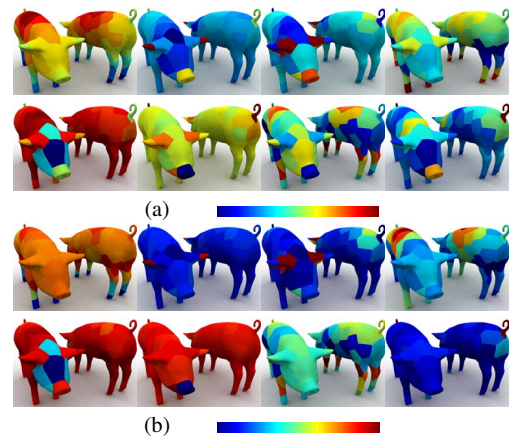


Figure 6: (a) The first eight SVD basis vectors from the first map in Figure 5, sorted by decreasing singular value, and (b) the same vectors “untangled” using [SBCBG11] to better show their support. Bases are colored using the scale below the images. These respect symmetries and are spread depending on the usefulness of ϕ for mapping each patch.

shows a plot of the singular values from our decomposition. These singular values have a relatively long tail, so low-rank approximations of A can be obtained by projecting onto a restricted basis; Figure 8 shows such a projection onto the first nine basis vectors.




Figure 8: A low-rank map using the basis from Figure 6; it is nearly indistinguishable from the original map.

5.2. Understanding Collections of Shapes

Using SVD bases as in Section 5.1 makes it possible to express a single map using a few basis vectors and a smaller correspondence matrix. In some sense, the compactness here is not surprising, since the bases are tailor-made for the map in question. The two bases, however, live on M_1 and M_2 *individually*, so in some sense the map is only expressed in the reduced correspondence matrix.

Consider now a *collection* of shapes M_1, \dots, M_n , each with its own patch decomposition. We seek a “probabilistic basis” on each shape that captures its maps to all the others. Computing such a basis is a simple extension of our previous method: we simply concatenate the outgoing maps to *all* other shapes, project out the uniform distribution, and perform SVD. Figure 9 shows the results of an experiment in which a database of twenty shapes is mapped pairwise using the method in Section 4 and analyzed using this technique. The resulting bases, illustrated for one shape in Figure 9(b), are more robust than those from Figure 6 since they are not subject to the particularities of a single map; they can be “untangled” using the method in [SBCBG11] (shown in Figure 9(c)) to illustrate their support more clearly, although this process does not affect their span. Figure 9(d) shows that approximation using these non-map-specific bases remains relatively effective, demonstrating their generality.

Some of the highest “errors” in Figure 9(d) are along the diagonal, which represents self maps. The fact that descriptors match exactly can make the identity map dominate soft mapping output regardless of symmetries, see Figure 10(a); thus we leave them out of the SVD computation. Projecting onto the shape’s mapping basis can alleviate this issue and makes for more symmetric self mappings as in Figure 10(b).




Figure 10: Maps from a shape to itself with source (a) using the optimization from Section 4 (b) and after projecting onto the reduced basis (c). The maps on the right are more symmetric at the cost of being more spread out.

6. Discussion and Limitations

As a new map representation, soft correspondences have properties and limitations that should be considered for future research and for the using the algorithms as stated. Most obviously, our method depends strongly on the choice of the dissimilarity measure ϕ , and as noted in [ASC11] much work

remains to be done to find effective descriptors in the non-isometric case. That said, as illustrated in Figure 11, our method is more robust to errors and inaccuracies in ϕ than methods that match the values directly, since the continuity term can help sharpen and disambiguate matchings.




Figure 11: A Gaussian in ϕ_{ij} from a map in Figure 1 (same source patches). Our optimization taking into account GEMD continuity greatly sharpens the map and corrects the incorrect tail-to-ear match shown in red.

Our optimization does not allow the shapes of the patches on M_1 and M_2 to change. Choosing patches that align with geometric features may make soft maps simpler or more sparse, and choosing an effective ϕ in this case may be easier. Even so, Figure 12 illustrates that our maps are fairly insensitive to the choice of patches by showing maps to a surface with different choices of Voronoi centers, except when they group salient features in different ways. Also, our optimization can handle surface decompositions on the order of hundreds of patches but cannot generate soft correspondences at the triangle level; while such dense mappings might be useful in some applications, we find that our decompositions are sufficient for capturing most geometric relationships.




Figure 12: Two maps from the source on the left to different discretizations of the same shape; both are qualitatively similar despite the differences in discretization of M_2 .

There are some applications for which point-to-point methods need considerable adjustment to make use of soft maps. For instance, it is not obvious how to transfer texture directly from one surface to another using a soft map, since using “expected” texture values likely will lead to smearing and blurring. More generally, some intuition for point-to-point mapping may stem from the diffeomorphism properties of a map, which no longer exist in the soft case.




Figure 9: Pairwise maps between the models in (a) yield bases for mapping on the shapes including those in (b), which can be untangled to yield bases (c). A subset of the models used in the paper is shown. The untangled basis reveals the power of the reduced-rank representation, directly coupling symmetric patches that cannot be disambiguated using ϕ or continuity. For instance, feet and legs are coupled in the untangled basis and thus always are mapped together. L^2 soft correspondence matrix approximation error using the truncated twenty-vector basis is shown in (d); the color bar scales between 0% and 25% error.

Test case	Soft maps			[KLF11]		
	\mathcal{E}_{av}	\mathcal{E}_{max}	E_S	\mathcal{E}_{av}	\mathcal{E}_{max}	E_S
Pigs	.67	.76	.017	.50	.74	.46
Humans	.64	.76	.027	.45	.73	.46
Bearings	.46	.60	.026	.30	.59	.48
Chairs	.77	.86	.019	.26	.40	.25
Cats	.72	.87	.020	.35	.63	1.87
Ants	.73	.86	.021	.52	.81	.46
Four-legged	.62	.72	.037	.42	.62	.47
Teddies	.62	.73	.017	.36	.67	.46

Table 1: Validation and scaled sharpness for maps in Figure 5 and comparison with maps computed using [KLF11].

6.1. Validation

Using traditional measures of map distortion and agreement with ground truth can be difficult for soft correspondences, since they encode maps in a novel way. Even so, given a sparse set of ground truth correspondences $(x_1, y_1), \dots, (x_m, y_m) \in M_1 \times M_2$ we generate a validation score as follows. Choose k_i and l_i such that $x_i \in U_{k_i}$ and $y_i \in V_{l_i}$. Then, we score a map using the metric:

$$\mathcal{E}(A) = \frac{1}{m} \sum_{i=1}^m \frac{1}{2} \left(\frac{A_{k_i l_i}}{\max_j A_{k_i j}} + \frac{A_{k_i l_i}}{\max_k A_{k j_i}} \right) \in [0, 1] \quad (5)$$

This metric rewards soft correspondences whose peaks align with the (x_i, y_i) pairs. We do not punish correspondences that contain symmetries, so maps with high sharpness alone are *not* rewarded. We show both average and maximum values of $\mathcal{E}(A)$ when multiple ground truth maps are available.

Table 1 contains validation scores \mathcal{E} and sharpness values E_S for the maps displayed in Figure 5, scaled so that patch-to-patch permutation maps have sharpness 1. We find that our optimization yields maps whose peaks align well with ground truth correspondences, even when WKS values do not match well. For comparison, we convert maps from [KLF11] to soft correspondences by tabulating how many points from each source patch are mapped to points in each target patch in forward and backward maps and re-normalizing; as expected, these maps have much higher sharpness E_S since they are converted from nearly continuous point-to-point maps, but their average symmetric *and*

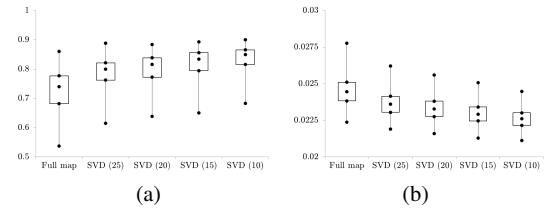


Figure 13: Box plot of validation scores (a) and scaled sharpness scores (b) for the pairwise human maps.

maximum validation scores are lower. Note that it is difficult to accomplish a reverse comparison in which our soft correspondences are converted to point-to-point maps, as the process for “sharpening” a soft map to one or many point-to-point maps is a topic for future research.

Figure 13 shows the results of a larger experiment in which all pairs of twenty human models from the SHREC 2007 database were mapped to each other; validation scores and scaled sharpness values E_S for the resulting maps as well as their SVD projections onto different numbers of bases are shown. Interestingly, the validation scores *increase* after SVD projection due to phenomena like that illustrated in Figure 10, at the cost of decreased sharpness.

7. Conclusions and Future Work

There are many ways to view soft correspondences within the larger context of mapping algorithms. Primarily, they serve as a new map representation acknowledging uncertainty in the mapping problem, improving upon dissimilarity matrices using continuity to cull false matches. They can also be viewed as superpositions of symmetric or slippage-prone point-to-point maps whose spread reflects potential mapping quality latent in a given descriptor. Regardless of interpretation, soft correspondences deal with local and global ambiguities gracefully, admit straightforward analysis, and can be computed using convex optimization.

Soft maps are a new way of representing and understand-

ing maps that brings many possible avenues for future inquiry. Soft maps as we have computed them can be used for tasks such as transfer of labels and other information between meshes using expectations. Allowing Voronoi centers to shift or be clustered might make soft maps more versatile, and replacing Voronoi regions with regions that better conform to salient surface geometry may provide more intuitive maps. Formalizing additional properties of soft maps such as the behavior of modes may help gain higher-level understanding of the underlying point-to-point maps. Methods for sharpening soft maps would make methods like [NBCW*11] applicable to cycle-based analysis of soft map collections despite the increased spread resulting from composition around long cycles. Our mapping basis computations might provide methods for expressing the possible maps between surfaces in a collection using few parameters. It also may be possible to combine probabilistic and functional approaches to mapping for alternative dense and compact map representations [OBCS*12].

Acknowledgments

This research was made with government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate Fellowship 32CFR168a. We also acknowledge the following grants: the Hertz Foundation Fellowship, the NSF GRF program, the Siebel Scholars program, NSF grants FODAVA 808515/CCF 1011228/CCF 1161480, a Google Research Award, and support from the Max Planck Center for Visual Computing and Communication.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV Workshops* (2011), pp. 1626–1633. 3, 4, 7
- [ASP*04] ANGUELOV D., SRINIVASAN P., PANG H.-C., KOLLER D., THRUN S., DAVIS J.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *NIPS* (2004). 3
- [BBK06] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 2
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 2
- [BBM05] BERG A., BERG T., MALIK J.: Shape matching and object recognition using low distortion correspondences. In *CVPR* (2005), vol. 1, IEEE, pp. 26–33. 3
- [BDM09] BURKARD R., DELL'AMICO M., MARTELLO S.: *Assignment problems*. SIAM, 2009. 4
- [CLM*11] CHANG W., LI H., MITRA N., PAULY M., RUSINKIEWICZ S., WAND M.: Computing correspondences in geometric data sets. *Eurographics tutorial* (2011). 2
- [GBAL09] GEBAL K., BÆRENTZEN J., AANÆS H., LARSEN R.: Shape analysis using the auto diffusion function. In *Comp. Graph. Forum* (2009), vol. 28, pp. 1405–1413. 2
- [HZTA04] HAKER S., ZHU L., TANNENBAUM A., ANGENENT S.: Optimal mass transport for registration and warping. *Int. J. Comp. Vision* 60, 3 (Dec. 2004), 225–240. 3
- [IBM] IBM, INC: IBM ILOG CPLEX optimizer. See <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>. 5
- [KLCF10] KIM V. G., LIPMAN Y., CHEN X., FUNKHOUSER T.: Möbius transformations for global intrinsic symmetry analysis. *Proc. SGP* (2010). 2
- [KLF11] KIM V., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *Proc. SIGGRAPH* (2011). 2, 8
- [LD11] LIPMAN Y., DAUBECHIES I.: Conformal wasserstein distances: Comparing surfaces in polynomial time. *Adv. Mathematics* 227, 3 (2011), 1047–1077. 3
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *Trans. Graph.* 28 (2009), 72:1–72:12. 2, 3
- [Mém07] MÉMOLI F.: On the use of Gromov-Hausdorff distances for shape comparison. In *Point Based Graphics* (2007), pp. 81–90. 3
- [Mém09] MÉMOLI F.: Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment* (2009). 3
- [Mém11] MÉMOLI F.: A spectral notion of Gromov-Wasserstein distance and related methods. *Applied and Computational Harmonic Analysis* 30, 3 (2011), 363–401. 3
- [MST06] MÉMOLI F., SAPIRO G., THOMPSON P.: Geometric surface and brain warping via geodesic minimizing lipschitz extensions. In *MFCA-2006 Int. Workshop on Math. Found. of Comp. Anatomy* (2006), pp. 58–67. 5
- [NBCW*11] NGUYEN A., BEN-CHEN M., WELNICKA K., YE Y., GUIBAS L.: An optimization approach to improving collections of shape maps. *Comp. Graph. Forum* 30, 5 (2011), 1481–1491. 3, 9
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *Proc. SIGGRAPH, to appear* (2012). 9
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *Comp. Graph. Forum* 29, 5 (2010), 1555–1564. 2
- [OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. In *Comp. Graph. Forum* (2008), vol. 27, pp. 1341–1348. 2
- [PW09] PELE O., WERMAN M.: Fast and robust earth mover's distances. In *ICCV* (2009). 4
- [RBBK10] RAVIV D., BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Full and partial symmetries of non-rigid shapes. *IJCV* 89, 1 (2010), 18–39. 2
- [RCB97] RANGARAJAN A., CHUI H., BOOKSTEIN F.: The soft-assign procrustes matching algorithm. In *Information Processing in Medical Imaging* (1997), Springer, pp. 29–42. 3
- [SBCBG11] SOLOMON J., BEN-CHEN M., BUTSCHER A., GUIBAS L.: Discovery of intrinsic primitives on triangle meshes. *Comp. Graph. Forum* 30, 2 (2011), 365–374. 6, 7
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP* (2009), pp. 1383–1392. 2
- [SU97] SCHEINERMAN E., ULLMAN D.: *Fractional graph theory*. Wiley, 1997. 3

- [TA77] TIKHONOV A., ARSEININ V.: *Solutions of ill-posed problems*. Scripta Series in Math. Winston, 1977. 5
- [Tak10] TAKANO Y.: Metric-preserving reduction of earth mover's distance. *Asia-Pac. J. Op. Res.* 27, 1 (2010), 39–54. 5
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.: Intrinsic shape matching by planned landmark sampling. In *Comp. Graph. Forum* (2011), vol. 30, pp. 543–552. 3
- [Vil03] VILLANI C.: *Topics in optimal transportation*, vol. 58. Amer Mathematical Society, 2003. 4
- [vKZHC010] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Eurographics STAR* (2010). 2
- [WL78] WITTE F. P., LUCAS D.: Probabilistic tracking in a multitarget environment. In *Decision and Control* (January 1978), vol. 17, pp. 1212–1216. 3
- [WSSC11] WINDHEUSER T., SCHLICKWEI U., SCHIMDT F., CREMERS D.: Large-scale integer linear programming for orientation preserving 3d shape matching. *Proc. SGP 30*, 5 (2011), 1471–1480. 3
- [ZS08] ZASS R., SHASHUA A.: Probabilistic graph and hypergraph matching. In *CVPR* (2008), pp. 1–8. 3

Appendix A: Parameter Choice

Let M_1, M_2 be two surfaces split into N_1, N_2 patches, respectively. Let $(\bar{U}, E_1), (\bar{V}, E_2)$ be the “segment graphs” on M_1, M_2 , where $\bar{U} = \{U_1, \dots, U_{N_1}\}, \bar{V} = \{V_1, \dots, V_{N_2}\}$, and E_1, E_2 encode the adjacency of patches on the two models. Take Φ to be the matrix of descriptor distances, and let A be the soft correspondence matrix we wish to compute.

Consider the optimization (4). For now assume $\beta = 0$, so that our problem becomes a linear program (LP). For convenience, we use $\phi, \mathbf{a} \in \mathbb{R}^{N_1 N_2}$ for the entries of Φ and A in vector form. Then, the primal linear program takes the form:

$$\begin{aligned} & \text{minimize} && \phi^T \mathbf{a} && + \lambda y \\ & \text{subject to} && && \\ (\alpha) & && \mathbf{1}^T \mathbf{a} && = 1 \\ (\gamma_e) & && C_e \mathbf{a} + D_e \mathbf{x} && = \mathbf{0} \quad \forall e \in E_1 \cup E_2 \\ (z_e) & && \mathbf{d}_e^T \mathbf{x} - y && \leq 0 \quad \forall e \in E_1 \cup E_2 \end{aligned}$$

Here, \mathbf{x} is a vector of auxiliary variables required to compute the GEMD, and C_e, D_e contain coefficients used to constrain the GEMD computation explained in Section 3.3. The vector \mathbf{d}_e contains unit coefficients for any variables containing GEMD values and zeros elsewhere; thus, the constraint (z_e) ensures that y contains the maximum GEMD. Variables in the above program are also constrained to be nonnegative. The dual of this LP is as follows:

$$\begin{aligned} & \text{maximize} && \alpha \\ & \text{subject to} && \\ (\mathbf{a}) & && \alpha \mathbf{1} + \sum_{e \in E_1 \cup E_2} C_e^T \gamma_e && \leq \phi \\ (\mathbf{x}) & && \sum_{e \in E_1 \cup E_2} D_e^T \gamma_e + \sum_{e \in E_1 \cup E_2} z_e \mathbf{d}_e && \leq \mathbf{0} \\ (y) & && - \sum_{e \in E_1 \cup E_2} z_e && \leq \lambda \end{aligned}$$

The variables z_e are constrained to be nonpositive and other variables are unconstrained. We modify this dual to find the value λ where the optimal soft correspondence A switches from being nonuniform to uniform. Since λ is isolated in the dual, we can turn it from a constant into a variable and turn the old dual objective into a constraint (σ) :

$$\begin{aligned} & \text{maximize} && && -\lambda \\ & \text{subject to} && && \\ (\mathbf{a}) & && \alpha \mathbf{1} + \sum_{e \in E_1 \cup E_2} C_e^T \gamma_e && \leq \phi \\ (\mathbf{x}) & && \sum_{e \in E_1 \cup E_2} D_e^T \gamma_e + \sum_{e \in E_1 \cup E_2} z_e \mathbf{d}_e && \leq \mathbf{0} \\ (y) & && && - \sum_{e \in E_1 \cup E_2} z_e - \lambda \leq 0 \\ (\sigma) & && \alpha && = \frac{\|\phi\|_1}{N_1 N_2} \end{aligned}$$

Intuitively, this LP seeks the minimal λ yielding a uniform soft correspondence. In particular, the new constraint (σ) applies strong duality to enforce a uniform optimal correspondence, since the optimal energy of the primal LP for the uniform map is $\|\phi\|_1 / N_1 N_2$, the sum of descriptor differences with no contribution from GEMD. Again, z_e is nonpositive. We simplify this program by substituting $\alpha = \|\phi\|_1 / N_1 N_2$ and then dual this modified dual to get a new primal LP:

$$\begin{aligned} & \text{minimize} && \phi^T \mathbf{a} - \frac{\|\phi\|_1}{N_1 N_2} \mathbf{1}^T \mathbf{a} \\ & \text{subject to} && \\ (\gamma_e) & && C_e \mathbf{a} + D_e \mathbf{x} && = \mathbf{0} \quad \forall e \in E_1 \cup E_2 \\ (z_e) & && \mathbf{d}_e^T \mathbf{x} - y && \leq 0 \quad \forall e \in E_1 \cup E_2 \\ (\lambda) & && && -y = -1 \end{aligned}$$

All variables again are constrained to be nonnegative. Finally, we eliminate y to obtain our final program:

$$\begin{aligned} & \text{minimize} && \phi^T \mathbf{a} - \frac{\|\phi\|_1}{N_1 N_2} \mathbf{1}^T \mathbf{a} \\ & \text{subject to} && \\ (\gamma_e) & && C_e \mathbf{a} + D_e \mathbf{x} && = \mathbf{0} \quad \forall e \in E_1 \cup E_2 \\ (z_e) & && \mathbf{d}_e^T \mathbf{x} && \leq 1 \quad \forall e \in E_1 \cup E_2 \end{aligned}$$

If we solve this LP, then the optimal objective is equal to the threshold $-\lambda$, where any larger λ would provide a uniform distribution. Additionally, this LP finds the soft correspondence $A / \|A\|_1$ whose matching score is as good as possible to make up for its EMD error at the λ threshold. To incorporate spread-based regularization, we simply add $\beta \|A\|_{F_{\text{row}}}^2$ to the objective above for small $\beta > 0$. This modification maintains the convexity of the program, so interior-point solvers do not suffer due to nonlinearity; in fact, the regularization increases the conditioning of the convex program, which in practice converges much faster with such a modification.