On reconstruction of non-rigid shapes with intrinsic regularization

Yohai S. Devir Guy Rosman Alexander M. Bronstein Michael M. Bronstein
Ron Kimmel
{yd|rosman|bron|mbron|ron}@cs.technion.ac.il
Department of Computer Science
Technion - Israel Institute of Technology

Abstract

Shape-from-X *is a generic type of inverse problems in computer vision, in which a shape is reconstructed from some measurements. A specially challenging setting of this problem is the case in which the reconstructed shapes are non-rigid. In this paper, we propose a framework for intrinsic regularization of such problems. The assumption is that we have the geometric structure of a shape which is intrinsically (up to bending) similar to the one we would like to reconstruct. For that goal, we formulate a variation with respect to vertex coordinates of a triangulated mesh approximating the continuous shape. The numerical core of the proposed method is based on differentiating the fast marching update step for geodesic distance computation.*

1. Introduction

In many tasks, both in human and computer vision, one tries to deduce the shape of an object given an observation or a measurement thereof. For example, in a shadow theater (Fig. 1) we recognize shapes of objects from silhouettes projected on a screen. Another example is *shape from stereo*, where a three-dimensional (3D) object is reconstructed from a pair of two-dimensional (2D) images of the object taken from different viewpoints. These and many other problems, in which an object is reconstructed based on some measurement, are known as *shape reconstruction problems*. They are a subset of what is called *inverse problems*. Most such inverse problems are underdetermined, in the sense that measuring different objects may yield similar measurements. Thus, in the above illustration, the essence of the shadow theater is that it is hard to distinguish between shadows cast by an animal and shadows cast by hands. Therefore prior knowledge about the unknown object is needed.

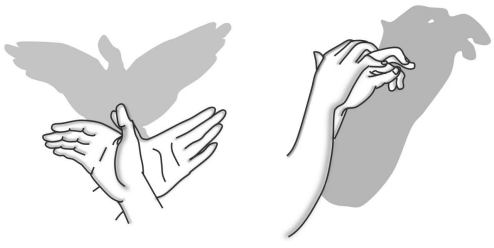

Figure 1. Shadow theater example - a shadow of a bird and a camel being cast by two hands.

Of particular interest are reconstruction problems involving non-rigid shapes. The world surrounding us is full with objects such as live bodies, paper products, plants, clothes etc., which may be deformed to different postures. These objects may be deformed to an infinite number of different postures. While bending, though, objects tends to preserve their internal geometric structure. Two objects differing by a bending are said to be *intrinsically similar*. In many cases, while we do not know the measured object, we have a prior on its intrinsic geometry. For example, in the shadow theater, though we do not know which exact posture of the hand casts the shadow, we know that the unknown object is a deformation of a hand. In other words, our prior is a shape of the hand, and the variability of the reconstructed object is the bending of the hand.

Many methods of shape-from-X reconstruct the shape by minimizing some functional. In *shape from shading (SFS)* [23], a 3D shape is recovered from one or more 2D light reflectance images. In *bundle adjustment* [21], a 3D shape is reconstructed from a set of point matches between 2D images of the shape from different viewpoints taken by possibly non-calibrated cameras. Certain functionals can be minimized by solving the corresponding Euler-Lagrange equations, for example, by a gradient descent. Using a functional measuring intrinsic similarity, knowledge about the intrinsic structure of the shape can be incorporated into shape reconstruction methods.

There are many papers that model distortion between an object and a known shape by means of local deformations or

use local shape properties as shape priors. Eckstein *et al*. [8] and Kilian *et al*. [12] use the *local* distortion of the mesh edges to define a global one. Salzmann *et al*. [18] propose a method for modeling deformations of planar surfaces with no holes from 2D images based on preserving the lengths of edges in the mesh. Local deformations functionals, even though usually contained in global deformations ones, may be deceiving. For example consider a plane, along with a corresponding second plane in which all internal samples were moved in random tangent directions. Locally, there is a strong distortion, although globally, the planes are very similar.

A few methods such as [15, 9] measure global intrinsic similarity by mapping the metric structure of 3D shapes to a simpler space. Those methods are difficult to incorporate into an optimization process which alters the shape, beside of their inherent inaccuracy due to the limitation of the embedding space to represent the intrinsic geometry.

This limitation can be solved by using the *Gromov-Hausdorff* (GH) distance as was suggested by Mémoli and Sapiro [16]. Bronstein *et al*. [5] proposed a method formulating the GH distance as a multidimensional scaling problem and embedding one surface into another while reducing the intrinsic dissimilarity between the surfaces. Since *this method* limits one surface to be embedded into another, it cannot be adopted for finding a general shape embedded in $\mathbb{R}^3$.

In a prior work [1], Anguelov *et al*. studied the problem of non-rigid shapes reconstruction from a small set of markers located on the shapes. This work is limited to a specific class of shapes (human bodies in [1]) and requires complex learning from many prior examples. In this work we use a much simpler model of a single body in a single position.

In this paper we use the *stress* [4] between a known prior shape and a given shape as an intrinsic similarity functional. In order to make this functional consistent and robust to triangulation, we calculate it using the *fast marching method (FMM)* [13]. The numerical core of our method is based on an extension to the FMM that enables us to calculate the derivatives of pairwise geodesic distances w.r.t. the spatial coordinates of the mesh vertices .

Our approach is similar in spirit to Bronstein *et al*. [6] who showed a combination of intrinsic and extrinsic similarities criteria. Their method is based on calculating the geodesic distances on a shape using Dijkstra's algorithm (as opposed to FMM used here), it is sensitive to the shape triangulation and is affected by metrication errors.

Benmansour *et al*. [3] proposed a method for differentiating geodesic distances calculated by FMM w.r.t. the Riemannian metric defined on a rectangular Cartesian mesh. They did not, however, explorer the link between the intrinsic geometry and the extrinsic geometry of an embedded surface. Our work uses a more general framework defined for triangulated surfaces in order to explorer this very relation.

The rest of the paper is organized as follows: Section 2 introduces the mathematical background. Section 3 reviews methods of calculating geodesic distances. Section 4 presents the proposed reconstruction framework. Section 5 demonstrates the proposed method on some examples. Finally, we provide concluding remarks in Section 6.

For implementation details and derivations, the reader is referred to [7].

2. Mathematical background

2.1. Shape space

We model a shape X as a *metric space* - a pair (X, d_X) where X is a 2D smooth, compact, connected and complete Riemannian manifold embedded into the Euclidean space $\mathbb{R}^3$, and $d_X : X \times X \to \mathbb{R}^+ \cup \{0\}$ is the *geodesic metric* measuring the length of the minimal curve contained in the surface connecting two points. We denote by $\mathbb{X}$ the shape space in which each element is a shape.

Given two shapes $Z, X \in \mathbb{X}$, we define the *correspondence* between them as $C \subset X \times Z$ such that $\forall x \in X, \exists z \in Z : (x, z) \in C$ and $\forall z \in Z, \exists x \in X : (x, z) \in C$. We say $x \in X$ and $z \in Z$ are in correspondence if $(x, z) \in C$. The *distortion* of a correspondence is given by

$$\operatorname{dis} C = \max_{\substack{(x,z)\in C \\ (x',z')\in C}} |d_X(x, x') - d_Z(z, z')|, \tag{1}$$

which measures how different the corresponding metric structures in X and Z are.

A subset $X' \subseteq X$ is said to be an *ε-net* in (or *ε-covering* of) X if

$$\max_{x\in X} \min_{x'\in X'} d_X(x, x') \leq \varepsilon. \tag{2}$$

If a correspondence C has $\operatorname{dis} C \leq \varepsilon$ and the set $X' = \{x : \exists z, (x, z) \in C\}$ of all points in X having correspondence in Z is an ε-net, then X and Z are *ε-isometric*.

The similarity of two shapes in $\mathbb{X}$ is formalized using a metric $d_{\mathbb{X}}$ on the shape space. Since the elements of $\mathbb{X}$ are metric spaces we can use the Gromov-Hausdorff [11] distance as $d_{\mathbb{X}}$,

$$d_{\mathrm{GH}} = \min_C \operatorname{dis} C \tag{3}$$

If $d_{\mathrm{GH}}(X, Z) \leq \varepsilon$, then X and Z are 2ε-isometric, and if X and Z are ε-isometric, $d_{\mathrm{GH}}(X, Z) < 2\varepsilon$. In particular, $d_{\mathrm{GH}}(X, Z) = 0$ iff X and Z are isometric. [1]

[1] Most polyhedral shapes without boundaries are in fact rigid, *i.e.* do not have incongruent isometries. Moreover, small wrinkles that appear in real world non-rigid objects are not represented in limited resolution models. As a result, two meshes approximating two isometric deformations of a non-rigid shape are only approximately isometric.

2.2. Measurement space

Let $\mathbb{Y}$ be a measurement space, and $d_{\mathbb{Y}} : \mathbb{Y} \times \mathbb{Y} \to \mathbb{R}^+ \cup \{0\}$ be a distance on $\mathbb{Y}$. The operator $P : \mathbb{X} \to \mathbb{Y}$ creates a measurement of a shape X. For example, in the SFS inverse problem, $\mathbb{Y}$ is the space of gray scale images, P is an operator assigning a patch on the shape to a pixel in the image whose intensity is determined by the average dot product between the surface normal at each point in this patch and the illumination direction.

2.3. Reconstruction problems

Given a measurement $Y = P(\tilde{X})$ of an unknown shape $\tilde{X}$, and a prior shape X_0 such that $d_{\mathbb{X}}(\tilde{X}, X_0)$ is small, we formulate a generic problem of shape reconstruction as follows:

$$\hat{X} = \underset{X \in \mathbb{X}}{\operatorname{argmin}}\, d_{\mathbb{X}}(X, X_0) + \lambda d_{\mathbb{Y}}(P(X), Y), \qquad (4)$$

where $\lambda > 0$ is some parameter.

We look for a shape X which is intrinsically similar to the prior X_0 and whose measurement is similar to Y. Our way of solving Problem (4) consists of deforming an initial shape and evaluating the result of this deformation (X) on $d_{\mathbb{X}}(X, X_0)$ and $d_{\mathbb{Y}}(P(X), Y)$.

Calculating $d_{\mathbb{X}}(X, X_0)$ requires computation of all geodesic distances in both X and X_0.

3. Geodesic distance computation

In the discrete settings we use triangulated mesh as an approximation to our shapes. Two common and simple ways of approximating the geodesic distances on a triangulated mesh are using Dijkstra's shortest paths algorithm, and using the FMM ([13, 19, 22]). Other algorithms such as of Surazhsky *et al*. [20] yields much accurate results, but are more complicated to implement, and make the derivative computation impractical.

Given a mesh X and a source vertex x_s on it as the input, a distance computation algorithm returns the *distance map* $d(x_i) \approx d_X(x_i, x_s)$, an approximation of the geodesic distance from all vertices x_i to x_s. For simplicity of notation, we use d_i as a shortening of $d(x_i)$.

In Dijkstra's algorithm, the shape is approximated as an undirected graph, and geodesics as shortest paths in it. Thus, the distance $d_X(x_s, x_i)$ is expressed as the minimum of $d_X(x_s, x_j) + E_{ij}$ over all x_j adjacent to x_i, and the computation is performed using a dynamic programming-based successive approximation procedure, in which distance values are propagated to adjacent vertices. While being easy to implement, the main disadvantage of Dijkstra's algorithm is its inconsistent to triangulation. Under general conditions, increasing the mesh resolution does not decrease the approximation error, a phenomenon known as *metrication error*.

A family of continuous variants of shortest path Dijkstra's algorithm known as *fast marching methods (FMM)* is virtually free of metrication errors. The major difference constitutes in the fact that the surface is approximated as a continuous two-dimensional structure (usually, a triangular mesh), and paths are allowed to pass anywhere on its faces rather than being restricted to the edges. The computation differs mainly in the update step, according to which distances are propagated. Given an acute triangle x_1, x_2, x_3 in which the current approximation of the distances $d_X(x_1, x_s)$ and $d_X(x_2, x_s)$ are d_1 and d_2, respectively, the approximation of d_3 of $d_X(x_3, x_s)$ is computed according to

$$\begin{aligned} d_3 &= \mathcal{F}_{\text{FMM}}(d_1, d_2, a, b, c) \\ &= \begin{cases} d_1 + \frac{rv}{c} + \frac{s}{c}\sqrt{c^2 - v^2} & \text{some condition holds} \\ \min\{d_1 + b, d_2 + a\} & \text{otherwise} \end{cases} \end{aligned} \qquad (5)$$

where a, b and c are the length of the triangle edges, $v = d_2 - d_1$, $r = \frac{b^2 + c^2 - a^2}{2c}$ and $s = \sqrt{b^2 - r^2}$.

Obtuse triangles are replaced by two new acute triangles, $x_1x_0x_3$ and $x_2x_0x_3$ contained in a flattening of $x_0x_1x_2$ and some of its adjacent triangles to a plane. We denote those triangles as *virtual triangles*.

4. Implementation Considerations

In order to solve Eq. 4 we first define a measure of the intrinsic dissimilarity. In the discrete case, we use the L_2 version of the stress to measure the intrinsic similarity between two surfaces Z and X, both with N samples, which is defined as

$$\sigma(X, Z) = \min_{C} \sum_{\substack{j,k \\ (x_j, z_j) \in C \\ (x_k, z_k) \in C}} \left(d_X(x_j, x_k) - d_Z(z_j, z_k)\right)^2, \qquad (6)$$

where C is a correspondence between X and Z. In case the correspondence C is predetermined, we denote the stress given a correspondence C by $\sigma(X, Z; C)$.

In this work, we assume the correspondence C to be known, which is reasonable in many cases. In those cases where C is unknown, an initial correspondence can usually be found from local shape descriptors (see [17, 14, 10] and references therein), and one can alternate optimization for shape deformation with fixed correspondence and optimization for optimal correspondence with fixed deformation.

Using σ as a version of the intrinsic dissimilarity and assuming the correspondence to be known, problem 4 reduces

to

$$\hat{X} = \underset{X \in \mathbb{X}}{\operatorname{argmin}}\, \sigma(X, X_0; C) + \lambda d_{\mathbb{Y}}(P(X), Y). \quad (7)$$

Given the derivatives of $d_{\mathbb{Y}}(P(X), Y_0)$ w.r.t. the coordinates of X's vertices, Eq. 7 can be solved as an optimization problem where the optimization variables are the coordinates of X's vertices. The initialization is the prior X_0 if no better guess is available. In order to find $\hat{X}$ of Eq. 7 using gradient-descent type optimization techniques, we need to differentiate $\sigma(X, X_0)$ w.r.t. the coordinates of the samples of X. Due to the ill-conditioning of $\sigma(X, X_0)$, analytic derivatives are needed.

Vertex-wise, the optimization has the form of

$$\overline{x}_i \leftarrow \overline{x}_i + \overline{\alpha}_i, \quad (8)$$

where $\overline{\alpha}_i = (\alpha_i^x, \alpha_i^y, \alpha_i^z)$ is the distance x_i moves in each coordinate direction.

Since the analysis is similar for all coordinates, we omit the superscript index for notation simplicity.

4.1. Gradient of the stress

Given two shapes X and X_0, and a bijective correspondence C, the stress between X and X_0 can be written as

$$\sigma(X, X_0; C) = \sum_{\substack{j,k \\ (x_j, x_{0j}) \in C \\ (x_k, x_{0k}) \in C}} \left(d_X(x_j, x_k) - d_{X_0}(x_{0j}, x_{0k})\right)^2. \quad (9)$$

Differentiating the stress w.r.t. α_i, we obtain

$$\frac{\partial \sigma(X, X_0; C)}{\partial \alpha_i} = 2 \sum_{j,k} \delta_{jk} \cdot \frac{\partial d_X(x_j, x_k)}{\partial \alpha_i} \quad (10)$$

where $\delta_{jk} = d_X(x_j, x_k) - d_{X_0}(x_{0j}, x_{0k})$. The pairwise distances $d_X(x_j, x_k)$ are computed after every iteration of the optimization by running the FMM algorithm N times, each time with a different source vertex. The pairwise distances in X_0 are pre-computed once in a similar way. The derivatives of the geodesic distances $\frac{\partial d_X(x_j, x_k)}{\partial \alpha_i}$ are calculated as follows.

Consider a vertex x_0, and a source vertex x_s. We denote the triangle used for the calculation of the $d_X(x_s, x_0)$ as $x_0x_1x_2$, whose edges lengths are$\{a, b, c\}$. Calculating $\frac{\partial a}{\partial \alpha_i}$, $\frac{\partial b}{\partial \alpha_i}$ and $\frac{\partial c}{\partial \alpha_i}$ is straightforward if $x_0x_1x_2$ is an acute non-virtual triangle. If $x_0x_1x_2$ is a virtual triangle, the calculation of those derivatives is based on calculating the derivatives for every intermediate value along the unfolding process.

Also given $d_X(x_s, x_1)$ and $d_X(x_s, x_2)$ the distances of x_1 and x_2 from the source and the derivatives of these estimated distances w.r.t. the X's vertices locations, $(\frac{\partial d_X(x_s, x_1)}{\partial \alpha_i}, \frac{\partial d_X(x_s, x_2)}{\partial \alpha_i})$. $\frac{\partial}{\partial \alpha_i} d_X(x_s, x_0)$, can be calculated by differentiating Eq. 5 using the chain rule. Note that the derivatives of the triangle edges' length are calculated only once for all source vertices.

In order to calculate the derivatives of all pairwise distances we run this modified FMM algorithm N times, once for every source vertex. In the initialization step, we set $\frac{\partial d_X(x_s, x_s)}{\partial \alpha_i} = 0, i \in [1 \dots N]$. Next, in every FMM iteration, we calculate the $\frac{\partial d_X(x_s, x_i)}{\partial \alpha_j}, j \in [1 \dots N]$ immediately after calculating $d_X(x_s, x_i)$.

4.2. Path fixing

Consider a vertex x_0 for which there are more than one shortest paths connecting it to the source. The gradient of the distance function at x_0 is not defined and $d_X(x_s, x)$ is not differentiable. If w.l.o.g. $d(x_0, x_s)$ is calculated in triangle $x_0x_1x_2$, increasing d_1 or d_2 will not lead to the increase of d_0, since d_0 will be updated through another path and another triangle. In order to overcome this, we calculate a pairwise distance map on the prior and then fix the order of vertices updated and also fix the triangles used for updating each vertex.

4.3. Smoothing the transition between calculations

$d_X(x_0, x_s)$ may be calculated according to triangle $x_0x_1x_2$ in a couple of different ways. If $\angle x_1x_0x_2$ is acute, it will be updated using Eq. 5 in triangle $x_0x_1x_2$. If $\angle x_1x_0x_2$ is obtuse, it will be updated using Eq. 5 in the virtual triangle $x_0x_1\hat{x}_3$ or $x_0x_2\hat{x}_3$. As a consequence, a small perturbation in the mesh changing $\angle x_1x_0x_2$ from an acute angle to obtuse angle (or vice versa), results in different update schemes which cause non-differentiability of the geodesic distance function. Such scenario is demonstrated in Fig. 2 (left).

Moreover, since $\angle x_1x_0\hat{x}_3$ and $\angle x_2x_0\hat{x}_3$ must be non-obtuse, if a small perturbation changes $\angle x_1x_0\hat{x}_3$ and / or $\angle x_2x_0\hat{x}_3$ from being acute angle to obtuse angle, more unfolding are needed resulting in different virtual triangles and thus different calculation scheme of $d_X(x_0, x_s)$. Such scenario is demonstrated in Fig. 2 (middle).

A combination of more than one such scenario is also possible as is demonstrated in Fig. 2 (right).

The discontinuities are treated using the following notion. In all of those scenarios, there is an angle (denoted by β) for which if a perturbation causes it to rise above a threshold L_1, a different calculation is performed. Therefore, for each such scenario we denote by $d_X^{\beta < L_1}(x_s, x_0)$ the distance of x_0 from x_s calculated using the current value of β and by $d_X^{\beta \geq L_1}(x_s, x_0)$ we denote a second value calculated as if β exceeded the threshold. We denote these calculations as *speculative calculations*.

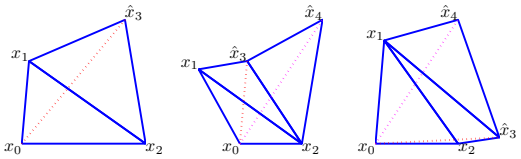

Figure 2. The cases in which the method of calculation is changing. In all images - updating x_0 comes from $x_0x_1x_2$ triangle. Left: The angle of the update triangle changes from acute to obtuse. Middle: A change of the virtual triangles. Right: A combination of the scenarios on the left and the middle.

Next, we set

$$d_X(x_s,x_0) = H\cdot h(\beta,L_0,L_1)d_X^{\beta\geq L_1}(x_s,x_0) + (1-H\cdot h(\beta,L_0,L_1))d_X^{\beta<L_1}(x_s,x_0), \quad (11)$$

where L_0 is a parameter, and $h(\beta,L_0,L_1)$ is a differentiable monotone increasing function for which $h(\beta,L_0,L_1) = 0\,\forall\beta < L_0$ and $h(\beta,L_0,L_1) = 1\,\forall\beta \geq L_1$, and H is $h(\frac{\theta}{2},L_0,L_1)$ if both $\angle\hat{x}_3x_0x_1 > L_0$ and $\angle\hat{x}_3x_0x_2 > L_0$ and 1 otherwise.

If $\min\{\angle\hat{x}_3x_0x_1,\angle\hat{x}_3x_0x_2\} > L_0$, $d_X^{\beta\geq L_1}(x_s,x_0)$ is calculated as a weighted sum between the case assuming $\angle\hat{x}_3x_0x_2 > L_1$ and the case where $\angle\hat{x}_3x_0x_1 > L_1$. The weight is calculated in a similar manner according to the value of $\angle\hat{x}_3x_0x_2$.

Another case is that a vertex is updated using Dijkstra scheme (the second case in Eq. 5), and $d_1+b = d_2+a$. This is treated by replacing $\min\{d_1+b,d_2+a\}$ in Eq. 5 with

$$d_0 = h_\varepsilon(K)(d_2+a) + (1-h_\varepsilon(K))(d_1+b) \quad (12)$$

where $K = (d_1+b)-(d_2+a)$, ε is some small constant and h_ε is a monotone increasing differentiable function for which $h_\varepsilon(e) = 0\,\forall e \leq -\varepsilon$, $h_\varepsilon(e) = 1\,\forall e \geq \varepsilon$.

Theoretically, since a speculative calculation may be based on another speculative calculation, there can be a $\mathcal{O}(2^N)$ speculative calculations where N is the number of vertices in the mesh. However, in our experiments, for small values of L_1 in non-pathological meshes, we found in most cases less than 2^5 speculative calculations.

Derivatives of h are calculated by the chain rule using the derivatives of the edges forming the triangles affecting the weighted sum.

It should be noted that by implementing the modifications described here and in Sec. 4.2, we replace the stress with a similar *smooth* concave function.

4.4. Limiting the paths affecting the stress

According to Eq. 6, a relative distortion in a single long distance has more effect on the stress than a few relatively similar distortions in short distances. In most applications, making local patches intrinsically similar to the corresponding local patches in a second shape is more important than keeping all pairwise distances similar. Therefore we ignore distances between pairs of vertices whose geodesic distance in the prior shape is above a predetermined threshold.

5. Results

In this section, we show a few examples of using the gradient of the metric for synthesizing a shape from a measurement given a prior of a non-rigid deformation of that shape. The animals and human body models were taken from TOSCA database. [2]

The values of weight variables λ shoud be set according to desired tradeoff between instrinsic similarity and the other criteria. In the following experiments these parameters were set experimentally.

5.1. Shape-preserving denoising

In the first experiment, two corresponding models of a palm are used. A prior model X_p and an unknown model X_u for which a random Gaussian noise was added to every vertex's coordinate yielding a noisy model X_n. The cost function used in this example is

$$\text{cost}(X) = \sigma(X,X_p;C) + \lambda_1 d_{\text{ext}}(X,X_n) + \lambda_2 D_{\text{vol}}(X,X_p,r_v) \quad (13)$$

where $D_{\text{vol}}(X,X',r) = (|\text{vol}(X)-\text{vol}(X')|-r)^2$ if $|\text{vol}(X)-\text{vol}(X')| > r$ where $\text{vol}(X)$ is the volume of X, and 0 otherwise. $d_{\text{ext}}(X,X_n) = \sum_i \|x_i-(X_n)_i\|^2$ for all corresponding vertices in X and X_n.

The volume penalty term is added to prevent the shape from being flattened, since a flatten shape and a non-flatten one are intrinsically nearly identical. This term is needed since due to the high magnitude of the noise, the extrinsic term will not keep the volume of the shape.

The results are shown in Figure. 3.

The reason that the stress between the prior and the unknown shape is higher than the stress between the prior and the final result is that current shape deformation tools do not fully preserve the internal structure.

5.2. Sparse model fitting

In the second experiment, we show how the stress can be reduced as a postprocessing of an arbitrary pair of nearly isometric shapes. Here we demonstrate it as a postprocessing for a shape completion. The inputs we use in this experiment are an initial shape X_0 (Fig. 4 left) and a set of 12 markers marked both on X_0 and on a target model (Fig. 4

[2] Data available from `http://tosca.cs.technion.ac.il`.

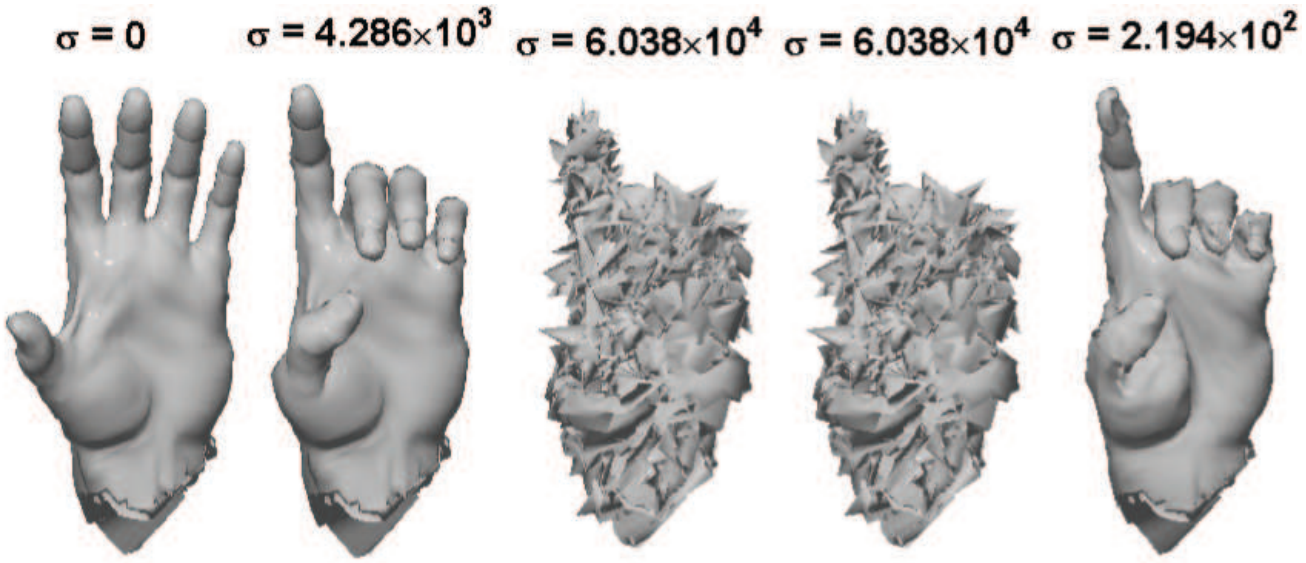

Figure 3. Shape denoising. From left to right: the given prior, the unknown shape, the measurement - a noisy shape, denoising using the prior's volume only, denoising using the prior's intrinsic geometry and volume. Stress value of a shape is the stress between this shape and the prior.

second from left). The deformation was made using the algorithm of [2] (Fig. 4 third from left), we denote its result as X_d.

We minimize the following cost function

$$\text{cost}(X) = \sigma(X, X_0; C) + \lambda_1 \sum_{x_i \in X} d_{\mathbb{R}^3}(x_i, X_d)^2 + \lambda_2 D_{\text{vol}}(X, X_0, r_v). \quad (14)$$

where $d_{\mathbb{R}^3}(x_i, X_d)$ is the minimal Euclidean distance from x_i to a point on the faces of X_d.

The stress between the prior and the unknown shape is 1.11×10^4. The stress between the prior and the deformation's result is 1.08×10^4. After 800 iterations, the stress between the prior and the optimization's result was 2.4×10^3. It should be noted that although the result is not visually better, it is less distorted, while still visually pleasing.

5.3. Bundle adjustment

In the bundle adjustment problem, an unknown object X_u is viewed by $K \geq 2$ cameras. The image positions of points projected unto these views are used to reconstruct the object geometry, and possibly, camera locations. Let $\{s_i^{(k)}\}_{i=1}^N \subseteq \mathbb{R}^2, k \in 1, \ldots, K$ denote a set of projected coordinates obtained by projecting the points $(x_i)_{i=1}^N \subseteq \mathbb{R}^3$ into a set of cameras with projection matrices $P_k, k = 1, .., K$, $P_k(\tilde{x}_i) = \tilde{s}_i^{(k)}$, where $\tilde{x}_i$ and $\tilde{s}_i^{(k)}$ are homogenous coordinates representation of x_i and $s_i^{(k)}$ respectively. We ignore in this experiment occlusions and mismatches as well as the estimation of camera parameters.

Next, let $\{r_i^{(k)}\}_{i=1}^N \subseteq \mathbb{R}^2, k \in 1, \ldots, K$ denote a set of noisy projected coordinates, $r_i^{(k)} = s_i^{(k)} + n$, where $n \sim N(0, \sigma)$ is Gaussian noise.

In this example, we reconstruct the unknown object X given a intrinsically similar prior shape X_p, assuming a known correspondences between vertices of X_p and $\{s_i^{(k)}\}_{i=1}^N$.

The cost function we minimize is

$$\text{cost}(X) = \sigma(X, X_p; C) + \lambda \sum_{k=1}^{K} \sum_{i=1}^{N} \|r_i^{(k)} - P_k(x_i)\|_2^2. \quad (15)$$

Results for the case of $K = 2$ are shown in Fig. 5.3.

6. Conclusions

We proposed a framework for regularization of inverse problems involving non-rigid shapes. Our approach assumes knowledge of a prior shape intrinsically similar to the unknown shape we try to reconstruct. At the numerical core of our method we compute the gradient of the geodesic distances matrix w.r.t. the coordinates of the shapes vertices. This computation is based on differentiating the update step of the fast marching method. It has been demonstrated on various applications.

Acknowledgment

The authors are grateful to Yuval Rabani for his fruitful comments, and to Ofir Weber and Mirela Ben-Chen for

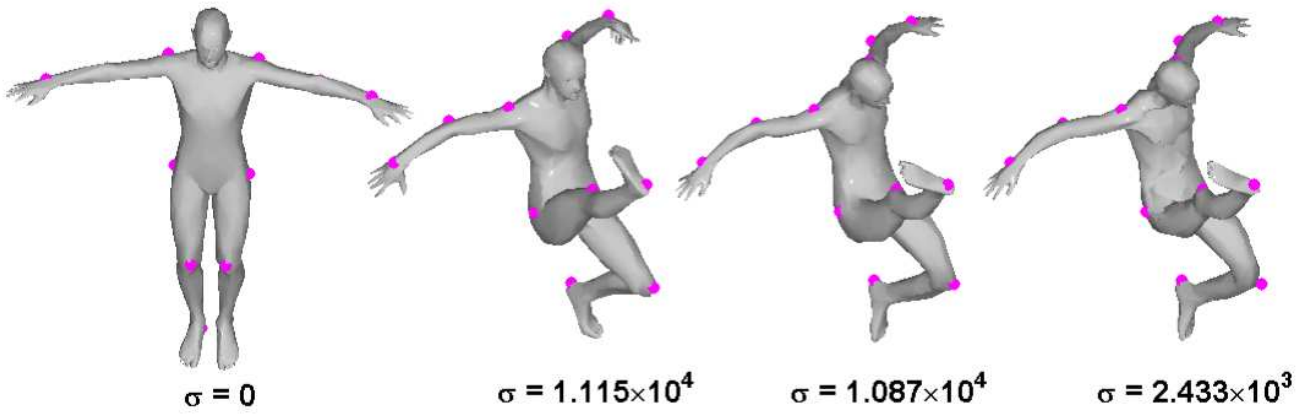

Figure 4. Stress reduction after sparse model fitting. From left to right: the prior shape, the unknown shape for which only the spatial locations of the markers are given, the result of a deformation applied to the prior shape in order to fit the marker points of the unknown shape, the result after reducing the stress. It can be noted that the final result is more intrinsically similar to the prior while preserving the deformation.

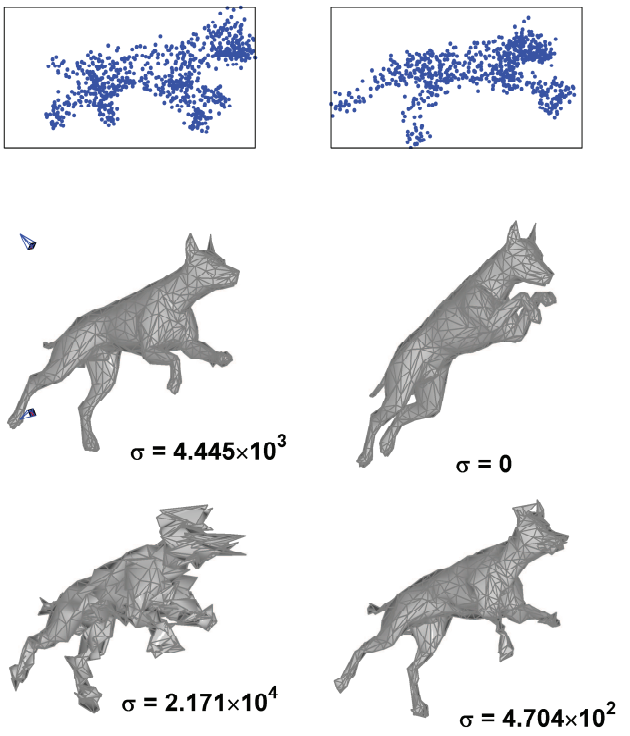

Figure 5. Top row: unknown shape's noisy projection. X_u (middle row, left) with cameras' locations and orientations are indicated by the blue pyramids. X_p (middle row, right). Reconstruction result without the intrinsic prior (bottom row, left), and with the intrinsic prior (bottom row, right). Stress values are relative to X_p.

their assistance with the 3D models. The model of the arm is courtesy of Autodesk. This research was supported by ONR grant No. N00014-06-1-0978, by The Israel Science Foundation grant No. 623/08 and by New York Metropolitan Research Fund.

References

[1] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: Shape completion and animation of people. *ACM Trans. Graph*, 24(3):408–416, 2005. 2

[2] M. Ben-Chen, O. Weber, and C. Gotsman. Variational harmonic maps for space deformation. *ACM Trans. Graph.*, 28(3), 2009. 6

[3] F. Benmansour, G. Carlier, G. Peyr, and F. Santambrogio. Derivatives with respect to metrics and applications: Subgradient marching algorithm. Technical report, INRIA, 2009. 2

[4] I. Borg and P. Groenen. *Modern multidimensional scaling*. Springer-Verlag, New York, 1997. Theory and applications. 2

[5] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *PNAS*, 103:1168–1172, Jan. 2006. 2

[6] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Topology-invariant similarity of nonrigid shapes. *IJCV*, 81(3):281–301, Mar. 2009. 2

[7] Y. Devir, G. Rosman, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Shape reconstruction with intrinsic priors. Technical Report CIS-2009-03, Department of Computer Science, Technion – Israel Institute of Technology, April 2009. 2

[8] I. Eckstein, J.-P. Pons, Y. Tong, C. C. J. Kuo, and M. Desbrun. Generalized surface flows for mesh processing. In *Proc. SGP*, pages 183–192, 2007. 2

[9] A. Elad (Elbaz) and R. Kimmel. On bending invariant signatures for surfaces. *PAMI*, 25(10):1285–1295, 2003. 2

[10] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Proc. Symp. Geom. Processing*, pages 197–206, 2005. 3

[11] M. Gromov. *Structures métriques pour les variétés riemanniennes*. 1981. 2

[12] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. *ACM Trans. Graph.*, 26(3):64, 2007. 2

[13] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *PNAS*, 95(15):8431–8435, 1998. 2, 3

[14] M. Körtgen, G. J. Park, M. Novotni, and R. Klein. 3d shape matching with 3d shape contexts. In *The 7th Central European Seminar on Computer Graphics*, Apr. 2003. 3

[15] N. Litke, M. Droske, M. Rumpf, and P. Schröder. An image processing approach to surface matching. In *Proc. SGP*, pages 207–216, July 2005. 2

[16] F. Mèmoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005. 2

[17] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Trans. Graph.*, 21(4):807–832, 2002. 3

[18] M. Salzmann, J. Pilet, S. Ilic, and P. Fua. Surface deformation models for nonrigid 3D shape recovery. *PAMI*, 29(8):1481–1487, Aug. 2007. 2

[19] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *PNAS*, 93(4):1591–1595, 1996. 3

[20] V. Surazhsky, T. Surazhsky, D. Kirsanov, S. J. Gortler, and H. Hoppe. Fast exact and approximate geodesics on meshes. *ACM Trans. Graph.*, 24(3):553–560, July 2005. 3

[21] B. Triggs, P. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment: A modern synthesis. In *Vision Algorithms Workshop: Theory and Practice*, pages 298–372, 1999. 1

[22] J. N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Trans. Automatic Control*, 40(9):1528–1538, 1995. 3

[23] R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah. Shape from shading: A survey. *PAMI*, 21(8):690–706, 1999. 1