

Eigen Deformation of 3D Models

Tamal K. Dey

Pawas Ranjan

Yusu Wang

April 6, 2012

Abstract

Recent advances in mesh deformations have been dominated by two techniques: one uses an intermediate structure like a cage which transfers the user intended moves to the mesh, the other lets the user to impart the moves to the mesh directly. The former one lets the user deform the model in real-time and also preserve the shape with sophisticated techniques like Green Coordinates. The direct techniques on the other hand free the user from the burden of creating an appropriate cage though they take more computing time to solve larger non-linear optimizations. It would be ideal to develop a cage-free technique that provides real-time deformation while respecting the local geometry. Using a simple eigen-framework we devise such a technique. Our framework creates an *implicit* skeleton automatically. The user only specifies the motion in a simple and intuitive manner, and our algorithm computes a deformation whose quality is similar to that of the cage-based scheme with Green Coordinates.

1 Introduction


The creation of deformed models from an existing one is a quintessential task in animations and geometric modeling. A user availing such a system would like to have the flexibility in controlling the deformation in real-time while preserving the isometry. In recent years, considerable progress has been made to meet these goals and a number of approaches have been suggested.

Some of the earliest techniques for mesh deformation involved using skeletons [33]. A user typically creates a skeletal shape which is bound to the mesh. The mesh is then deformed by deforming the skeleton and transforming the changes back to the mesh. This approach puts a burden on the user to create an appropriate skeleton and bind the mesh to it. Later work [9, 1, 32] sought to reduce this burden by automatically creating a skeleton. Automatic generation of good skeletons and accurate transformation of deformations from skeleton to mesh remain challenging till today.

Later approaches replaced the skeletons with a sparse cage surrounding the mesh and then controlled the deformation through the movement of the cage. The use of a cage is akin to the concept of control polyhedron that is used for free-form deformations. The authors in [27] introduced the concept of control polyhedron and others refined

it later [18, 22]. It is well recognized that a control polyhedron does not provide sufficient flexibility to deform meshes with complicated topology and geometry [15, 21]. More recent techniques increase this flexibility by introducing sophisticated coordinate functions that bind the cage to the mesh. In general, each vertex of the mesh is associated with weights called *coordinates* for each vertex of the cage. This allows the mesh vertices to be represented as a linear combination of the vertices of the cage. Cage based techniques vary in how the weights of the vertices are computed. Early attempts include extending the notion of barycentric coordinates to polyhedra [16, 23, 31]. More recently, Mean Value Coordinates [10, 11, 19], Harmonic [15] and Green [21] Coordinates have been proposed for the purpose. The authors in [21] pointed out that the Mean Value and Harmonic Coordinates do not necessarily preserve shapes though they provide affine invariant deformations. They overcome this difficulty by providing a real-time deformation tool that preserves shapes. Recently, in [14], Sorkine et al. use biharmonic coordinates to integrate cages and skeletons under a single framework. This allows the user to use different types of techniques simultaneously based on the result desired. Also, in [3], Ben-Chen et al. use a small set of control points on the original mesh to guide the deformation of the cage. Nevertheless, the limitation of creating pseudo structures like cages and skeletons by users still persists.

Creating pseudo structures, especially cages, can be time-consuming and tricky, as Figure 1 illustrates. The cage on the left, for example, fails to envelop the mesh correctly. The cage on the right envelops the entire mesh, but has self-intersections leading to incorrect calculation of coordinates. It falls upon the user to manually move the cage vertices to



(a) cage intersects with mesh (b) self-intersecting cage

Figure 1: Creating correct cages for meshes




Figure 2: Comparing with green coordinates

rectify the cage, which can become time-consuming. A user typically spends more time creating a good cage, than deforming the mesh. The state-of-the-art would be enhanced if one can have a tool that has the capabilities of Green Coordinates but without the need for the cage. Our approach is geared towards that. Figure 2 illustrates the point by showing how our method produces similar quality deformations as Green Coordinates but without any cage.

There are other approaches that impart the deformation directly to the surface mesh and thus eliminate the need for intermediate structures like cages or skeletons; see e.g. [4, 5, 13, 30, 29, 35]. These techniques usually optimize an energy function tied to the deformation and user control to achieve high quality deformations. However, they either require non-linear solvers or multiple iterations of linear solvers to compute new vertex positions, making them slower for large meshes. Also see [6] for a survey on various deformation techniques that use the Laplacian operator to formulate the energy.

Our work. We introduce a novel approach that allows the user to apply the deformation directly to the mesh but without solving any non-linear system and thus improving both time and numerical accuracy. The method uses a skeleton but without explicitly constructing one. It computes the eigenvectors of the Laplace-Beltrami operator to provide a low frequency harmonic functional basis which helps creating an *implicit* skeleton. The skeleton is a high-level abstraction of the shape of the mesh, lacking small features and details. It deforms this skeleton by computing a new

set of eigen coefficients. These coefficients are solutions of a *linear* system which can be computed in real-time. Finally, it adds the details back to the skeleton to get the deformed mesh. We point out that, unlike other skeleton-based approaches, our implicit skeleton is simply the original mesh whose vertex coordinates are derived from a truncated set of eigenvectors.

Our eigen-framework retains the advantages of both the cage-based and cage-less approaches. Figures 2,10,13 illustrate this point. Our deformation software is easy to use and efficient. We are also able to handle both isometric as well as non-isometric deformations like stretching gracefully, as shown in Figure 3. More examples can be found in the video submitted with this paper.

Comparison with previous work on spectral deformation. Recently, Rong et al. [25, 26] proposed a deformation framework using the eigenvectors of the Laplace operator. Although this seems similar to our approach, the reasons why we use eigenvectors are different. In particular, Rong et al. perform as-rigid-as-possible [29] deformations by trying to preserve the Laplace operator. However, they use the eigenvectors to change the problem domain from spatial to spectral, thereby reducing the size of the optimization problem to the number of eigenvectors used. Usually, they need at least 100 eigenvectors, and more eigenvectors make their final deformation look better.

Our method, on the other hand, just needs a functional basis of low frequency harmonic functions in which the meshes are represented. We use low frequency eigenvectors in order to get a smooth fit for a target deformation since our goal is to guarantee a *smoothly varying deformation* rather than *isometry*. Hence, we can guide deformations by using as few as 8 eigenvectors. Furthermore, since we do not try to preserve the Laplace operator, and hence isometry, we can handle stretching better than [25, 26], as Figure 4 illustrates.

Also, [25, 26] use deformation transfer based techniques to recover details, and sometimes produce artifacts in the deformed mesh, as shown in Figure 12. We develop a more sophisticated iterative technique to recover the details with greater accuracy. Finally, in [25, 26], the positions of the constrained vertices need to be changed each time the user wish to change the scope of the deformation, increasing users' burden to specify the intended deformation.

2 Eigen-framework

2.1 Laplace-Beltrami Operator

Consider a smooth, compact surface M isometrically embedded in \mathbb{R}^3 . Given a twice differentiable function $f : M \rightarrow \mathbb{R}$, the Laplace-Beltrami operator Δ_M of f is defined as the divergence of the gradient of f .

The Laplace-Beltrami operator has many useful properties and has been widely used in many geometric processing

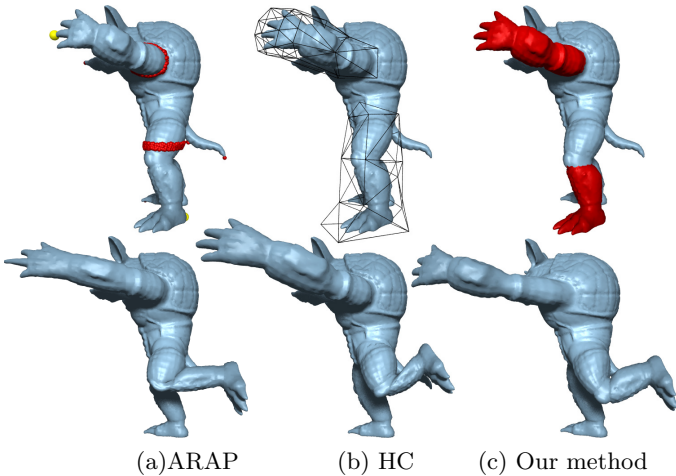


Figure 3: Comparisons when we stretch the arm and bend the leg of the armadillo. Note that our method handles stretching better than as-rigid-as-possible (ARAP), and extreme bending better than harmonic coordinates (HC).

applications; see [20, 34] for surveys on Laplacian mesh processing applications. For example, it is well known that the Laplace operator uniquely decides the intrinsic geometry of the input manifold M . Hence, isometric manifolds share the same Laplacian, which makes the Laplace operator a natural tool to capture or describe isometric deformation. Indeed, this idea has been used to build local coordinates for mesh editing and deformation to help produce as-rigid-as possible type of deformation [28, 35]. The eigenfunctions of the Laplace operator form a natural basis for square integrable functions defined on M . Analogous to Fourier harmonics for functions on a circle, Laplacian eigenfunctions with lower eigenvalues correspond to low frequency modes, while those with higher eigenvalues correspond to high frequency modes that describe the details of the input manifold M .

In our problem, the input is a triangular mesh approximating a hidden surface M . In such case, we need a discrete version of the Laplace operator computed from this mesh. Several choices are available in the literature [2, 7, 12, 23, 24]. In this paper, we use the mesh-Laplacian developed in [2], although other discretizations of the Laplace operator should also be fine. Both the mesh-Laplace operator itself and its eigenvalues have been shown to converge to those of the hidden manifold as the mesh approximates the manifold better [2, 8]. See [6] for a discussion of the effects that the various discretizations have on Laplacian based surface deformation techniques.

2.2 Eigen-skeleton

Given the Laplace operator Δ_M of an input manifold, let ϕ_1, ϕ_2, \dots denote the eigenfunctions of Δ_M . These eigenfunctions form a basis for $\mathcal{L}^2(M)$, the family of square integrable functions on M . Hence we can re-write any function $f \in \mathcal{L}^2(M)$ as $f = \sum_{i=1}^{\infty} \alpha^i \phi_i$, where $\alpha^i = \langle f, \phi_i \rangle$ and $\langle \cdot, \cdot \rangle$ is the inner product in the functional space $\mathcal{L}^2(M)$. Un-




Figure 4: Stretching the arm of the armadillo. Note that spectral mesh deformation (SMD), causes the entire mesh to deform in order to preserve the mesh volume.

der this view, the function f can be considered as a vector $\alpha = [\alpha^1, \alpha^2, \dots]$ in the infinite-dimensional *eigenspace* spanned by the Laplacian eigenfunctions.

Now, consider the *coordinate functions* (f_x, f_y, f_z) defined on M whose values at each point are simply the x, y and z -coordinate values of the point, respectively. By re-writing these coordinate functions, we can represent a surface by three vectors $(\alpha_x, \alpha_y, \alpha_z)$ in its eigenspace. We call these the *coordinate weights* of M . The embedding of a manifold is fully decided by its coordinate weights once the eigenfunctions are given.

Finally, since higher eigenfunctions have higher frequencies and hence capture smaller details, we can truncate the number of eigenfunctions (i.e. use only the top few coordinate weights) for reconstructing the surface to get varying levels of detail.

Eigen-skeleton. Given a surface mesh, also denoted by M , with n vertices, we compute the eigenvectors of the mesh-Laplacian computed from M , denoted still by ϕ_1, \dots, ϕ_n . We now restrict ourselves to the first few, say $m < n$, eigenvectors of the shape. This gives us a higher level abstraction of the surface that captures its coarser features. Specifically, let $P = \{p_1, p_2, \dots, p_m\}$ be the set of points reconstructed from the vertex set $V = \{v_1, v_2, \dots, v_n\}$ of M using only the first m eigenvectors $\phi_1, \phi_2, \dots, \phi_m$ of the mesh-Laplacian. That is, if

$$\hat{f}_x = \sum_{i=1}^m \alpha_x^i \phi_i, \quad \hat{f}_y = \sum_{i=1}^m \alpha_y^i \phi_i, \quad \hat{f}_z = \sum_{i=1}^m \alpha_z^i \phi_i,$$

then $p_i = \{\hat{f}_x(v_i), \hat{f}_y(v_i), \hat{f}_z(v_i)\}$ for $i = 1, \dots, m$. Consider the mesh $K = K_m$ with vertices p_i and the connectivity same as that of M . We call this mesh the *eigen-skeleton*¹ of M . For different values of m , the eigen-skeleton K_m abstracts the input surface M at different levels of detail.

3 Algorithm

Our algorithm will compute the target configuration by deforming the eigen-skeleton. In particular, the eigen-skeleton K_m is decided by the $3m$ coordinate weights

¹The concept of eigen-skeletons is not new and has been used for mesh compression in [17]. For more applications, please refer to the survey papers [20, 34].

$\alpha_x^1, \dots, \alpha_x^m; \alpha_y^1, \dots, \alpha_y^m$ and $\alpha_z^1, \dots, \alpha_z^m$. We will deform the eigen-skeleton by computing a new set of coordinate weights by solving only a linear system. Since the number of eigen-vectors used is typically much smaller than the number of vertices involved in deformation, a solution can be obtained in real-time. We will see later that, other than being efficient, the use of coordinate weights also has the advantage that the deformation tends to be smooth across the entire shape. Since the new eigen-skeleton lacks smaller features, we design a novel and effective algorithm to add back details using the one-to-one correspondence between the vertices of the eigen-skeleton and the original mesh. The high level framework for our algorithm is described in Algorithm 1. Next, we describe each step in detail.

Algorithm 1: Deformation Framework

Input : Input mesh M

Output: Deformed mesh M^*

```

1 begin
2   Compute eigen-skeleton  $K$  for  $M$ 
3   While (user initiates deformation) {
4     Step 1: Interpret user-specified deformation and
5       compute a coarse target configuration  $\tilde{K}$ 
6     Step 2: Obtain  $K^*$ , a smooth approximation to  $\tilde{K}$ 
7     Step 3: Add shape details to obtain  $M^*$ 
8   }
```

3.1 Step 1: Coarse Guess-Target Configuration

Our software uses a standard and simple interface for the user to specify the intended target configuration. First, the user selects a mesh region that he wishes to deform. We call it the *region of interest*, R , and let $V_R \subseteq V$ denote the set of vertices in this region. Next, the user specifies the type of transformation desired for the region of interest, which can be either a **translation-type** or a **rotation-type**. The user then indicates the target configuration by simply dragging some point, say $\mathbf{v} \in V_R$, to its target position $\tilde{\mathbf{v}}$.

From the type of transformation combined with the position of \mathbf{v} and $\tilde{\mathbf{v}}$, our algorithm computes either a translational vector $\mathbf{t} = \tilde{\mathbf{v}} - \mathbf{v}$ if the desired transformation is a translation-type, or a rotational pivot \mathbf{p} and a rotation matrix \mathbf{r} if the desired transformation is a rotation-type. We then compute a coarse target configuration \tilde{K} for the eigen-skeleton K using the following simple procedure: For all points $v_i \notin V_R$, the target position for the corresponding point p_i in the eigen-skeleton is simply $\tilde{p}_i = p_i$. For each point $v_i \in V_R$, if the type of transformation is translation, then the target position is $\tilde{p}_i = p_i + \mathbf{t}$. If the type of transformation is rotation, then the target position is $\tilde{p}_i = \mathbf{r}(p_i - \mathbf{p}) + \mathbf{p}$.

In other words, we simply cut the region of interest and apply to it the target transformation indicated by the user,





Figure 5: The dragon model with its eigen-skeleton created using 8 eigen-vectors

while the rest of the shape remains intact. Such an initial guess of target configuration is of course rather unsatisfactory. In fact, the deformation is not even continuous (along the boundary of the region of influence R , there is a dramatic, non-continuous change in the deformation). However, we will see later that in Step 2, our algorithm takes this initial target configuration and produces a much better, smoothly bent eigen-skeleton. See Figures 5 and 6 for an example: in order to bend the body of the dragon, we specify a rotation on the back half of the dragon. We then apply this rotation on the entire region of interest in the eigen-skeleton to obtain a target configuration \tilde{K} , as shown in the left image in Figure 6. Note that Step 2 will produce a nice, global deformation for the eigen-skeleton, as shown on the right in Figure 6.

Observe that the translation-type and rotation-type motions are only high-level guidance for producing the final deformation in Step 2. The final deformation is of course not necessarily rigid. A stretching effect, for example, can be achieved by a simple translation-type motion. From the user's point of view, the amount of work to specify the deformation is very little and rather intuitive, while the algorithm reconstructs a more complex deformation from the user's coarse input.

3.2 Step 2: Eigen-skeleton Deformation

After Step 1, we have a guess-target configuration \tilde{K} for the eigen-skeleton K . In this step, we wish to compute an improved target deformed eigen-skeleton K^* from the guess-target configuration \tilde{K} . In Step 3 described in next section, we will add details back to K^* to obtain a deformed surface M^* for the input surface M . The following diagram illustrates successive structures.



Recall that \tilde{p}_i is the position of the i th vertex v_i in the guess-target skeleton \tilde{K} . Now consider the coordinate func-

tions $(\tilde{f}_x, \tilde{f}_y, \tilde{f}_z)$ of the guess-target skeleton \tilde{K} . Note, each function \tilde{f}_a , where $a \in \{x, y, z\}$, is a function $M \rightarrow \mathbb{R}$ on the input surface M . The guess configuration \tilde{K} is often far from being satisfactory. In particular, by cutting the region of influence and simply translating and rotating this part, a discontinuity exists at the boundary of region of influence. In other words, there is no smooth transition across the cut. See the enlarged picture in Figure 6 left image. This means that the coordinate functions \tilde{f}_a are not smooth across the cut. To get a smooth deformed skeleton, we wish to find a smooth approximation f_a^* for each \tilde{f}_a . This will give rise to an improved deformed skeleton K^* with the i th vertex $p_i^* = (f_x^*[i], f_y^*[i], f_z^*[i])$.

To this end, note that since the eigenfunctions ϕ_j s of M form a basis for the family of square-integrable functions on M , each \tilde{f}_a can be written as a linear combination of all eigenfunctions ϕ_i s for $i = 1, \dots, n$. Furthermore, eigenfunctions with low eigenvalues are analogous to modes with low-frequency while those with high eigenvalues correspond to high-frequency modes. Since we aim to obtain a smooth reconstruction of \tilde{f}_a , we want to ignore high frequency modes. Hence we find a smooth reconstruction of \tilde{f}_a using only the top m low-frequency eigenfunctions ϕ_1, \dots, ϕ_m of M . This is achieved as follows: Suppose $f_a^* = \sum_{j=1}^m \tilde{\alpha}_a^j \phi_j$, and $A_j = (\tilde{\alpha}_x^j, \tilde{\alpha}_y^j, \tilde{\alpha}_z^j)$ for $j = 1, \dots, m$. We want to find weights $(\tilde{\alpha}_x, \tilde{\alpha}_y, \tilde{\alpha}_z)$ that minimize the following energy function where $\phi_j[i]$ is the value of the j -th eigenfunction ϕ_j on the vertex v_i :

$$E = \sum_{i=1}^n \left\| \sum_{j=1}^m A_j \phi_j[i] - \tilde{p}_i \right\|^2. \quad (1)$$

Intuitively, the discontinuity in the coordinates of guess-target configuration \tilde{K} requires high frequency eigenfunctions to reconstruct it, and using only low-frequency modes produces smoother f_a^* s, which induces better deformed skeleton K^* . See the right picture of Figure 6 — the skeleton reconstructed from new coordinate weights $(\tilde{\alpha}_x, \tilde{\alpha}_y, \tilde{\alpha}_z)$ after Step 2 shows a smooth transition from the region of interest to the rest.

An alternative interpretation. Before we describe how we minimize the above energy function, we provide an alternative interpretation for the formulation of our energy function. Recall after Step 1, we have a guess-target configuration \tilde{K} for the eigen-skeleton K , and \tilde{p}_i is the position of vertex v_i in this skeleton \tilde{K} . Intuitively, if \tilde{K} turns out to be the eigen-skeleton of an isometric deformation \tilde{M} of M , then there exist new coordinate weights $(\tilde{\alpha}_x, \tilde{\alpha}_y, \tilde{\alpha}_z)$, such that

$$\left\| \sum_{j=1}^m A_j \phi_j[i] - \tilde{p}_i \right\|^2 = 0 \quad \forall v_i \in V$$

where $A_j = (\tilde{\alpha}_x^j, \tilde{\alpha}_y^j, \tilde{\alpha}_z^j)$ represents the j -th entry of each $\tilde{\alpha}_a$. This is true because the manifold \tilde{M} has the same eigenfunctions as M , and its corresponding coordinate functions can be written as a linear combination of the eigenfunctions of




Figure 6: Left picture: A coarse discontinuous initial guess. Rotating the entire region of interest (colored red) causes the discontinuity at its boundary. We use the mesh connectivity information from the original mesh to further emphasize this point. Right picture: After step 2, we obtain a smooth transition across the boundary.

\tilde{M} (i.e. ϕ_1, \dots, ϕ_m). The new coordinate weights $\tilde{\alpha}_a$ are simply the first m coefficients for ϕ_1, \dots, ϕ_m in this linear combination.

If the deformation is not isometric, then we can try to find the best fit $(\tilde{\alpha}_x^j, \tilde{\alpha}_y^j, \tilde{\alpha}_z^j)$ for $j \in [1, m]$ by minimizing the above quantity over all vertices in V , that is, minimizing the energy function as defined in Eqn (1). Experimental results show that our method tends to preserve isometry in practice when such a deformation is possible; see for example Figure 10 and Table 2. At the same time, since we do not try to preserve the Laplace operator, we can handle non-isometric deformations like stretching in a more natural manner, compared to [29, 26, 25]; see for example Figure 4.

Minimizing the Energy function E . There are $3m$ variables in the energy function in Eqn (1). To minimize E , we compute its gradient with respect to A_k

$$\begin{aligned} \frac{\partial E}{\partial A_k} &= 2 \sum_{i=1}^n \phi_k[i] \left(\sum_{j=1}^m A_j \phi_j[i] - \tilde{p}_i \right) \\ &= 2 \left(\sum_{j=1}^m A_j \langle \phi_k \cdot \phi_j \rangle - (\langle \phi_k \cdot \tilde{f}_x \rangle, \langle \phi_k \cdot \tilde{f}_y \rangle, \langle \phi_k \cdot \tilde{f}_z \rangle) \right) \end{aligned}$$

where $(\tilde{f}_x, \tilde{f}_y, \tilde{f}_z)$ are the coordinate functions of the guess-target skeleton. Now, setting the partial derivatives to zero for all A_k , we get

$$\sum_{j=1}^m \langle \phi_k \cdot \phi_j \rangle A_j = (\langle \phi_k \cdot \tilde{f}_x \rangle, \langle \phi_k \cdot \tilde{f}_y \rangle, \langle \phi_k \cdot \tilde{f}_z \rangle)$$

which leads to a linear system of equations in the following form: $\Phi A^* = b$, where Φ is an m by m matrix with $\Phi_{i,j} = \langle \phi_i \cdot \phi_j \rangle^2$, A^* is an m by 3 matrix with $A_{i,\cdot}^* = A_i$ and b is also an m by 3 matrix with the i th row as $(\langle \phi_i \cdot \tilde{f}_x \rangle, \langle \phi_i \cdot \tilde{f}_y \rangle, \langle \phi_i \cdot \tilde{f}_z \rangle)$. Using A^* as coordinate weights, we reconstruct the new deformed eigen-skeleton K^* .

²In the ideal case, the Laplace-Beltrami operator is symmetric, which makes its eigenvectors orthonormal, and Φ the identity matrix. However, we use area weights when building the Laplace-Beltrami operator, which makes it asymmetric, and Φ a matrix with non-zero off diagonal entries.

3.3 Step 3: Shape Recovery

We now have the deformed eigen-skeleton K^* . Since we use only the top few eigenvectors for deformation, this skeleton lacks small features and fine details of the original mesh. To obtain the deformed mesh M^* , we need to add appropriate details back to K^* .

In order to keep track of all the shape details, when creating the original eigen-skeleton K , we also keep track of the difference between v_i and its reconstruction p_i . We call it the *detail vector* which is given by $d_{v_i} = v_i - p_i$. However, since the mesh is deforming, we cannot simply add d_{v_i} back to \tilde{p}_i .

To address this issue, we keep track of d_{v_i} in a *local coordinate frame* around p_i . In particular, for each p_i , we compute three axes that are given by: (i) the normal at p_i , (ii) projection of an edge incident at p_i onto a tangent plane at p_i , and (iii) a third vector orthogonal to the previous two. This frame remains consistent with the local orientation of the vertex. For each detail vector d_{v_i} , we record its coordinates in this local frame, which is the projection of d_{v_i} onto the three axes. After the eigen-skeleton is deformed to a new configuration K^* , we compute the new frames, and reconstruct \tilde{d}_{v_i} using the same coordinates but in the new frame. We then obtain the deformed location \tilde{v}_i for vertex v_i by adding the new detail vector back to the skeleton point \tilde{p}_i ; that is, $\tilde{v}_i = \tilde{p}_i + \tilde{d}_{v_i}$.

A drawback of this local-frame based scheme is that the resulting eigen-skeleton becomes very thin near the extremities, with a lot of small features collapsing together, when very few eigenvectors are used. This leads to poor normal estimation around sharp features. See the dragon foot in Figure 7(a). To overcome this difficulty, we compute two sets of new detail vector \tilde{d}_{v_i} : one is obtained by using the local-frames as described above, denoted by $\tilde{d}_{v_i}^{(1)}$; and the other, denoted by $\tilde{d}_{v_i}^{(2)}$, is obtained by simply applying the target deformation transformation computed in Step 1 to the original d_{v_i} . The advantage of $\tilde{d}_{v_i}^{(2)}$ is that it tends to preserve local details. However, just using $\tilde{d}_{v_i}^{(2)}$ alone has the problem that the changes around boundary of R are often dramatic. See the body of the dragon in Figure 7(b).

To get the best out of both strategies, we obtain a final detail vector \tilde{d}_{v_i} by interpolating between the two detail vectors $\tilde{d}_{v_i}^{(1)}$ and $\tilde{d}_{v_i}^{(2)}$. In particular, we assign a large weight to the local-frame based detail vector (i.e., $\tilde{d}_{v_i}^{(1)}$) near the boundary of R and diminish away from the boundary both inside and outside R . We do so because we observed that the normal estimation (and hence $\tilde{d}_{v_i}^{(1)}$) is reliable away from the extremities and near the boundary of the region of interest and that $\tilde{d}_{v_i}^{(2)}$ provides accurate recovery of the sharp features. This works for most models commonly used in the real world, although it is theoretically possible for a mesh and its corresponding skeleton to be very thin even in regions away from the extremities. The details of this interpolation scheme are described in Section 4.1. Figure 7 shows results for recovering the details of the deformed dragon using each

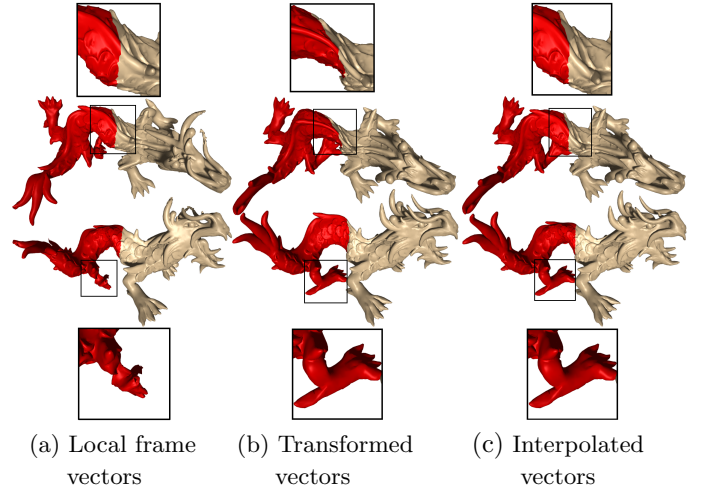


Figure 7: Adding details back to the dragon

individual strategies (a, b) and using the integrated strategy (c).

4 Implementation

4.1 Recovery details

For interpolating the detail vectors, we need to assign a weight to each vertex which should depend on how far it is from the boundary of the region of interest. To do this, we need to (1) identify the boundary ∂R of the region of interest R ; (2) compute a per-vertex function denoting the distance from the boundary ∂R ; and (3) use this function to assign the interpolation weights.

The boundary vertices are identified by considering all vertices in R and simply choosing the ones whose one-ring neighborhood contains vertices that are not in R . We precompute the one-ring neighborhoods on the original mesh just once to reduce computation time during actual deformation of the mesh.

Next, we first compute the following function for each vertex v_i : $g(v_i) = \min_{v \in \partial R} d_g(v_i, v)$, where ∂R is the set of boundary vertices of the region of interest R , and $d_g(v_i, v)$ denotes the geodesic distance between two vertices. Again, we precompute the all-pair geodesic distance matrix once for the original mesh and use it subsequently for all deformations.

Once we have g , we find the approximate diameter of R as $\text{diam}R = \max_{v \in R} g(v)$. We use the diameter to compute two cutoff values $\delta_1 = \frac{\text{diam}R}{8}$, and $\delta_2 = \frac{\text{diam}R}{4}$. The interpolation weights are then computed as:

$$w(v_i) = \begin{cases} 1 & \text{if } g(v_i) < \delta_1 \\ 0 & \text{if } g(v_i) > \delta_2 \\ \frac{\delta_2 - g(v_i)}{\delta_2 - \delta_1} & \text{otherwise} \end{cases}$$

The final detail vector at each vertex v_i is then

$$\tilde{d}_{v_i} = \left(w(v_i) \cdot \frac{\tilde{d}_{v_i}^{(1)}}{\|\tilde{d}_{v_i}^{(1)}\|} + (1 - w(v_i)) \cdot \frac{\tilde{d}_{v_i}^{(2)}}{\|\tilde{d}_{v_i}^{(2)}\|} \right) \cdot \|\tilde{d}_{v_i}^{(1)}\|$$

Note that the two detail vectors $\tilde{d}_{v_i}^{(1)}$ and $\tilde{d}_{v_i}^{(2)}$ have the same length. The above formula simply interpolates their directions to obtain \tilde{d}_{v_i} . Intuitively, the closer a point is to the boundary ∂R of the region of interest, the larger role the local-frame detail vector $\tilde{d}_{v_i}^{(1)}$ plays to guarantee smooth transition. When a point is far from ∂R , the skeleton tends to be much thinner, and in this case we rely more on the transformation-based detail vector $\tilde{d}_{v_i}^{(2)}$ to reconstruct \tilde{d}_{v_i} .

4.2 Choice of number of eigenvectors

The eigenvectors capture details at different scales. Consequently, the use of different number of eigenvectors for deformation causes changes at different scales.




Figure 8: Far Left: head of camel. Right: Eigen-skeleton of the head of the camel constructed using 8, 50 and 300 eigenvectors, respectively.

In general, the eigen-skeleton created with only the top few eigenvectors causes shape changes at a global level. To capture local changes, we need a larger number of eigenvectors. Specifically, if the region of interest R is small, then we need more eigenvectors to build the skeleton so that R is reconstructed reasonably well in this skeleton and the change of the corresponding coordinate weights are sufficient to deform R . See Figure 8 where if we choose too few eigenvectors, the eigen-skeleton of the ear collapses into roughly a point, and cannot represent the ear at all. Since deformation is computed for the eigen-skeleton, the deformation of the ear cannot be described by such a skeleton. Using more eigenvectors we can capture the ear in the skeleton and further deform it.

On the other hand, if the region of interest is large, the change usually needs to be spread over a large area. If we now choose too many eigenvectors, minimizing the energy function in Step-2 tries to preserve local details of the eigen-skeleton (as there are more terms, i.e. A_j s with large j , describing them). Roughly speaking, the optimization of the weights of the lower eigenvectors is overwhelmed by the large number of higher eigenvectors. Hence, the deformation of the eigen-skeleton returned in Step 2 tends to have some dramatic changes for a few points while trying to preserve local details elsewhere. Therefore in the case of a large region of interest, we need to choose a small number of eigenvectors to build the eigen-skeleton so that the weight for global deformation is emphasized.




Figure 9: Adding details back to the dragon; Left: Directly from eigen-skeleton, Right: After iterative improvement

In summary, the number of eigenvectors n_{ev} to be used to reconstruct the eigen-skeleton should be chosen based on the size of R , the region of interest. At the same time, it turns out that the deformation returned by our algorithm is rather robust with respect to n_{ev} , as long as n_{ev} is within a reasonable range. We thus use the following simple strategy to decide n_{ev} . First, compute $\delta_3 = \frac{\text{diam}R}{\text{diam}(V \setminus R)}$, where $\text{diam}(V \setminus R) = \max_{v \notin R} g(v)$ is the approximate diameter of the complement of the region of interest. Now choose n_{ev} as:

$$n_{ev} = \begin{cases} 8 & \text{if } \delta_3 \geq 0.75 \\ 50 & \text{if } 0.75 > \delta_3 \geq 0.1 \\ 300 & \text{otherwise} \end{cases}$$

This simple strategy works well for all the models we experimented with. However, the user can easily override these defaults to choose their own value for n_{ev} .

4.3 Additional modifications

Finally, we observe that since the eigen-skeleton can be rather coarse when n_{ev} is small (8 or 50), the local frame estimation sometimes simply becomes too error-prone on $K_{n_{ev}}^*$ to recover a smooth shape through the interpolation strategy.

For this reason, we iteratively improve the quality of the eigen-skeleton based on the algorithm introduced in Section 3 as follows: Recall that K_m^* denotes the deformed eigen-skeleton reconstructed using m eigenvectors. Instead of directly recovering the deformed mesh M^* from K_m^* , we first recover another intermediate eigen-skeleton $\tilde{K}_{n'}$ from $K_{n_{ev}}^*$, with $n' > n_{ev}$ using Algorithm 1. This is achieved by using the detail vectors to record the change from $\tilde{K}_{n_{ev}}$ to $\tilde{K}_{n'}$, instead of $\tilde{K}_{n_{ev}}$ to the original mesh M . In particular, in our software, $n_{ev} = 8$ or 50 , and $n' = 300$ (this iterative approach is not needed if $n_{ev} = 300$). The result is an eigen-skeleton that already captures the main deformation, and that also contains sufficient details.

Next, we feed $\tilde{K}_{n'}$ as the coarse-guess configuration to the linear solver in Step 2 to obtain a new deformed eigen-skeleton $K_{n'}^*$. This is done to smooth out any errors that may have been introduced due to poor local frame estimation on $K_{n_{ev}}^*$. We then use this new eigen-skeleton $K_{n'}^*$ and local-frame based detail estimation (instead of the interpolation method) to recover the shape-detail of the

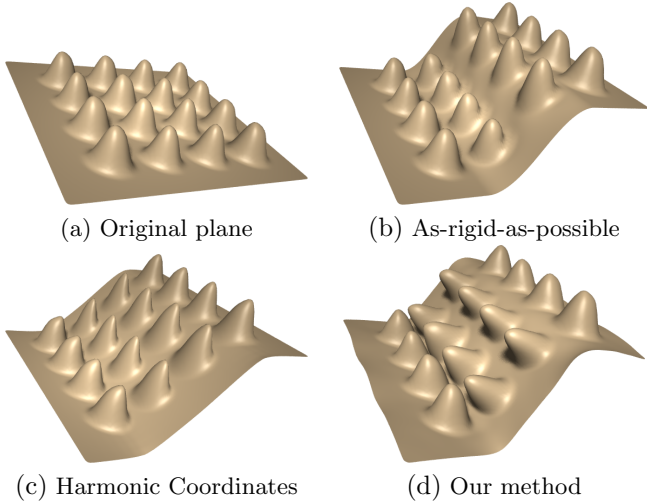


Figure 10: Bending a bumpy plane (dense mesh)




Figure 11: Bending a bumpy plane (coarse mesh)

deformed mesh M^* . See Figure 9 for an example. The final deformation algorithm for the case that $n_{ev} = 8$ or 50 is summarized in the following diagram.

Iteration 1:

$$K_{n_{ev}} \xrightarrow{\text{Step 1}} \tilde{K}_{n_{ev}} \xrightarrow{\text{Step 2}} K_{n_{ev}}^* \xrightarrow[\text{Interpolation based}]{\text{Step 3}} \tilde{K}_{n'}$$

Iteration 2:

$$\tilde{K}_{n'} \xrightarrow{\text{Step 2}} K_{n'}^* \xrightarrow[\text{local-frame based}]{\text{Step 3}} M^*$$

For the case where $n_{ev} = 300$, the original algorithm 1 is applied as before. We remark that potentially one can perform more iterations to improve the deformation quality. However, we observe in practice that two iterations provide a good trade-off between quality and simplicity / efficiency.

4.4 Interactivity

To make the software interactive, we precompute the eigenvectors for the mesh along with the matrix Φ since it depends on the original mesh only. Notice that Φ is symmetric and hence can be factored using Cholesky decomposition. We also precompute the all-pairs geodesic distance matrix used for interpolating detail vectors. To maintain interactive rates, we only deform the eigen-skeleton. Once the user is




Figure 12: Moving the arm of Neptune using our method and spectral mesh deformation (SMD)




Figure 13: Twisting a bar using our method

satisfied with the shape of the eigen-skeleton, the details are added. When deforming the eigen-skeleton, the right-hand side (b) for our linear-solver can be quickly computed by multiplying the matrix of eigenvectors with a matrix containing the coarse guess. We can then compute the new coordinate weights by performing simple backward and forward substitutions. The entire process is simple and can be computed in real-time. Adding details can be a little slow (see Table 1) since we need to compute the normal for each vertex and hence is separated from the interactive part.

5 Results

We implemented our deformation algorithm using C, OpenGL and MATLAB. For comparisons, we wrote our own code for as-rigid-as-possible deformations [29] and used the implementation of cage-based deformation using harmonic coordinates provided in open-source software called BLENDER. For spectral surface deformation, we used the code provided by the authors.

Figure 3 compares our method with harmonic coordinates and as-rigid-as-possible deformations. For as-rigid-as-possible deformation, red dots denote the fixed vertices, while yellow dots represent the vertices that are moved. The partial cages used for deforming using harmonic coordinates are depicted using black edges. For our method, the red portions are the regions of interest. Figures 10 and 11 show the results of bending a plane with smooth bumps using different techniques. Harmonic coordinates are not able to orient the details correctly while for as-rigid-as-possible, the quality of the deformation seems to depend on mesh density.

The timing data for different stages of our algorithm are presented in Table 1. Timings of Step 1 and 2 are coupled together since they are used in each step of interactive de-



Figure 14: Editing the dancing children

formation. Step 3 is used after the user is satisfied with the shape of the skeleton. Table 2 compares the root mean square error in edge lengths. Our method introduces very little error in edge lengths, similar to as-rigid-as-possible approach which aims to optimize such error. Figure 13 shows the result of twisting a bar using our method, while Figure 14 shows that we can handle meshes of arbitrary genus. More deformations using our method can be found in the video submitted with this paper.

We also present comparisons with spectral mesh deformation in Figures 4, 12. The results of spectral mesh deformation are global and cannot be constrained to small regions. For example, in Figure 4, even though only the arm was stretched, the entire mesh got deformed in an attempt to preserve the volume and the Laplace operator of the mesh. Our method is able to handle such deformations more naturally. Also, the detail recovery method used in spectral mesh deformation can introduce artifacts into the deformed mesh. For example, in Figure 12, although only the arm was moved, the staff of Neptune got slightly deformed as well. Even the hand looks unnatural after the deformation. This happens because 100 eigenvectors are not enough to capture the finer details of the model.

Acknowledgments

We would like to thank the anonymous reviewers for their comments, and also the authors of [26] for providing the software implementation of their work. Meshes used in this paper were obtained from AIM@Shape Shape Repository. This work is supported by the National Science Foundation

Model (# vertices)	n_{ev}	Step1 & 2	Step 3
Armadillo (25k)	50	0.018	0.125
dragon (22.5k)	8	0.013	0.161
camel (7k)	8	0.002	0.049
	300	0.012	0.013
plane (10k)	50	0.016	0.061
bar (13.5k)	50	0.017	0.074
children (20k)	50	0.017	0.095

Table 1: Timing data (in seconds) for our algorithm

Model	ARAP	ED
Armadillo (Stretch Arm)	0.0023	0.0022
Armadillo (Bend Knee)	0.001	0.0007
Armadillo (Combined)	0.0052	0.0021
plane	0.0028	0.0022

Table 2: Comparison of relative RMS errors in deformations using as-rigid-as-possible (ARAP) and our method (ED)

grants CCF-0830467 and CCF-0747082.

References

- [1] I. Baran and J. Popović. Automatic rigging and animation of 3D characters. In *Proc. SIGGRAPH '07*, pages 72:1–72:8.
- [2] M. Belkin, J. Sun, and Y. Wang. Discrete laplace operator on meshed surfaces. In *Proc. SCG '08*, pages 278–287.
- [3] M. Ben-Chen, O. Weber, and C. Gotsman. Variational harmonic maps for space deformation. In *Proc. SIGGRAPH '09*, pages 34:1–34:11.
- [4] M. Botsch and L. Kobbelt. Real-time shape editing using radial basis functions. *Comput. Graph. Forum*, 24(3):611–621, 2005.
- [5] M. Botsch, M. Pauly, M. Gross, and L. Kobbelt. Primo: coupled prisms for intuitive surface modeling. In *Proc. SGP '06*, pages 11–20.
- [6] M. Botsch and O. Sorkine. On linear variational surface deformation methods. *IEEE Trans. on Visualization and Comput. Graph.*, 14(1):213–230, Jan. 2008.
- [7] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. SIGGRAPH '99*, pages 317–324.
- [8] T. K. Dey, P. Ranjan, and Y. Wang. Convergence, stability, and discrete approximation of Laplace spectra. In *Proc. SODA '10*, pages 650–663.
- [9] H. Du and H. Qin. Medial axis extraction and shape manipulation of solid objects using parabolic PDEs. In *Proc. ACM Sympos. Solid Modeling Appl. '04*, pages 25–35.
- [10] M. S. Floater. Mean value coordinates. *Comput. Aided Design*, 20(1):19–27, 2003.
- [11] M. S. Floater, G. Kos, and M. Reimers. Mean value coordinates in 3D. *Comput. Aided Design*, 22(7):623–631, 2005.
- [12] K. Hildebrandt and K. Polthier. On approximation of the laplacebeltrami operator and the willmore energy of surfaces. In *Proc. SGP '11*, pages 1513–1520.
- [13] K. Hildebrandt, C. Schulz, C. V. Tycowicz, and K. Polthier. Interactive surface modeling using modal analysis. *ACM Trans. Graph.*, 30(5):119:1–119:11, Oct 2011.
- [14] A. Jacobson, I. Baran, J. Popović, and O. Sorkine.

- Bounded biharmonic weights for real-time deformation. *Proc. SIGGRAPH '11*, pages 78:1–78:8.
- [15] P. Joshi, M. Meyer, T. DeRose, B. Green, and T. Sanocki. Harmonic coordinates for character articulation. In *Proc. SIGGRAPH '07*, pages 71:1–71:10.
- [16] T. Ju, S. Schaefer, J. Warren, and M. Desbrun. A geometric construction of coordinates for convex polyhedra using polar duals. In *Proc. SGP '05*, pages 181–186.
- [17] Z. Karni and C. Gotsman. Spectral compression of mesh geometry. In *Proc. SIGGRAPH '00*, pages 279–286.
- [18] K. G. Kobayashi and K. Ootsubo. T-ffd:free-form deformation by using triangular mesh. In *Proc. Sympos. Solid Modeling Appl. '03*, pages 226–234.
- [19] T. Langer, A. Belyaev, and H.-P. Seidel. Spherical barycentric coordinates. In *Proc. SGP '06*, pages 81–88.
- [20] B. Levy. Laplace-beltrami eigenfunctions: Towards an algorithm that understands geometry. In *IEEE Internat. Conf. on Shape Modeling Appl. '06, Invited Talk*.
- [21] Y. Lipman, D. Levin, and D. Cohen-Or. Green coordinates. In *Proc. SIGGRAPH '08*, pages 78:1–78:10.
- [22] R. MacCracken and K. I. Joy. Free-form deformations with lattices of arbitrary topology. In *Proc. SIGGRAPH '96*, pages 181–188.
- [23] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Math.*, 2(1):15–36, 1993.
- [24] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-beltrami spectra as "shape-dna" of surfaces and solids. *Comput.-Aided Design*, 38(4):342–366, 2006.
- [25] G. Rong, Y. Cao, and X. Guo. Spectral surface deformation with dual mesh. In *Proc. Internat. Conf. on Comput. Animation and Social Agents '08*, pages 17–24.
- [26] G. Rong, Y. Cao, and X. Guo. Spectral mesh deformation. *The Visual Comput.*, 24(7-9):787–796, 2008.
- [27] T. W. Sederberg and S. R. Parry. Free-form deformation of solid geometric models. In *SIGGRAPH '86*, pages 151–160.
- [28] O. Sorkine. Differential representations for mesh processing. *Comput. Graph. Forum*, 25(4):789–807, 2006.
- [29] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Proc. SGP '07*, pages 109–116.
- [30] O. Sorkine, Y. Lipman, D. Cohen-Or, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Proc. SGP '04*, pages 179–188.
- [31] J. Warren. Barycentric coordinates for convex polytopes. *Advances in Computational Math.*, 6(2):97–108, 1996.
- [32] O. Weber, O. Sorkine, Y. Lipman, and C. Gotsman. Context-aware skeletal shape deformation. pages 265–274.
- [33] S. Yoshizawa, A. G. Belyaev, and H.-P. Seidel. Free-form skeleton-driven mesh deformations. In *Proc. ACM Sympos. Solid Modeling Appl. '03*, pages 247–253.
- [34] H. Zhang, O. van Kaick, and R. Dyer. Spectral mesh processing. *Comput. Graph. Forum*, 29(6):1865–1894, 2010.
- [35] K. Zhou, J. Huang, J. Snyder, X. Liu, H. Bao, B. Guo, and H.-Y. Shum. Large mesh deformation using the volumetric graph laplacian. *ACM Trans. Graph.*, 24(3):496–503, 2005.