

# Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching

Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel<sup>†</sup>

Department of Computer Science, Technion Israel Institute of Technology, Haifa 32000, Israel

Edited by Alexandre J. Chorin, University of California, Berkeley, CA, and approved December 5, 2005 (received for review October 3, 2005)

An efficient algorithm for isometry-invariant matching of surfaces is presented. The key idea is computing the minimum-distortion mapping between two surfaces. For this purpose, we introduce the generalized multidimensional scaling, a computationally efficient continuous optimization algorithm for finding the least distortion embedding of one surface into another. The generalized multidimensional scaling algorithm allows for both full and partial surface matching. As an example, it is applied to the problem of expression-invariant three-dimensional face recognition.

Gromov–Hausdorff distance | isometric embedding | iterative-closest-point | partial embedding

The problem of comparison between deformable surfaces has recently gained a lot of interest in pattern recognition, especially in the fields of three-dimensional (3D) data analysis and synthesis. It arises, for example, in medical imaging and 3D face recognition (1). The fundamental question is how to efficiently yet accurately quantify the similarity between a given reference surface (“model”) and some other surface (“probe”), usually a bent version of the model. We would like to capture the distinction of the model and the probe intrinsic properties associated with the metric structure of the surface, while ignoring the extrinsic properties that describe the way the surface is immersed into the ambient space and that often change while the surface bends. A deformation that preserves the intrinsic structure of the surface is called an “isometry.” Our goal is thus to define a computable isometry-invariant measure of similarity between surfaces. Often, the problem is even more complicated, when having the probe only partially available (see Fig. 1). We refer to the latter setting as “partial surface matching” and to the case where the whole probe is available as “full surface matching.”

One of the earliest attempts of isometry-invariant surface matching is the classical “iterative closest point” (ICP) algorithm (2). It addresses a particular case of the partial matching problem where only rigid (Euclidean) isometries are allowed. An efficient method for the construction of near-isometry-invariant representations of surfaces [called “canonical forms” (CFs)] based on Euclidean embeddings was presented in ref. 3, as a generalization of ref. 4. This approach used a multidimensional scaling (MDS) algorithm (5). MDS is closely related to dimensionality reduction (6, 7) and can be performed in a computationally efficient manner. Euclidean embeddings are used in theoretical computer science for representing metric spaces usually arising from geometry of graphs (8).

Recently, Mémoli and Sapiro (9) used a discrete probabilistic approximation of the Gromov–Hausdorff (GH) distance (10) to compare surfaces up to isometries. Their method is based on evaluating all of the permutations between two sets of surface samples, which is computationally expensive. Moreover, the GH distance is inappropriate for partial matching.

One of the main contributions of this work is the “partial embedding (PE) distance,” motivated by Gromov’s theory. For its efficient computation, we introduce an extension of MDS, which we call generalized MDS (GMDS). The key idea is to measure the minimum possible distortion when trying to iso-

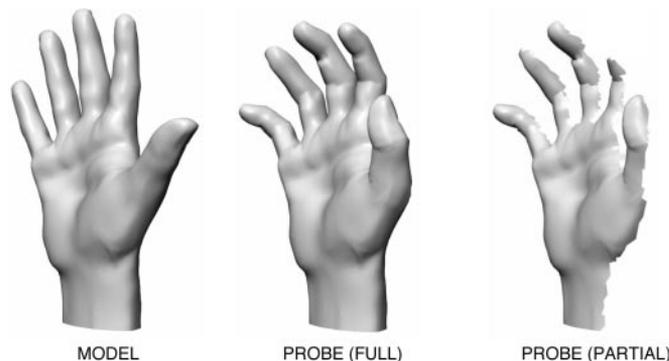


Fig. 1. Illustration of the deformable surface matching problem.

metrically embed one surface into another. By using GMDS, we can handle full as well as partial surface matching; this ability is one advantage over a straightforward use of the GH distance.

The paper includes six sections. *Theoretical Foundations* reviews the theoretical background of isometry-invariant surface matching. We define our PE distance and show its relation to the GH distance. The next section, *GMDS*, introduces the GMDS and addresses implementation and optimization considerations. The following section focuses on partial matching of surfaces. In *Results*, we show some experimental results. Finally, in *Conclusions*, we summarize the main results.

## Theoretical Foundations

Let  $\mathcal{S}$  be a smooth connected and compact Riemannian surface (2-manifold) with minimal geodesics  $C_{\mathcal{S}}^*(s, s')$  between every two points  $s, s'$  on  $\mathcal{S}$ . We define the geodesic distance function  $d_{\mathcal{S}}: \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}^+$  as  $d_{\mathcal{S}}(s, s') = \text{length}\{C_{\mathcal{S}}^*(s, s')\}$ . For simplicity, we say “surface  $\mathcal{S}$ ,” implying the underlying metric space  $(\mathcal{S}, d_{\mathcal{S}})$  induced by the Riemannian geometry of  $\mathcal{S}$ . In our context, surfaces are perceived from a point of view of metric rather than Riemannian geometry; in other words, we do not consider the Riemannian metric tensor, but the geodesic distances that it induces.

In practice, we work with sampled surfaces represented by finite metric spaces. We use the following definitions: a subset  $\mathcal{S}' \subset \mathcal{S}$  is called an “ $r$ -covering” (or “ $r$ -net”) of  $\mathcal{S}$  if  $\mathcal{S} = \cup_{s \in \mathcal{S}'} B_{\mathcal{S}}(s, r)$ , where  $B_{\mathcal{S}}(s_0, r) = \{s \in \mathcal{S} : d_{\mathcal{S}}(s, s_0) < r\}$  is a ball of radius  $r$  around  $s_0$  in  $\mathcal{S}$ . A finite  $r$ -covering of  $\mathcal{S}$  consisting of  $N$  points is denoted by  $\mathcal{S}'_N$ . The metric on  $\mathcal{S}'_N$  is assumed to be the restricted metric  $d_{\mathcal{S}'}(s, s') = d_{\mathcal{S}}(s, s')$  for all  $s, s'$  in  $\mathcal{S}'_N$ . We refer to  $r$  as the “sampling radius.” An arbitrary finite sampling of  $\mathcal{S}$  consisting of  $N$  points is denoted by  $\mathcal{S}_N$ , where  $\mathcal{S}$  is called “continuous” surface and  $\mathcal{S}_N$  a “discrete” one.

Conflict of interest statement: No conflicts declared.

This paper was submitted directly (Track II) to the PNAS office.

Abbreviations: CF, canonical form; FMM, fast marching method; GH, Gromov–Hausdorff; MDS, multidimensional scaling; GMDS, generalized MDS; PE, partial embedding.

<sup>†</sup>To whom correspondence should be addressed. E-mail: ron@cs.technion.ac.il.

© 2006 by The National Academy of Sciences of the USA

Given two surfaces  $\mathcal{S}$  and  $\mathcal{Q}$ , a transformation  $\psi: \mathcal{Q} \rightarrow \mathcal{S}$  is said to have “distortion”  $\varepsilon$  if

$$\text{dis } \psi \equiv \sup_{q, q' \in \mathcal{Q}} |d_{\mathcal{Q}}(q, q') - d_{\mathcal{S}}(\psi(q), \psi(q'))| = \varepsilon. \quad [1]$$

If in addition  $\psi$  is “ $\varepsilon$ -surjective” (i.e.,  $d_{\mathcal{S}}(s, \psi(\mathcal{Q})) \leq \varepsilon$  for all  $s \in \mathcal{S}$ ), it is called an “ $\varepsilon$ -isometry,” and  $\mathcal{S}$  and  $\mathcal{Q}$  are called “ $\varepsilon$ -isometric.” A transformation  $\psi: \mathcal{Q} \rightarrow \mathcal{S}$  with  $\text{dis } \psi = 0$  is called an “isometry,” and  $\mathcal{S}$  and  $\mathcal{Q}$  admitting such a transformation are called “isometric.”

**Desiderata.** Let  $\mathbb{M}$  denote the set of all compact surfaces, and  $\tilde{\mathbb{M}}$  denote the space of the equivalence classes of all isometries in  $\mathbb{M}$ , i.e., a space in which a point represents all of the isometries of a compact surface. The similarity between two surfaces can be measured by a nonnegative function  $d: \mathbb{M} \times \mathbb{M} \rightarrow \mathbb{R}^+$ . Slightly abusing common terminology, we call  $d$  a “distance,” even if it is not a metric (often, the terms distance and metric are used as synonyms). Ideally,  $d$  should have the following properties:

**Desideratum 1: Metric properties.**  $d$  should be a metric on  $\mathbb{M}$ , i.e., satisfy the following conditions: (i)  $d(\mathcal{S}, \mathcal{Q}) \geq 0$ ; (ii)  $d(\mathcal{S}, \mathcal{Q}) = 0$  if and only if  $\mathcal{S}$  and  $\mathcal{Q}$  are isometric; (iii)  $d(\mathcal{S}, \mathcal{Q}) = d(\mathcal{Q}, \mathcal{S})$ ; and (iv)  $d(\mathcal{R}, \mathcal{Q}) \leq d(\mathcal{R}, \mathcal{S}) + d(\mathcal{S}, \mathcal{Q})$ .

**Desideratum 2: Good similarity measure.** (i) If  $d(\mathcal{S}, \mathcal{Q}) \leq \varepsilon$ , then  $\mathcal{S}$  and  $\mathcal{Q}$  are  $c\varepsilon$ -isometric, where  $c > 0$ ; (ii) if  $\mathcal{S}$  and  $\mathcal{Q}$  are  $\varepsilon$ -isometric, then  $d(\mathcal{S}, \mathcal{Q}) \leq c\varepsilon$ . Note that this property is an extension of the second metric axiom to  $\varepsilon$ -isometries.

**Desideratum 3: Partial matching.** Given  $\mathcal{S}'$  a patch of  $\mathcal{S}$ ,  $d(\mathcal{S}, \mathcal{S}') = 0$ .

**Desideratum 4: Sampling consistency.** Given  $\mathcal{S}'$  a finite  $r$ -covering of  $\mathcal{S}$ ,  $\lim_{r \rightarrow 0} d(\mathcal{S}', \mathcal{Q}) = d(\mathcal{S}, \mathcal{Q})$ .

**Desideratum 5: Efficiency.**  $d$  should be efficiently computable, or at least efficiently approximated. Practically, given  $\mathcal{S}_N$  and  $\mathcal{Q}_{N'}$ ,  $d(\mathcal{S}_N, \mathcal{Q}_{N'})$  should be computable in polynomial time of  $\max\{N, N'\}$ .

Unfortunately, *Desiderata 1–3* cannot be satisfied simultaneously. In practical applications, it appears that the ability to perform partial matching (*Desideratum 3*) is the most important. In the following, we show a distance that satisfies *Desiderata 3–5* and partially satisfies *Desiderata 1* and *2*.

It is also important to note the following potential danger of partial matching. Consider an extreme case of  $\mathcal{S}'$  a smooth curve on  $\mathcal{S}$ . Obviously, if *Desideratum 3* holds, then  $d(\mathcal{S}, \mathcal{S}') = 0$ , but because a curve can be isometrically embedded into any surface, we are liable to obtain  $d(\mathcal{Q}, \mathcal{S}') = 0$  for an arbitrary  $\mathcal{Q}$  which is not necessarily an isometry of  $\mathcal{S}$ . In other words, when the patch  $\mathcal{S}'$  is too small, partial matching can be meaningless.

**GH Distance.** Our starting point is the GH distance introduced by Gromov (10). For two surfaces  $\mathcal{S}$  and  $\mathcal{Q}$ , it is defined by

$$d_{\text{GH}}(\mathcal{S}, \mathcal{Q}) \equiv \inf_{\mathcal{Z}, \phi, \psi} d_{\text{H}}^{\mathcal{Z}}(\phi(\mathcal{S}), \psi(\mathcal{Q})), \quad [2]$$

where  $\phi: \mathcal{S} \rightarrow \mathcal{Z}$  and  $\psi: \mathcal{Q} \rightarrow \mathcal{Z}$  are isometric embeddings into the metric space  $\mathcal{Z}$ , and  $d_{\text{H}}^{\mathcal{Z}}(\mathcal{A}, \mathcal{B}) \equiv \max\{\sup_{a \in \mathcal{A}} d_{\mathcal{Z}}(a, \mathcal{B}), \sup_{b \in \mathcal{B}} d_{\mathcal{Z}}(b, \mathcal{A})\}$  is the Hausdorff distance between two subsets  $\mathcal{A}$  and  $\mathcal{B}$  of  $\mathcal{Z}$ . For bounded metric spaces  $\mathcal{S}$  and  $\mathcal{Q}$ , the GH distance can be defined alternatively (11) as

$$d_{\text{GH}}(\mathcal{S}, \mathcal{Q}) = \frac{1}{2} \inf_{\substack{\phi: \mathcal{S} \rightarrow \mathcal{Q} \\ \psi: \mathcal{Q} \rightarrow \mathcal{S}}} \max\{\text{dis } \phi, \text{dis } \psi, \text{dis}(\phi, \psi)\}, \quad [3]$$

where  $\text{dis}(\phi, \psi) \equiv \sup_{s \in \mathcal{S}, q \in \mathcal{Q}} |d_{\mathcal{S}}(s, \psi(q)) - d_{\mathcal{Q}}(q, \phi(s))|$ .

**Proposition 1.** (i)  $d_{\text{GH}}$  is a finite metric on  $\tilde{\mathbb{M}}$ ; (ii) if  $d_{\text{GH}}(\mathcal{S}, \mathcal{Q}) \leq \varepsilon$ , then  $\mathcal{S}$  and  $\mathcal{Q}$  are  $2\varepsilon$ -isometric; (iib) if  $\mathcal{S}$  and  $\mathcal{Q}$  are  $\varepsilon$ -isometric, then  $d_{\text{GH}}(\mathcal{S}, \mathcal{Q}) \leq 2\varepsilon$ ; (iii)  $|d_{\text{GH}}(\mathcal{S}', \mathcal{Q}) - d_{\text{GH}}(\mathcal{S}, \mathcal{Q})| \leq r$ .

For proof, see ref. 11. The GH distance is a good candidate for

our “ideal” similarity measure between surfaces; from *Proposition 1*, it follows that it satisfies *Desiderata 1, 2, and 4*. However, there are two main disadvantages. First, the definition in Eq. 2 is hard to compute (*Desideratum 5*). Secondly,  $d_{\text{GH}}$  does not allow partial matching between surfaces (*Desideratum 3*). Given a surface  $\mathcal{S}$  and a patch  $\mathcal{S}' \subset \mathcal{S}$  such that  $\text{diam}(\mathcal{S} \setminus \mathcal{S}') \leq R$  (i.e.,  $\mathcal{S}'$  is an  $R$ -covering of  $\mathcal{S}$ ), we can have  $d_{\text{GH}}(\mathcal{S}, \mathcal{S}') = R$ , although we would have liked the distance between the surface and its patch to be zero.

**PE Distance.** Our assumption was that we have a model surface  $\mathcal{S}$  and a probe surface  $\mathcal{Q}$ , which can be an isometrically bent patch of  $\mathcal{S}$ . We have shown the disadvantage of the GH distance in such a setup. To allow for partial matching, we define the PE distance

$$d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) \equiv \frac{1}{2} \inf_{\psi: \mathcal{Q} \rightarrow \mathcal{S}} \text{dis } \psi.$$

Intuitively,  $d_{\text{PE}}$  measures the metric distortion obtained trying to embed  $\mathcal{Q}$  into  $\mathcal{S}$  in the “most isometric” way. It can be considered as only part of the GH distance, hence the name “partial embedding.”

**Proposition 2.** Given compact surfaces  $\mathcal{S}$ ,  $\mathcal{Q}$ , and  $\mathcal{R}$ ,  $d_{\text{PE}}$  satisfies the following: (i)  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) \geq 0$ ; (ii)  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) = 0$  if and only if  $\mathcal{Q}$  is isometrically embeddable into  $\mathcal{S}$ ; and (iii)  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) \leq d_{\text{PE}}(\mathcal{S}, \mathcal{R}) + d_{\text{PE}}(\mathcal{R}, \mathcal{Q})$ .

The proof closely follows the proof of theorem 7.3.30 in ref. 11. Still,  $d_{\text{PE}}$  is not symmetric, and furthermore, not a metric on  $\tilde{\mathbb{M}}$ . Yet, it satisfies some important metric properties, namely: isometry invariance (ii) and a nonsymmetric version of the triangle inequality (iii). Moreover, because  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) \leq d_{\text{GH}}(\mathcal{S}, \mathcal{Q})$ , it trivially follows that if  $\mathcal{S}$  and  $\mathcal{Q}$  are  $\varepsilon$ -isometric, then  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) \leq 2\varepsilon$ . The converse appears to be true under some more restrictive assumptions.

Given  $\mathcal{Q}'$  an  $r'$ -covering of  $\mathcal{Q}$  and  $\mathcal{S}'$  an  $r$ -covering of  $\mathcal{S}$ , the distance  $d_{\text{PE}}$  satisfies the following important properties.

**Proposition 3.** (i)  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) - r' \leq d_{\text{PE}}(\mathcal{S}, \mathcal{Q}') \leq d_{\text{PE}}(\mathcal{S}, \mathcal{Q})$ ; (ii)  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) \leq d_{\text{PE}}(\mathcal{S}', \mathcal{Q}) \leq d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) + r$ .

Note that a particular case of *Proposition 3i* is  $d_{\text{PE}}(\mathcal{Q}, \mathcal{Q}') = 0$ . By virtue of this property [unlike the corresponding property  $d_{\text{GH}}(\mathcal{Q}, \mathcal{Q}') \leq r$  of the GH distance]  $d_{\text{PE}}$  can be used for partial matching. From *Proposition 3* it also follows that  $d_{\text{PE}}$  is consistent with sampling, namely, given  $\mathcal{S}_N$  and  $\mathcal{Q}_{N'}$ , two finite coverings of  $\mathcal{S}$  and  $\mathcal{Q}$ , respectively,  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) - r' \leq d_{\text{PE}}(\mathcal{S}_N, \mathcal{Q}_{N'}) \leq d_{\text{PE}}(\mathcal{S}, \mathcal{Q}) + r$ . In other words, if we have a sufficiently dense finite sampling of  $\mathcal{S}$  and  $\mathcal{Q}$ ,  $d_{\text{PE}}(\mathcal{S}_N, \mathcal{Q}_{N'})$  is a sufficiently good approximation of the continuous  $d_{\text{PE}}(\mathcal{S}, \mathcal{Q})$ .

### GMDS

We now address a practical question of computing the PE distance  $d_{\text{PE}}$ , starting from the CF approach (3). As an input, we assume to be given two sampled surfaces  $\mathcal{S}_N$  and  $\mathcal{Q}_{N'}$  (representing the continuous surfaces  $\mathcal{S}$  and  $\mathcal{Q}$ ) and the geodesic distances between the samples  $\{s_i\}_{i=1}^N$  and  $\{q_i\}_{i=1}^{N'}$ , represented by an  $N \times N$  matrix  $\mathbf{D}_{\mathcal{S}_N} = (d_{\mathcal{S}_N}(s_i, s_j))$  and an  $N' \times N'$  matrix  $\mathbf{D}_{\mathcal{Q}_{N'}}$ , respectively.

**CF and MDS.** In ref. 3, a three-stage algorithm for near-isometry-invariant surface matching was proposed. First, each of the sampled surfaces is mapped into an  $m$ -dimensional Euclidean space by a near-isometric embedding (“flattening”), obtained by minimization of the stress function

$$\sigma(\mathbf{X}; \mathbf{D}_{\mathcal{S}_N}) = \frac{1}{N} \sum_{i > j} (d_{\mathbb{R}^m}(x_i, x_j) - d_{\mathcal{S}}(s_i, s_j))^2, \quad [4]$$

with respect to  $\mathbf{X}$ . Here  $\mathbf{X}$  denotes an  $N \times m$  matrix of coordinates in  $\mathbb{R}^m$  and  $d_{\mathbb{R}^m}$  denotes the Euclidean metric. The stress can be thought of as an  $L_2$  measure of the distance distortion caused by such an embedding. Algorithms minimizing the stress  $\sigma$  are known as MDS (5).  $\mathbf{X}_{S_N} = \operatorname{argmin}_{\mathbf{X}} \sigma(\mathbf{X}; \mathbf{D}_{S_N})$  was called the CF of  $\mathcal{S}$  and  $\min_{\mathbf{X}} \sigma(\mathbf{X}; \mathbf{D}_{S_N})$  the “embedding error.”

Note that both CFs  $\mathbf{X}_{S_N}$  and  $\mathbf{X}_{Q_N}$  are defined up to a Euclidean isometry (translation, rotation, and reflection). The second stage of the CF algorithm is an alignment of the CFs in  $\mathbb{R}^m$ . It is obtained by diagonalizing the matrix of the second-order moments of each CF and reordering the axes such that the diagonal values are decreasing. At the third stage, a distance based on high-dimensional moments (12) is used to compare between  $\mathbf{X}_{S_N}$  and  $\mathbf{X}_{Q_N}$  after alignment. We denote the distance computed as rigid matching between the CFs of  $\mathcal{S}$  and  $\mathcal{Q}$  by  $d_{\text{CF}}(\mathcal{S}, \mathcal{Q})$ .

Strictly speaking, the CF algorithm satisfies only *Desideratum 5*: it is very efficient computationally. Unlike alternative combinatorial optimization, MDS is a continuous optimization problem that can be solved by using standard optimization methods. Recently, a very fast multigrid implementation of the CF algorithm was presented (13). However, the CF algorithm lacks true isometry invariance, because, in general, a surface cannot be isometrically embedded into  $\mathbb{R}^m$  (8). Therefore, this “flattening” introduces a distortion. Partial matching is also complicated using this method.

**GMDS.** The choice of a Euclidean embedding space was followed by other choices, like spherical (14, 15) and hyperbolic (16) spaces, which proved to be advantageous for certain types of surfaces. The leitmotif of these approaches was choosing an embedding space in which geodesic distances can be expressed analytically. In practice, this possibility exists only in very few spaces.

So far, we have a good distance  $d_{\text{PE}}$ , which satisfies our *Desiderata* on one hand and an efficient MDS framework for computing Euclidean embeddings on the other. Next, we unite theory and practice. We consider a generalization of the CF algorithm, where surface  $\mathcal{Q}$  is embedded into another surface  $\mathcal{S}$ . We call this procedure GMDS.

Similarly to the Euclidean case, we define the “generalized stress” as

$$\sigma_p(\mathbf{U}; \mathbf{D}_{Q_N}, d_{\mathcal{S}}, \mathbf{W}) = \left( \frac{1}{\sum_{j>i} w_{ij}} \sum_{j>i} (w_{ij}(d_{\mathcal{S}}(\mathbf{u}_i, \mathbf{u}_j) - d_{\mathcal{Q}}(q_i, q_j)))^p \right)^{\frac{1}{p}}, \quad [5]$$

for  $1 \leq p < \infty$ , and

$$\sigma_{\infty}(\mathbf{U}; \mathbf{D}_{Q_N}, d_{\mathcal{S}}, \mathbf{W}) = \max_{i,j=1, \dots, N'} w_{ij} |d_{\mathcal{S}}(\mathbf{u}_i, \mathbf{u}_j) - d_{\mathcal{Q}}(q_i, q_j)|, \quad [6]$$

for  $p = \infty$ , where the matrix  $\mathbf{U}$  represents the positions of  $N'$  points on  $\mathcal{S}$  in some local or global parametric coordinates  $\mathbf{u}_i$ , and  $\mathbf{W} = (w_{ij})$  is a symmetric matrix of nonnegative weights. Note two major differences between the generalized stress and the stress defined in Eq. 4. First, the metric in the embedding space is not given analytically and is approximated numerically. Second, we replaced the  $L_2$  norm with a general  $L_p$  norm. Furthermore, the embedding space is 2D, whereas in the CF algorithm the embedding space is usually at least 3D.

Choosing  $p = \infty$  and  $\mathbf{W} = \mathbf{1}_{N' \times N'}$  we easily obtain

$$d_{\text{PE}}(\mathcal{S}, Q_{N'}) = \frac{1}{2} \min_{\mathbf{U}} \sigma_{\infty}(\mathbf{U}; \mathbf{D}_{Q_{N'}}, d_{\mathcal{S}}, \mathbf{1}_{N' \times N'}). \quad [7]$$

In other words, minimization of the generalized stress allows computing  $d_{\text{PE}}$ . We discuss numerical optimization algorithms in the next section. Moreover, it is possible to use a similar procedure to compute  $d_{\text{GH}}$  (A.M.B., M.M.B., and R.K., unpublished results). In practice, it may appear that using other norms (e.g.,  $p = 1$  or 2) is advantageous. In *Results*, it is demonstrated that we can use  $\sigma_p$  instead of  $\sigma_{\infty}$ . We denote the distances computed in this way by  $d_{\text{PE}}^p$ .

An important issue is the choice of the parametric coordinate system. For objects acquired by a range sensor, a global parameterization  $U \rightarrow \mathcal{S}$  (e.g.,  $U = [0, 1]^2$ ), is usually readily available. Global parameterization for objects with simple topology (e.g., homeomorphic to a disk or a sphere) is also very natural and can be computed from a general mesh by a variety of efficient algorithms (17). For general triangulated meshes with more sophisticated topology, a single global coordinate system may not suffice. In this case, we propose to resort to a local coordinate system, in which each point on the polyhedral surface is represented by an index of the triangle to which it belongs and a vector in barycentric coordinates of that triangle.

**Computational Aspects.** For  $1 \leq p < \infty$ , the generalized stress (5) is minimized by iterative gradient-type optimization methods. Here, for simplicity, we show the projected gradient descent algorithm for a global parameterization  $U \rightarrow \mathcal{S}$ ; in practice, more efficient optimization methods such as conjugate gradients or quasi-Newton can be used (18). Note that because the problem is nonconvex, such methods do not necessarily converge to the global minimum. Local convergence is a common property of MDS algorithms in general (5). There are standard techniques to overcome this problem, one of which is multiscale or multigrid optimization (13, 19).

For globally parameterized surfaces, the basic projected gradient descent step has the form

$$\mathbf{U}^{(k+1)} = P_U \left( \mathbf{U}^{(k)} - \alpha^{(k)} \nabla_{\mathbf{U}} \sigma_p(\mathbf{U}^{(k)}) \right), \quad [8]$$

where  $k$  is the iteration number.  $\alpha^{(k)}$  is the step size on iteration  $k$ ; it is computed by using line search (18).  $P_U$  is a projection operator onto the parameterization domain  $U$ .  $\nabla_{\mathbf{U}} \sigma_p$  denotes the gradient of  $\sigma_p$  with respect to  $\mathbf{U}$ ; its computation involves the distance function  $d_{\mathcal{S}}$  and its derivatives. When local parameterization is used, the optimization algorithm becomes more elaborate because it requires unfolding of triangular faces of the mesh and transformations from barycentric coordinate systems of adjacent triangles to compute the stress along a polylinear path.

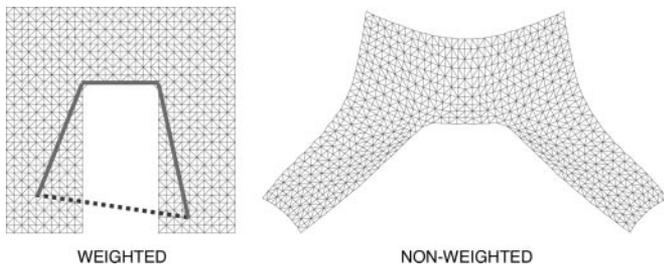
For  $p = \infty$ , we rewrite the problem as constrained optimization with  $2N' + 1$  variables, a linear objective, and  $N'(N' - 1)$  nonlinear inequality constraints using an artificial variable  $\tau$ ,

$$\min_{\tau \geq 0} \tau \text{ s.t. } |w_{ij}(d_{\mathcal{S}}(\mathbf{u}_i, \mathbf{u}_j) - d_{\mathcal{Q}}(q_i, q_j))| \leq \tau, \quad [9]$$

for all  $j > i$ . The new problem can be solved by using standard constrained optimization methods, e.g., penalty-barrier or augmented Lagrangian (18).

A crucial part of the GMDS problem is the computation of geodesic distances. We use the “fast marching method” (FMM), which computes geodesic distances on surfaces by solving the eikonal equation on general triangulated meshes (20). For surfaces with global parameterization, an efficient parametric version of FMM can be used (21).

The geodesic distances  $d_{\mathcal{Q}}(q_i, q_j)$  between the fixed sample points of  $\mathcal{Q}$  can be precomputed by using FMM. However, the  $\mathbf{u}_i$  coordinates representing  $\psi(q_i)$  on  $\mathcal{S}$  change during the iterations of the numerical minimization algorithm. Thus, the distances



**Fig. 2.** Partial matching of a planar patch with a plane. (Left) Correct use of weights (shown in solid red is an inconsistent distance that must be excluded from the stress computation and in dashed blue is the corresponding correct distance). (Right) The embedding result when inconsistent distances are not excluded.

$d_S(\mathbf{u}_i, \mathbf{u}_j)$  have to be reevaluated at every iteration. This computation is critical for the GMDS. For that goal, we developed a numerical procedure we termed “three-point geodesic distance approximation” (detailed in A.M.B., M.M.B., and R.K., unpublished results). The idea is to produce a  $C^1$ -approximation for  $d_S(\mathbf{u}_i, \mathbf{u}_j)$  and its derivatives, interpolating their values from the matrix  $\mathbf{D}_S$ .

**Partial Matching of Surfaces**

Let  $Q'$  be a patch of the probe surface  $Q$  and  $Q'_N$  be a finite sampling of  $Q'$ . We consider the distance  $d_{PE}(\mathcal{S}, Q'_N)$ . By our assumption, the metric on  $Q'_N$  is given by restricting  $d_Q$  to  $Q'_N$ , i.e.,  $d_{Q'_N} = d_Q|_{Q'_N}$ . However,  $d_{Q'_N}$  is computed numerically on  $Q'_N$  and can be inconsistent with  $d_Q|_{Q'_N}$ . Specifically, geodesics  $C_Q^*(q_i, q_j)$  that pass through points not in  $Q'$  give rise to inconsistent distances, for which  $d_{Q'_N}(q_i, q_j) \neq d_Q(q_i, q_j)$  (see Fig. 2 Left). In pathological cases, the inconsistency can be arbitrarily large, such that the partial matching makes little sense (see example in Fig. 2 Right). To overcome this problem, the inconsistent distances  $d_{Q'_N}(q_i, q_j)$  are excluded defining the corresponding weights  $w_{ij} = 0$  in the generalized stress (5).

A practical problem is how to locate the inconsistent distances. Here, we address two possible scenarios: (i) the original surface  $Q$  or its sampled version is available, and (ii) only the patch  $Q'_N$  is available. The first scenario occurs when one intentionally wishes to exclude parts of  $\mathcal{S}$ . In this case, the FMM is used first to compute the distances  $d_Q(q_i, q_j)$  on  $Q$  and then to compute the distances  $d_{Q'_N}(q_i, q_j)$  on  $Q'$ . The weights are defined by

$$w_{ij} = \begin{cases} 1 & |d_Q(q_i, q_j) - d_{Q'_N}(q_i, q_j)| < \delta \\ 0 & |d_Q(q_i, q_j) - d_{Q'_N}(q_i, q_j)| \geq \delta \end{cases} \quad [10]$$

for all  $i, j = 1, \dots, N'$ , where  $\delta$  stands for the FMM accuracy.

The second scenario is more general, yet, also a more challenging one. When the original surface  $Q$  is not given, we need to rely on the geodesics of  $Q'$  [in practice, the geodesics are found by backtracking (22)]. By using the above argument, we define the weights by

$$w_{ij} = \begin{cases} 0 & C_{Q'}^*(q_i, q_j) \cap \partial Q' \neq \emptyset \\ 1 & \text{else} \end{cases}, \quad [11]$$

for all  $i, j = 1, \dots, N'$ , i.e., excluding the geodesics that touch the boundary  $\partial Q'$ .

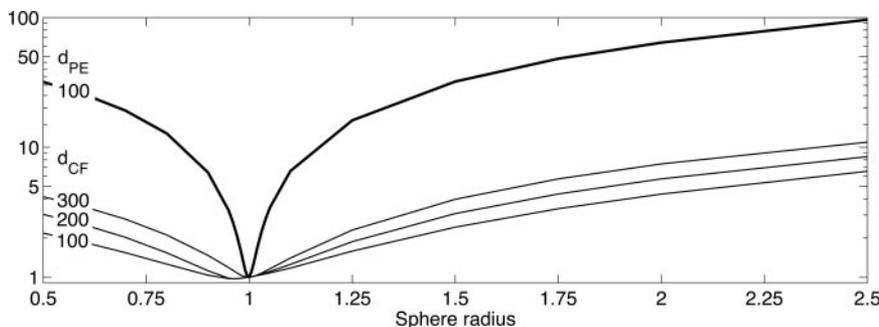
**Results**

**Embedding of Spherical Surfaces.** In the first experiment, we measured the distance between a unit 2D sphere sampled at 3,200 points, and spheres with radii in the range  $0.5 \div 2.5$  sampled at a smaller number of points according to the “farthest point sampling” strategy (23) with random seed. Ten random samplings were used for each radius. Two distance measures were compared as follows:  $d_{PE}^2$  (obtained by the minimization of  $\sigma_2$ ) with 100 points and  $d_{CF}$  with 100, 250, and 500 points.  $d_{PE}^2$  was computed by using local representation of points in barycentric coordinates.

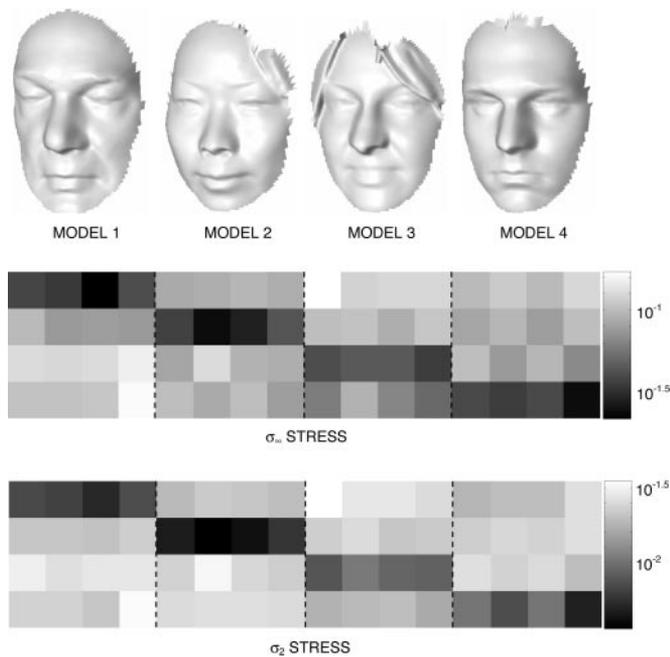
Fig. 3 presents the normalized distances as a function of the sphere radius. The PE distance appears to be extremely sensitive to the geometry of the surfaces. A change as small as 0.1% in the spherical patch radius (from 1.000 to 0.999) increases  $d_{PE}^2$  by a value exceeding the variance due to sampling. Similar results are achieved by using  $d_{PE}^2$  and slightly inferior with  $d_{PE}$ . For comparison, with the same number of points (100),  $d_{CF}$  is unable to discern between spheres differing in radius by  $>10\%$ . Increasing the number of points to 500 makes  $d_{CF}$  sensitive to radius differences of  $\approx 2.64\%$ , which is still one order of magnitude below the sensitivity of  $d_{PE}^2$ .

**Face Recognition Experiment.** In ref. 1, we argued that human facial expressions can be approximated as isometric deformations of the face. According to this model, an isometry-invariant surface matching results in nearly expression-invariant face recognition.

In the second experiment, we applied our GMDS approach to a toy 3D face recognition problem of 20 faces taken from the Notre Dame database (24, 25). This experiment is a proof of concept rather than a real face recognition system, in which tune-up plays a crucial role. As models, four different faces were used. Each model contained  $\approx 4,000$  points and was represented using global parameterization. As probes, faces of the same subjects with different facial expressions were used (four probes per subject). The probes were cropped by using geodesic mask (1) (leaving the nose and the forehead region) and subsampled by using farthest point sampling to 53 points. The distances were weighted as described in *Partial Matching of Surfaces*. All dis-



**Fig. 3.** Normalized distance between a densely sampled unit sphere and spheres of different radii measured using  $d_{PE}^2$  (top heavy line) and  $d_{CF}$  (three bottom lines). The numbers on the plots stand for the number of samples in the embedded surface.

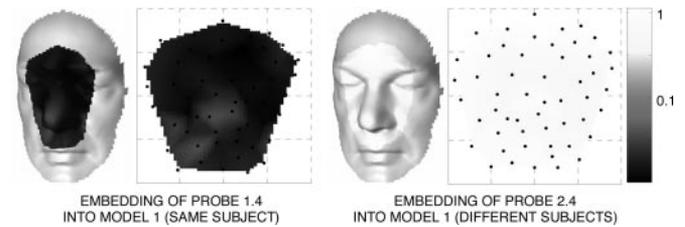


**Fig. 4.** Toy face recognition experiment. (Top) Model facial surfaces. (Middle and Bottom) Similarity matrices obtained using  $d_{PE}$  (Middle) and  $d_{PE}^2$  (Bottom), respectively. All values are shown in logarithmic scale.

tances were computed numerically by a version of parametric FMM (21). We tested two PE distances:  $d_{PE}$  and  $d_{PE}^2$ . The generalized stress function  $\sigma_p$  and its gradient were implemented in C. For 53 points, the computation time of  $\sigma_p$  and  $\nabla\sigma_p$  is  $\approx 20$  msec on a Pentium IV PC.

Fig. 4 presents the similarity matrix using  $d_{PE}$  (second row) and  $d_{PE}^2$  (third row). Note that in both cases a perfect separation between different subjects is possible; the recognition error (measured as “equal error rate,” or EER; see ref. 26 for definition) in this experiment is zero. It is worthwhile noting the small number of points used in our approach, compared with the thousands of points required in the CF algorithm (1, 3) for similar matching accuracy.

An example of embedding probe surfaces of two different



**Fig. 5.** Local stress map shown on the facial surface and in the parameterization domain in logarithmic scale. (Left) Embedding of probe 1.4 (subject 1, probe 4) into model 1. (Right) Embedding of probe 2.4 into model 1.

subjects into a model surface are shown in Fig. 5. The figure depicts the “local stress” value  $\sigma_\infty(k) = \max_i |d_S(\mathbf{u}_k, \mathbf{u}_i) - d_Q(q_k, q_i)|$  at each point  $k$ , i.e., the contribution of the  $k$ th point to  $\sigma_\infty$ . We observe that the stress is significantly higher when a probe of a different subject is embedded into the model surface.

## Conclusions

We introduced a previously undescribed framework for isometry-invariant surface matching. The core of our computational framework is GMDS, a generalization of standard MDS algorithms. We exemplified the use of GMDS on the problem of expression-invariant 3D face recognition.

Our approach favorably compares with previous attempts to perform isometry-invariant surface matching. First, our PE distance naturally allows for isometry-invariant matching of partially missing surfaces. Secondly, the properties of our distance and its computation are completely deterministic. Thirdly, GMDS used for the PE distance computation is a continuous optimization problem and can be solved very efficiently by using standard optimization methods. In light of recent results (13), we expect a significant performance boosting by using multigrid methods.

Finally, we note that the same numerical framework can be applied for computing the exact GH distance between surfaces (A.M.B., M.M.B., and R.K., unpublished results).

We thank Alexander Brook and Roman Goldenberg for thorough reading of the manuscript, Facundo Mémoli and Vladimir Lin for fruitful discussions, and our reviewers for valuable comments. This work was supported in part by United States-Israel Science Foundation Grant 2004274, Israeli Science Foundation Grant 738/04, and the Charles Krown Research Fund.

- Bronstein, A. M., Bronstein, M. M. & Kimmel, R. (2005) *Int. J. Computer Vision* **64**, 5–30.
- Zhang, Z. Y. (1994) *Int. J. Computer Vision* **13**, 119–152.
- Elad, A. & Kimmel, R. (2003) *IEEE Trans. PAMI* **25**, 1285–1295.
- Schwartz, E. L., Shaw, A. & Wolfson, E. (1989) *IEEE Trans. PAMI* **11**, 1005–1008.
- Borg, I. & Groenen, P. (1997) *Modern Multidimensional Scaling: Theory and Applications* (Springer, Berlin).
- Roweis, S. T. & Saul, L. K. (2000) *Science* **290**, 2323–2326.
- Donoho, D. & Grimes, C. (2003) *Proc. Natl. Acad. Sci. USA* **100**, 5591–5596.
- Linial, N., London, E. & Rabinovich, Y. (1995) *Combinatorica* **15**, 333–344.
- Mémoli, F. & Sapiro, S. (2005) *Foundations Computational Math*, 10.1007/s10208-004-0145-y.
- Gromov, M. (1981) *Structures Métriques pour les Variétés Riemanniennes*, Textes Mathématiques (CEDIC, Paris), Vol. 1.
- Burago, D., Burago, Y. & Ivanov, S. (2001) *A Course in Metric Geometry*, Graduate Studies in Mathematics (Am. Math. Soc., Providence, RI), Vol. 33.
- Tal, A., Elad, M. & Ar, S. (2001) in *Proceedings of the 6th Eurographics Workshop on Multimedia*, eds. Jorge, J. A., Correia, N. M., Jones, H. & Kamegai, M. B. (Springer, New York), pp. 97–108.
- Bronstein, M. M., Bronstein, A. M., Kimmel, R. & Yavneh, I. (2006) *Numerical Linear Algebra Applications*, in press.
- Elad, A. & Kimmel, R. (2002) *Geometric Methods in Bio-medical Image Processing* (Springer, Berlin), Vol. 2191, pp. 77–89.
- Bronstein, A. M., Bronstein, M. M. & Kimmel, R. (2005) *Lecture Notes Computer Sci.* **3459**, 622–631.
- Walter, J. & Ritter, H. (2002) in *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, eds. Hand, D., Keim, D. & Ng, R. (Assoc. Computing Machinery, New York), pp. 123–131.
- Floater, M. S. & Hormann, K. (2004) in *Advances in Multiresolution for Geometric Modelling*, eds. Dodgson, N. A., Floater, M. S. & Sabin, M. A. (Springer, Berlin), pp. 157–186.
- Bertsekas, D. (1999) *Nonlinear Programming* (Atlanta Scientific, Atlanta), 2nd Ed.
- Nash, S. (2000) *J. Optimization Methods Software* **14**, 99–116.
- Kimmel, R. & Sethian, J. A. (1998) *Proc. Natl. Acad. Sci. USA* **95**, 8431–8435.
- Spira, A. & Kimmel, R. (2004) *Interfaces Free Boundaries* **6**, 315–327.
- Kimmel, R. (2003) *Numerical Geometry of Images* (Springer, Berlin).
- Eldar, Y., Lindenbaum, M., Porat, M. & Zeevi, Y. Y. (1997) *IEEE Trans. Image Processing* **6**, 1305–1315.
- Chang, K. I., Bowyer, K. W. & Flynn, P. J. (2005) *IEEE Trans. PAMI* **27**, 619–624.
- Flynn, P. J., Bowyer, K. W. & Phillips, P. J. (2003) in *Audio- and Video-Based Biometric Person Authentication*, eds. Kittler, J. & Nixon, M. S. (Springer, Berlin), pp. 44–51.
- Ashbourn, J. (2002) *Biometrics: Advanced Identity Verification* (Springer, Berlin).