Detail-Preserving Mesh Unfolding for Nonrigid Shape Retrieval

YUSUF SAHILLIOĞLU
Computer Engineering Department, Middle East Technical University, Turkey
and
LADISLAV KAVAN
School of Computing, University of Utah

We present a shape deformation algorithm that unfolds any given 3D shape into a canonical pose that is invariant to nonrigid transformations. Unlike classical approaches, such as least-squares multidimensional scaling, we preserve the geometric details of the input shape in the resulting shape, which in turn leads to a content-based nonrigid shape retrieval application with higher accuracy. Our optimization framework, fed with a triangular or a tetrahedral mesh in 3D, tries to move each vertex as far away from each other as possible subject to finite element regularization constraints. Intuitively this effort minimizes the bending over the shape while preserving the details. Avoiding geodesic distances in our computation renders the method robust to topological noise. Compared to state-of-the-art approaches, our method is simpler to implement, faster, more accurate in shape retrieval, and less sensitive to topological errors.

Categories and Subject Descriptors: I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling—*Curve, surface, solid, and object representations*

General Terms: Shape Representation, Shape Retrieval

Additional Key Words and Phrases: Canonical pose, detail preservation, nonrigid shape retrieval

ACM Reference Format:

Yusuf Sahillioğlu and Ladislav Kavan. 2016. Detail-preserving mesh unfolding for nonrigid shape retrieval. ACM Trans. Graph. 35, 3, Article 27 (May 2016), 11 pages.
DOI: http://dx.doi.org/10.1145/2893477

1. INTRODUCTION

Three-dimensional (3D) shapes are the building blocks of many computer graphics and vision applications ranging from interpolation [Freifeld and Black 2012] to deformation [Nealen et al. 2006]. Consequently, it is important to find the best representation of these shapes tailored to the specific application in mind. In this article, we will introduce a new canonical representation that is shown to be more advantageous for the shape retrieval application.

Thanks to the increasing sizes and numbers of 3D model repositories, the task of fast and robust retrieval of shapes never loses attention [Li et al. 2014]. Our new canonical shape representation facilitates a scalable search algorithm that is shown to be effective on 3D shape databases of articulated, or equivalently nonrigid, objects. Existing approaches essentially address this problem by performing 3D nonrigid shape comparisons using (i) descriptors, (ii) global intrinsic properties, or (iii) canonical forms, where the last option is the most promising one as the canonical form can easily be integrated into a simpler and well-studied rigid shape retrieval process. Despite this advantage, there has been few canonical forms proposed in the literature. Moreover, all but the most recent one [Lian et al. 2013] are based on embedding procedures that introduce serious distortions (Figure 1). The feature-preserving canonical form of Lian et al. [2013] has obtained significantly better retrieval accuracy than the state-of-the-art canonical methods. This is due to the fact that semantically similar objects with varying geometric details could not be well distinguished by a canonical pose that distorts details. Our algorithm, compared to this closest work [Lian et al. 2013], is simpler to implement, faster, more accurate in shape retrieval, and less sensitive to topological noise. In addition to this main contribution, our canonical pose simplifies texture mapping and allows geodesic distance approximation by revealing the extrinsic properties of the shapes.

In order to compute the detail-preserving canonical pose from a given arbitrary pose, we formulate an optimization problem whose solution separates each mesh vertex as far as possible while preserving the geometric details of the original shape such as the bumps and cavities on the surface. Hence we promote a pose with minimum bending and maximum geometric detail.

The source code and the executables for the method that we present in this article are publicly available at http://www.ceng.metu.edu.tr/∼ys/pubs.

2. RELATED WORK

Shape retrieval is an important problem coming in various scenarios based on the query paradigms as well as the supported deformations between the query and the database models [Tangelder and Veltkamp 2008]. The former includes query-by-keyword, query-by-sketch, query-by-example, or a combined usage, whereas the latter determines the similarity measure to be considered while

This work was supported by TUBITAK under the project EEEAG-115E471, and the NSF awards IIS-1622360 and IIS-1350330.
Authors' addresses: Y. Sahillioğlu (Current address), CENG METU, Universiteler Mah. Dumlupinar Blv. No:1, Ankara, Turkey; email: ys@ceng.metu.edu.tr; L. Kavan, School of Computing, 50 S. Central Campus Drive, University of Utah, Salt Lake City, USA; email: ladislav.kavan@gmail.com.

DOI: http://dx.doi.org/10.1145/2893477

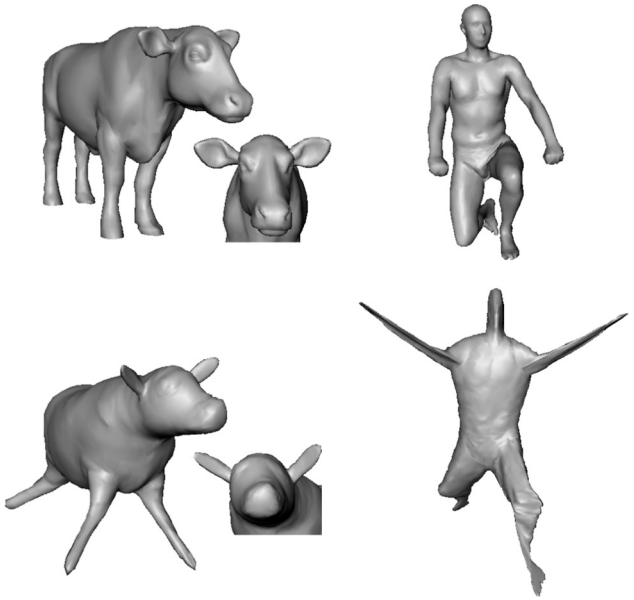

Fig. 1. Original shapes (top row) unfolded by minimization of Equation (2), or equivalently, by the least-squares MDS method [Elad and Kimmel 2003] (bottom row). Note the loss of details.

matching the query with the database models; popular ones being rigid matching and nonrigid matching. We address 3D shape retrieval with nonrigid matching and query-by-example setting, meaning that an example shape in arbitrary pose will be nonrigidly matched to the database models with the hope of returning the semantically similar objects to the user.

Shape representation issue is closely related to the shape retrieval problem. One may represent the exemplar query shape and the database shapes with global descriptors to convert the problem into a lower-dimensional and hence simpler descriptor matching problem. The popular descriptors that are invariant to rigid transformations, that is, translations and rotations, for rigid shape retrieval are spin images [Johnson and Hebert 1999; Assfalg et al. 2007], wavelets [Paquet et al. 2000], statistical moments [Zhang and Chen 2001; Novotni and Klein 2003], shape contexts [Belongie et al. 2002], shape distributions [Osada et al. 2002], spherical harmonics [Funkhouser et al. 2003; Kazhdan et al. 2003], and lightfield descriptor [Chen et al. 2003]. Note that Funkhouser et al. [2003] and Chen et al. [2003] allow sketch queries as part of their example-based retrieval engines for rigid shapes. Alternatively, rigid shape retrieval problem may be solved by efficiently registering the query to each database models in the low-dimensional space of rigid transformations [Besl and Mckay 1992]. The initial guess for Besl and Mckay [1992] can always be improved by the fast Principal Component Analysis (PCA)-based alignment [Kazhdan 2007]. A simple correspondence based on the closest points of the registered pair would then yield a cost descriptor that is useful for retrieval purposes.

For the nonrigid shape retrieval, which is more generic than the rigid counterpart, the existing methods are based on (i) articulation-invariant descriptors, (ii) global intrinsic properties, or (iii) articulation-invariant postures (a.k.a. the canonical poses). The descriptors here are invariant to nonrigid transformations (i.e., rigid transformations plus bending). Laplace-Beltrami differential operator is used extensively in the formulation of these descriptors [Reuter et al. 2006; Sun et al. 2009; Zaharescu et al. 2009; Bronstein and Kokkinos 2010; Raviv et al. 2010; Bronstein et al. 2011]. A global intrinsic property for all nonrigid shapes is the geodesic distance, which is consequently a good clue for nonrigid matching. Many algorithms exploit the preservation of geodesic distances under nonrigid transformations in order to utilize descriptors [Mahmoudi and Sapiro 2009] or graphs [Hilaga et al. 2001; Barra and Biasotti 2013] for nonrigid shape retrieval. Another important usage of the geodesic distances is the creation of canonical poses that are invariant to nonrigid transformations. In the canonical pose, the Euclidean distance between any two surface points approximates the geodesic distance between the corresponding vertices in the original pose. This idea is realized with the classical multidimensional scaling (MDS) method [Gower 1966; Jain and Zhang 2007] where the leading eigenvectors of the geodesic affinity matrix define the canonical pose. A similar pose is achieved with the least-squares MDS where a similar objective function is minimized [Elad and Kimmel 2003; Au et al. 2010]. Both of these MDS procedures require dense geodesic distance computations [Crane et al. 2013] which render them inappropriate for high-resolution meshes. de Silva and Tenenbaum [2002] and Panozzo et al. [2013] address this issue by finding geodesics only for the few landmark points and interpolating the result for the remaining data points. Another problem with the geodesic-based MDS methods is their vulnerability to the topological noise as a small modification on shape connectivity may change geodesics, that is, the shortest paths, drastically. To remedy this problem, Rustamov [2007] proposes a canonical pose based on Laplace-Beltrami eigenfunctions, whose sign and order switch issue, however, poses another problem. Lipman and Funkhouser [2009] also avoids the use of geodesic distances by conformally mapping shapes with sphere topology into the extended complex plane via Möbius embedding. All these canonical poses exhibit significant distortions, as it is impossible to satisfy all pairwise distances simultaneously in the restricted flat Euclidean space [Bronstein et al. 2008]. In order to eliminate the distortion, Bronstein et al. [2006] propose to embed one of the shapes to be matched into the surface of the other via the generalized MDS, which requires minimization of a nonconvex stress function. Due to the heavy computational complexity involved, generalized MDS is presented as a solution for the shape correspondence problem that deals with only two shapes, rather than the problem of shape retrieval from large databases. Given plenty of execution time, generalized MDS, as well as the other nonrigid shape correspondence algorithms [Jain and Zhang 2006; Ovsjanikov et al. 2010; Sahillioğlu and Yemez 2012, 2013], may be utilized to address the retrieval problem as the correspondence cost defines a plausible similarity measure.

Among the three types of nonrigid shape retrieval methods discussed above, namely (i) descriptor-based, (ii) intrinsic-based, and (iii) canonical, the last one is generally promoted as the most promising approach because by definition all shapes of the same class converges to the same canonical pose up to rigid transformations. This makes all the simpler and well-acknowledged rigid retrieval techniques available for the nonrigid retrieval problem, which is a big win. It is then natural to seek for improvements on the detail-oblivious canonical representation as objects of the same class with varying geometric details could not be distinguished sufficiently in the original setting. The first work addressing this issue to some extent is Lian et al. [2013], which deforms near-rigid mesh segments toward the corresponding components on the distorted least-squares MDS canonical pose. Although details are preserved to some extent, segmentation errors on the surface cause problems on the output shape (Figure 12). Besides the computational cost is too expensive for real search engines and the method is sensitive to topological

errors. Their resulting detail-preserving canonical pose, on the other hand, performs much better than the detail-oblivious canonical poses for nonrigid shape retrieval purposes, as expected. We are inspired by this improvement and introduce our detail-preserving 3D canonical posing algorithm, which produces a smoother and more accurate unfolding in a smaller amount of time than that of Lian et al. [2013], hence improving their good retrieval performance even further. We also note that we avoid geodesic distances in our computations for a scalable method that is also robust to topological noises. Thanks to the simple closed form of our deformation model, our algorithm is not only efficient but also easy to implement. We finally note that in a concurrent work [Sahillioglu 2015] obtains a detail-preserving canonical form by treating the vertices of the Landmark MDS embedding [de Silva and Tenenbaum 2002] as handles to deform the original shape. However, their work employs the problematic geodesic distances and a simpler yet faster deformation regularization energy. Instead, we do not use geodesics and achieve regularization with a more sophisticated scheme based on springs and the finite element method, which in turn achieves more accurate results in terms of element inversions and retrieval performance.

Our method is based on shape deformation, which allows us to explicitly control the geometric details of the shapes, a property that is lacking in all the embedding-based canonical pose creation methods. There are different deformation energies that guide shape deformation by measuring the difference between the current deformed configuration and the initial fixed rest-pose [Nealen et al. 2006; Botsch and Sorkine 2008]. The most basic one is the Dirichlet energy [Bonet and Wood 1997] that penalizes rotations and consequently allows only small deviations from the initial pose. Green's strain energy [Bonet and Wood 1997] allows not only the proper rotations but also the improper ones (i.e., reflections), which in turn enables a better separation from the initial pose while permitting undesired element inversions. This problem is handled in the state-of-the-art deformation energies that facilitate large deformations with no or very few inverted elements [Irving et al. 2004; Muller et al. 2005; Sorkine and Alexa 2007; Liu et al. 2008; Chao et al. 2010; Stomakhin et al. 2012].

Deformation energies become more interesting when they are coupled with positioning and regularization constraints. In interactive shape deformation, for instance, user specifies the position of the manipulation handles as the positioning constraints that work against the deformation energy [Sorkine et al. 2004]. Regularization constraints, on the other hand, try to keep the mesh in good shape during the whole deformation process. They essentially help the deformation energy term by additional soft or hard specifications on some regional properties such as angle, area, and volume [Muller et al. 2004; Irving et al. 2007; Adams et al. 2008; Schuller et al. 2013]. We adapt the composite element approach of Irving et al. [2007], which enforces constant volume in one ring of each vertex.

Finally, we mention the special poses in the literature that are used for tasks other than shape retrieval. A commonality of all these poses is that the geometric details are preserved. Mitra et al. [2007], for instance, strives to bring a given arbitrary pose into a special pose, which is in their case the symmetrized pose that enhances Euclidean symmetries present in the input. They find the symmetrizing transformation based on a set of corresponding symmetric point pairs computed by a curvature-based search algorithm. Twigg and Kačić-Alesić [2011] approximately compensates the mesh sagging effect under gravity by estimating rest-length parameters of a spring-mass system. Inverse design methodsin Skouras et al. [2012] and Chen et al. [2014] aim to compute a rest-shape for 3D printing, which can deform into the desired target shape under specified forces when fabricated. Based on the concepts of electrostatics, Wang et al. [2013] achieve spherical surface parameterization for shapes with arbitrary topology.

3. ALGORITHM

In this section, we describe our shape deformation algorithm for canonical pose optimization. *Deformation* is defined as a change in the shape or pose of an object due to applied forces. In other words, deformation is a mapping from $\mathbb{R}^m$ to $\mathbb{R}^m$, where we consider $m = 3$ in this work.

The input to our system is an elastic 2D triangular surface embedded in $\mathbb{R}^3$. For more realistic results, we consider volumetric elasticity, which requires tetrahedralizing the entire volume of the input surface [Jacobson et al. 2013]. After this preprocessing, we denote with $V = \{1, 2, \ldots, n\}$ and $E = \{1, 2, \ldots, e\}$ the set of vertex and edge indices, respectively. The position of vertex i of our tetrahedral mesh M is given by $\mathbf{v}_i$.

Our task is to bring M into its canonical pose, which can be realized by minimizing the bending on it while conserving the shape details. The problem can then be formulated as a search over all possible vertex locations so as to minimize a convenient energy functional $E(\mathbf{v})$:

$$\mathbf{v}^* = \arg\min_{\mathbf{v}} \; E(\mathbf{v}). \tag{1}$$

The search space is reduced by exploiting the fact that the optimal locations $\mathbf{v}^*$ for a bending-free pose move every vertex as far away from each other as possible while keeping M in good shape with full original details. This is equivalent to detail-preserving unfolding of the input mesh. In the following, we derive our $E(\mathbf{v})$ that addresses this requirement.

3.1 Canonical Pose Optimization Energy

We start with the simplest energy function:

$$E(\mathbf{v}) = \sum_{i<j} (||\mathbf{v}_i - \mathbf{v}_j|| - g(i, j))^2, \tag{2}$$

where $g(.,.)$ is the geodesic distance between two vertices on a given surface. This is a mass-spring system, arguably the simplest deformation model consisting of point masses connected together by massless springs. Spring forces are governed by Hooke's law (see Nealen et al. [2006] for a more detailed discussion). Our mass-spring system in Equation (2) consists of V as masses and the so-called geodesic springs that run between all pairs of vertices $\{(i, j)|i < j\} \in V$. Minimization of this energy makes the lengths of the geodesic springs close to the corresponding pairwise geodesic distances, which effectively separates any given point by the amount of the geodesic distance in between. This is indeed a well-known Euclidean embedding method, Multidimensional Scaling (MDS), that comes in least-squares form [Elad and Kimmel 2003]. The primary purpose of the MDS methods is to represent the pairwise (dis)similarity data (e.g., geodesic distances) as Euclidean distances in a low-dimensional space in order to make it accessible to visual inspection and further exploration. Since MDS methods do not utilize any information other than the pairwise data, they are very likely to lose the geometric details of the input in the resulting embedding, as exemplified in Figure 1 (see also Figures 7–9 and 12 for outputs of different types of MDS method).

In order to alleviate the detail-preservation problem of Equation (2), we add new springs to our existing mass-spring system:

$$E(\mathbf{v}) = \frac{1}{2} \sum_{(i,j)\in\mathcal{G}\cup\mathcal{E}} k_{ij}(||\mathbf{v}_i - \mathbf{v}_j|| - r_{ij})^2, \tag{3}$$

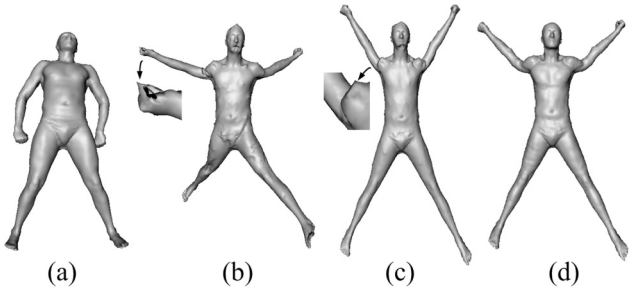

Fig. 2. Original mesh (a) and its canonical pose due to Equation (3) (b), Equation (4) (c), and Equation (4) with constraints (d). Inverted or highly-distorted parts are zoomed in.

where the spring set is the union of the geodesic springs $\mathcal{G}$ (between all pairs of vertices $\{(i,j)|i < j\} \in V$) and the so-called edge springs $\mathcal{E}$ (between connected pairs of vertices $(i,j) \in E$), and k_{ij} is the spring stiffness. The rest length r_{ij} takes the value of geodesic distance (for geodesic springs) or original edge length (for edge springs).

The problem with Equation (3) is its dependence on the original geodesic distances and the pure spring-based approach to capture volumetric elasticity (see Figure 2(b)). Geodesic distances require significant time and storage complexity for high-resolution meshes. More importantly, since we expect some stretching in the unfolded pose that we aim for, original geodesic distances should alter in an unpredictable manner. Therefore, we remove geodesic distances entirely from our system and employ a charge-based method (Equation (4)). With springs, on the other hand, it is very difficult to control the compressibility of a tetrahedron, and also impossible to detect inversions. Setting stiffness parameters for nonuniformly tesselated meshes is yet another issue. Therefore, we add the support for the finite element method, representing a principled approach to address these problems (Equations (5) and (6)).

Our proposed energy function for the canonical pose computation is as follows:

$$E(\mathbf{v}) = \frac{1}{2}\left(\sum_{(i,j)\in\mathcal{C}} -k_{ij}||\mathbf{v}_i - \mathbf{v}_j||^2 + \alpha \sum_{(i,j)\in\mathcal{E}} k_{ij}(||\mathbf{v}_i - \mathbf{v}_j|| - r_{ij})^2 \right), \tag{4}$$

where $\mathcal{C}$ denotes the charge springs (between all pairs of vertices $\{(i,j)|i < j\} \in V$), $\mathcal{E}$ is the set of edge springs as before, k_{ij} is the spring stiffness, and α is a tuning parameter controlling the effect of the edge springs. Spring stiffness value k_{ij} is always 1 in this paper. We control the regularization effect with α parameter, and elaborate on the choice of α in Section 4.2.2. The rest length r_{ij} is the original edge length for edge springs and equal to 0 for charge springs. The charge springs essentially maximize the distance between vertex pairs by minimizing the negative distance between them. The name charge comes from the fact that identically charged particles repulse each other in electric fields, just like our vertices V repulsing each other. To summarize, edge springs help detail-preservation, whereas charge springs are responsible for the unfolding operation. The main reason of the success of the proposed energy function (Equation (4)) over the original one (Equation (2)) is the elimination of unpredictable geodesic distances in the unfolded pose. Another advantage is time and storage efficiency as well as topological noise robustness (Figure 13) obtained by ruling out geodesic computations. We visualize the charge and edge springs on a didactic example in Figure 3. See also Figure 2(c) and Figure 6 (right) to see an input pose and the resulting canonical pose based on these springs.

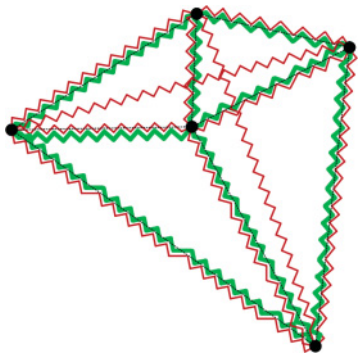

Fig. 3. A 2D mesh with 5 vertices and 8 edges has $\binom{5}{2} = 10$ charge springs (red) and 8 edge springs (bold green). Edge springs are generally stiffer to better control the repulsing effect of the charge springs, for example, edge springs prevent the vertical charge spring (the longest one) from separating the top and the bottom vertices by infinite distance.

Although our final energy in Equation (4) is free of geodesic distances and hence the associated problems, it still involves regularization problems due to springs, as shown in Figure 2(c) (see also Figure 6 (right)). We address this issue by minimizing Equation (4) subject to finite element constraints. The first set of constraints is responsible for preserving local volume around each vertex:

$$c_i(\mathbf{v}) : \sum_{t\in\eta(i)} \text{vol}(t) = l_i \quad \forall i \in V, \tag{5}$$

where $\eta(i)$ are the tetrahedra indices in the one-ring neighborhood of vertex i, function $\text{vol}(t)$ measures the volume of t, and l_i is the initial local volume around i in the original pose. Note that this is the one-ring composite element constraint in Irving et al. [2007], which avoids locking, that is, the inability of a given overconstrained finite element space to approximate solutions. The second set of constraints are inequalities to prevent inversions:

$$c_t(\mathbf{v}) : \text{vol}(t) > \epsilon \quad \forall t \in T, \tag{6}$$

where T is the set all tetrahedra indices and ϵ is a small constant for numerical inaccuracies, being one-hundredth of the minimum volume of all the original tetrahedra. These two sets of constraints along with the edge springs achieve a nice regularization effect, which we visualize in Figure 2(d) (see also Figure 6 (middle)). Note that edge springs are still a fundamental part of our formulation, as they prevent constraints from being satisfied with arbitrary edge lengths. In addition to or in place of Equations (5) and (6), we have also tried another constraint candidate, which preserves the local edge length summation around each vertex, that is, sum of lengths of the edges incident to a vertex is preserved, and yet not observed any improvement. For an efficient solution of our nonlinear constrained optimization problem, we have employed the interior-point method implemented in a state-of-the-art KNITRO solver [Byrd et al. 2006]. Specifically, we solve the following constrained optimization problem:

$$\begin{aligned} &\underset{\mathbf{v}}{\text{minimize}} \quad E(\mathbf{v}) \\ &\text{subject to} \quad c_i(\mathbf{v}),\ \forall i \in V \quad \text{and} \quad c_t(\mathbf{v}),\ \forall t \in T. \end{aligned}$$

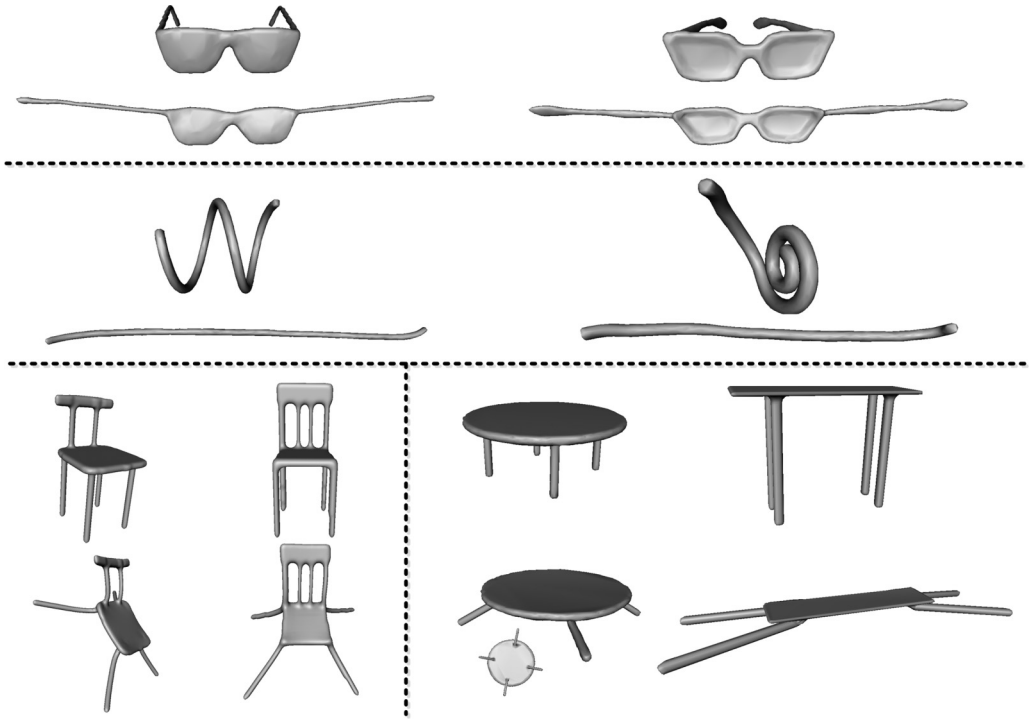

Fig. 4. Two examples of input mesh (top) and its resulting canonical pose (bottom) for each nonanimal class of Watertight dataset.

4. EXPERIMENTAL RESULTS

4.1 Datasets

We demonstrate the potential of our detail-preserving mesh unfolding algorithm in a content-based nonrigid shape retrieval application. To this end, we use Watertight [Giorgi et al. 2007], McGill [Siddiqi et al. 2008], and SCAPE [Anguelov et al. 2005] datasets, union of which forms a large database of 507 models with articulating parts. We query this database using random individual models selected from the database itself. The part that we have used from the Watertight benchmark consists of glasses, ant, chair, octopus, table, fish, rope, armadillo, and fourleg classes of cardinality 20 each. The remaining 11 classes have been omitted due to uninteresting type and/or pose of their constituent elements. We have used all the 256 articulated models in the McGill database consisting of various classes ranging from humans to spiders. SCAPE dataset, on the other hand, is a reconstructed sequence of a human actor in 71 different fixed-connectivity poses. Models in both datasets are nonuniformly sampled. We have also used the hand model from Panozzo et al. [2013] to compare with their embedding result. As a side application, we also demonstrate our geodesic distance approximation capability on the SCAPE and Watertight models.

4.2 Evaluation Criteria

We evaluate the quality of our resulting canonical poses visually (Figures 4–8) and quantitatively, namely in terms of the amount of stretching (Table I (first row)), geodesic distance approximation accuracy (Figure 9 and Table I (second row)), and nonrigid retrieval performance (Figures 10 and 11 and Tables II–IV). We also compare our canonical poses and their retrieval performance with the state-of-the-art approaches in Section 4.2.3.

4.2.1 *Visual Evaluation.* In Figures 4 and 5, we see the detail-preserving unfolding of example models from various Watertight classes, whereas in Figure 6 we show the benefit of our local volume preservation and inversion prevention constraints. The embedding step in the framework of Panozzo et al. [2013] is based on Metric MDS [Cox and Cox 2000] and least-squares meshes [Sorkine and Cohen-Or 2004], and aims to unfold the input mesh without distorting the original triangles significantly. Since this is exactly what we want to achieve in our canonical poses, we find it meaningful to compare our result with their embedding. Our result (Figure 7 (fourth row)) respects the original shape details much more than [Panozzo et al. 2013] does (Figure 7 (third row)), which implies that their algorithm can be improved by employing our unfolding technique. We also show another unfolding by the well-known Landmark MDS method [de Silva and Tenenbaum 2002], which led to the worst result in terms of shape preserving (Figure 7 (second row)).

4.2.2 *Quantitative Evaluation.* We measure the percentage of stretching ($\mathcal{S}$) by comparing the maximum geodesic distances before (g_0) and after (g_1) our canonical pose computation:

$$\mathcal{S} = \left\lceil 100 \frac{g_1 - g_0}{g_0} \right\rceil. \tag{7}$$

Due to the nature of the unfolding task, we observe stretching to a certain extent that is controlled by parameter α in Equation (4) (see Table I). We fix this value to 5K in most of the experiments while noting that the lower values allow further stretches. Models of the glasses class, for instance, overstretch with the usual value of $\alpha = 5$K. We consequently obtain a more plausible unfolding with more powerful edge springs imposed by $\alpha = 20$K, as illustrated in Figure 8. Note that the executable of our algorithm takes this tuning

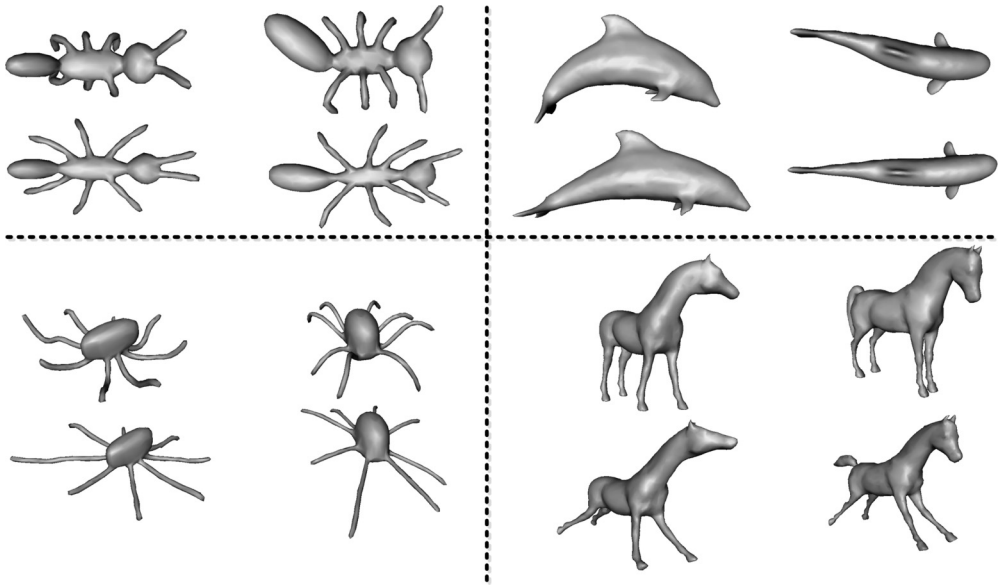

Fig. 5. Two examples of input mesh (top) and its computed canonical pose (bottom) for each animal class of Watertight dataset.

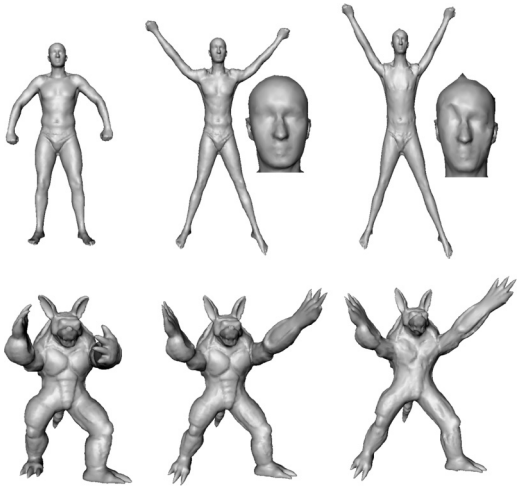

Fig. 6. SCAPE (top) and armadillo (meshes) at left are brought to their canonical poses with (middle) and without (right) the finite element constraints in Equations (5) and (6). Absence of constraints causes detail loss (chests) and distortion (heads) as well as additional stretching.

parameter α from the command prompt, which enables users to progressively update the value if necessary.

We also measure the similarity between the pairwise Euclidean distances at the computed canonical pose (**E**) and the corresponding geodesic distances on the manifold (**G**) via:

$$\mathcal{A} = \left\| \frac{\mathbf{E}}{||\mathbf{E}||_F} - \frac{\mathbf{G}}{||\mathbf{G}||_F} \right\|_F, \tag{8}$$

where values close to 0 imply a good geodesic approximation as the difference is taken between unit matrices. Visual inspection of this approximation is made in Figure 9 by running [de Silva and Tenenbaum 2002] based on the standard affinity matrix **G**, that

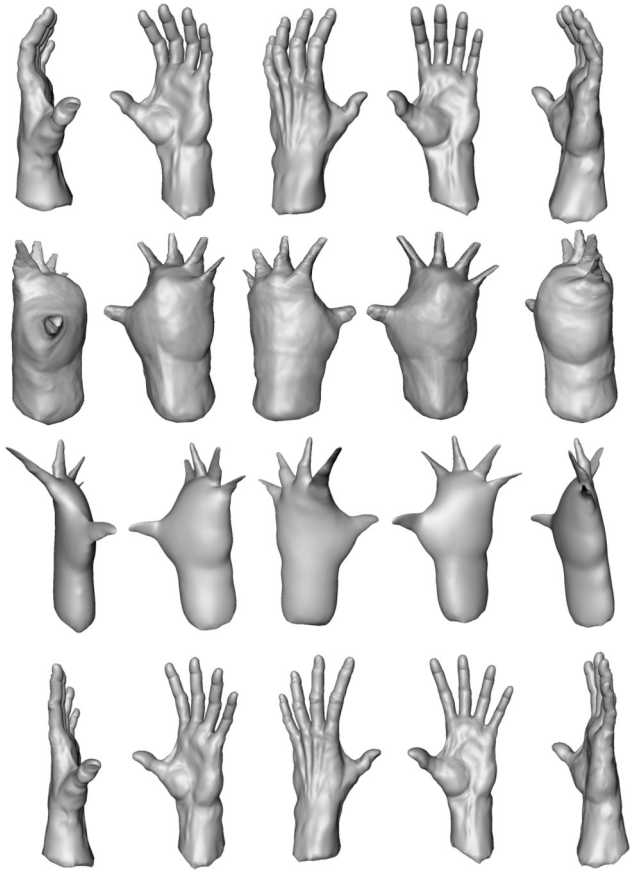

Fig. 7. Original hand model (top) is unfolded with, from top to bottom, de Silva and Tenenbaum [2002]; Panozzo et al. [2013], and our method.

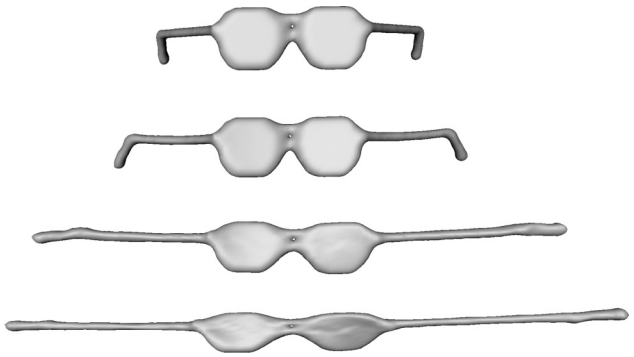

Fig. 8. Original glasses model (top) understretches with $\alpha = 100$K (second row), overstretches with $\alpha = 5$K (fourth row), and unfolds well with $\alpha = 20$K (third row).

Table I. Stretching Percentages $\mathcal{S}$ and Geodesic Approximation Error $\mathcal{A}$ after Canonical Pose Optimization

	Aa	B	C	D	E	F	G	H	I	J	K
$\mathcal{S}$	11	13	4	7	7	12	14	13	11	14	6
$\mathcal{A}$	.55	.46	.71	.74	.66	.50	.40	.51	.57	.44	.62

A: Glasses, B: Ant, C: Chair, D: Octopus, E: Table, F: Fish, G: Rope, H: Armadillo, I: Fourleg, J: Scape, K: Hand.

is, pairwise geodesic distances, and our affinity matrix $\mathbf{E}$, that is, pairwise Euclidean approximations of geodesic distances. Namely, $\mathbf{G}_{i,j} = g(i, j)$ and $\mathbf{E}_{i,j} = ||\mathbf{v}_i - \mathbf{v}_j||$. Recall that de Silva and Tenenbaum [2002] use M leading eigenvectors of the affinity matrix to obtain an M-dimensional spectral embedding of landmark points, and then computes embedding coordinates for the remaining data points based on their distances from the landmarks. Visual similarity between the middle and left columns of Figure 9 is therefore another sign that Euclidean distances in our resulting canonical pose efficiently approximate geodesic distances locally on the manifold. Finally, the correlation between $\mathcal{S}$ and $\mathcal{A}$ suggests that the more we allow stretching, the better the geodesic distances are approximated.

Our resulting canonical poses are invariant to nonrigid transformations, that is, all articulated shapes are expected to converge to the poses that may differ by purely rigid transformations. This fact encourages us to implement a simple shape correspondence method that takes two shapes S and T, unfolds them with our method, resolves the rigid transformation ambiguity by PCA alignment [Kazhdan 2007] followed by standard ICP [Besl and Mckay 1992], and matches the closest points between the aligned shapes. We demonstrate the moderate quality of a resulting correspondence in Figure 10 (top row), and note that one may infer more correspondences from the unfolded pairs of Figures 4 and 5. We also define the following ground-truth distortion $\mathcal{D}$ to measure the deviation of a correspondence $\phi : S \to T$ from the ground truth:

$$\mathcal{D}(\phi) = \frac{1}{|\phi|} \sum_{(i,j)\in\phi} g(\vartheta(i), j), \tag{9}$$

where $\vartheta(i)$ stands for the ground-truth correspondence of source vertex $i \in S$ on target T as known a priori for the SCAPE dataset. The maximum geodesic distance on the target model is normalized to 1 in order to simplify the interpretation of this measure. We observe average $\mathcal{D} = 0.178$ for 20 random pairs from the SCAPE, which is worse than the state-of-the-art nonrigid correspondence algorithms [Ovsjanikov et al. 2010; Sahillioğlu and Yemez 2013], yet accurate enough to initiate a pose-independent shape retrieval

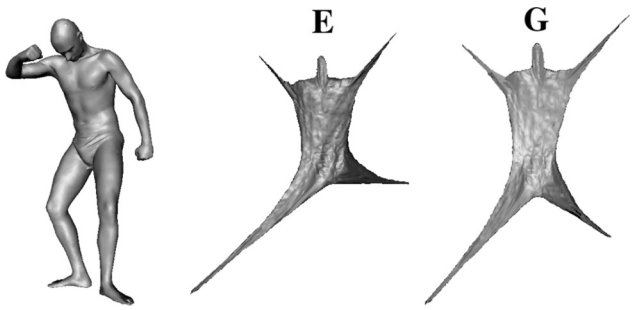

Fig. 9. Landmark MDS embeddings [de Silva and Tenenbaum 2002] of the input pose (left) based on affinity matrices $\mathbf{E}$ (our geodesic approximation, middle) and $\mathbf{G}$ (original geodesics, right).

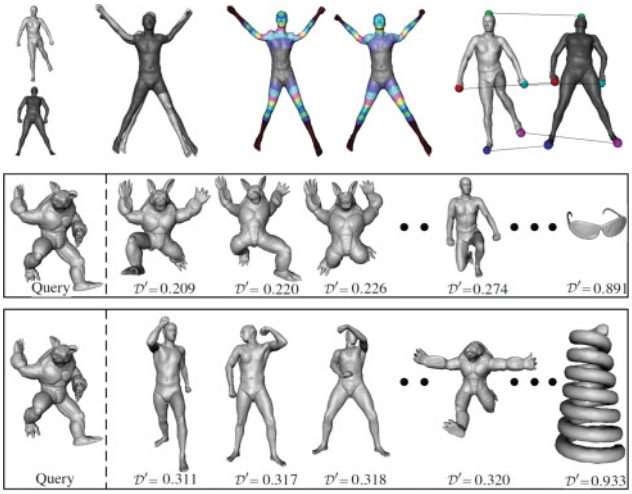

Fig. 10. Shape correspondence (top row) and retrieval (middle row) applications based on the computed canonical poses. Notice the distortion in the transfer of colors based on the computed correspondences. This inaccuracy can, however, be compensated in a retrieval application. Notice also the advantage of using our detail-preserving canonical pose (middle row) over a detail-oblivious canonical pose [Elad and Kimmel 2003] (bottom row) in discriminating an armadillo and a man.

technique. After representing the query and database shapes with their canonical poses, we perform pairwise matchings between query and database, as described above. We then measure shape similarities using the correspondence cost descriptor of Jain and Zhang [2007], which is,

$$\mathcal{D}'(\phi) = \frac{1}{|\phi|} \sum_{(i,j)\in\phi} ||\mathbf{v}_i - \mathbf{v}_j||, \tag{10}$$

where query vertex i maps to the database model vertex j after pairwise matching, and the maximum term in the summation is normalized to 1. These similarity scores, as shown under Figure 10, intuitively take smaller values when matching a query to semantically equivalent articulated shapes since the canonical pose representations of both shapes in this scenario will be very similar. Ranking based on $\mathcal{D}'$ yields the result set for pose-independent shape retrieval (Figure 10). By comparing the middle and bottom rows of Figure 10, we verify the main motivation of this article. As we claimed in Section 1, similar objects with varying geometric

Table II. Retrieval Performances in Terms of NN, 1-Tier, and 2-Tier Metrics on Our Large Database of 507 Models

Canonical Pose Creation	NN	1-tier	2-tier
Our method	99.1	90.2	97.3
Least-squares MDS [Elad and Kimmel 2003]	96.4	82.8	87.5
Classical MDS [Elad and Kimmel 2003]	91.4	75.8	79.9

Query and database models are compared Via Equation (10).

Table III. Retrieval Performances in Terms of NN, 1-Tier, and 2-Tier Metrics on Our Large Database of 507 Models

Canonical Pose Creation	NN	1-tier	2-tier
Our method	99.2	89.9	97.5
Least-squares MDS [Elad and Kimmel 2003]	96.1	82.4	87.6
Classical MDS [Elad and Kimmel 2003]	91.2	75.1	79.7

Query and database models are compared via CM-BOF.

Table IV. Retrieval Performances in Terms of NN, 1-Tier, and 2-Tier Metrics on the McGill Database

Canonical Pose Creation	NN	1-tier	2-tier
Our method	99.7	92.1	98.6
Lian et al. [2013]'s method	99.6	86.6	95.3
Least-squares MDS [Elad and Kimmel 2003]	98.1	80.9	86.9
Classical MDS [Elad and Kimmel 2003]	95.2	75.1	80.9

Query and database models are compared via CM-BOF.

details, such as a human and an armadillo, could not be well distinguished by a canonical pose that distorts details, such as least-squares MDS [Elad and Kimmel 2003] (Figure 10 (bottom row)).

In addition to the correspondence cost descriptor (Equation (10)) for solving the pose-independent retrieval problem, we employ a compact shape descriptor extracted from the resulting canonical form. Specifically, we implement the visual-similarity based Clock Matching Bag-of-Features (CM-BOF) approach [Lian et al. 2010] in order to (i) make the comparisons in Section 4.2.3 fair and (ii) achieve faster retrieval performance than the original correspondence-based approach. For shape similarity comparison with CM-BOF, we begin with the same PCA-based pose normalization [Kazhdan 2007] between the canonically posed query and the canonically posed database model. Since each database model is already equipped with an histogram of visual word occurrences in an offline process, the only task remaining is to extract the corresponding histogram on the query shape. To this end, SIFT features [Lowe 2004] computed on each of the several different views of the query is quantized against the visual vocabulary, which in turn represents the query as a histogram of visual word occurrences, and enables comparison with the database histogram. Once canonical forms are computed, a pair of shapes with 9K tetrahedra is compared with CM-BOF approach in only a few milliseconds, whereas the correspondence-based approach takes about a second. Both approaches achieve similar and satisfactory performances in retrieval accuracy (Tables II and III).

The average execution time of our canonical pose optimization algorithm on a 2.53GHz PC for SCAPE dataset, whose models have 45K tetrahedra, is 985 seconds. The hand model with 51K tets takes 1684 seconds to get unfolded into its canonical pose. When the number of tets drops to 35K, 20K, 9K, and 3K for various Watertight models, our method demands 575, 280, 21, and 4 seconds, respectively. Note that the number of iterations, and hence the execution times, depend highly on the spatial distance between the original pose and the canonical pose to be computed.

4.2.3 *Visual and Quantitative Comparisons.* We first evaluate our nonrigid retrieval accuracy on our large database of 507 models using three commonly used performance metrics, namely the

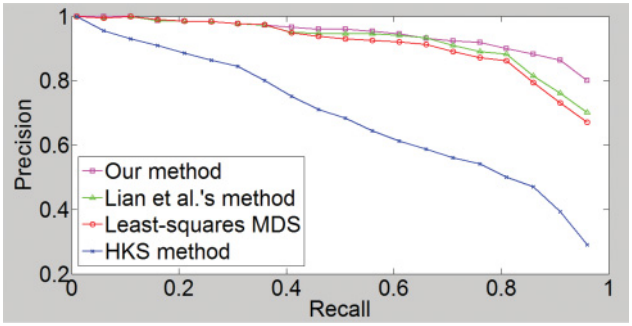

Fig. 11. Precision-recall plot of the retrieval methods.

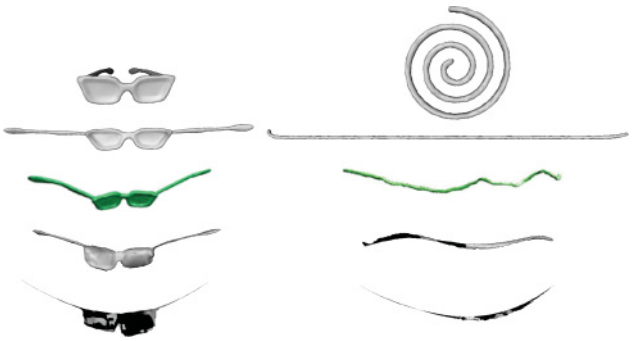

Fig. 12. Unfolding of two objects (top rows of the left and right columns) using our method (second row), Lian et al. [2013] (third row), least-squares MDS (fourth row), and classical MDS (fifth row). Third row images are taken from Lian et al. [2013].

Nearest Neighbor (NN), which is the percentage of the first-closest matches that belong to the query class, and the First Tier (1-tier), which is the ratio of the relevant matches to the size of the query class C when the number of retrieved models, that is, top K matches, is $|C|$. We relax $K = 2|C|$ to obtain the Second Tier (2-tier) metric. The retrieval method leading to the results in Table II converts each shape to its detail-preserving (Table II (first row)) or detail-oblivious (Table II (second and third rows)) canonical poses and then compares query with the database rigidly via Equation (10). Our high NN value indicates the potential of our algorithm in a classification application. Table III is created the same way as Table II except CM-BOF is used instead of Equation (10) for shape comparison.

We next pass to the McGill database in order to make fair comparisons with the only other detail-preserving canonical pose method in the literature [Lian et al. 2013]. Since we are unable to reproduce the complicated algorithm described in Lian et al. [2013], we directly import their visuals and quantifications obtained on the McGill database. In their comprehensive comparisons with the existing methods, they outperform the state of the art of the nonrigid retrieval approaches that utilize MDS canonical poses, which implies that performing better than [Lian et al. 2013] would put us in front of those methods as well. We employ the same performance metrics, which are NN, 1-tier, 2-tier, and Precision-Recall, and use the same shape comparison protocol, namely the CM-BOF approach (Table IV). For the Precision-Recall plots in Figure 11, the vertical axis is the Precision, which is the ratio of the relevant matches to the number of retrieved models, whereas the horizontal axis is the Recall, which is the ratio of relevant matches to the size

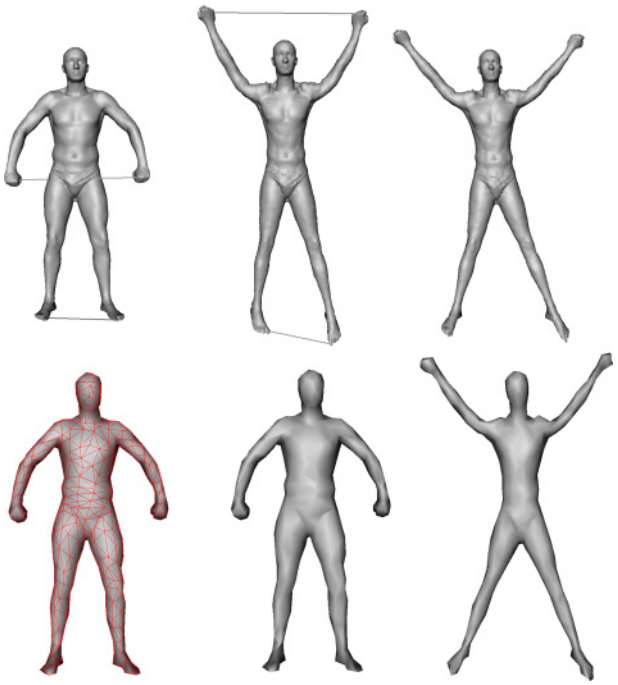

Fig. 13. (Top row) SCAPE model with topological noise on hands and feet (left) is unfolded to a model (middle) that is close to the unfolding of a noise-free input (right). (Bottom row) Low-resolution mesh with irregular tetrahedra (left and right) is unfolded to a model (right) that is close to the unfolding of a high-resolution input (top row-right).

of the query class. Ideally, this curve should be a horizontal line at unit precision. In addition to the plot associated with our method, we provide plots for the other canonical-based retrieval methods, namely Lian et al. [2013] and least-squares MDS. These retrieval methods all apply the same protocol described in the first paragraph of this section. We also provide the plot of a descriptor-based method, which directly compares the low-dimensional heat-kernel signatures [Sun et al. 2009] of the query and database models, as done in Bronstein et al. [2011]. The lower performance of the descriptor-based approach supports our claim in the Section 1 that canonical-based methods are more promising than descriptor-based counterparts as the canonical form can easily be integrated into a simpler and well-studied rigid shape retrieval process.

The main reason of our success over the closest work [Lian et al. 2013] is our resulting canonical pose that is more smooth and accurate in terms of unfolding the object, as shown in Figure 12. We also note that the reported execution time of Lian et al. [2013] for canonical pose creation of a model with 10K triangles is 243 seconds, whereas we perform the same task in about 10 seconds on a similar computer, which effectively makes our method more scalable. Finally, we are more robust to topological noise, for example, a "short circuit" in the feet, than [Lian et al. 2013] as we avoid the use of geodesic distances. We empirically show this robustness by adding 10 virtual edges between random point pairs of each McGill model and rerunning our retrieval method, which achieved almost the same performance in Table IV. Besides, we display the effect of topological noise on a SCAPE model in Figure 13, where we also demonstrate the unfolding result on a low-resolution and nonuniformly sampled mesh (1.5K irregular tetrahedra) in an attempt to show that different versions of a 3D mesh with different

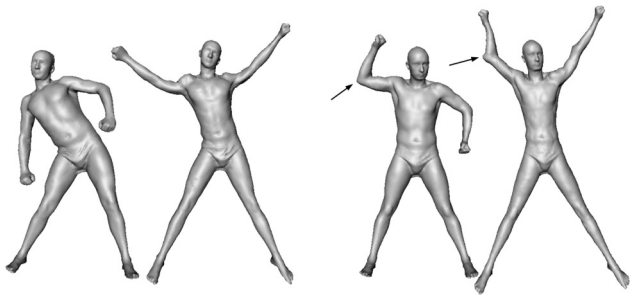

Fig. 14. Global (left pair) and local (right pair) similarities between original and canonical poses.

resolutions can have similar canonical poses using the proposed method with the same parameters and settings, for example, $\alpha = 5\text{K}$ in Equation (4). As verified in Figure 13 (bottom), neither the amount (resolution) nor the quality (sampling) of the tetrahedra in the input mesh degrades the performance of our method.

4.3 Limitations

Regardless of the original pose we start with, our algorithm always converges to an unfolded canonical pose that may facilitate applications such as pose-independent shape retrieval, geodesic distance approximation, and texture mapping [Zigelman et al. 2002], to name a few. However, we have to note that our regularization scheme makes the method slightly sensitive to the original pose in that one may observe global and/or local effects of the input on the output (Figure 14), which decreases the performance of a potential shape correspondence algorithm. We, however, show that such a loose correspondence can still be very useful in a shape retrieval framework.

5. CONCLUSION

We have introduced a shape deformation algorithm that aims to unfold a given deformable model without distorting its geometric details. This detail-preserving unfolding leads to a canonical pose that is invariant to nonrigid transformations and consequently suitable for nonrigid shape retrieval. We compute this canonical pose by minimizing an energy functional whose solution moves each mesh point as far away from each other as possible with respect to some regularization constraints. Our method proves useful, not only for detail-preserving mesh unfolding but also for pose-independent shape retrieval as well as efficient approximation of geodesic distances. It is also less sensitive to topological noise and works for models of arbitrary genus. A potential application for future work may be handling different scenarios of the shape retrieval problem, such as real-time or partial retrieval. As another future work, we plan to incorporate semantic parameters into our framework so as to create more specific poses such as the Vitruvian Man pose that suggests, for instance, foot to be one seventh of the height of a man.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their constructive comments, and Norm Badler, James O'Brien, Tiantian Liu, Mark Pauly and Lifeng Zhu for fruitful discussions.

REFERENCES

B. Adams, M. Ovsjanikov, M. Wand, H.-P. Seidel, and L. Guibas. 2008. Meshless modeling of deformable shapes and their motion. In *Proceedings of the SCA* (2008).

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

J. Assfalg, M. Bertini, and A. Delbimbo. 2007. Content-based retrieval of 3d objects using spin image signatures. *Trans. Multimedia* 9, 3 (2007), 589–599.

O. K.-C. Au, D. Cohen-Or, C.-L. Tai, H. Fu, and Y. Zheng. 2010. Electors voting for fast automatic shape correspondence. In *Computer Graphics Forum (Proc. Eurographics)* 29, 2 (2010), 645–654.

V. Barra and S. Biasotti. 2013. 3D shape retrieval using kernels on extended reeb graphs. *Pattern Recognit.* 46, 11 (2013), 2985–2999.

S. Belongie, J. Malik, and J. Puzicha. 2002. Shape matching and object recognition using shape contexts. *IEEE Trans. PAMI* 24, 4 (2002), 509–522.

P. J. Besl and N. D. Mckay. 1992. A method for registration of 3D shapes. *IEEE Trans. PAMI* 14, 2 (1992), 239–256.

J. Bonet and R. D. Wood. 1997. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge University Press, Cambridge, England.

M. Botsch and O. Sorkine. 2008. On linear variational surface deformation methods. *IEEE Trans. Visual. Comp. Graphics* 14, 1 (2008), 213–230.

A. M. Bronstein, M. M. Bronstein, and R. Kimmel. 2006. Generalized multidimensional scaling: A framework for isometry invariant partial surface matching. *Proc. Natl. Acad. Sci. USA* 103, 5 (2006), 1168–1172.

A. M. Bronstein, M. M. Bronstein, and R. Kimmel. 2008. *Numerical Geometry of Non-Rigid Shapes*. Springer, New York, USA.

M. Bronstein, M. Bronstein, L. Guibas, and M. Ovsjanikov. 2011. Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. Graph.* 30, 1 (2011), Article No. 1.

M. M. Bronstein and I. Kokkinos. 2010. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proceedings of the Computer Vision and Pattern Recognition (CVPR).*

R. Byrd, J. Nocedal, and R. Waltz. 2006. An integrated package for nonlinear optimization. In *Large-Scale Nonlinear Optimization* 83 (2006), 35–59.

I. Chao, U. Pinkall, P. Sanan, and P. Schroder. 2010. A simple geometric model for elastic deformations. *ACM Trans. Graph.* 29, 4 (2010), 38.

D. Chen, X. Tian, Y. Shen, and M. Ouhyoung. 2003. On visual similarity based 3D model retrieval. *Comput. Graphics Forum* 22, 3 (2003), 223–232.

X. Chen, C. Zheng, W. Xu, and K. Zhou. 2014. An asymptotic numerical method for inverse elastic shape design. *ACM Trans. Graph.* 33, 4 (2014), Article No. 95.

T. Cox and M. Cox. 2000. *Multidimensional Scaling*, 2nd ed. Chapman & Hall/CRC.

K. Crane, C. Weischedel, and M. Wardetzky. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. In *Proceedings of the SIGGRAPH.*

V. de Silva and J. B. Tenenbaum. 2002. Global versus local methods for nonlinear dimensionality reduction. In *Proceedings of the NIPS.* 705–712.

A. Elad and R. Kimmel. 2003. On bending invariant signatures for surfaces. *IEEE Trans. PAMI* 25, 10 (2003), 1285–1295.

O. Freifeld and M. J. Black. 2012. Lie bodies: A manifold representation of 3D human shape. In *Proceedings of the European Conference on Computer Vision (ECCV).* 1–14.

T. Funkhouser, P. Min, M. Kazhdan, J. Chen, A. Halderman, D. Dobkin, and D. Jacobs. 2003. A search engine for 3D models. *ACM Trans. Graph.* 22, 1 (2003), 83–105.

D. Giorgi, S. Biasotti, and L. Paraboschi. 2007. SHREC: Shape retrieval contest: Watertight models track.

J. C. Gower. 1966. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika* 53, 3–4 (1966), 325–338.

M. Hilaga, Y. Shinagawa, T. Kohmura, and T. Kunii. 2001. Topology matching for fully automatic similarity estimation of 3D shapes. In *Proceedings of the SIGGRAPH.*

G. Irving, C. Schroeder, and R. Fedkiw. 2007. Volume conserving finite element simulations of deformable models. In *Proceedings of the SIGGRAPH.*

G. Irving, J. Teran, and R. Fedkiw. 2004. Invertible finite elements for robust simulation of large deformation. In *Proceedings of the SCA.* 131–140.

A. Jacobson, L. Kavan, and O. Sorkine. 2013. Robust inside-outside segmentation using generalized winding numbers. *Proc. SIGGRAPH* 32, 4 (2013), 33:1–33:12.

V. Jain and H. Zhang. 2006. Robust 3D shape correspondence in the spectral domain. In *Proceedings of the IEEE Int. Conf. on Shape Modeling and Applications (SMI).* 118–129.

V. Jain and H. Zhang. 2007. A spectral approach to shape-based retrieval of articulated 3D models. *Comput.-Aided Des.* 39, 5 (2007), 398–407.

A. Johnson and M. Hebert. 1999. Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Trans. PAMI* 21, 5 (1999), 433–449.

M. Kazhdan. 2007. An approximate and efficient method for optimal rotation alignment of 3D models. *IEEE Trans. PAMI* 29, 7 (2007), 1221–1229.

M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. 2003. Rotation invariant spherical harmonic representation of 3D shape descriptors. *Comput. Graphics Forum (Proc. SGP)* 6 (2003), 156–164.

B. Li, Y. Lu, C. Li, A. Godil, et al. 2014. SHREC'14 track: Large scale comprehensive 3D shape retrieval. In *Proceedings of the Eurographics Workshop on 3D Object Retrieval.* 131–140.

Z. Lian, A. Godil, and X. Sun. 2010. Visual similarity based 3D shape retrieval using bag-of-features. In *Proceedings of the IEEE Int. Conf. on Shape Modeling and Applications (SMI).* 25–36.

Z. Lian, A. Godil, and J. Xiao. 2013. Feature-preserved 3D canonical form. *Int. J. Comput. Vision (IJCV)* 102, 1–3 (2013), 221–238.

Y. Lipman and T. Funkhouser. 2009. Möbius voting for surface correspondence. *ACM Trans. Graph.* 28, 3 (2009), Article No. 72.

L. Liu, L. Zhang, Y. Xu, C. Gotsman, and S. Gortler. 2008. A local/global approach to mesh parameterization. *Comput. Graphics Forum* 27, 5 (2008), 1495–1504.

D. G. Lowe. 2004. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision (IJCV)* 60, 2 (2004), 91–110.

M. Mahmoudi and G. Sapiro. 2009. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models* 71, 1 (2009), 22–31.

N. Mitra, L. Guibas, and M. Pauly. 2007. Symmetrization. *ACM Trans. Graph*. 26, 3 (2007), 63.

M. Muller, B. Heidelberger, M. Teschner, and M. Gross. 2005. Meshless deformations based on shape matching. *ACM Trans. Graph.* 24, 3 (2005), 471–478.

M. Muller, R. Keiser, A. Nealen, M. Pauly, M. Gross, and M. Alexa. 2004. Point based animation of elastic, plastic and melting objects. *Proc. SCA* (2004), 141–151.

A. Nealen, M. Muller, R. Keiser, E. Boxerman, and M. Carlson. 2006. Physically based deformable models in computer graphics. *Comput. Graphics Forum* 25, 4 (2006), 809–836.

M. Novotni and R. Klein. 2003. 3D Zernike descriptors for content based shape retrieval. In *Proceedings of the ACM Symposium on Solid Modeling and Applications.* 216–225.

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

M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. 2010. One point isometric matching with the heat kernel. *Comput. Graphics Forum* 29, 5 (2010), 1555–1564.

D. Panozzo, I. Baran, O. Diamanti, and O. Sorkine-Hornung. 2013. Weighted averages on surfaces. In *Proceedings of the SIGGRAPH.*

E. Paquet, M. Rioux, A. Murching, T. Naveen, and A. Tabatabai. 2000. Description of shape information for 2-D and 3-D objects. *Signal Proc. Image Comm.* 16, 1–2 (2000), 103–122.

D. Raviv, M. Bronstein, A. Bronstein, and R. Kimmel. 2010. Volumetric heat kernel signatures. In *Proceedings of the ACM Multimedia Workshop on 3D Object Retrieval.*

M. Reuter, F. Wolter, and N. Peinecke. 2006. Laplace-Beltrami spectra as 'shape-DNA' of surfaces and solids. *Comput.-Aided Des.* 38, 4 (2006), 342–366.

R. Rustamov. 2007. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the Symposium on Geometry Processing.* 225–233.

Y. Sahillioglu. 2015. A shape deformation algorithm for constrained multidimensional scaling. *Comput. Graphics* 53 (2015), 156–165.

Y. Sahillioğlu and Y. Yemez. 2012. Minimum-distortion isometric shape correspondence using EM algorithm. *IEEE Trans. PAMI* 34, 11 (2012), 2203–2215.

Y. Sahillioğlu and Y. Yemez. 2013. Coarse-to-fine isometric shape correspondence by tracking symmetric flips. *Comput. Graphics Forum* 32, 1 (2013), 177–189.

C. Schuller, L. Kavan, D. Panozzo, and O. Sorkine-Hornung. 2013. Locally injective mappings. *Comput. Graphics Forum (Proc. SGP)* 32, 5 (2013), 125–135.

K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. Dickinson. 2008. Retrieving articulated 3-D models using medial surfaces. *Mach. Vision and Applications* 19, 4 (2008), 261–275.

M. Skouras, B. Thomaszewski, B. Bickel, and M. Gross. 2012. Computational design of rubber balloons. *Comput. Graphics Forum* 31, 2 (2012), 835–844.

O. Sorkine and M. Alexa. 2007. As-rigid-as-possible surface modeling. In *Proceedings of the Symposium on Geometry Processing.* 109–116.

O. Sorkine and D. Cohen-Or. 2004. Least-squares meshes. In *Proceedings of the Shape Modeling International.* 191–199.

O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rossl, and H.-P. Seidel. 2004. Laplacian surface editing. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing.* 179–188.

A. Stomakhin, R. Howes, C. Schroeder, and J. Teran. 2012. Energetically consistent invertible elasticity. In *Proceedings of the SCA.*

J. Sun, M. Ovsjanikov, and L. Guibas. 2009. A concise and provably informative multi-scale signature based on heat diffusion. *Comput. Graphics Forum* 28, 5 (2009), 1383–1392.

J. Tangelder and R. Veltkamp. 2008. A survey of content based 3D shape retrieval methods. *Multimedia Tools Appl.* 39, 3 (2008), 441–471.

C. Twigg and Z. Kačić-Alesić. 2011. Optimization for sag-free simulations. In *Proceedings of the SCA.*

H. Wang, K. Sidorov, P. Sandilands, and T. Komura. 2013. Harmonic parameterization by electrostatics. *ACM Trans. Graph.* 32, 5 (2013), Article No. 155.

A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. 2009. Surface feature detection and description with applications to mesh matching. In *Proceedings of the Computer Vision and Pattern Recognition (CVPR).* 373–380.

C. Zhang and T. Chen. 2001. Efficient feature extraction for 2D/3D objects in mesh representation. In *Proceedings of the IEEE International Conference on Image Processing.* 3.

G. Zigelman, R. Kimmel, and N. Kiryati. 2002. Texture mapping using surface flattening via multidimensional scaling. *IEEE Trans. Vis. Comp. Graphics* 8, 2 (2002), 198–207.

Received November 2014; revised December 2015; accepted February 2016