Lie Bodies: A Manifold Representation of 3D Human Shape

[1]Oren Freifeld and [2]Michael J. Black

[1]Division of Applied Mathematics, Brown University, Providence, RI 02912, USA
[2]Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany
freifeld@dam.brown.edu, black@is.mpg.de

Abstract. Three-dimensional object shape is commonly represented in terms of deformations of a triangular mesh from an exemplar shape. Existing models, however, are based on a Euclidean representation of shape deformations. In contrast, we argue that shape has a manifold structure: For example, summing the shape deformations for two people does not necessarily yield a deformation corresponding to a valid human shape, nor does the Euclidean difference of these two deformations provide a meaningful measure of shape dissimilarity. Consequently, we define a novel manifold for shape representation, with emphasis on body shapes, using a new Lie group of deformations. This has several advantages. **First** we define triangle deformations exactly, removing non-physical deformations and redundant degrees of freedom common to previous methods. **Second**, the Riemannian structure of *Lie Bodies* enables a more meaningful definition of body shape similarity by measuring distance between bodies on the manifold of body shape deformations. **Third**, the group structure allows the valid composition of deformations. This is important for models that factor body shape deformations into multiple causes or represent shape as a linear combination of basis shapes. **Finally**, body shape variation is modeled using *statistics on manifolds*. Instead of modeling Euclidean shape variation with Principal Component Analysis we capture shape variation on the manifold using Principal Geodesic Analysis. Our experiments show consistent visual and quantitative advantages of Lie Bodies over traditional Euclidean models of shape deformation and our representation can be easily incorporated into existing methods.

Key words: Shape deformation, Lie group, Statistics on manifolds

1 Introduction

Three dimensional mesh models of objects play a central role in many computer vision algorithms that perform *analysis-by-synthesis* [1–4]. Capturing the variability of 3D meshes for an object class is critical and the increasing availability of 3D mesh data enables statistical learning methods to be used to build such models. In particular, for capturing human shape variation, *deformable template* models are popular for representing non-rigid deformations and articulations [1, 2, 5]. Such models have wide application in vision, graphics, virtual reality, shape

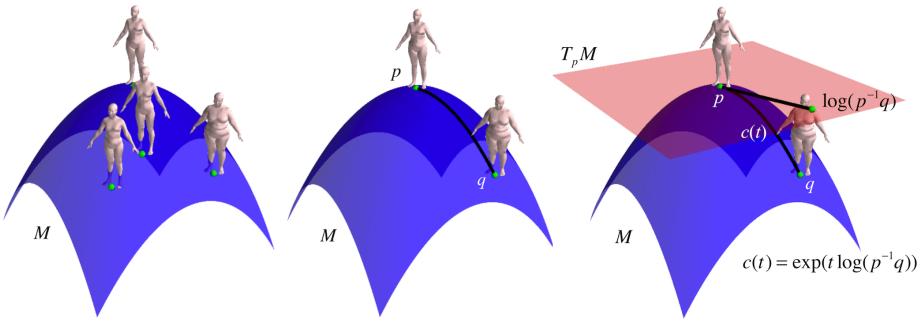

Fig. 1. Body manifold. Left: Human shapes as points on a manifold, M. Every point represents a deformation from a template. Center: Distance between shapes p and q is measured via a geodesic distance; i.e. the length of the path between them along M. Right: The tangent space at p, denoted by T_pM, is a vector space

compression, biometrics, and the fashion industry. Current methods typically use a Euclidean representation of deformations and measure distance in a Euclidean space, ignoring the geometry of the space of deformations. These methods model triangle deformations as elements of $\mathbb{R}^{3\times3}$ while the true deformations are only 6 dimensional. Despite the use of heuristics to remove excessive degrees of freedom (DoF) their deformations might still be noisy or have negative determinant. The latter is physically impossible (e.g. reflections). In contrast, we propose a novel manifold representation of shape deformations (Fig. 1) that eliminates the above problems and has many other benefits, both practical and theoretical. In particular, respecting the underlying geometry enables better statistical learning methods, distance computation, shape interpolation, and the valid composition of multiple causes of shape deformation.

Here we consider surfaces represented as triangulated meshes. While we illustrate our model with human body shapes, the formulation is completely general and applies to any triangular mesh model (*Lie Shapes*). We assume a dataset of registered meshes with the same graph topology of N_t triangles. The *deformable mesh statistical modeling problem* has two parts: (I) Given a template mesh $\mathcal{T}$ and a set of training meshes $\{\mathcal{M}_i\}$, for each $\mathcal{M}_i$, quantify the variation between $\mathcal{T}$ and $\mathcal{M}_i$. (II) Learn a statistical model of these variations. The model should also support sampling (or synthesis) of new shapes.

For effective statistical learning, it is crucial to have an appropriate shape representation. The simplest approach is to work directly with the points in $\{\mathcal{M}_i\}$ or their displacements from $\mathcal{T}$. While this may work well for rigid objects, it is a poor choice for non-rigid and/or articulated objects such as the human body where the deformations are more complex and can result from a composition of multiple causes. The common approach represents shape in terms of deformation matrices acting on the triangles of $\mathcal{T}$ [1, 2, 6, 5, 7, 8]. Modeling transformations, rather than the transformed objects themselves, has a long history [9, 10].

We define an appropriate mathematical representation for mesh deformations in terms of a manifold. This *Lie Body* representation is based on a simple new 6D Lie group of triangle deformations (Fig. 2) that eliminates redundant DoF. The deformations can be computed exactly, in closed-form, without heuristics. The $6N_t$ dimensional manifold of Lie Bodies, denoted by M, has a Riemannian structure inducing a *left-invariant* metric between shapes that can be computed in closed-form. This metric defines distances between body shapes in a principled way using geodesic distances (Fig. 1). It also allows us to compute *statistics on manifolds* using methods such as Principal Geodesic Analysis (PGA) [11, 12]. These better capture the statistics of human body shape deformations than do standard Euclidean methods such as Principal Component Analysis (PCA).

We focus here on human shape and show that our formulation results in a better, more parsimonious, and more accurate model of shape deformation than the traditional Euclidean representation. We evaluate performance in several ways: 1) Our group structure results in meshes that are better behaved and exhibit lower variance across a database of registered body shapes. 2) For a fixed number of low-dimensional shape vectors we find that our model is better able to predict biometric measurements. 3) For a fixed number of shape vectors, our reconstruction of Euclidean shape is even better than the Euclidean model. 4) Finally we show that our representation better captures properties of body shape related to human perception.

In summary, our main contributions are: 1) A novel non-linear manifold representation for deformations of triangular meshes. This manifold has the minimal number of DoF required for arbitrary triangle deformations, provides a heuristic-free way to compute deformation, and eliminates non-physical deformations. 2) This representation is consistent in the sense it has a group structure: deformations can be composed or inverted in a meaningful way. 3) We provide a principled way to measure distances and interpolate between shapes using geodesic distances and geodesic paths. 4) We show how statistics on the manifold capture shape variation over a database of aligned triangular meshes. 5) Closed-form formulas, or efficient approximations, are given for all computations.

2 Related Work

Sumner and Popović [6] define shape in terms of affine deformations of triangles from $\mathcal{T}$. A triangle deformation is represented by a 3×3 deformation matrix and a 3D displacement vector. The 9 dimensional space of deformations is under-constrained as deformations outside the plane of the triangle are undefined. They deal with this heuristically by adding a fourth virtual vertex defined by the cross product of two of the triangle edges. Anguelov et al. [5] use these matrices to define the SCAPE model, which factors deformations due to body shape from those due to pose. Bălan [13] builds a SCAPE model from the CAEASR dataset [14] and regularizes the ambiguity in the 3×3 deformations using a smoothness constraint that penalizes difference in deformations between neighboring triangles. Hasler et al. [2] learn a multilinear model of affine deformations. Like all the

above models, distances between deformations are measured in a Euclidean space with redundant DoF. Hasler et al. [7] use an even higher-dimensional representation of deformations. They model deformations with 15 DoF and a non-linear encoding of triangle deformations that captures dependencies between pose and shape. Again, Euclidean distance still plays a central role.

Note that in fact, triangle deformations lie in a 6 dimensional non-linear manifold. Unlike previous methods, our deformations live on this manifold, have positive determinant by construction, and thus exclude non-physical deformations such as reflections.

Chao et al. [8] represent deformations, using an analogy to elasticity, as a rotation plus an affine residual deformation closest (in a Euclidean sense) to an isometry. Their affine deformations have positive determinant but have more than 6 DoF. Additionally, to compute an approximation of their shape distance and the path between shapes, they require an expensive optimization scheme. In contrast, we provide accurate closed-form formulas for paths and distances. Also, in [8] there is no notion of a Lie group or shape deformation statistics.

Grenander and Miller [15] explore the group of scaled-rotations for open 2D contour deformations. However, in their approach, the group structure is lost when applied to closed contours as the contour-closure imposes a linear constraint. In [16], this restriction is removed but the method is limited to 2D.

The idea of employing Lie groups to model 3D deformations is not new and is widely used in computational anatomy [10–12, 17–19]. However, these methods work on volumetric representations rather than triangulated meshes.

Alexa [20] uses a Lie group for triangle deformations but both his motivation (addressing noncommutativity) and solution (defining a new group operation) are quite different. Essentially, his approach operates in the tangent space at the identity of standard matrix groups while ignoring a non-vanishing Lie bracket; this approximation can be justified only if matrices are close to each other. Additionally, the set of matrices comprising his group is chosen ad-hoc, and so, depending on the case, suffer from either excessive DoF or are not expressive enough to accurately capture arbitrary mesh deformations. Of note, graphics applications described in [20] (or [21]) can benefit from our representation.

Note that our work is not directly related to the classical work of Kendall [22], despite the fact we use a manifold representation for shapes. In [22], shapes are represented by a set of landmarks, and their quotient spaces are studied. In contrast, our manifold represents a group of *transformations* acting on 3D triangular surfaces. For a more recent work on geometric modeling based on Kendall's theory, see [21]. Note that since we provide closed-form formulas for geodesic paths, interpolation (or extrapolation) on our manifold is considerably simpler than the algorithms presented in [21].

Our use of the words manifold and Riemannian metric should not confuse the reader with vast literature on non-linear manifold learning in a Euclidean space. Unlike these approaches, we change the actual representation of the shape deformations and the manifold is then a natural property of Lie Bodies. These approaches, however, can still benefit from our representation as they can be

employed in a Euclidean tangent space to our manifold (to be described later). Another possible source of confusion might arise due to the term geodesic distance. Many mesh registration methods try to align meshes so as to preserve geodesic distances on the surface of the mesh. Here when we refer to a geodesic distance, it is the distance between two *shapes* living on a $6N_t$ dimensional manifold, and not between *points on the 3D surface of a particular shape.*

In summary, we introduce a manifold representation for accurate arbitrary non-rigid deformations using a finite-dimensional Lie group acting directly on the 3D triangular surfaces. Our mathematical representation can be easily utilized in existing methods such as [2, 5], and thus has a wider scope than any such particular method.

3 Mathematical Model

We begin the mathematical definition of Lie Bodies with the basic definitions of triangle deformations that represent how $\mathcal{T}$ is deformed to a new shape. Next, we show how such deformations can be described by 6 DoF and how they form a group that is also a smooth manifold. Finally, we describe the computation of geodesic distances, geodesic paths and statistics on the (Riemannian) manifold.

3.1 Triangle deformation

Let $(v_0, v_1, v_2) \subset \mathbb{R}^3$ be an ordered triplet defining a (non-degenerate) triangle. Without loss of generality, we assume $v_0 = [0, 0, 0]^T$ so we can identify a triangle with its edge matrix: $[v_1 - v_0, v_2 - v_0] = [v_1, v_2] \in \mathbb{R}^{3\times 2}$.

Definition 1. *Let X and Y be a pair of triangles. $Q \in \mathbb{R}^{3\times 3}$ is called a deformation matrix (acting on X by deforming it to Y) if $Y = QX$.*

This defines the standard Euclidean deformations used in previous work. Consider the collection of all triangles $\{X_t\}_{t=1}^{N_t}$ in a template mesh $\mathcal{T}$. Applying deformations to all X_t independently will typically result in a mesh that is disconnected. To create a consistent mesh, we take the standard approach and solve for a valid mesh where the edges are as close as possible to the edges of the deformed triangles in a Least Squares (LS) sense [1, 5].

3.2 Imposing a group structure

If Y, X and Q are as in Definition 1, then Q is not unique, as $Y = QX$ gives only six constraints. Previous approaches have resorted to ad-hoc methods to constrain or regularize the additional DoF [6, 13]. Instead we explicitly work in an appropriate non-linear 6D space. The intuition is that X can be deformed to Y by a combination of isotropic scaling, a particular in-plane-deformation, and a 3D rotation. See Fig. 2 for an illustration. We now make this precise.

Definition 2. GL*(n), the general linear group of degree n, is the set of $n \times n$ real non-singular matrices, together with the operation of matrix multiplication.*

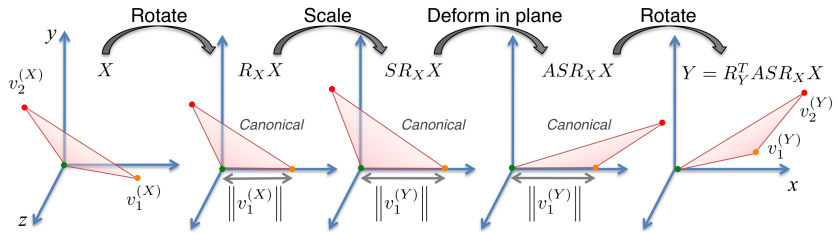

Fig. 2. Triangle deformation. Deforming $X = [v_1^X, v_2^X]$ to $Y = [v_1^Y, v_2^Y]$

GL(n) is a *matrix group*: A group whose elements are square matrices with matrix multiplication as its binary operation. Every matrix group is a *Lie group*.

Definition 3. *A Lie group G is a group that is also a smooth manifold and whose operations, composition ($G \times G \to G$) and inversion ($G \to G$), are smooth.*

See [23] for formal definitions of a smooth manifold and for the smoothness of these operations. Since G is a group it is closed under composition and inversion. Note that standard Euclidean deformations, as defined in Definition 1, are not closed under linear combinations (as used, e.g., in PCA). Moreover, such deformations (let alone their linear combinations) might be singular or have negative determinant. Thus, they do not form a group under either matrix multiplication or addition, and might represent non-physical deformations. We will show that a group structure is important for the statistical analysis of shape deformations and to enable the principled combination of these deformations. We will also show that the Euclidean distance is not suitable for measuring differences between deformations; this too makes statistical analysis with PCA inappropriate.

We argue for a new type of deformation to appropriately model 3D shape in triangulated meshes. These deformations do not suffer from the problems of Euclidean deformations, have a group structure, and give rise to a meaningful distance. The simplest component of our representation is isotropic scaling, which can be defined by a group: Let G_S denote the group of $\mathbb{R}^+$ together with the operation of standard multiplication. A second component models in-plane deformations. Before we define these, we need the following

Definition 4. *A triangle (v_0, v_1, v_2) is said to be canonical (or in a canonical position) if $(v_0, v_1, v_2) = ((0,0,0), (x_1, 0, 0), (x_2, y_2, 0))$, such that $x_1 > 0$, $y_2 > 0$ and $x_2 \in \mathbb{R}$. In effect, v_1 lies on the positive x-axis, and v_2 is in the upper open half of the xy-plane (Fig. 2). Whenever is convenient, we will regard such points as 2D, dropping the third coordinate.*

We now define our in-plane deformations, acting on canonical triangles.

Definition 5. $G_A \triangleq \left\{A \in \mathrm{GL}(2) : A[1,0]^T = [1,0]^T, \det A > 0\right\}$. *Equivalently,*

$$G_A \triangleq \left\{A = \left(\begin{smallmatrix} 1 & U \\ 0 & V \end{smallmatrix}\right) : U \in \mathbb{R}, V > 0\right\} . \tag{1}$$

Proposition 1. *G_A is a subgroup of* GL*(2). (see [24] for proof)*

By a slight abuse of notation, we can also regard G_A as a subgroup of GL(3), using the bijection $\left(\begin{smallmatrix} 1 & U \\ 0 & V \end{smallmatrix}\right) \leftrightarrow \left(\begin{smallmatrix} 1 & U & 0 \\ 0 & V & 0 \\ 0 & 0 & 1 \end{smallmatrix}\right)$. Remark: G_A is not Abelian; if A, B are in G_A, then usually $AB \neq BA$. Geometrically, if $A \in G_A$ and X is canonical, then AX is canonical too. Note that the first edge (column) of X equals the first edge of AX (Fig. 2). Moreover, we have the following

Proposition 2. *If X and Y are two canonical triangles, then there exists a unique $(A, S) \in G_S \times G_A$ such that $Y = ASX$. (see [24] for proof and formulas)*

This is illustrated in Fig. 2. Of course, not all triangles are canonical, hence we have the third and final component: SO(3), the rotation group:

Definition 6. SO*(3), the special orthogonal group of degree 3, is the subgroup of* GL*(3) given by* $\mathrm{SO}(3) = \left\{R \in \mathrm{GL}(3) : R^T R = I, \det(R) = +1\right\}$.

Fact: If X is a triangle, then we can always find a rotation matrix R_X, that depends on X, such that $R_X X$ is canonical. See [24] for details.

Now the entire story in Fig. 2 is complete: If X and Y are two arbitrary triangles, then the fact above and Proposition 2 imply we can always find R_X, R_Y in SO(3) such that $R_X X$ and $R_Y Y$ are canonical, and a unique $(A, S) \in G_A \times G_S$ such that $R_Y Y = ASR_X X$. Equivalently, $Y = R_Y^T ASR_X X$. Setting $R \triangleq R_Y^T R_X \in \mathrm{SO}(3)$, yields our *triangle deformation equation*:

$$\boxed{Y = RR_X^T ASR_X X}\,. \tag{2}$$

Consider X as belonging to the template shape $\mathcal{T}$. Thus, X and R_X are fixed. For varying Y, the triplet (R, A, S) has six DoF: 3 for R, 2 for A and 1 for S. By construction, $\det(RR_X^T ASR_X) = VS > 0$. We have thus found an invertible 3×3 matrix deforming one triangle to another and this matrix has only 6 DoF. Note that unlike the traditional Euclidean approach, with 9 DoF, we do not resort to any heuristic regularization to constrain the excess 3 DoF. Furthermore, the fact that the determinant is strictly positive eliminates deformations that have no physical meaning such as reflections.

There still remains one problem. The set $\left\{RR_X^T ASR_X\right\}$ such that $R \in \mathrm{SO}(3)$, $A \in G_A$ and $S \in G_S$ does not form a matrix subgroup. Fortunately, this can be easily fixed, by defining a new group:

Definition 7. *The triangle deformation group, denoted by G_T, is the set*

$$\{(R, A, S) : R \in \mathrm{SO}(3), A \in G_A, S \in G_S\} \tag{3}$$

together with the composition map

$$G_T \times G_T \to G_T, ((R_1, A_1, S_1), (R_2, A_2, S_3)) \mapsto (R_1 R_2, A_1 A_2, S_1 S_2)\,. \tag{4}$$

G_T is a direct product of SO(3), G_A, and G_S and so it too is a group. G_T is not Abelian and has a natural identification with a matrix group using the bijection:

$$(R, A, S) \leftrightarrow \begin{pmatrix} R & 0_{3\times2} & 0_{3\times1} \\ 0_{2\times3} & A & 0_{2\times1} \\ 0_{1\times3} & 0_{1\times2} & S \end{pmatrix} \in \mathbb{R}^{6\times6}\,. \tag{5}$$

Fig. 3. Smooth Manifold. (left) T_pM is a tangent space to M at p. T_pM and M are connected via exp and log. (right) Interpolation (light blue) along a geodesic path between two human shapes p and q (green)

So far we have discussed the deformation of a single triangle. For a mesh of N_t triangles, we use a direct product to represent the Lie group of triangular mesh deformations: $M \triangleq G_T^{N_t}$, rendering M a Lie group (and thus, a smooth manifold) of dimension $6N_t$. We use the term M-valued to describe elements in M. The structure of M gives us group closure (ensuring a consistent representation) as well as a meaningful measure of distance between shapes. Both distance and addition are crucial for classical statistical analysis in Euclidean spaces. On a manifold, which is also a Lie group, these are replaced by geodesic distance and composition (i.e. the group binary operation). Before we approach the topic of geodesic distances, we must introduce the concept of a tangent space.

3.3 Lie groups and Lie algebras

Standard Euclidean statistical methods such as PCA do not apply to non-linear manifolds. However, for Lie groups, many statistical tools can be applied with minor changes by using the concept of a tangent space, which is Euclidean.

We let T_pM denote the tangent space to M at $p \in M$ (Fig. 3 (left)); e.g., T_IM is the tangent space at the identity. Unlike M, T_pM is a vector space, of the same dimension as M. The connection between M and T_pM is via the exponential and logarithmic maps. For matrix groups, these are the matrix exponential and logarithm. If A is a square matrix, then $\exp(A)$ is defined by $\exp(A) = \sum_{n=0}^{\infty} \frac{A^n}{n!}$, while $\log(A)$ is a matrix B satisfying $\exp(B) = A$. T_pM can be identified with T_IM (being Euclidean spaces of the same dimension), while exp connects T_pM and M by $a \mapsto p\exp(a)$; e.g., for $p, q \in M$, set $a = \log(p^{-1}q)$ to get $q = p\exp(\log(p^{-1}q))$ (see Fig. 3 (left)). For $p = I$, this is $q = \exp(\log(q))$.

Every Lie group has a corresponding Lie algebra whose general definition is quite technical [23]. Luckily for $\mathbb{R}$-valued matrix groups we can use the following

Definition 8. *Let G be an $\mathbb{R}$-valued matrix group. Its Lie algebra $\mathfrak{g}$ is given by the vector space (viewed as the tangent space at I) $\mathfrak{g} = \{A : \exp(A) \in G\} = \exp^{-1}(G)$ and $[\cdot,\cdot] : \mathfrak{g}\times\mathfrak{g} \to \mathfrak{g}$, the Lie bracket of $\mathfrak{g}$, is given by $A, B \mapsto AB - BA$.*

For example, the Lie algebra of GL(n), denoted by $\mathfrak{gl}(n)$, is given by the set of all $\mathbb{R}$-valued $n \times n$ matrices. Let A, B be in $\mathfrak{g}$. In general, $\exp(A)\exp(B) \neq \exp(A+$

B), with equality if and only if $[A, B] = 0$. We now describe $\mathfrak{g}_T \triangleq \exp^{-1}(G_T)$, the Lie algebra of G_T. Naturally, it is defined using the direct product of the following three Lie algebras: $\mathfrak{g}_T = \mathrm{so}(3) \times \mathfrak{g}_A \times \mathfrak{g}_S$, where 1) $\mathfrak{g}_S \triangleq \exp^{-1}(G_S)$ is $\mathbb{R}$. 2) $\mathrm{so}(3) \triangleq \exp^{-1}(\mathrm{SO}(3))$ is well known to coincide with $\left\{A \in \mathfrak{gl}(3) : A = -A^T\right\}$. 3) $\mathfrak{g}_A \triangleq \exp^{-1}(G_A)$ is to be described shortly. The computation of the matrix exponential (or logarithm) involves an infinite sum (unless A is diagonalizable or nilpotent). In several cases, however, it is possible to derive closed-form formulas; e.g., for $\exp : \mathrm{so}(3) \to \mathrm{SO}(3)$ and $\log : \mathrm{SO}(3) \to \mathrm{so}(3)$, computations are given by the well known Rodrigues' formula [25]. The first map is surjective, unlike the second. The pair does, however, form a local bijection around the zero matrix in so(3) and the identity in SO(3). For $\mathfrak{g}_A$ and G_A, we have the following

Proposition 3. *$\mathfrak{g}_A \triangleq \exp^{-1}(G_A)$, is given by: $\mathfrak{g}_A = \{A \in \mathfrak{gl}(2) : g = \left(\begin{smallmatrix} 0 & u \\ 0 & v \end{smallmatrix}\right)\}$. The maps $\exp : \mathfrak{g}_A \to G_A$ and $\log : G_A \to \mathfrak{g}_A$ form a bijection and can be computed in a closed-form. (see [24] for proof and formulas)*

Proposition 3 is important not only because of the bijectivity; it also provides simple and exact closed-form formulas; e.g., generic expm and logm functions in Matlab or SciPy use the Padé approximation [26], are slower, and are not easily vectorized. Since such computations are needed frequently and apply to all N_t triangles, this makes their use impractical here.

The map $\exp : \mathfrak{g}_T \to G_T$ is defined by the product map ($\exp : \mathrm{so}(3) \to \mathrm{SO}(3)$, $\exp : \mathfrak{g}_A \to G_A$, $\exp : \mathfrak{g}_S \to G_S$). Similarly, the Lie algebra of M is given by $\mathfrak{m} \triangleq \mathfrak{g}_T^{N_t}$ and the product map $\exp : \mathfrak{m} \to M$ is given by an N_t-tuple of $\exp : \mathfrak{g}_T \to G_T$ maps. See Fig. 1 for an illustration.

3.4 Statistical analysis, geodesic paths, and geodesic distances

Geodesic distance. We first need a way to define distances between points in M. To this aim, we endow M with a Riemannian structure: For every $p \in M$, we define the same inner product $\langle \cdot, \cdot \rangle : T_pM \times T_pM \to \mathbb{R}$ independently of p, by taking the standard inner product in $\mathbb{R}^{6N_t}$. It can be shown that this induces a *geodesic distance* on M, of the form $d(p, q) = \left\|\log(p^{-1}q)\right\|_F$ for $p, q \in M$, where $\|\cdot\|_F$ is the Frobenius norm (it is a general result for connected matrix groups). The block diagonal structure in Eq. (5) enables easy computation: If $p_i = \{g_{i,j}\}_{j=1}^{N_t} \in M$, $i \in \{1, 2\}$ and $g_{i,j} \triangleq (R_i, G_i, S_i)_j \in G_T$, then

$$d(p_1, p_2)^2 = \sum_{j=1}^{N_t} d_{G_T}(g_{1,j}, g_{2,j})^2 \tag{6}$$

where, for $(g_1, g_2) \in G_T \times G_T$ (dropping the j notation),

$$d_{G_T}(g_1, g_2)^2 \triangleq \left\|\log(R_1^T R_2)\right\|_F^2 + \left\|\log(A_1^{-1} A_2)\right\|_F^2 + \left|\log(S_2/S_1)\right|^2 . \tag{7}$$

One shortcoming of Euclidean deformations and the associated distance is the lack of left-invariance; e.g., if Q_1 and Q_2 are two deformation matrices, one

would like to have $\|Q_1 - I\|_F = \|Q_2 Q_1 - Q_2 I\|_F$ but this rarely holds (consider the case where $Q_1 \in \mathrm{SO}(3)$ and $Q_2 = 2I$). In contrast, note that d is left-invariant: $d(p_3 p_1, p_3 p_2) = d(p_1, p_2)$ for every p_1, p_2, p_3 in M. This makes our representation especially useful for methods that factor deformation to several causes such as shape and pose. Finally, a geodesic path interpolating between $p_1, p_2 \in M$ is given in closed-form by $p(t) = p_1 \exp(t \log(p_1^{-1} p_2))$ where $t \in [0, 1]$, $p(0) = p_1$, $p(1) = p_2$. See Fig. 3 (right). Taking $t < 0$ or $t > 1$ is extrapolation.

Interior mean and PGA. Given a set $\{p_i\} \subset M$, the sample mean is a linear combination of the data, and might not be in M. A typical replacement is the sample Interior Mean [27], defined as $\mu = \arg\min_\mu \sum d(p_i, \mu)^2$, which can be computed by an efficient iterative algorithm [11] that can be applied to each mesh triangle in parallel. In our experiments convergence is typically reached with few iterations. This is the only non-closed-form computation required in this work. However, one can also use a closed-form approximation: $\mu \approx \exp(1/n \sum \log p_i)$.

We also replace standard PCA by PGA [11, 12]. This amounts to PCA at $T_\mu M$. In effect, we first set $g_i = \mu^{-1} p_i$, and then compute regular PCA on $\{\log g_i\}$, to get a K-dimensional subspace of $T_\mu M$. To synthesize from the PGA subspace, we *compose* μ with the exponent of a linear combination of eigenvectors $\{V_k\}_{k=1}^K \subset T_\mu M$: $\mu \exp(\sum_{k=1}^K \alpha_k V_k)$. As M is not Abelian, this is not the same as $\exp(\log \mu + \sum_{k=1}^K \alpha_k V_k)$. Note that PGA, so defined, is sometimes viewed as a linearized PGA; see discussion in [28].

4 Experiments

We model a dataset of 986 body scans of adult women standing in a similar pose [14]. A template mesh with N_t=16218 triangles is aligned to each of the scans and we then work with these aligned meshes. Because here we are interested in body shape, we remove the head and hands for the analysis but they can trivially be included in the model. Given aligned triangles we compute the M-valued deformations (and encode them as $RR_X^T A S R_X$ as in Eq. (2)) and Euclidean deformations (using the method in [13]). We seek a parsimonious statistical representation of body shape variation where parsimony has two components: it implies that the model has low variance and low reconstruction error.

Variance comparison. To fairly compare representations, we measure the variance in the same, *Euclidean*, space - $(\mathbb{R}^{3\times 3})^{N_T}$. While this should give Euclidean deformations an advantage, we find that the total variance for the Euclidean case is 1.68 times the total variance of the new deformations. We attribute this to the fact that our method does not admit non-physical deformations. This suggests that, even if one avoids a manifold approach, one is still better off computing M-valued deformations, encoding them as $RR_X^T A S R_X$, and then working with these in a Euclidean space.

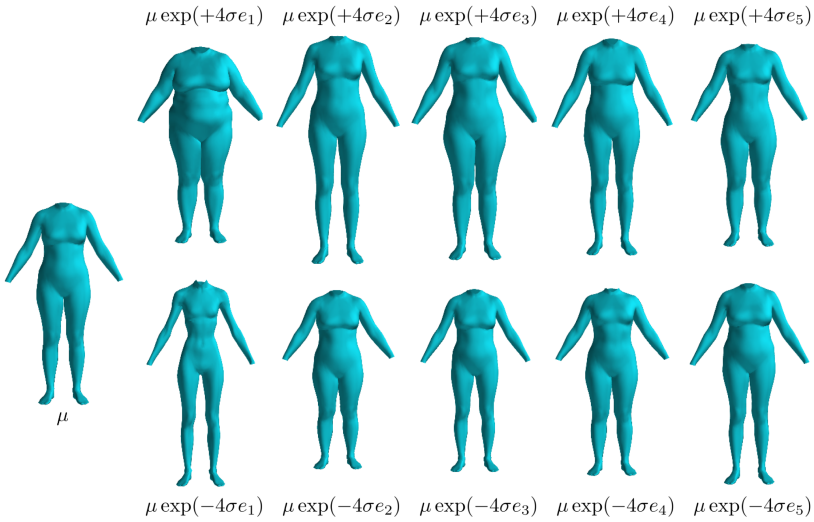

Fig. 4. PGA. The single shape in the leftmost column is μ. The other columns present, from left to right, the first five eigenvectors by showing $\mu \exp(+4\sigma e_i)$ (top) and $\mu \exp(-4\sigma e_i)$ (bottom), where e_i is the i^{th} eigenvector, $i \in \{1, 2, 3, 4, 5\}$

Table 1. Mean edge RMS for mesh reconstruction using a subspace

#PC's	5	10	15	20	25	30	35	40	45	50	100
Euclidean method. Ave. RMS [mm]	2.71	2.53	2.43	2.34	2.28	2.23	2.19	2.15	2.11	2.08	1.91
Our method. Ave. RMS [mm]	2.57	2.43	2.32	2.26	2.21	2.17	2.12	2.09	2.06	2.03	1.88

PCA, PGA, and reconstruction. In addition to low variance, a good representation must model the data. In particular, we seek a low-dimensional approximation of body shape variation that captures as much of the variance as possible. Consequently we evaluate our ability to reconstruct the data meshes using the Euclidean and our Manifold approaches. We compute the interior mean and PGA and show the first few eigenvectors in Fig. 3. We also compute PCA for the Euclidean deformations.

We reconstructed all the meshes using the same number of components computed with PCA and PGA. For each mesh $\mathcal{M}_i$ we compute the root mean squared (RMS) reconstruction error across all edges in M_i and then average the result over all $\{\mathcal{M}_i\}$. The results in Table 1 show a lower error for every number of basis vectors used. Note this is despite the fact that RMS is a Euclidean error.

Human shape perception. We also evaluate the perceived quality of the reconstructions using a two-alternative forced choice perceptual experiment. This evaluates how well each representation captures features of body shape that are important to human perception. Mechanical Turk workers were shown images like those in Fig. 5 drawn from a set of 300. The center shape always showed the original mesh (head and hands were removed to focus people on body shape). On either side, in the same pose, we showed the PGA and PCA

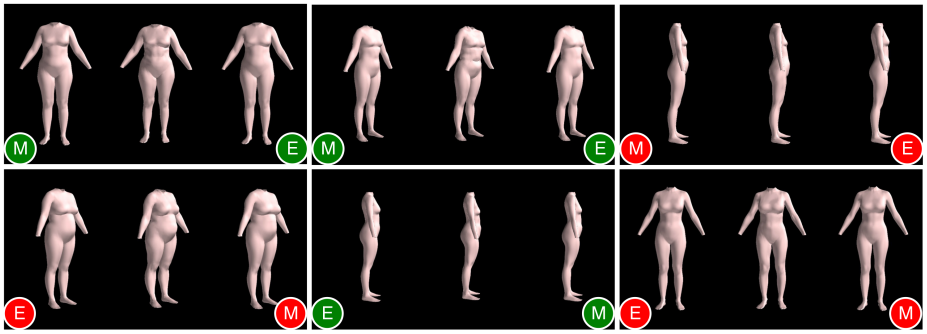

Fig. 5. Perceptual task. Examples images shown in the shape perception experiment (workers did not see the "M" and "E" notations). The true mesh is always in the center. On either side, at random, is either the manifold representation (M) or the Euclidean representation (E) generated with 20 coefficients. Workers had to indicate which body looked more like the center. Green indicates examples that were "easy" (i.e. the manifold was selected as better more than 90% of the time) and red indicates "hard" examples where workers were roughly split in their decisions

reconstructions. The left/right location was randomly varied and the reconstructed bodies were shown in a random order in one of three viewing directions: profile, frontal, and oblique. Each comparison was presented 10 times and each worker could try as many of the 300 examples as they wished. Of the 3000 answers, 7 were disqualified for technical reasons. M-valued reconstruction was preferred in 1670 out of 2993 answers (55.8%). For each of the 300 test images, the portion (usually out of 10) of the workers who preferred the M-valued reconstruction to the Euclidean was computed. A tie was reached in 20 % of the cases (59/300) while in 80%, a majority was achieved. When a majority was achieved, the manifold approach was selected as better significantly more often: $\Pr(M \text{ won}|\text{majority was achieved}) \approx 55/80 = 0.69$ while $\Pr(\text{Euclidean won}|\text{ majority was achieved}) \approx 25/80 = 0.31$.

Predicting biometric measurements. An important quantitative measure of body mesh quality is the accuracy with which body shape can be used to compute anthropometric measurements. We compute a simple linear regression from subspace coefficients to body measurements [4]. We compare three subspaces: PCA on Euclidean matrices, PCA on $RR_X^T ASR_X$ matrices, and PGA on M-valued deformations. The data was split into two halves for training and test sets. Results are summarized in Fig. 6. There is a clear advantage to the manifold approach, especially for a small number of coefficients.

5 Conclusions

With the availability of devices like Microsoft's Kinect, there is increasing interest in modeling the 3D shape of non-rigid and articulated objects – particularly,

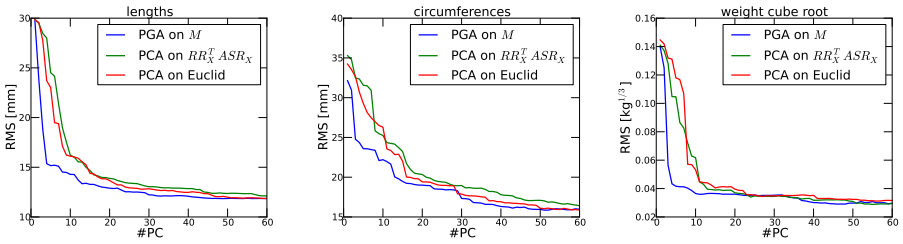

Fig. 6. Linear prediction of body measurements from shape coefficients. RMS error as a function of the number of coefficients. Left: Average results for length measurements ('spine to elbow', 'shoulder breadth', 'stature', 'knee height', 'spine to shoulder', 'arm'). Center: Average results for circumference measurements ('chest', 'thigh', 'ankle', 'under bust'). Right: Cube root of weight. See [24] for additional plots

human body shape [4]. To model object shape variation we need an appropriate representation and we propose a new Lie group for representing shape as a deformation from a template mesh. The approach has many nice properties. Representing triangle deformations with our Lie group, gives an exact solution for 6 DoF triangle deformations, unlike previous Euclidean methods. The distance between shapes is properly defined as a geodesic distance on the non-linear manifold of deformations. Statistics of shape variation are represented using Principal Geodesic Analysis. These benefits come with little additional overhead since the main equations are efficiently computed in closed-form. In addition to theoretical benefits, we have shown that the Lie representation of shape deformation consistently outperforms the Euclidean approach.

In future work we will incorporate our representation into models such as SCAPE [5], which factors body shape variation into separate pose, shape, and pose-dependent deformations. Here our group structure is ideal for composition of deformations. While focusing on human bodies, we emphasize the generality of the approach and envision "Lie Shapes" providing an improved foundation for shape representation, analysis, and synthesis.

This work was supported in part by NIH-NINDS EUREKA (R01-NS066311).

References

1. Balan, A., Sigal, L., Black, M., Davis, J., Haussecker, H.: Detailed human shape and pose from images. CVPR (2007) 1–8
2. Hasler, N., Ackermann, H., Rosenhahn, B., Thormahlen, T., Seidel, H.: Multilinear pose and body shape estimation of dressed subjects from image sets. CVPR (2010) 1823–1830
3. Salzmann, M., Fua, P.: Linear local models for monocular reconstruction of deformable surfaces. PAMI **33** (2011) 931–944
4. Weiss, A., Hirshberg, D., Black, M.: Home 3D body scans from noisy image and range data. ICCV (2011) 1951–1958

5. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: SCAPE: Shape completion and animation of people. ACM ToG **24** (2005) 408–416
6. Sumner, R., Popović, J.: Deformation transfer for triangle meshes. ACM ToG **23** (2004) 399–405
7. Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.: A statistical model of human pose and body shape. Computer Graphics Forum **28** (2009) 337–346
8. Chao, I., Pinkall, U., Sanan, P., Schröder, P.: A simple geometric model for elastic deformations. ACM ToG **29** (2010) 38:1–38:6
9. Grenander, U.: General pattern theory: A mathematical study of regular structures. Clarendon Press (1993)
10. Grenander, U., Miller, M.: Pattern theory: From representation to inference. Oxford University Press, USA (2007)
11. Fletcher, P., Lu, C., Joshi, S.: Statistics of shape via principal geodesic analysis on Lie groups. CVPR **1** (2003) 95–101
12. Fletcher, P., Joshi, S.: Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis (2004) 87–98
13. Balan, A.: Detailed Human Shape and Pose from Images. PhD thesis, Brown Univ., Providence, RI (2010)
14. Robinette, K., Blackwell, S., Daanen, H., Boehmer, M., Fleming, S., Brill, T., Hoeferlin, D., Burnsides, D.: Civilian American and European Surface Anthropometry Resource (CAESAR) final report. AFRL-HE-WP-TR-2002-0169, US AFRL (2002)
15. Grenander, U., Miller, M.: Representations of knowledge in complex systems. J. Royal Stat. Soc. B (Methodological) (1994) 549–603
16. Freifeld, O., Weiss, A., Zuffi, S., Black, M.: Contour people: A parameterized model of 2D articulated human shape. CVPR (2010) 639–646
17. Miller, M., Christensen, G., Amit, Y., Grenander, U.: Mathematical textbook of deformable neuroanatomies. PNAS **90** (1993) 11944–11948
18. Grenander, U., Miller, M.: Computational anatomy: An emerging discipline. Quarterly of Applied Math. **56** (1998) 617–694
19. Fletcher, P., Joshi, S., Lu, C., Pizer, S.: Gaussian distributions on Lie groups and their application to statistical shape analysis. Inf. Proc. in Medical Imaging, Springer (2003) 450–462
20. Alexa, M.: Linear combination of transformations. ACM ToG **21**-3 (2002) 380–387
21. Kilian, M., Mitra, N., Pottmann, H.: Geometric modeling in shape space. ACM ToG **26**-3 (2007) 64:1–64:8
22. Kendall, D.: Shape manifolds, Procrustean metrics, and complex projective spaces. Bull. London Math. Soc. **16** (1984) 81–121
23. Lee, J.: Introduction to smooth manifolds. Graduate Texts in Mathematics **218**. Springer Verlag (2003)
24. Freifeld, O., Black, M.: Lie bodies: Supplemental material. MPI-IS-TR-005, Max Planck Institute for Intelligent Systems (2012)
25. Murray, R., Li, Z., Sastry, S.: A mathematical introduction to robotic manipulation. CRC Press (1994)
26. Higham, N.: The scaling and squaring method for the matrix exponential revisited. SIAM J. Matrix Analysis and Applications **26**-4 (2005) 1179–1193
27. Karcher, H.: Riemannian center of mass and mollifier smoothing. Comm. Pure and Applied Math. **30** (1977) 509–541
28. Sommer, S., Lauze, F., Hauberg, S., Nielsen, M.: Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations. ECCV (2010) 43–56