Regularized Harmonic Surface Deformation

Yeara Kozlov, Janick Martinez Esturo, Hans-Peter Seidel, and Tino Weinkauf

Max Planck Institute for Informatics, Germany

{ykozlov,janick,hpseidel,weinkauf}@mpi-inf.mpg.de

Abstract

Harmonic surface deformation is a well-known geometric modeling method that creates plausible deformations in an interactive manner. However, this method is susceptible to artifacts, in particular close to the deformation handles. These artifacts often correlate with strong gradients of the deformation energy. In this work, we propose a novel formulation of harmonic surface deformation, which incorporates a regularization of the deformation energy. To do so, we build on and extend a recently introduced generic linear regularization approach. It can be expressed as a change of norm for the linear optimization problem, i.e., the regularization is baked into the optimization. This minimizes the implementation complexity and has only a small impact on runtime. Our results show that a moderate use of regularization suppresses many deformation artifacts common to the well-known harmonic surface deformation method, without introducing new artifacts.

1 Introduction

Surface deformation is an important task in geometry processing. Deforming models involves interactive modeling sessions driven by a user, who deforms an object by manipulating a subset of the surface vertices. Linear deformation methods [BS08] have proven effective in this context as they often create plausible and realistic-looking deformations, while still allowing for interactive runtimes. Deformations are usually modeled as the minimizers of specific deformation energies that are defined locally at each point of the domain and measure a specific distortion property.

Harmonic surface deformation is a well-known linear gradient-domain method introduced by Zayer et al. [ZRKS05] that uses a differential surface representation to perform global mesh deformations. Deformation constraints are smoothly distributed over the entire mesh using harmonic functions and surface details are preserved in the reconstructed deformations. Furthermore, the method is parameter-free.

Linear methods such as harmonic deformations are prone to artifacts, as linear energy terms are inadequate to accurately model the non-linear physical forces and processes involved in a deformation. Common deformation artifacts include flipped triangles (in the case of planar surfaces), protruding triangles, degenerate elements, volume loss as well as local and global shape distortion. Many of those artifacts occur close to the deformation handles. Figure 1 shows an example.

A number of non-linear correction methods [Lip12, AL13, SKPSH13, KABL14] exists that suppress these artifacts in planar or volumetric settings. These methods are very powerful as they guarantee artifact-free deformations, but they typically have a big impact on the runtime. Most importantly, they are not applicable to deformations of surfaces. Martinez Esturo et al. [MRT14] introduce an alternative linear energy regularization method: a quadratic regularization term is proposed that is strongly coupled to the problem-specific deformation energy. For a number of problems, this regularization yields to artifact-free results, albeit it cannot be guaranteed. Technically, this energy regularization requires only minor modifications to the algorithm with little impact on the runtime. The amount of regularization can be adjusted using a single parameter.

In this work, we apply linear energy regularization to the harmonic surface deformation method of Zayer et al. [ZRKS05]. Hereby, we follow the general ideas of Martinez Esturo et al. [MRT14]. We demonstrate that energy regularization enhances harmonic deformation results and suppresses a variety of artifacts. Our main contributions are:

- We provide an energy-regularized formulation of harmonic surface deformation.

- We refine the discretization of the energy differential operator of Martinez Esturo et al. [MRT14] for better estimates in high curvature regions.

- We evaluate the effectiveness of our approach. In particular, we demonstrate that moderate use of energy regularization improves deformation results by resolving artifacts without introducing new ones.

This paper is structured as follows: we discuss related shape editing techniques and correction methods (Section 2) and review harmonic deformations as well as energy regularization (Section 3). Then we
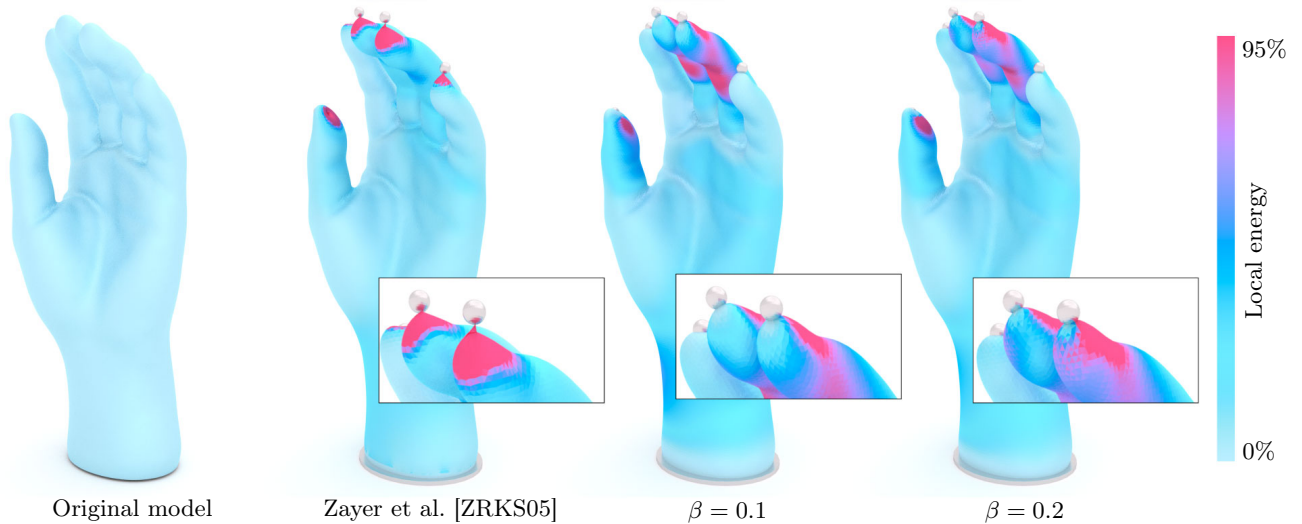
Figure 1: *Hand*. The base of the model is fixed, and each sphere represents a deformation handle, which was rotated. Harmonic surface deformation is susceptible to protruding triangles and surface self-intersection near small deformation handles. Low amounts of regularization (β) suppress these artifacts.

introduce our approach to linear energy regularization for harmonic surface deformation (Section 4). We perform a qualitative analysis of our results (Section 5), followed by quantitative evaluation and discussion (Section 6). Lastly, we present our conclusions and outlook for future work (Section 7).

2   Related Work

The goal of interactive surface deformation is to create meaningful deformations while preserving surface properties such as local details and curvature. Linear deformation methods play a major role in this area, since they provide the interactivity and often produce plausible deformations. Most often, linear methods represent the surface using its differential properties [Sor06]. One can distinguish these methods with respect to their sensitivity regarding rotation and translation. Rotation sensitive methods such as Yu et al. [YZX+04] and Zayer et al. [ZRKS05] use the gradients of affine transformations to construct a deformation guidance field, and solve a Poisson problem for geometry reconstruction. Since translations introduce local changes to the tangent plane of the surface, these methods are not suitable for shape deformations that involve large translations. On the other hand, translation sensitive methods such as [SCL+04] can handle large translations but not rotations. We refer to the survey of Botsch and Sorkine [BS08] for a detailed review of linear deformation methods.

Linear techniques often cannot guarantee that the used deformations are smooth everywhere [JBPS11]. This leads to deformation artifacts such as flipped triangles, protruding elements and volume loss due to rotations.

Artifacts can be avoided by improving the smoothness of the transformation interpolation field. However, bi- or tri-harmonic weights create additional local extrema in the interpolation field that lead to unintuitive deformations results. Jacobson et al. [JWS12] tackle this problem by forcing a desired topology for the interpolation field. This requires solving a non-linear conic problem.

Serveral correction methods have been explored as another means to reducing artifacts in various geometry processing tasks. Lipman [Lip12] presents a generic tool for constructing orientation preserving (i.e., no triangle flips allowed) triangle mesh mappings, while limiting worst-case conformal distortion. This method has non-interactive run times and is only defined for planar meshes. Schüller et al. [SKPSH13] propose a specialized optimization based on a barrier energy function to repress zero-area elements and flipped triangles at interactive rates. Their iterative scheme solves for an injective mapping to a new mesh configuration. They guarantee inversion-free mappings of planar triangular and volumetric tetrahedral meshes. Aigerman and Lipman [AL13] extend [Lip12] to volumetric meshes. Their algorithm takes a deformation created by common deformation techniques and returns a similar deformation that is injective and minimizes the distortion of the mesh volumetric elements. The method is not interactive. Most recently, Kovalsky et al. [KABL14] present a method based on linear matrix inequalities for restricting the range of singular values. It enables, e.g., bounded distortion mappings of planar or volumetric domains, but is also computationally to expensive for interactive applications.

These correction methods guarantee that deformations are inversion-free, and in some cases even protrusion-free. However, they all require solving non-linear systems, which results in loss of interac-

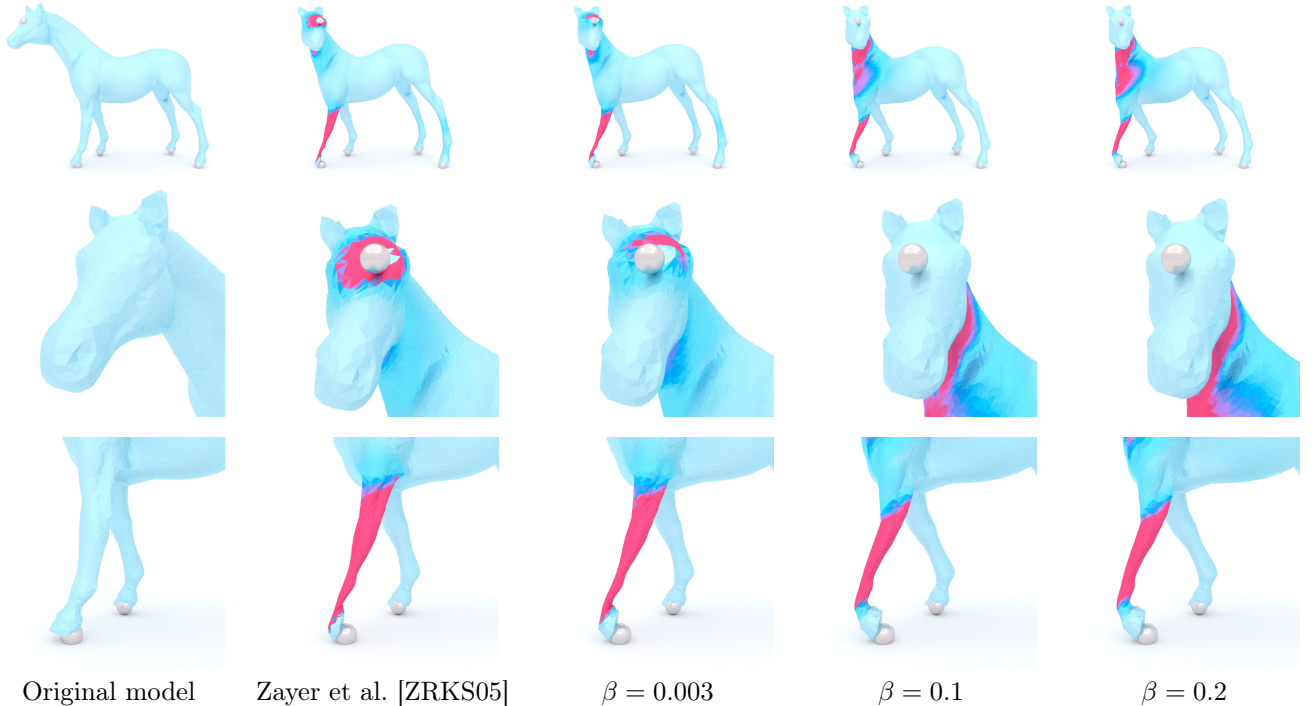| Original model | Zayer et al. [ZRKS05] | $\beta = 0.003$ | $\beta = 0.1$ | $\beta = 0.2$ |

Figure 2: *Horse*. Each sphere corresponds to a handle region. The front foot and the head is rotated, while the rest of the handles are kept fixed. Harmonic surface deformation results in strong distortions close to the head, which can be seen in the closeups (middle row), and severe volume loss at the leg (bottom row). With increasing values of β, regularization gradually improves volume conservation close to the foot.

tivity for moderate to large meshes. In contrast, we follow the recent linear approach to regularization by Martinez Esturo et al. [MRT14], which has no significant impact on runtime. While we cannot guarantee artifacts-free deformations, our method successfully suppresses usual deformation artifacts. Furthermore, the correction methods mentioned above are not applicable to surface meshes embedded in $\mathbb{R}^3$. In contrast, our method is well-define for surface meshes.

3 Background

In this Section, we continue to review the formal details required in our work. We consider triangulated surface meshes $\mathcal{M} = (\mathcal{T}, \mathcal{V}, \mathcal{E})$ defined by sets of vertices $i \in \mathcal{V}$, oriented edges $\mathcal{E} \subset \mathcal{V}^2$, and triangles $\mathcal{T} \subset \mathcal{V}^3$. Coordinates of vertices $i \in \mathcal{V}$ are denoted by $\mathbf{x}_i \in \mathbb{R}^3$. A missing subscript either indicates a vector of stacked coefficients, e.g., $\mathbf{x} \in \mathbb{R}^{3|\mathcal{V}|}$, the vector of stacked vertex coordinates $\mathbf{x}_i$, or a matrix of component-wise coefficients, e.g., $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times 3}$. For a triangle $t \in \mathcal{T}$, $\mathbf{X}_t \in \mathbb{R}^{3 \times 3}$ denotes the column-wise concatenation of the coefficients of its vertices. Using the notations

$$||\mathbf{y}||_{\mathbf{N}}^2 = \mathbf{y}^{\mathrm{T}} \mathbf{N} \mathbf{y} \quad \text{and} \quad ||\mathbf{Y}||_{\mathbf{N}}^2 = \mathrm{Tr}(\mathbf{Y}^{\mathrm{T}} \mathbf{N} \mathbf{Y}) \ , \quad (1)$$

we denote (squared) vector and matrix norms that are induced by symmetric and positive definite matrices $\mathbf{N}$. ($\mathrm{Tr}(\cdot)$ denotes the trace of a matrix.)

For the piecewise linear functions on $\mathcal{M}$ a discrete gradient operator $\mathbf{G} \in \mathbb{R}^{3|\mathcal{T}| \times |\mathcal{V}|}$ can be assembled from local per-triangle gradient operators $\mathbf{G}_t$: for triangles $t = (i, j, k) \in \mathcal{T}$ with normalized normals $\mathbf{n}_t$, the local gradient operators are given by

$$\mathbf{G}_t = \begin{bmatrix} (\mathbf{x}_j - \mathbf{x}_i)^{\mathrm{T}} \\ (\mathbf{x}_k - \mathbf{x}_i)^{\mathrm{T}} \\ \mathbf{n}_t^{\mathrm{T}} \end{bmatrix}^{-1} \begin{pmatrix} -1 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \ , \quad (2)$$

see, e.g., [BS08]. Given a scalar function on $\mathcal{M}$ defined by the vertex-based coefficients $\mathbf{u} \in \mathbb{R}^{|\mathcal{V}|}$, $\mathbf{G}\,\mathbf{u}$ is the vector of stacked and constant per-triangle gradients. Note that gradients computed by $\mathbf{G}$ are defined in a common coordinate system.

3.1 Harmonic Guidance for Surface Deformation

Zayer et al. [ZRKS05] propose a variant of gradient-domain deformations in which local deformation constraints are propagated using harmonic functions. Deformed surfaces are reconstructed from manipulated surface gradients by minimizing the global deformation energy

$$E(\mathbf{x}) = \sum_{t \in \mathcal{T}} A_t \, ||\mathbf{G}_t \, \mathbf{X}_t - \mathbf{Z}_t||_F^2 \quad (3)$$

subject so suitable boundary constraints. Here, A_t denotes the area of triangle t, $||\mathbf{M}||_F^2 = \mathrm{Tr}(\mathbf{M}^{\mathrm{T}} \mathbf{M})$ is the (squared) Frobenius norm of $\mathbf{M}$, and $\mathbf{Z}_t \in$

$\mathbb{R}^{3\times 3}$ are prescribed component-wise *guidance gradients* that are constant per triangle. Dense guidance gradients are computed from user-specified transformations associated to a set of handle regions. In [ZRKS05], harmonic functions $h(\mathbf{x})$ given by solutions of the Poisson equation $\mathbf{L}\,\mathbf{h} = \mathbf{0}$ are used for the global propagation of the sparse set of given transformations. Here, $\mathbf{L} = \mathbf{G}^{\mathrm{T}}\mathbf{A}\,\mathbf{G}$ is a discretization of the Laplace-Beltrami operator [BS08], in which $\mathbf{A}$ is a diagonal matrix of replicated triangle areas. Please, see Zayer et al. [ZRKS05] for further details on the quaternion-based propagation of transformations using harmonic functions. As minimizers of (3) are characterized by the same differential operator, in practice a single factorization of $\mathbf{L}$ can be used to perform both transformation propagation and energy minimization.

Drawbacks. Linear deformation methods are susceptible to various artifacts. This is because the physical energies involved in deformations are non-linear by nature, and are only approximated by linear methods [BS08]. Specifically, harmonic surface deformation is susceptible various deformation artifacts, which are in particular related to the size of the handle regions. Small deformation handles are likely to cause protruding and intruding triangles. Large deformation handles cause local shape distortion on the boundary between the constrained and free mesh areas. Other artifacts include volume loss close to deformation handles when the deformation includes large rotations and local surface intersection. It was observed [MRT14] that these artifacts correlate with large spatial variation in the optimized energy on the domain.

3.2 Linear Energy Regularization

Martinez Esturo et al. [MRT14] propose a generic linear energy regularization scheme that suppresses geometric artifacts in a number of different applications. It is applicable to regularize problem-specific squared energies of the general form
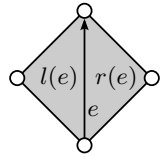
$$E_{\mathcal{P}}(\mathbf{u}) = ||\mathbf{E}\,\mathbf{u} - \mathbf{c}||_{\mathbf{A}}^2 \qquad (4)$$

over d-dimensional piecewise linear functions defined by vertex-based coefficients $\mathbf{u} \in \mathbb{R}^{d|\mathcal{V}|}$. Here, $\mathbf{E} \in \mathbb{R}^{n|\mathcal{T}|\times d|\mathcal{V}|}$ is a problem-specific linear energy operator that maps unknown functions $\mathbf{u}$ to triangle-constant *local energies* of dimension n, $\mathbf{c}$ are problem-specific energy constants, and $\mathbf{A}$ is a diagonal matrix of replicated triangle areas that performs domain-wide integration of triangle-constant quantities. Energies $E_{\mathcal{P}}$ are regularized by introducing a regularization term

$$E_{\mathcal{R}}(\mathbf{u}) = ||\mathbf{D}\,(\mathbf{E}\,\mathbf{u} - \mathbf{c})||_{\mathbf{B}}^2 \qquad (5)$$

that measures squared *variations* of local energies. For piecewise constant local energies, pointwise energy variations are estimated by the sparse differential operator $\mathbf{D}$.

For each pair of neighboring triangles, it can be discretized along all internal edges $e \in \mathcal{E}_i$ from the set of non-boundary edges $\mathcal{E}_i \subseteq \mathcal{E}$: let $l(e)$ and $r(e)$ denote the left and right triangle at e, respectively. Then, for scalar local energies ($n = 1$), the nonzero coefficients of $\mathbf{D}$ are given by

$$D_{et} = \begin{cases} 1 & \text{if } l(e) = t \\ -1 & \text{if } r(e) = t \end{cases}, \text{ for all internal edges } e \in \mathcal{E}_i$$

and triangles $t \in \mathcal{T}$. For vector-valued local energies ($n > 1$), the differential operator is given by a component-wise replication, which can be expressed as $\mathbf{D}\otimes\mathbf{I}_n$ using the Kronecker product $\otimes$ and the $n\times n$ identity matrix $\mathbf{I}_n$. The constant estimates of pointwise local energy variations are integrated using the diagonal matrix $\mathbf{B}$ of replicated internal edge lengths.

The total regularized energy is given by a weighted combination of both terms

$$E_{\beta}(\mathbf{u}) = (1 - \beta)\,E_{\mathcal{P}}(\mathbf{u}) + \beta\,E_{\mathcal{R}}(\mathbf{u}) = ||\mathbf{E}\,\mathbf{u} - \mathbf{c}||_{\mathbf{W}_{\beta}}^2 \qquad (6)$$

that can be expressed compactly using the β-weighted norm

$$\mathbf{W}_{\beta} = (1 - \beta)\,\mathbf{A} + \beta\,\mathbf{D}^{\mathrm{T}}\mathbf{B}\,\mathbf{D} . \qquad (7)$$

The amount of regularization is steered by $\beta \in [0, 1)$. Note that this formulation of energy regularization is also valid for energies in the components of the unknown functions $\mathbf{u}$, which are then given by $E_{\beta}(\mathbf{U}) = ||\mathbf{E}\,\mathbf{U} - \mathbf{C}||_{\mathbf{W}_{\beta}}^2$. Please, see Martinez Esturo et al. [MRT14] for further details and applications on this energy regularization scheme.

4 Enhancing Harmonic Surface Regularization

We continue to show that the concept of energy regularization is applicable to harmonic surface deformation in a straightforward way. For this, we rewrite the component-wise deformation energy (3) to the equivalent formulation

$$E(\mathbf{X}) = ||\mathbf{G}\,\mathbf{X} - \mathbf{Z}||_{\mathbf{A}}^2 \qquad (8)$$

using the global gradient operator $\mathbf{G}$, the matrix $\mathbf{Z}$ of all stacked prescribed gradients, and diagonal matrix $\mathbf{A}$ of replicated triangle areas. Note that the local energies correspond to the to the summed terms of (3). Comparing our problem-specific energy (8) and the regularizable generic energy (4), we obtain the correspondences that the generic energy operator $\mathbf{E}$ is given by the gradient operator $\mathbf{G}$, the constant energy term $\mathbf{C}$ is given by the gradient field $\mathbf{Z}$, and the dimension of the local energies is $n = 3$. Hence, the energy-regularized version of the deformation energy (8) is given by

$$E_{\beta}(\mathbf{X}) = ||\mathbf{G}\,\mathbf{X} - \mathbf{Z}||_{\mathbf{W}_{\beta}}^2 \qquad (9)$$

by simply applying the weighted norm $\mathbf{W}_\beta$ for energy integration and smoothness estimation. Technically, to apply regularization, this substitution of $\mathbf{A}$ for $\mathbf{W}_\beta$ allows for a straightforward implementation. Specifically, the remainder of the original harmonic surface deformation approach is unchanged, in particular the harmonic function-based transformation propagation for the guidance gradients $\mathbf{Z}$.

Our experiments demonstrate that this simple energy modification suppresses a variety of deformation artifacts of the original energy formulation (see Section 5). Still, the original discretization of the energy differential operator $\mathbf{D}$ is defined independently of the local surface curvature, which leads to poor estimates of energy variation in high curvature regions. We continue to provide a refined differential operator discretization that is based on surface curvature and yields better estimates of energy variation.
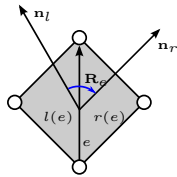
Curvature-based Energy Differential Operator. Martinez Esturo et al. [MRT14] estimate the variation of local energies by finite differences of their respective local energy residuals (Section 3.2). In our application of gradient-domain deformations, the energy residual $\mathbf{E}_t$ of a triangle t is given by the deviation between the deformed mesh gradients and the gradients of the guidance field:

$$\mathbf{E}_t = \mathbf{G}_t\,\mathbf{X}_t - \mathbf{Z}_t\ . \tag{10}$$

The (squared) energy variation between two neighboring triangles t_1 and t_2 is now simply given by $||\mathbf{E}_{t_1} - \mathbf{E}_{t_2}||_F^2$.

This estimation works well if both triangles are aligned and therefore also the residuals in the columns of $\mathbf{E}_{t_1}$ and $\mathbf{E}_{t_2}$ live in the same tangent space, e.g., for planar meshes. However, if the tangent spaces are not aligned, this estimation is likely to break. The effect is particularly noticeable at sharp edges of a curved surface, where areas of low energy values along the edge are surrounded by higher energy values. This results in a less smooth energy distribution across edges of high curvature.

We adjust the discretization of the operator $\mathbf{D}$ to also be applicable to high curvature regions by locally compensating for curvature. Similar local alignments of tangent spaces is used, e.g., by Crane et al. [CDS10], for the computation of connections relating neighboring tangent spaces. For the $n = 3$ dimensional local energy residuals (10), our refined energy differential operator $\mathbf{D}^R$ is a $|\mathcal{E}_i| \times |\mathcal{T}|$ matrix of 3×3 block matrices. For each internal edge $e \in \mathcal{E}_i$, we use the rotation $\mathbf{R}_e$ that aligns the normals of its left and right triangles $\mathbf{n}_r, \mathbf{n}_l$ such that $\mathbf{n}_r = \mathbf{R}_e \mathbf{n}_l$. Then, the nonzero blocks of $\mathbf{D}^R$ are given by $D_{et}^R = \begin{cases} \mathbf{R}_e & \text{if } l(e) = t \\ -\mathbf{I}_3 & \text{if } r(e) = t \end{cases}$ for all internal edges $e \in \mathcal{E}_i$ and

triangles $t \in \mathcal{T}$. This way, $\mathbf{D}^R$ compensates for local curvatures: residuals that are similar relative to their respective tangent spaces are not estimated to be different anymore. Note that this operator refinement simplifies to the original formulation for planar meshes. It reduces the sparseness of the resulting linear system only by a constant factor, which does not impair resulting runtimes. However, it is only applicable to this particular case of $n = 3$ dimensional local energy residuals. For different values of n the discretization should therefore fall back to the $\mathbf{D}$ operator of [MRT14].

Implementation. The operators $\mathbf{G}$, $\mathbf{D}^R$, $\mathbf{A}$, and $\mathbf{B}$ as well as $\mathbf{L}$ are assembled once when the surface mesh is loaded. For each deformation of the model the guidance field $\mathbf{Z}$ is computed and the corresponding normal equations

$$\mathbf{G}^\top \mathbf{W}_\beta\, \mathbf{G}\,\mathbf{X} = \mathbf{G}^\top \mathbf{W}_\beta\, \mathbf{Z} \tag{11}$$

of (9) are solved for the coordinates $\mathbf{X}$ of the deformed mesh. The norm $\mathbf{W}_\beta$ is assembled for a given β value using the refined differential operator $\mathbf{D}^R$. After elimination of positional hard boundary constraints, the system (11) is symmetric positive definite and it is solved using a Cholesky factorization with fill-in reducing reordering [GJ+10]. Similar to [ZRKS05], this factorization can be reused for different guidance fields $\mathbf{Z}$ as long as the handle configuration and β values are unchanged.

5   Results

We evaluate our approach qualitatively on a number of different models. The models presented have $3 - 36.6k$ vertices. Deformations are performed using single and multiple handles, varying sizes of the handle regions and the regularization weight β.

Energy Visualization. We use color to visualize the squared magnitude of local triangle-constant energies $||\mathbf{E}_t||_F^2$. Energies are linearly mapped to the color space (shown in Figure 1) by setting the maximum of the color interval to correspond to the 95th percentile of the energy values. The top 5% of the energy values is clipped to allow visualization of the variation of lower energies with higher contrast.

Small Deformation Handles. For small handle regions, harmonic surface deformation tends to create artifacts such as protruding triangles and local surface intersections close to the handles. We show two examples of these artifacts: the *Hand* deformation (Figure 1) is created by fixing the base of the model, and a single vertex on each finger acts as the deformation handle. All handles are rotated inwards,
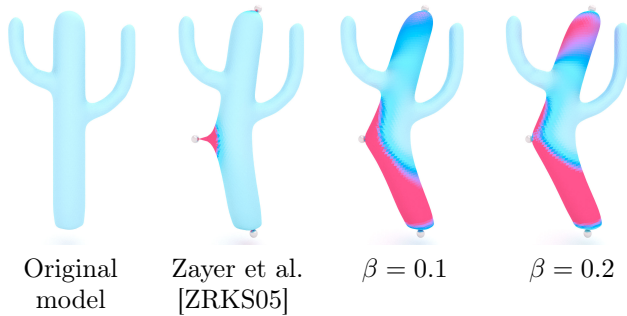
5

Figure 3: *Cactus 2*. Single vertices at the bottom and at the top of the cactus are fixed, and a single vertex is pulled away from the cactus body. Regularization results in deformations free of protruding triangles.

Figure 4: *Cow*. The handle on face is fixed, and the horns are rotated upwards. Regularization helps preserve the volume near the base of the horns.

resulting in protruding triangles and local surface intersections near the deformation handles. Regularization ($\beta > 0$) suppresses these artifacts. Similar artifacts to protruding triangles, albeit of smaller magnitude, occur if handle regions include up to several dozen vertices. In the *Cactus 2* example shown in Figure 3, single vertices on the base and the top are fixed, and a single vertex on the side of the cactus is used as a deformation handle. The deformation suffers from large protruding triangles. Introducing regularization corrects this artifact. For both deformations, even low amounts of regularization result in a smoother energy distribution, affecting more triangles in the mesh. This way the optimization favors more global changes to the mesh instead of local concentrations of energies, which result in the observed local artifacts.

Large Deformation Handles. Harmonic surface deformation is commonly used with large deformation handles, as careful design of these constraints can lead to pleasing deformations. However, the boundary of large handle regions is also susceptible to artifacts. The *Cactus 1* model in Figure 5 is based on a benchmark deformation from [BS08]. The base of the model is fixed, and the top is rotated and translated. Using the original method by Zayer et al. [ZRKS05], the boundary between the constrained and free mesh regions exhibits strong changes of the directions of mesh normals, distorting the local shape. Our regularization reduces distortions of local geometry close to this

boundary, creating a smooth transition between the constrained and free mesh regions.

Strong Rotations. Strong rotations, especially on elongated limbs, are a challenge for a number of deformation methods [KS12]. Using harmonic surface deformation, strong rotations usually cause loss of volume. This effect is illustrated in Figures 4 and 2. The foot of the *Horse* is rotated. In the resulting deformation, most of the lower leg suffers from volume loss, creating a "candy wrapper"-like artifact. Low values of regularization are effective at preserving some of the limb's volume to create more plausible results. In the *Cow* deformation, the horns are rotated upwards while the front of the face is fixed. This deformation is even more challenging, as the region deformed is small and the mesh geometry is rather coarse. Regularization helps to restore some of the lost volume.

6 Evaluation and Discussion

In this section, we perfom a quantitative evaluation of our approach. Deformations are created for $\beta \in [0, 1)$, using a higher sampling rate for lower β values, for which we usually observe the strongest changes of the deformation. Please, see the accompanying **video** for the deformations at these β values.

Maximal Isometric and Conformal Errors. Two error measures proved to be most useful for evaluating deformation quality (see, e.g., [LZX$^+$08]): the local per triangle isometric error e_t^{iso} is given by the sum of squared deviations from 1 of the singular values of the deformation gradient. The local per triangle conformal error e_t^{conf} is computed by half of the squared sum of the pairwise differences between the singular values of the deformation gradient. The maximal isometric (conformal) error, $\max(e_t^{\text{iso}})$ ($\max(e_t^{\text{conf}})$), is given by the maximal value of the isometric (conformal) error over all triangles. Both of these error measures indicate strong distortions of the mesh geometry. They are also loosely related to artifacts such as protruding or intruding elements, degenerate triangles, and surface self-intersections. We also examined the total (integrated and normalized) isometric and conformal errors as possible indicators for deformation quality evaluation, but these proved ineffective.

Figure 6 shows the behavior of the maximal isometric and conformal errors for different values of β. In almost all tested deformations, regularization strongly decreases the maximal isometric and conformal errors. The behavior of these error measurements is different for the *Cactus 1* example of Figure 5, as they slightly increase with regularization. The reason is that deformations defined using large handle regions usually don't result in artifacts associated with

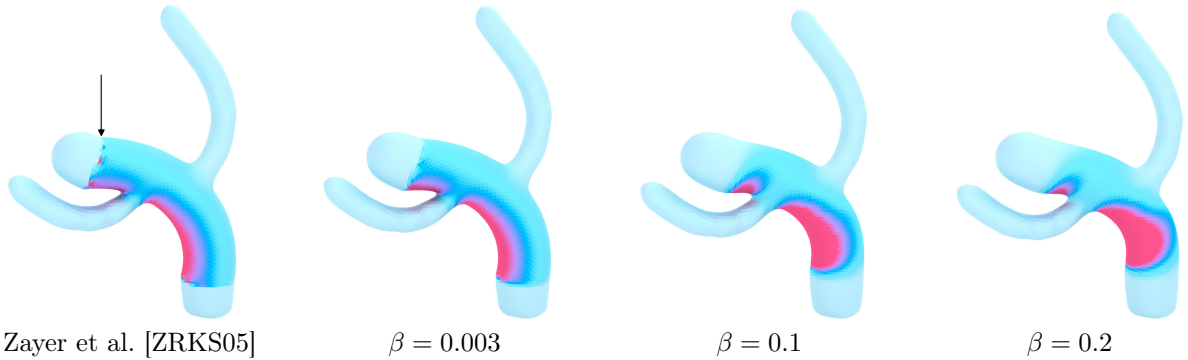| Zayer et al. [ZRKS05] | $\beta = 0.003$ | $\beta = 0.1$ | $\beta = 0.2$ |

Figure 5: *Cactus 1.* The entire base and top regions are constrained. The base is fixed and the top is translated and rotated. The transition between the constrained and free mesh regions is apparent both in the energy domain and in the irregularities of mesh normal directions. Low regularization weights create smoother transitions and less distortions on the handle boundary, while preserving the global shape of the deformation.

large local isometric and conformal errors. Additionally, we note that no new local maximal errors are introduced for moderate β values, meaning that deformation quality does not derogate for higher β values. We also confirm this behavior in all other tested examples, as the total deformation energy of the regularized deformation defined in (3) stays within the same order of magnitude as the original deformation.

Space of Regularized Deformations. We observe that the initial introduction of regularization ($\beta > 0$) has strong effects on deformations suffering from strong artifacts for $\beta = 0$. After this initial reaction interval, regularization creates gradual changes to the mesh geometry, indicated both by our deformation results (e.g., in Figures 2 and 5) and by the behavior of the error measures $\max(e_t^{\mathrm{iso}}), \max(e_t^{\mathrm{conf}})$ in Figure 6. As the regularization energy becomes more dominant with increasing β values, $\max(e_t^{\mathrm{iso}}), \max(e_t^{\mathrm{conf}})$ also slightly increase. For very high values of $\beta > 0.9$, the regularized energy formulation strongly deviates from the original problem, which can create new artifacts. In our experience, choosing a fixed value of $\beta \approx 0.2$ usually suppresses artifacts effectively without negatively affecting the mesh geometry. This means that a constant regularization can simply be added to existing implementations without exposing users to a new parameter. In addition, although the linear system (11) has to be refactored when β changes, examining different β values can usually be done at interactive rates due to the high performance of sparse linear solvers. Hence, the space of regularized deformations can also be explored interactively.

Curvature-based Differential Operator. Figure 7 demonstrates the benefits of using our curvature-based energy differential operator on the *Accordion* mesh with highly curved edges. The method by Zayer et al. [ZRKS05] suffers from local shape distortion near the deformation handle. The differential operator by Martinez Esturo et al.

[MRT14] estimates the energy variation between neighboring triangles ineffectively, resulting in global shape distortions. Our curvature-based operator estimates the energy differential more effectively, resulting in a deformation which has the same geometry as the solution by Zayer et al. [ZRKS05], but suppresses the local artifact. The deviation of Martinez Esturo et al. [MRT14] from our results increases with higher regularization weight. For smooth meshes, using our curvature-based differential operator has negligible effects, as vertex coordinates were only affected marginally.

Performance. Regularization affects the runtime in two ways: it reduces the sparsity of the linear system (11) being solved and it requires a one-time operation to setup of the curvature-based differential operator $\mathbf{D}^R$. Our measurements on an Intel Core i7 $2.2GHz$ system indicates that the effect of regularization on interactive runtime is insignificant for all tested meshes: For example, factorization time of the linear system for the *Cactus* model ($|\mathcal{V}| = 10k$) is ≈ 0.04 seconds for *both* $\beta = 0$ and $\beta > 0$, its solution time is ≈ 0.003 seconds. Similarly, factorization time for the *Hand* model ($|\mathcal{V}| = 36.6k$) is ≈ 1.13 seconds for $\beta = 0$ and ≈ 1.14 seconds for $\beta > 0$ with a solution time of ≈ 0.1 seconds. Hence, changing the regularization weight to examine different regularization weight can be done interactively.

7 Conclusions

In this work, we have provided an energy regularized formulation of harmonic surface deformation. Our approach expands the capabilities of the original method, allowing the creation of artifact-free deformations for a wider range of deformation constraints and handle configurations. Our formulation of a curvature-based differential energy operator improves the estimation of energy differentials in high-curvature mesh regions. This reduces geometric dis-
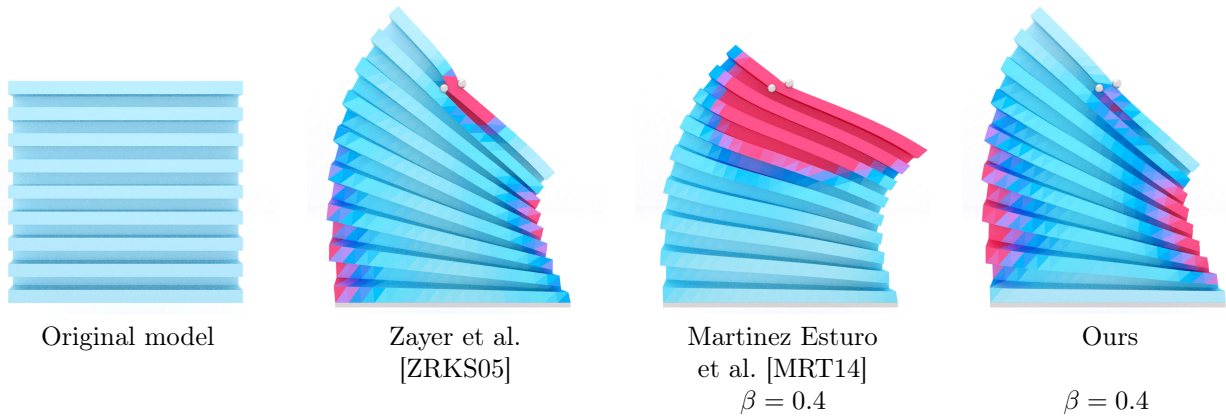
| | Original model | Zayer et al. [ZRKS05] | Martinez Esturo et al. [MRT14] $\beta = 0.4$ | Ours $\beta = 0.4$ |

Figure 7: *Accordion.* The base is fixed and two vertices on the top row are rotated. Harmonic surface regularization [ZRKS05] creates an artifact near the deformation handles. Martinez Esturo et al. [MRT14] suppress the artifact, but distorts the global mesh shape. Our curvature-based differential operator creates a smooth, artifact-free deformation without distorting the mesh geometry.

tortions introduced by the original energy differential operator around these regions. The evaluation of our results demonstrates that even low regularization weights can effectively suppress many deformation artifacts without negatively affecting the performance of the original method. In addition, no new artifacts are created,

Future work. An interesting direction for future work is the application of energy regularization to other (non-linear) surface deformation approaches, e.g., to the work of Jacobson et al. [JBK+12], who observe similar artifacts. In addition, since artifacts are usually localized close to handle regions, a local energy smoothness formulation could achieve better control on deformation artifacts.

Acknowledgment. The *Horse*, *Hand*, and *Cow* models are provided by the AIM @ Shape project. The *Cactus* model courtesy of Botsch and Sorkine [BS08].

References

[AL13]   Noam Aigerman and Yaron Lipman. Injective and bounded distortion mappings in 3d. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 32(4):106:1–106:14, 2013.

[BS08]   Mario Botsch and Olga Sorkine. On linear variational surface deformation methods. *IEEE TVCG*, 14(1):213–230, 2008.

[CDS10] Keenan Crane, Mathieu Desbrun, and Peter Schröder. Trivial connections on discrete surfaces. *Comput. Graph. Forum (Proc. SGP)*, 29(5):1525–1533, 2010.

[GJ+10] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. http://eigen.tuxfamily.org, 2010.

[JBK+12] Alec Jacobson, Ilya Baran, Ladislav Kavan, Jovan Popović, and Olga Sorkine. Fast automatic skinning transformations. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 31(4):77:1–77:10, 2012.

[JBPS11] Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 30(4):78:1–78:8, 2011.

[JWS12] A. Jacobson, T. Weinkauf, and O. Sorkine. Smooth shape-aware functions with controlled extrema. *Comput. Graph. Forum (Proc. SGP)*, 31(5):1577–1586, 2012.

[KABL14] Shahar Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. Controlling singular values with semidefinite programming. *ACM Trans. Graph. (Proc. SIGGRAPH)*, page (to appear), 2014.

[KS12]   Ladislav Kavan and Olga Sorkine. Elasticity-inspired deformers for character articulation. *ACM Trans. Graph. (Proc. SIGGRAPH ASIA)*, 31(6):196:1–196:8, 2012.

[Lip12] Yaron Lipman. Bounded distortion mapping spaces for triangular meshes. *ACM Trans. Graph.*, 31(4):108:1–108:13, 2012.

[LZX+08] Ligang Liu, Lei Zhang, Yin Xu, Craig Gotsman, and Steven J. Gortler. A local/global approach to mesh parameterization. *Comput. Graph. Forum (Proc. SGP)*, 27(5):1495–1504, 2008.
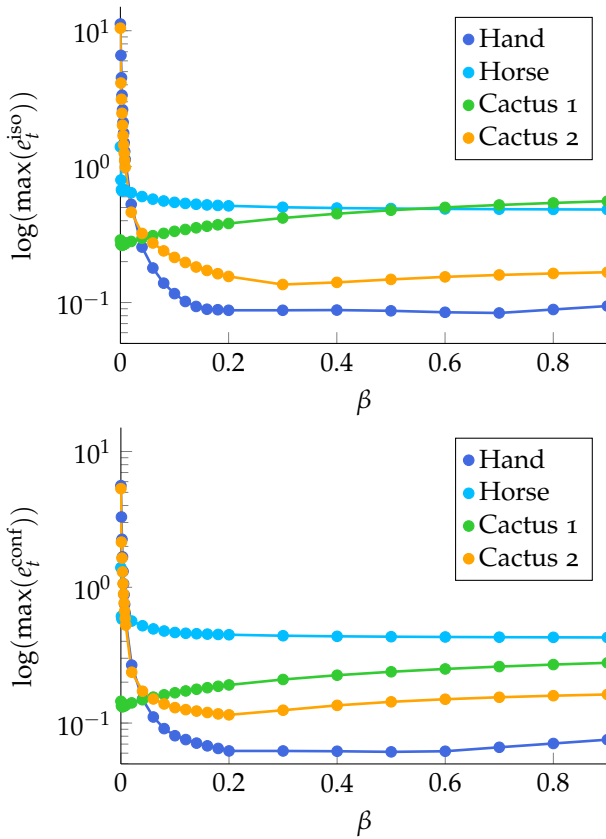
Figure 6: Graphs of the maximal isometric and conformal errors over the regularization weight β. Regularization typically decreases these errors drastically already for small values of β. Note the logarithmic scale on the y-axes.

[MRT14]   Janick Martinez Esturo, Christian Rössl, and Holger Theisel. Smoothed quadratic energies on meshes. *ACM Trans. Graph.*, page (to appear), 2014.

[SCL+04]   O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Proc. SGP*, pages 175–184, 2004.

[SKPSH13] Christian Schüller, Ladislav Kavan, Daniele Panozzo, and Olga Sorkine-Hornung. Locally injective mappings. *Comput. Graph. Forum (Proc. SGP)*, 32(5):125–135, 2013.

[Sor06]   Olga Sorkine. Differential representations for mesh processing. *Comput. Graph. Forum*, 25(4):789–807, 2006.

[YZX+04]   Yizhou Yu, Kun Zhou, Dong Xu, Xiaohan Shi, Hujun Bao, Baining Guo, and Heung-Yeung Shum. Mesh editing with poisson-based gradient field manipulation. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 23(3):644–651, 2004.

[ZRKS05]   Rhaleb Zayer, Christian Rössl, Zachi Karni, and Hans-Peter Seidel. Harmonic