

Scalable Locally Injective Mappings

MICHAEL RABINOVICH and ROI PORANNE

ETH Zurich

and

DANIELE PANOZZO

ETH Zurich and New York University

and

OLGA SORKINE-HORNUNG

ETH Zurich

We present a scalable approach for the optimization of flip-preventing energies in the general context of simplicial mappings, and specifically for mesh parameterization. Our iterative minimization is based on the observation that many distortion energies can be optimized indirectly by minimizing a simpler proxy energy and compensating for the difference with a reweighting scheme. Our algorithm is simple to implement and scales to datasets with millions of faces. We demonstrate our approach for the computation of maps that minimize a conformal or isometric distortion energy, both in two and three dimensions. In addition to mesh parameterization, we show that our algorithm can be applied to mesh deformation and mesh quality improvement.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—*Geometric algorithms, languages, and systems*

General Terms: Mesh Parameterization, Optimization

Additional Key Words and Phrases: parameterization, bijectivity, scalability

1. INTRODUCTION

Mappings are an essential tool in computer graphics and geometry processing. One of the most basic uses, and the main focus of this paper, is *mesh parameterization*. Many practical applications such as texture mapping, remeshing, shape correspondence and attribute transfer rely on the computation of a low-distortion parameterization. The problem has been extensively studied, and a plethora of algorithms have been devised. Linear methods were proposed first, providing efficient ways to compute parameterizations, but only able to ensure injectivity of the map when the mesh boundary is fixed a priori, which induces a high distortion. As more powerful processors became available, nonlinear optimization became tractable, allowing one to compute free boundary, injective or bijective maps of a very high quality. Still, current nonlinear approaches typically require long computation times and do not scale well to large datasets, such as detailed scanned surfaces like the one in Fig. 1.

In this paper, we propose a simple algorithm that combines the benefits of the two approaches: it scales well to large datasets with millions of elements (Fig. 1) and minimizes state-of-the-art nonlinear energies (Fig. 2). In particular, we focus on minimizing *flip-preventing* energies and we propose an algorithm that is guaranteed to produce optimized maps without any flipped elements. On medium-sized meshes, our algorithm is two orders of magnitude faster than competing methods, while being able to minimize many different energies with minimal code changes.




Fig. 1. We compare the behavior of our algorithm on progressively simplified versions of the Lucy model. We observe that the number of iterations required is not dependent on the resolution of the mesh.

The key idea of our method is to minimize the nonlinear energy using much simpler proxy functions that permit the use of a local/global approach. Our algorithm scales well to large datasets even using a single core, and it can take advantage of the recent developments in parallel solution of linear systems to be deployed on multi-core architectures. While we are unable to provide a strict bound on the convergence rate, we experimentally found that the number of iterations required by our method is related to the geometric surface complexity and is not affected by the tessellation density (Fig. 1). Since each iteration merely requires solving a sparse linear system, massive datasets can be parameterized quickly.

2. PREVIOUS WORK

Mappings are one the most researched subjects in computer graphics, and specifically the problem of generating locally injective 2D and 3D mappings has garnered a lot of attention in the past decades. In this section, we mention only the most closely related work on the topic of large scale mesh parameterization and we refer to [Floater and Hormann 2005; Sheffer et al. 2006] for in-depth surveys.




Fig. 2. A locally injective parameterization obtained by minimizing the symmetric Dirichlet energy [Smith and Schaefer 2015] on a mesh with over 25 million triangles, computed with our algorithm in 80 minutes. The algorithm starts from a highly distorted locally injective initialization and in only 40 iterations, each requiring to solve a sparse linear system, it converges to a map with low isometric distortion that is guaranteed to be free of inverted elements.

Linear methods. Linear methods compute a mesh parameterization by solving a linear system, where each mesh vertex is represented as a weighted average of its neighbors. They have been proposed to parameterize topological disk patches [Tutte 1963] or topological spheres [Aigerman and Lipman 2015]. For topological disks, linear methods can be guaranteed to produce bijective parameterizations if the patch boundary is fixed to a convex shape and the weights are positive [Floater 2003]. Free-boundary methods exist that minimize a measure of conformal distortion [Desbrun et al. 2002; Lévy et al. 2002; Zayer et al. 2007; Ben-Chen et al. 2008; Mullen et al. 2008], but they are not guaranteed to produce a bijective map.

Nonlinear methods. Many nonlinear deformation energies have been proposed in the literature for both conformal and isometric distortion. They are typically minimized using standard optimization methods, such as Newton [Sheffer and de Sturler 2001; Chao et al. 2010; Sanan 2014], Gauss-Newton [Sanan 2014], quasi-Newton [Smith and Schaefer 2015] and second-order cone programming [Aigerman et al. 2014]. To simplify implementation and reduce memory usage, several works [Hormann and Greiner 2000; Labsik et al. 2000; Schreiner et al. 2004] opt for a block descent optimization, where in each iteration only a single vertex is free to move. Similarly, [Levi and Zorin 2014; Fu et al. 2015] optimize an independent subset of vertices in parallel. These methods do not scale to large datasets, since the number of iterations they need grows quickly with the size of the mesh. In contrast, we observe that the number of iterations required by our method is related to the geometric complexity and not the dataset size (Fig. 1): even datasets with millions of elements can be parameterized within a few iterations.

The local/global minimization of isometric distortion [Sorkine and Alexa 2007; Liu et al. 2008] iteratively alternates between a local step and a global step. In the local step, each element is individually perfectly mapped (without any distortion), and in the global step, a linear system is solved to stitch all elements back together. This process very quickly recovers from a bad initialization, but it is slow to converge to a local minimum when it is close to it. The decoupling of a local condition from a global “stitching” has been successfully used in other parameterization algorithms [Weber et al. 2012] and to enforce complex constraints [Bouaziz et al. 2012; Poranne et al. 2013]. Our method uses the local/global paradigm and enriches it with a reweighting scheme to efficiently minimize nonlinear, flip-preventing energies.

Non-flipping invariant. A recent series of works [Schüller et al. 2013; Fu et al. 2015; Smith and Schaefer 2015] have proposed

parameterization energies with a term that goes to infinity when an element inverts. These *flip-preventing* energies are minimized starting from a flipless initialization (e.g., [Tutte 1963]) using line search to ensure that they never leave the feasible region. This approach is guaranteed to create a locally injective map given a feasible starting point, but the energies are numerically difficult to optimize, leading to high running times. Our algorithm is specifically designed to optimize these energies and, differently from all existing methods, quickly recovers from the highly distorted starting point.

Bounding, projections. Another approach to the creation of locally injective maps is based on directly *bounding* the distortion. Similar to the above, a parameterization algorithm is first employed to generate an initial locally injective map. The energy of the map is then optimized while adhering to a specified distortion bound [Lipman 2012; Kovalsky et al. 2014; Sanan 2014; Chen and Weber 2015]. A similar, more recent approach projects any map, possibly with flips, to the closest map that has no flips [Aigerman and Lipman 2013; Kovalsky et al. 2015]. The major limitation with these approaches is that the elements in the solution they generate have suboptimal distortion. In fact, their distortion tends to be the highest possible without violating the bound (Fig. 3). These methods are not guaranteed to find a valid solution [Myles et al. 2014] and they often need thousands of iterations to find one. An interesting special case is the convex energy used in ex-rotated elasticity simulations, for which Liu et al. [2016] propose a numerical approach that can both minimize the energy and robustly recover from flipped elements. However, the ex-rotated energy requires to specify a fixed rotation per element as input, inhibiting its use for parametrization tasks and limiting its applicability to a very specific isometric distortion measure [Liu et al. 2016].

Stiffening. Another strategy is the so-called *stiffening*, where the idea is to try and coerce inverted elements to reorient correctly. Examples include [Irving et al. 2004], where an added force acts on inverted elements in a simulation, and [Martin et al. 2011], where a volume term is added to the energy of each element. Another method related to ours is found in [Bommes et al. 2009]: there, the inverted elements are progressively reweighted in the parameterization energy, as opposed to minimizing the original energy in the space of locally injective maps. Finally, the IRLS algorithm [Pighin and Lewis 2007] can be similarly considered as a special form of stiffening. All these methods can be enriched with a line search to ensure the generation of locally injective maps, but they get stuck before reaching a local minimum and thus produce maps that are considerably inferior in quality when compared with the results of




Fig. 3. The bounded ARAP energy (left, result taken from [Fu et al. 2015]) pushes the triangles to the distortion bound, while a direct minimization of the AMIPS energy [Fu et al. 2015] evenly distributes the distortion over the surface (middle). Our approach can also optimize the AMIPS energy (right), further reducing the energy to a local minimum. This model has 11K faces and took our method less than 20 iterations to converge. Bottom row: the distortion magnitude is visualized by the saturation of the red color.

our algorithm (Fig. 4). We discuss these methods in more details in Sec. 4.

Large scale. Computing locally injective maps with a large number of elements is a challenging numerical problem that has been tackled in surprisingly few research papers in graphics. Apart from linear methods that scale well since they only involve the solution of a linear system, the only other works we are aware of are ABF++ [Sheffer et al. 2005] and [Kovalsky et al. 2015]. The former uses a multiresolution hierarchy to make the problem tractable at the expense of a higher distortion compared to the original ABF [Sheffer and de Sturler 2001], and it can only optimize for angle preservation. The algorithm of Kovalsky et al. [2015] is a projection method to the space of conformal maps. It does not minimize map distortion (Fig. 3), it cannot generate isometric parameterizations [Kovalsky 2016] and it often fails to find a valid map, making it impractical for many applications. Our method is the first algorithm that can robustly and efficiently compute locally injective maps that contain millions of elements and minimize a conformal or isometric distortion measure.

Other applications. In addition to parameterization, our algorithm can be used to deform 3D objects, and we can also adapt our method to minimize mesh improvement energies, as suggested in [Lipman 2012; Aigerman and Lipman 2013; Fu et al. 2015]. “Seamless” rigid and conformal parameterizations (i.e., parameterizations whose gradients match across seams, see [Myles and Zorin 2012; Myles et al. 2014; Diamanti et al. 2015]) are also supported, although we cannot enforce the integer translations that are necessary for remeshing applications [Bommes et al. 2012].

3. ALGORITHM

Denote the input triangle or tetrahedral mesh by $M = (V, F)$, where V is the set of vertices and F is the set of elements. A common way to define a distortion energy of a mapping is via a function of its Jacobians, measuring the distance of each Jacobian to its closest rotation. We denote the Jacobian of element $f \in F$ by \mathbf{J}_f . Then the




Fig. 4. We compare different numerical methods to minimize the symmetric Dirichlet energy (Eq. (11)) starting from a Tutte’s initialization (top left): the greedy stiffening [Bommes et al. 2012] (top right), Iteratively Reweighted Least Squares (bottom left) and our approach (bottom right). See Sec. 4 for a detailed explanation of these techniques.

energy we wish to minimize is

$$\min_{\mathbf{x}} E(\mathbf{x}) = \sum_{f \in F} A_f \mathcal{D}(\mathbf{J}_f(\mathbf{x})), \quad (1)$$

where $\mathcal{D}(\cdot)$ is the distortion measure, $\mathbf{x} \in \mathbb{R}^{|V| \times d}$ contains the mapping coordinates of the vertices ($d = 2, 3$) and A_f is the area (or volume) of element f . A popular choice is the As-Rigid-As-Possible (ARAP) measure [Liu et al. 2008], which is defined by

$$\mathcal{D}_{\text{ARAP}}(\mathbf{J}_f(\mathbf{x})) = \|\mathbf{J}_f(\mathbf{x}) - \mathbf{R}(\mathbf{J}_f(\mathbf{x}))\|_F^2, \quad (2)$$

where $\mathbf{R}(\mathbf{J}_f(\mathbf{x}))$ is the closest rotation to $\mathbf{J}_f(\mathbf{x})$ in the Frobenius norm, and $\|\cdot\|_F$ denotes the Frobenius norm.

Local/Global optimization. Liu *et al.* [2008] proposed to minimize the ARAP energy using a *local/global* algorithm, an iterative process that alternates between two steps. At iteration k , the local step greedily pushes the Jacobian of each element towards the closest minimum, which is a rotation. That is,

$$\mathbf{R}_f^k := \mathbf{R}(\mathbf{J}_f(\mathbf{x}^{k-1})) = \mathbf{U}\mathbf{V}^\top, \quad (3)$$

where $\mathbf{J}_f(\mathbf{x}^{k-1}) = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ is the *signed* SVD of $\mathbf{J}_f(\mathbf{x}^{k-1})$. The global step that follows minimizes the deviation from those rotations,

essentially averaging the distortion among all elements in order to stitch them together. This is done by solving

$$\mathbf{x}^k = \arg \min_{\mathbf{x}} \sum_{f \in F} A_f \|\mathbf{J}_f(\mathbf{x}) - \mathbf{R}_f^k\|_F^2 + \lambda \|\mathbf{x} - \mathbf{x}^{k-1}\|^2. \quad (4)$$

The second term is a proximal term that makes the solution unique, removing the invariance of the energy to rigid transformations, and it is equivalent to adding hard constraints to remove these degrees of freedom. We used $\lambda = 10^{-4}$ in this paper. Since the Jacobians are linear functions of \mathbf{x} , the above is a simple quadratic minimization that can be computed by solving a linear system (see [Liu et al. 2008]), which is always full rank by construction due to the proximal term. A notable property of the ARAP formulation is that the local/global scheme decreases the energy in each iteration, both in the local and the global step [Sorkine and Alexa 2007; Liu et al. 2008].

While the ARAP energy softly penalizes degenerate and inverted elements, they commonly appear in practice [Civit-Flores and Susin 2014; Martinez Esturo et al. 2014], making the resulting parameterizations unusable in many geometry processing tasks, such as texture mapping and remeshing (Fig. 5). However, the local/global optimization strategy itself is general and has the very interesting property of making large steps in the minimization when initialized with a point far from the solution. This well-known property makes the local/global approach attractive, since a few iterations are sufficient to obtain a result that is good enough for many practical purposes. In addition to parametrization, this technique has been successfully applied to shape deformation [Sorkine and Alexa 2007; Jacobson et al. 2012], architectural design [Bouaziz et al. 2012], and physical simulation [Bouaziz et al. 2014], where its fast progress in the initial iterations is paramount to achieving interactive performance.

To better understand how the local/global algorithm works and why it progresses quickly in the first iterations, let us first trivially rewrite it as a descent algorithm. The local step remains the same (Eq. (3)), and then the global step generates a search direction \mathbf{d}^k :

$$\mathbf{d}^k = \tilde{\mathbf{x}}^k - \mathbf{x}^{k-1}, \quad (5)$$

where $\tilde{\mathbf{x}}^k$ is computed as in Eq. (4) and \mathbf{x}^{k-1} is the previous iterate. Then the algorithm simply takes a full step in the search direction:

$$\mathbf{x}^k = \mathbf{x}^{k-1} + \alpha \mathbf{d}^k, \quad \alpha = 1, \quad (6)$$

and the iterations are continued for the next value $k := k + 1$.




Fig. 5. The ARAP energy (left) introduces 6K inverted triangles in the parametrization (highlighted in magenta), which cause highly distorted regions around the neck and the wing of the Gargoyle. Our algorithm avoids this problem by minimizing a flip-preventing symmetric Dirichlet energy (right). This model has 99K faces and took our algorithm 20 iterations and 3 seconds to optimize.

We observe that Eq. (4) minimizes the energy (1) with the ARAP distortion in (2) for a specific set of rotations. We suggest another point of view by introducing a family of proxy functions

$$\mathcal{P}^{\mathbf{R}_f}(\mathbf{J}_f) = \|\mathbf{J}_f - \mathbf{R}_f\|_F^2 \quad (7)$$

and equivalently rewriting Eq. (4) as

$$\mathbf{x}^k = \arg \min_{\mathbf{x}} \sum_{f \in F} A_f \mathcal{P}^{\mathbf{R}_f^k}(\mathbf{J}_f(\mathbf{x})) + \lambda \|\mathbf{x} - \mathbf{x}^{k-1}\|. \quad (8)$$

Note that the proxy in Eq. (7) is almost identical to the ARAP energy itself in Eq. (2), and they are in fact equal at the current iterate. Furthermore, we observe that

$$\nabla_{\mathbf{J}} \mathcal{P}^{\mathbf{R}_f^k}(\mathbf{J}) = \nabla_{\mathbf{J}} \mathcal{D}_{\text{ARAP}}(\mathbf{J}), \quad (9)$$

that is, the proxy and the energy have matching gradients w.r.t. the Jacobian. Since $\mathbf{J}(\mathbf{x})$ is a linear function of \mathbf{x} , the following also holds:

$$\nabla_{\mathbf{x}} \mathcal{P}^{\mathbf{R}_f^k}(\mathbf{J}(\mathbf{x})) = \nabla_{\mathbf{x}} \mathcal{D}_{\text{ARAP}}(\mathbf{J}(\mathbf{x})). \quad (10)$$

Since the proxy is quadratic and the gradients match, \mathbf{d}^k must point towards a descent direction for Eq. (2).

The reason for the local/global algorithm’s superior performance when compared with gradient descent or Newton iterations is that the quadratic proxy used by the local/global algorithm is customized to be a close fit to the original energy in a more *global* sense, while generic optimization methods only fit locally. The rotations \mathbf{R}_f^k are the parameters of the proxy and are strongly tied to the problem at hand, instead of being coefficients of a generic polynomial.

Flip-preventing energies. We use the observation above to extend the local/global approach to minimize energies other than ARAP, such as the challenging flip-preventing distortion energies, i.e., energies that become infinite when an element degenerates. We show in Sec. 5 that our algorithm can be easily used to minimize any rotation invariant distortion energy, but, for the sake of clarity, we first restrict our discussion to isometric distortion energies, such as the *symmetric* Dirichlet energy [Schreiner et al. 2004; Smith and Schaefer 2015]:

$$\mathcal{D}(\mathbf{J}_f(\mathbf{x})) = \|\mathbf{J}_f(\mathbf{x})\|_F^2 + \|\mathbf{J}_f^{-1}(\mathbf{x})\|_F^2. \quad (11)$$

The key difference of our approach to all existing methods is that we rely on an ad-hoc proxy function to find the descent direction at every step, instead of using the gradient descent direction or relying on the generic quadratic approximation used in Newton iterations.

Weighted proxy functions. The challenge in extending the local/global method to a more general energy \mathcal{D} is to define a proxy function \mathcal{P} that satisfies a similar condition to Eq. (9), that is, it has the same gradient as the energy itself. More concisely, that

$$\nabla_{\mathbf{J}} \mathcal{P}(\mathbf{J}) = \nabla_{\mathbf{J}} \mathcal{D}(\mathbf{J}). \quad (12)$$

This condition is important, since it directly implies that every descent direction obtained using the proxy is also a descent direction for \mathcal{D} . We propose the following family of proxy functions:

$$\mathcal{P}_{\mathbf{W}}^{\mathbf{R}}(\mathbf{J}) = \|\mathbf{W}(\mathbf{J} - \mathbf{R})\|_F^2, \quad (13)$$

where \mathbf{W} is a 2×2 matrix. The idea is to use something that is similar to the ARAP proxy (Eq. (7)), with the only difference being the term \mathbf{W} , which is the affine transformation required to warp the proxy to locally match the gradient of the distortion energy. This enhanced proxy energy is quadratic (with fixed \mathbf{R} and \mathbf{W}) and can




Fig. 6. Minimization of the symmetric Dirichlet energy on the Octopus model. Note how quickly the boundary of the UV map recovers from the distorted starting point. This model has 299K faces and took 20 iteration and 5.6 seconds to optimize.

be rewritten in matrix form and minimized (see details in Appendix A).

Given \mathbf{x}^{k-1} and \mathbf{R}_f^k , we can find a weight matrix \mathbf{W}_f^k , such that Eq. (12) holds for $\mathcal{P}_{\mathbf{W}}^{\mathbf{R}}$:

$$\nabla_{\mathbf{J}_f} \|\mathbf{W}_f^k (\mathbf{J}_f - \mathbf{R}_f^k)\|_F^2 = \nabla_{\mathbf{J}_f} \mathcal{D}(\mathbf{J}_f), \text{ at } \mathbf{J}_f = \mathbf{J}_f(\mathbf{x}^{k-1}). \quad (14)$$

For brevity, we drop the indices and simply write Eq. (14) as $\nabla_{\mathbf{J}} \|\mathbf{W}(\mathbf{J} - \mathbf{R})\|_F^2 = \nabla_{\mathbf{J}} \mathcal{D}(\mathbf{J})$. We can then expand the l.h.s:

$$\begin{aligned} \nabla_{\mathbf{J}} \|\mathbf{W}(\mathbf{J} - \mathbf{R})\|_F^2 &= \\ \nabla_{\mathbf{J}} \text{tr}(\mathbf{W}^\top \mathbf{W}(\mathbf{J} - \mathbf{R})(\mathbf{J} - \mathbf{R})^\top) &= \\ (\mathbf{W}^\top \mathbf{W} + \mathbf{W} \mathbf{W}^\top)(\mathbf{J} - \mathbf{R}). \end{aligned}$$

Assuming $\mathbf{J} - \mathbf{R}$ is invertible, we can manipulate this equation and find that \mathbf{W} must satisfy

$$\mathbf{W}^\top \mathbf{W} + \mathbf{W} \mathbf{W}^\top = \nabla_{\mathbf{J}} \mathcal{D}(\mathbf{J}) (\mathbf{J} - \mathbf{R})^{-1}. \quad (15)$$

If $(\mathbf{J} - \mathbf{R})^{-1}$ does not exist, we take the pseudo-inverse instead. Since the left-hand side is semi-positive definite, a solution to Eq. (15) exists if and only if the right hand-side is semi-positive definite, which is always the case for isometric energies such as the one in Eq. (11) (see Sec. 5 for more details). The unique solution can be found explicitly:

$$\mathbf{W} = \left(\frac{1}{2} \nabla_{\mathbf{J}} \mathcal{D}(\mathbf{J}) (\mathbf{J} - \mathbf{R})^{-1} \right)^{1/2}, \quad (16)$$

where the square root is a matrix principal root, which we compute via diagonalization. Alternatively, this equation can be rewritten to sidestep the diagonalization to improve performance, as detailed in Sec. 5 (see Eq. (23)).

Note that $\nabla_{\mathbf{J}} \mathcal{D}_{\text{ARAP}}(\mathbf{J}) = 2(\mathbf{J} - \mathbf{R})$, reducing the previous equation to $\mathbf{W} = \mathbf{I}$. Thus, our algorithm reduces to [Liu et al. 2008] when used on $\mathcal{D}_{\text{ARAP}}$.

Line search. The proxy function \mathcal{P} allows us to find a good descent direction during the minimization: at each step our algorithm alternates the computation of optimal rotations (local step) with the minimization of the proxy function to find a descent direction. Similarly to Eq. (8), we use the weighted proxy and solve,

$$\mathbf{p}^k := \arg \min_{\mathbf{x}} \sum_{f \in F} A_f \mathcal{P}_{\mathbf{W}_f^k}^{\mathbf{R}_f^k}(\mathbf{J}_f(\mathbf{x}^{k-1})) + \lambda \|\mathbf{x} - \mathbf{x}^{k-1}\|. \quad (17)$$

Then we use \mathbf{p}^k to define the search direction,

$$\mathbf{d}^k := \mathbf{p}^k - \mathbf{x}^{k-1} \quad (18)$$

and we use this direction for the next step,

$$\mathbf{x}^k = \mathbf{x}^{k-1} + \alpha \mathbf{d}^k,$$

where α is the step size to be determined by a line search minimizing the *original* energy (1), *without* the proximal term. To ensure that the generated map is flip-free, we start from a locally injective map (Tutte’s embedding) and we use a line-search strategy similar to [Smith and Schaefer 2015]. The idea is to choose α such that we never cross the point where an element flips. We thus initially set $\alpha = \min\{0.8\alpha_{\max}, 1\}$, where α_{\max} is the maximal step size with no foldovers. We then proceed with a bisection line search, where the initial interval is $[0, \alpha_{\max}]$, until we find a point that satisfies Wolfe’s conditions (see [Nocedal and Wright 2006], section 3.4). Our algorithm is guaranteed to generate a descent direction at every iteration, and therefore decrease the energy at each iteration. We have found experimentally that the cosine of the angle between \mathbf{d}^k and the gradient of E is always bounded away from zero, although we lack a formal proof. Assuming this is always true, then by the Zoutendijk theorem ([Nocedal and Wright 2006] section 3.2) global convergence could be guaranteed.

Reweighted local/global for isometric energies. Enriching the local/global algorithm with the matrix reweighting scheme and line search leads to our algorithm, which is simple to implement and scales to datasets with millions of elements. Algorithm 1 sketches the pseudocode of our algorithm, and we provide the source code of our reference implementation on GitHub [Rabinovich 2016].

We analyze the method’s properties in Sec. 4 and we show how it can be easily extended to various distortion energies and to 3D locally injective maps in Sec. 5.

Algorithm 1: Reweighted local/global

Input:

A mesh M with a set of vertices V and elements F

Output:

A set of mapping coordinates \mathbf{x} minimizing Eq. (1)

Initialization:

$\mathbf{x}^0 = \text{Tutte}(V, F)$

Optimization:

for $k = 1$ **to** max_iterations **do**

 Compute closest rotation \mathbf{R}_f^k for each Jacobian $\mathbf{J}_f(\mathbf{x}^{k-1})$

 Update the weights \mathbf{W}_f^k for each $f \in F$ using Eq. (16)

 Solve Eq. (17) to find \mathbf{p}^k

 Set \mathbf{d}^k using Eq. (18)

 Find α_{\max} , as in [Smith and Schaefer 2015](Section 3.3)

 Define $\alpha = \min\{1, 0.8\alpha_{\max}\}$ and interval $[0, \alpha_{\max}]$

 Perform bisection line search in the interval to find α

$\mathbf{x}^k = \mathbf{x}^{k-1} + \alpha \mathbf{d}^k$




Fig. 7. Minimization of the symmetric Dirichlet energy on the Buddha model. Despite the massive the dataset (470K faces), our algorithm produces an optimized locally injective parametrization in 14 seconds.

4. ANALYSIS

We apply our algorithm to a series of challenging parametrization problems (using the symmetric Dirichlet energy) to observe its properties and scalability to large datasets. We also provide a comparison against Newton descent and L-BFGS [Nocedal and Wright 2006] to demonstrate that our algorithm outperforms classical optimization techniques.

Qualitative convergence behavior. We observed that the convergence behavior in our case is similar to local/global optimization of the ARAP energy, that is, the algorithm progresses very rapidly at the beginning, but then slows down when close to a minimum (see Figures 6, 7, 8). Uneven meshing does not alter the performance of our method (Fig. 9). We demonstrate this by parametrizing two copies of the same surface, one with an homogeneous triangulation and another with an uneven triangulation with hundreds of slivers.




Fig. 8. Our algorithm quickly progresses towards the minimum of the energy. If sufficiently many iterations are executed, our algorithm numerically converges to a minimum, but the differences in quality compared to earlier iterations are negligible. This figure shows the result after 10, 20, and finally, 196 iterations, where the optimization has converged to a minimum.

Our algorithm is extremely robust, and it succeeds even when initialized with a map made of slivers, as demonstrated in Fig. 16 — the resulting parameterization after 20 iterations is very close to being isometric, as can be seen in the texture, even if the algorithm needs many more iterations to reach a numerical minimum. Newton’s method is much slower, and only after 2000 iterations does the quality of the parameterization start to resemble ours.

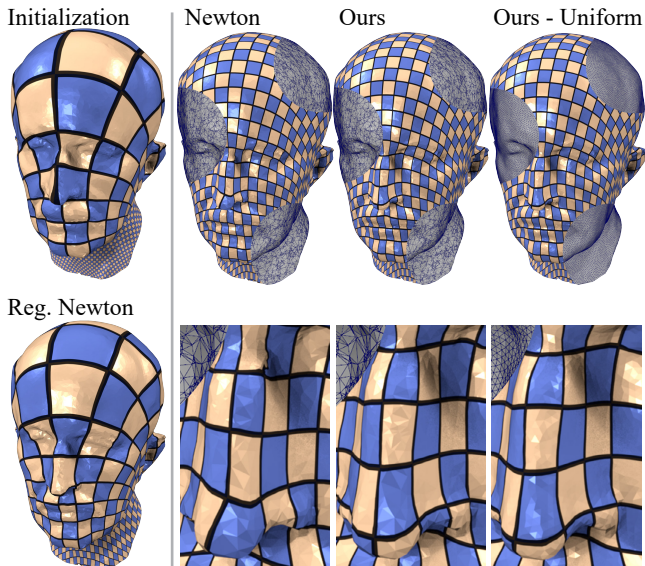


Fig. 9. Our algorithm is robust and consistently produces high quality locally-injective maps in as little as 20 iterations, even in presence of low-quality triangulations. Note that our algorithm reaches similar minima regardless of the mesh quality. In this case, Newton’s method stopped before reaching a minimum, while the regularized Newton did very little progress after 500 iterations.

Comparison with standard numerical methods. In Fig. 10 (top) we compare our iterations against L-BFGS. Since the L-BFGS iterations are faster than ours, we plot the time on the x-axis to ensure a fair comparison. Our method is considerably faster and produces a high quality map in a fraction of the time. In Fig. 10 (bottom), we compare against Newton’s method. This method can get stuck if the Hessian is not positive semidefinite (PSD). A common remedy in simulation applications is to regularize the Hessian by projecting the Hessian of each face on the set of PSD matrices (see e.g. [Teran et al. 2005]). We observe that for parameterization it is not clear which method performs better, as cases vary. We report the results of both the normal and the regularized version of Newton’s method. In both cases, our algorithm is considerably faster. Another example of this behavior is shown in Fig. 16.

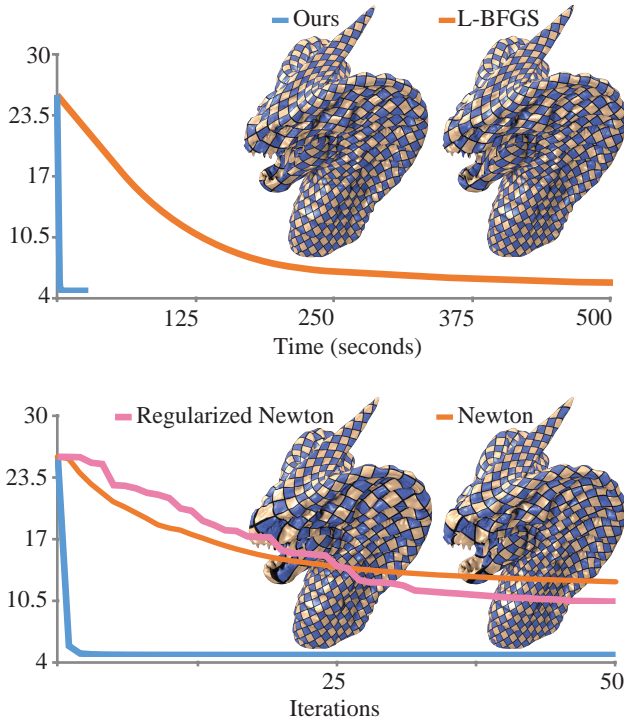



Fig. 10. Our iterations are slower than those of L-BFGS, since we have to solve a sparse linear system, but overall, they progress much faster (top). Our iterations are also making considerably more progress than Newton iterations when we are far from the minimum (bottom). This model has 386K faces, and in both cases we minimize the symmetric Dirichlet energy. See Fig. 16 for another comparison with Newton’s method.

Scalability. The efficiency of our method stems from the experimental observation that the number of iterations is related to the geometric complexity of the model, instead of depending on the density of the tessellation. We show evidence supporting this claim in Fig. 1, where we plot the energy on a progressive input mesh (Lucy), sampled with 0.6-6 million faces. Note that the majority of the competing methods become impractically slow for models larger than 100K vertices. The only exception is [Kovalsky et al. 2015], but this method does not minimize a distortion energy, often fails to find a solution and cannot be used to create isometric parameterizations [Kovalsky 2016].

Why not ARAP with Line search?. Our approach relies on the line search proposed by [Smith and Schaefer 2015] to ensure the local injectivity of the map. It is tempting to simply add the line search to e.g. ARAP parametrization algorithm, obtaining a method that is also guaranteed to produce locally injective maps. However, the quality of the map will be poor because the line-search will lock the algorithm whenever one triangle degenerates, and further progress will not be allowed.

In the inset, we repeat the experiment shown in Fig. 4 using the ARAP parametrization algorithm with the line search proposed by [Smith and Schaefer 2015]. The algorithm is not able to make much progress, producing a result with a much higher distortion than ours (Fig. 4).



Relationship to IRLS and greedy stiffening. Our approach is loosely related to the stiffening methods, which are often used in global parameterization [Bommes et al. 2009], where Eq. (4) is modified to include a per-element scalar weight w_f , i.e.,

$$\min_{\mathbf{x}} \sum_{f \in F} A_f w_f^k \|\mathbf{J}_f(\mathbf{x}) - \mathbf{R}_f^k\|_F^2. \quad (19)$$

The rationale is to increase this weight for the elements that flip, hoping that this will lead to a locally injective map. This approach is a heuristic that is not guaranteed to succeed, since increasing the weight of an element to remove a flip could cause a flip in a neighboring element, and this could (and does in practice) repeat indefinitely. We tried using this approach in our formulation, replacing our matrix weights with a scalar weight and updating it using the update rule proposed in [Bommes et al. 2009], but this is not guaranteed to generate a descent direction, and the algorithm gets stuck, as shown in Fig. 4.

Another related algorithm is iterative reweighted least squares (IRLS), which also uses scalar weights for the purpose of locally matching the energy value. This is sufficient to ensure convergence for the specific case of L^p norms. We tried to adapt this idea to our framework, using scalar weights and the update rule proposed in [Yoshizawa et al. 2004; Pighin and Lewis 2007]:

$$w_f^k = \frac{\mathcal{D}(\mathbf{J}_f(\mathbf{x}^{k-1}))}{\|\mathbf{J}_f(\mathbf{x}^{k-1}) - \mathbf{R}_f^k\|_F^2}. \quad (20)$$

If, in a given iteration, an element is mapped to an almost degenerate one, the symmetric Dirichlet energy in the numerator becomes huge, while the proxy in the denominator does not, and so in the next iteration the IRLS scalar weight strategy penalizes more on this element by giving it a larger weight. This could prevent it from flipping in the next iteration, and one could expect it to produce a result that is closer to the actual energy that we would like to minimize. If an element is scaled, we could expect the weight to be lower than 1, since the proxy tends to penalize it too much. However, our experiments show that this update rule (which turned out to be conceptually similar to the stiffening of [Bommes et al. 2009]) is not sufficient to guarantee convergence, and in fact eventually often fails to find a descent direction (Fig. 4).

Intuitively, the expressivity of scalar weights is limited, for instance when a triangle is stretched very heavily along just one axis, while not being stretched along the other axis. In this case, a single scalar weight per-triangle cannot distinguish between the axes, and both of them are penalized the same. A more proper choice is to penalize each axis separately, and this can only be expressed by using matrix weights. We elaborate on this in the next section.

5. GENERAL DISTORTION ENERGIES

We proceed with the treatment of specific isometric distortion measures, as well as generalizing to other types of distortions. From now on, we assume that $\mathcal{D}(\mathbf{J})$ is any rotation invariant distortion measure, as we define below. Additionally, we fill the gap that remains from the previous section, that is, ensuring that Eq. (15) is indeed valid, and that its right hand side is positive semidefinite. We start with the latter, as it naturally leads to the derivation of the relevant formulas.

SVD viewpoint. A common property of many useful distortion measures is that they are *rotation invariant*:

Definition 1. A distortion measure $\mathcal{D}(\mathbf{J})$ is *rotation invariant* if

$$\mathcal{D}(\mathbf{J}) = \mathcal{D}(\mathbf{U}\mathbf{J}\mathbf{V}^T)$$

for any rotation matrices \mathbf{U} and \mathbf{V} .

All rotation invariant distortion measures can be written solely in terms of the singular values of \mathbf{J} , as shown in the following lemma.

LEMMA 2. Let $\mathbf{J} = \mathbf{U}\mathbf{S}_\mathbf{J}\mathbf{V}^\top$ be the singular value decomposition of \mathbf{J} (we use $\mathbf{S}_\mathbf{A}$ to denote the diagonal matrix containing the singular values of \mathbf{A}). Then,

$$\mathcal{D}(\mathbf{J}) = \mathcal{D}(\mathbf{S}_\mathbf{J}), \quad (21)$$

$$\nabla_{\mathbf{J}}\mathcal{D}(\mathbf{J}) = \mathbf{U}\nabla_{\mathbf{S}_\mathbf{J}}\mathcal{D}(\mathbf{S}_\mathbf{J})\mathbf{V}^\top. \quad (22)$$

Proof. See Appendix B.1.

For example, the ARAP distortion measure, which is rotation invariant, can be written as

$$\|\mathbf{J} - \mathbf{R}\|_F^2 = \|\mathbf{S}_\mathbf{J} - \mathbf{I}\|_F^2 = \sum_{i=1}^d (\sigma_i - 1)^2,$$

where σ_i are the singular values of \mathbf{J} .

Using Lemma 2, we can write Eq. (16) in SVD form:

$$\mathbf{W} = \mathbf{U} \left(\frac{1}{2} \nabla_{\mathbf{S}_\mathbf{J}} \mathcal{D}(\mathbf{S}_\mathbf{J}) (\mathbf{S}_\mathbf{J} - \mathbf{I})^{-1} \right)^{1/2} \mathbf{U}^\top = \mathbf{U}\mathbf{S}_\mathbf{W}\mathbf{U}^\top. \quad (23)$$

This expression shows that the matrix weights in fact act as scalar weights per singular value of each Jacobian, that is, a single weight-per-axis. Since the matrix expression inside the square root is diagonal, it is trivial to take the root. From Eq. (23) we can clearly see that

$$\nabla_{\mathbf{S}}\mathcal{D}(\mathbf{S}_\mathbf{J})(\mathbf{S}_\mathbf{J} - \mathbf{I}) \succeq 0 \quad (24)$$

is required in order to take the root. In other words, $\nabla_{\mathbf{S}_\mathbf{J}}\mathcal{D}(\mathbf{S}_\mathbf{J})(\mathbf{S}_\mathbf{J} - \mathbf{I})$ must be PSD, which is always the case for isometric distortion measures, as we show in the following.

Isometric energies. We applied our algorithm to a variety of isometric distortion measures. In our context, a *true* isometric distortion measure is one that is rotationally invariant, minimal *only* for rotations, and is separable in terms of the singular values. In other words, a true isometric distortion measure can be written as

$$\mathcal{D}(\mathbf{S}_\mathbf{J}) = \mathcal{D}(\sigma_1, \dots, \sigma_d) = \sum_i f_i(\sigma_i), \quad (25)$$

where $f_i(\sigma_i)$ convex and minimal when $\sigma_i = 1$. This condition makes it is easy to show that Eq. (24) holds.

Without loss of generality, we demonstrate the solution to Eq. (23) using the symmetric Dirichlet energy. The same steps can be used for any other isometric distortion measures (see Table I and Figure 11). The symmetric Dirichlet energy in terms of the singular values is

$$\mathcal{D}(\mathbf{J}) = \|\mathbf{J}\|_F^2 + \|\mathbf{J}^{-1}\|_F^2 = \sum_{i=1}^d (\sigma_i^2 + \sigma_i^{-2}).$$

Hence,

$$(\nabla_{\mathbf{S}_\mathbf{J}}\mathcal{D}(\mathbf{S}_\mathbf{J}))_i = 2(\sigma_i - \sigma_i^{-3}), \quad (26)$$

where we introduce the notation $(\mathbf{D})_i$ to refer to the i 'th diagonal entry of the diagonal matrix \mathbf{D} . Thus, plugging Eq. (26) in Eq. (23) we obtain

$$(\mathbf{S}_\mathbf{W})_i = \sqrt{\frac{\sigma_i - \sigma_i^{-3}}{\sigma_i - 1}} \quad (27)$$




Fig. 11. Minimizing isometric distortions for the Bear model. Our approach is general and supports many distortion energies such as Symmetric Dirichlet (left), Hencky strain (middle) and Exponential Dirichlet (right). This model has 296K faces and required an average of 9 seconds to optimize.

when $\sigma_i \neq 1$, and $(\mathbf{S}_\mathbf{W})_i = 4$, which is the limit, otherwise. The expression under the root is always nonnegative, so we can always find weights to match the gradients.

General distortion measures. To extend this construction to energies that are rotation invariant, but not true isometric distortions (i.e., do not satisfy Eq. (25)), we need to slightly change our algorithm. Consider for example the AMIPS energy E_{iso}^* from [Fu et al. 2015] (see the paper for the reasoning behind this definition), that is defined using the distortion measure

$$\mathcal{D}_{iso}^*(\mathbf{J}) = \exp(s \cdot \mathcal{D}_{iso}(\mathbf{J})) \quad (28)$$

where

$$\mathcal{D}_{iso}(\mathbf{J}) = \frac{1}{2} \left[\left(\frac{\text{tr}(\mathbf{J}^\top \mathbf{J})}{\det(\mathbf{J})} \right) + \frac{1}{2} (\det(\mathbf{J}) + \det(\mathbf{J}^{-1})) \right]. \quad (29)$$

Algorithm 1 fails to minimize this energy, because (24) might not hold, and it thus may be impossible to compute weights that locally match the gradients. To address this issue, we generalize the local step (Eq. (3)) by replacing the closest rotation \mathbf{R} with another matrix Λ , which depends on the energy that we wish to minimize. Note that the local step is only used in the proxy energy, which becomes

$$\mathcal{P}_{\mathbf{W}}^\Lambda(\mathbf{J}) = \|\mathbf{W}(\mathbf{J} - \Lambda)\|_F^2. \quad (30)$$

The modification of the local step changes the formula for \mathbf{W} in Eq. (23), replacing the identity with the matrix \mathbf{S}_Λ , and the requirement in Eq. (24) becomes

$$\nabla_{\mathbf{S}_\mathbf{J}}\mathcal{D}(\mathbf{S}_\mathbf{J})(\mathbf{S}_\mathbf{J} - \mathbf{S}_\Lambda) \succeq 0. \quad (31)$$

The condition on the local step is now clear: The sign of each element of $\mathbf{S}_\mathbf{J} - \mathbf{S}_\Lambda$ must match the sign of each element of $\nabla_{\mathbf{S}_\mathbf{J}}\mathcal{D}(\mathbf{S}_\mathbf{J})$. The computation in the local step can then be easily adapted on a case-by-case basis. Note that this construction is a generalization of [Liu et al. 2008].

General construction. Given an arbitrary distortion energy, it is easy to find a matrix Λ that satisfies Eq. (31). For the special (and very general) case of *separably strictly convex* energies, we found a simple construction.

Definition 3. A function $\mathcal{D}(x_1, \dots, x_n)$ is *separably strictly convex*, if for each i the single variable function $\mathcal{D}(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)(x_i)$, which is constructed by freezing all of the other variables, is *strictly convex*.

In this case, we can set each $(\mathbf{S}_\Lambda)_i$ to be the solution of

$$\frac{\partial}{\partial \sigma_i} \mathcal{D}(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_d)(\sigma_i) = 0. \quad (32)$$

In other words, we set \mathbf{S}_Λ such that each entry of \mathbf{S}_Λ minimizes the corresponding entry of $\mathcal{D}(\mathbf{S}_\mathbf{J})$, assuming the other singular values are fixed. We show in Appendix B.2 that this choice always satisfies Eq. (31). The reasoning behind this choice is to always push each singular value towards its closest minimum. It is also a direct generalization of the isometric case shown above, as the choice of the closest rotation exactly satisfies Eq. (32).

As an example, we derive the expressions for the AMIPS energy (Eq. (28)) in Table I and show the resulting parameterization in Fig. 3. The complete algorithm is summarized in Algorithm 2.

Conformal distortions. While the above construction is general and guaranteed to create weights that match the gradients, it might be suboptimal for certain energies. For example, if we consider the conformal distortion energy proposed (for arbitrary dimension d) in [Fu et al. 2015]

$$\mathcal{D}(\mathbf{J}) = \frac{\text{tr}(\mathbf{J}^\top \mathbf{J})}{\det(\mathbf{J})^{2/d}}, \quad (33)$$

our general construction produces a matrix in the local step that is not a similarity. We experimentally observed an increase in performance by defining $\Lambda = \bar{\sigma}UV^\top$, where $\bar{\sigma}$ is a scalar, i.e., by finding the closest conformal transformation. We show in Appendix B.3 how to compute $\bar{\sigma}$ for the 2D and 3D case. Similarly to [Fu et al. 2015], we can also minimize the exponential of this energy (Fig. 14).

Algorithm 2: Generalized Reweighted local/global

Input:

A mesh M with a set of vertices V and elements F

Output:

A set of mapping coordinates \mathbf{x} minimizing Eq. (1)

Initialization:

$\mathbf{x}^0 = \text{Tutte}(V, F)$

Optimization:

for $k = 1$ **to** *max_iterations* **do**

 Compute the SVD $\mathbf{US}_\mathbf{J}\mathbf{V}^\top = \mathbf{J}_f(\mathbf{x}^{k-1})$ for each face f

 Update $\mathbf{W}_f^k = \mathbf{US}_\Lambda\mathbf{U}^\top$, taking \mathbf{S}_Λ from Table I

 Solve Eq. (17) to find \mathbf{p}^k

 Set \mathbf{d}^k using (18)

 Find α_{\max} , as in [Smith and Schaefer 2015](Section 3.3)

 Define $\alpha = \min\{1, 0.8\alpha_{\max}\}$ and interval $[0, \alpha_{\max}]$

 Perform bisection line search in the interval to find α

$\mathbf{x}^k = \mathbf{x}^{k-1} + \alpha \mathbf{d}^k$

6. APPLICATIONS

We ran our experiments on a 12-core Xeon clocked at 2.7 GHz, using the PARDISO solver [Schenk et al. 2007; Schenk et al. 2008; Kuzmin et al. 2013] for the linear system solve. We report the running times in Table II, using both a single- and multi-core implementation. The sparsity pattern of the linear system in every iteration never changes, allowing us to reuse the symbolic factorization between iterations. Our method requires a feasible, i.e., inversion free starting point: for 2D, we use Tutte’s parameterization with cotangent weights; if they produce a flipped element, we resort to uniform weights which are guaranteed to give us a valid starting point. For the mesh improvement and deformation examples in 3D, the rest pose is used as the starting point.




Fig. 12. Minimization of a conformal energy using our method. Note that visually the difference between 10 and 20 iterations is already quite small, and after 20 it becomes negligible.




Fig. 13. An example of “seamless” global parameterization computed with our method. Starting from Tutte’s embedding (left), we minimize the symmetric Dirichlet energy first, and then activate seamless soft constraints to make the parametrization’s derivatives match on the seams, up to permutation.

Single-patch 2D parametrization. Single-patch 2D parametrization is ubiquitously used in modeling software to generate UV maps given a predefined set of cuts. Our approach improves upon existing algorithms, providing superior quality and higher efficiency, in addition to supporting extremely detailed models with millions of elements. We show the quality of our results on several meshes with different energies throughout the paper. We highlight that our method is not tied to a specific parametrization energy, and can minimize all rotation invariant energies, provided that the local step satisfies Eq. (31).

Seamless 2D parametrization. “Seamless” parameterizations, i.e., parameterizations whose derivatives agree on the seams up to a rotation of multiples of $\pi/2$, are commonly used for remeshing purposes [Bommes et al. 2012]. Our algorithm can generate them by adding compatibility constraints [Bommes et al. 2009] to the seams, as shown in Fig. 13. Note that we currently do not support integer optimization in our framework and we thus cannot produce integer-grid maps [Bommes et al. 2013].

Table I. Energies we used in this paper, expressed also in terms of the singular values, with their derivatives and our choice for the local step. The entries of this table are used in Algorithm 2.

Name	$\mathcal{D}(\mathbf{J})$	$\mathcal{D}(\sigma)$	$(\nabla_{\mathbf{S}}\mathcal{D}(\mathbf{S}))_i$	$(\mathbf{S}_{\Lambda})_i$
Symmetric Dirichlet	$\ \mathbf{J}\ _F^2 + \ \mathbf{J}^{-1}\ _F^2$	$\sum_{i=1}^n (\sigma_i^2 + \sigma_i^{-2})$	$2(\sigma_i - \sigma_i^{-3})$	1
Exponential Symmetric Dirichlet	$\exp(s(\ \mathbf{J}\ _F^2 + \ \mathbf{J}^{-1}\ _F^2))$	$\exp(s \sum_{i=1}^n (\sigma_i^2 + \sigma_i^{-2}))$	$2s(\sigma_i - \sigma_i^{-3}) \exp(s(\sigma_i^2 + \sigma_i^{-2}))$	1
Hencky strain	$\ \log \mathbf{J}^T \mathbf{J}\ _F^2$	$\sum_{i=1}^n (\log^2 \sigma_i)$	$2(\frac{\log \sigma_i}{\sigma_i})$	1
AMIPS	$\exp(s \cdot \frac{1}{2} (\frac{\text{tr}(\mathbf{J}^T \mathbf{J})}{\det(\mathbf{J})} + \frac{1}{2} (\det(\mathbf{J}) + \det(\mathbf{J}^{-1}))))$	$\exp(s(\frac{1}{2}(\frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1}) + \frac{1}{4}(\sigma_1\sigma_2 + \frac{1}{\sigma_1\sigma_2})))$	$s \cdot \exp(s \cdot (\frac{1}{4}(\sigma_{i+1} - \frac{1}{\sigma_{i+1}\sigma_i^2}) + \frac{1}{2}(\frac{1}{\sigma_{i+1}} - \frac{\sigma_{i+1}}{\sigma_i^2})))$	$\sqrt{\frac{2\sigma_{i+1}^2+1}{\sigma_{i+1}^2+2}}$
Conformal AMIPS 2D	$\frac{\text{tr}(\mathbf{J}^T \mathbf{J})}{\det(\mathbf{J})}$	$\frac{\sigma_1^2 + \sigma_2^2}{\sigma_1\sigma_2}$	$\frac{1}{\sigma_{i+1}} - \frac{\sigma_{i+1}}{\sigma_i^2}$	$\sqrt{\sigma_1\sigma_2}$
Conformal AMIPS 3D	$\frac{\text{tr}(\mathbf{J}^T \mathbf{J})}{\det(\mathbf{J})^{\frac{2}{3}}}$	$\frac{\sigma_1^2 + \sigma_2^2 + \sigma_3^2}{(\sigma_1\sigma_2\sigma_3)^{\frac{2}{3}}}$	$\frac{-2\sigma_{i+1}\sigma_{i+2}(\sigma_{i+1}^2 + \sigma_{i+2}^2 - 2\sigma_i^2)}{(3\sigma_i\sigma_{i+1}\sigma_{i+2})^{\frac{2}{3}}}$	$\sqrt{\frac{\sigma_1^2 + \sigma_3^2}{2}}$

Table II. Size of the input datasets and running times with and without multi-core parallelization.

Model (Figure)	#Faces	#Vertices	Time (s)	Time (s)
			Multi core	Single core
Bimba (3)	11K	5.6K	< 1	< 1
Bunny (4)	108K	54K	3.7	9.8
Gargoyle (5)	99K	49K	2.8	6.5
Octopus (6)	300K	151K	5.6	18.4
Superman (8)	190K	95K	5.2	14.6
Buddha (7)	470K	235K	13.9	43.2
Max Planck (9)	84K	42K	2.2	6.2
Dragon head (10)	387K	194K	12.3	36.8
Bear (11)	296K	148K	9.1	25.2
Vase Lion (12)	99K	49K	2.6	6.2




Fig. 14. Our method can be used to deform tetrahedral meshes, minimizing a conformal energy (middle) or the exponential symmetric Dirichlet isometric energy (right). The cubes have 48K (top) and 350K (bottom) tetrahedra, and our algorithm took 5 and 80 seconds, respectively. We picked 4 edges of the cube (shown as cylinders) and manipulated them. The right image in each pair shows the interior of the deformed cube.

3D deformation. Volumetric deformation energies can be minimized with our method, benefiting in a similar way as 2D parametrization. In Fig. 14, we demonstrate an example of a cube deformation with two different discretization resolutions (48K and

250K tetrahedra) and two different distortion energies (exponential Dirichlet and exponential of the conformal AMIP). The running time of our algorithm is 0.5 and 8 seconds per iteration, respectively, and we used 10 iterations.

Mesh improvement. Maps can be used to improve the quality of meshes, and we compare our algorithm against two recent methods [Aigerman and Lipman 2013; Fu et al. 2015] in Fig. 15. We use the exponential Dirichlet energy, which is inspired by the exponential AMIPS energy proposed in [Fu et al. 2015]. As can be seen in Table III, our approach outperforms the competing methods in all cases except one. We performed 10 iterations for all models, which on average took 3 seconds each.

Table III. Comparison of mesh improvement.

Name	Init. Dihed.	BD	AMIPS	Our Method
Duck	(10,163)	(16,148)	(19.6,161.5)	(19, 138.16)
Elephant	(8,167)	(16,148)	(13.7,161.3)	(20.2,141.2)
Elephant2	(15,157)	(18,147)	(19,150)	(23.1,142.3)
Hand	(9,162)	(16,148)	(18,156)	(21.1, 143.3)
Max	(21,151)	(14,153)	(27.2,137.3)	(29,141.5)
Rocker	(10,163)	(16,148)	(21.6,148.7)	(22.8,139.4)
Skull	(0.8,178)	(14,153)	(21.1,147.6)	(17.4,157.8)
Dragon	(31,140)	(28,139)	(27.8,139.37)	(31.8, 137.6)

Comparison of mesh improvement achieved by running [Aigerman and Lipman 2013] (BD) and [Fu et al. 2015] (AMIPS). In each entry in the table we show the minimal and maximal dihedral angle, where the second column shows the initial values. As can be seen, in all cases except for the Skull dataset, our method outperforms the competing methods.

7. LIMITATIONS AND FUTURE WORK

We presented a general approach to simply and efficiently minimize practical distortion energies commonly used in geometry processing. The major theoretical limitation and advantage of our approach are both inherited from the local/global method: our algorithm is extremely fast while approaching a local minimum but requires many iterations to converge to a numerical minimum. This problem stems from the slow propagation of the rotations in the local step, making a small rotation over a large part of the parameterization hard to recover from. This is however not a practical limitation, since the difference between the result we obtain after 20 iterations and the converged result is negligible (Figures 8, 12). On the other




Fig. 15. An example of 3D mesh improvement computed with our method minimizing the Exponential Dirichlet energy. The elephant has 33.5K tetrahedra and the rocker arm 35.5K. Our entire optimization took 3 and 3.2 seconds, respectively.

hand, the fast progress of our method in the first few iterations is an extremely valuable property for many applications (Fig. 10).

We exemplify this property in a stress test similar to [Smith and Schaefer 2015] (see Fig. 16). The challenge in this test is to recover from Tutte’s embedding of a Hilbert-curve-shaped developable surface, which is flipless, but highly distorted. We observed that our first iteration makes the same progress as around 2000 Newton iterations — by combining the two algorithms, we reconstruct the Hilbert curve in less than 200 iterations.

Our algorithm requires a locally injective initialization, which we construct using Tutte’s embedding in 2D, but we are not aware of any general construction for 3D maps. This limitation is ameliorated by the fact that maps are mostly used for deformations in 3D, where the rest shape itself is a perfectly valid initial map which we directly use as a locally injective initialization.

A final limitation is that our proxy energy definition only works for rotation invariant distortion energies. It is an interesting venue for future work to extend our local/global method to a wider family of energies, such as those used in finite element simulations.

REFERENCES

- AIGERMAN, N. AND LIPMAN, Y. 2013. Injective and bounded distortion mappings in 3D. *ACM Trans. Graph.* 32, 4.
- AIGERMAN, N. AND LIPMAN, Y. 2015. Orbifold Tutte embeddings. *ACM Trans. Graph.* 34, 6.
- AIGERMAN, N., PORANNE, R., AND LIPMAN, Y. 2014. Lifted bijections for low distortion surface mappings. *ACM Trans. Graph.* 33, 4, 69:1–69:12.
- BEN-CHEN, M., GOTSMAN, C., AND BUNIN, G. 2008. Conformal flattening by curvature prescription and metric scaling. *Computer Graphics Forum* 27, 2, 449–458.
- BOMMES, D., CAMPEN, M., EBKE, H.-C., ALLIEZ, P., AND KOBBELT, L. 2013. Integer-grid maps for reliable quad meshing. *ACM Trans. Graph.* 32, 4 (July), 98:1–98:12.
- BOMMES, D., LÉVY, B., PIETRONI, N., PUPPO, E., SILVA, C., TARINI, M., AND ZORIN, D. 2012. State of the art in quad meshing. In *Eurographics STARs*.
- BOMMES, D., ZIMMER, H., AND KOBBELT, L. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3, 77:1–77:10.
- BOUAZIZ, S., DEUSS, M., SCHWARTZBURG, Y., WEISE, T., AND PAULY, M. 2012. Shape-Up: Shaping discrete geometry with projections. *Comput. Graph. Forum* 31, 5, 1657–1667.
- BOUAZIZ, S., MARTIN, S., LIU, T., KAVAN, L., AND PAULY, M. 2014. Projective dynamics: Fusing constraint projections for fast simulation. *ACM Trans. Graph.* 33, 4 (July), 154:1–154:11.
- CHAO, I., PINKALL, U., SANAN, P., AND SCHRÖDER, P. 2010. A simple geometric model for elastic deformations. *ACM Trans. Graph.* 29, 4, 38:1–38:6.
- CHEN, R. AND WEBER, O. 2015. Bounded distortion harmonic mappings in the plane. *ACM Trans. Graph.* 34, 4.
- CIVIT-FLORES, O. AND SUSIN, A. 2014. Robust treatment of degenerate elements in interactive corotational FEM simulations. *Computer Graphics Forum* 33, 6, 298–309.
- DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic parameterizations of surface meshes. *Comput. Graph. Forum* 21, 3.
- DIAMANTI, O., VAXMAN, A., PANOZZO, D., AND SORKINE-HORNUNG, O. 2015. Integrable PolyVector fields. *ACM Trans. Graph.* 34, 4, 38:1–38:12.
- FLOATER, M. S. 2003. One-to-one piecewise linear mappings over triangulations. *Mathematics of Computation* 72, 242, 685–696.
- FLOATER, M. S. AND HORMANN, K. 2005. Surface parameterization: a tutorial and survey. In *Advances in Multiresolution for Geometric Modelling*. Mathematics and Visualization. 157–186.
- FU, X.-M., LIU, Y., AND GUO, B. 2015. Computing locally injective mappings by advanced MIPS. *ACM Trans. Graph.* 34, 4.
- GILES, M. 2008. An extended collection of matrix derivative results for forward and reverse mode automatic differentiation. Tech. rep., Oxford University Computing Laboratory. Jan.
- HORMANN, K. AND GREINER, G. 2000. MIPS: An efficient global parameterization method. In *Curve and Surface Design*. Innovations in Applied Mathematics. 153–162.
- IRVING, G., TERAN, J., AND FEDKIW, R. 2004. Invertible finite elements for robust simulation of large deformation. In *Proc. Eurographics Symposium on Computer Animation*. 131–140.
- JACOBSON, A., BARAN, I., KAVAN, L., POPOVIĆ, J., AND SORKINE, O. 2012. Fast automatic skinning transformations. *ACM Trans. Graph.* 31, 4, 77:1–77:10.
- KOVALSKY, S. 2016. Private communication.
- KOVALSKY, S. Z., AIGERMAN, N., BASRI, R., AND LIPMAN, Y. 2014. Controlling singular values with semidefinite programming. *ACM Trans. Graph.* 33, 4.
- KOVALSKY, S. Z., AIGERMAN, N., BASRI, R., AND LIPMAN, Y. 2015. Large-scale bounded distortion mappings. *ACM Trans. Graph.* 34, 6, 191:1–191:10.
- KUZMIN, A., LUISIER, M., AND SCHENK, O. 2013. Fast methods for computing selected elements of the Green’s function in massively parallel nanoelectronic device simulations. In *Proc. Euro-Par*. Lecture Notes in Computer Science. 533–544.
- LABSIK, U., HORMANN, K., AND GREINER, G. 2000. Using most isometric parametrizations for remeshing polygonal surfaces. In *Proc. Geometric Modeling and Processing*. 220–228.

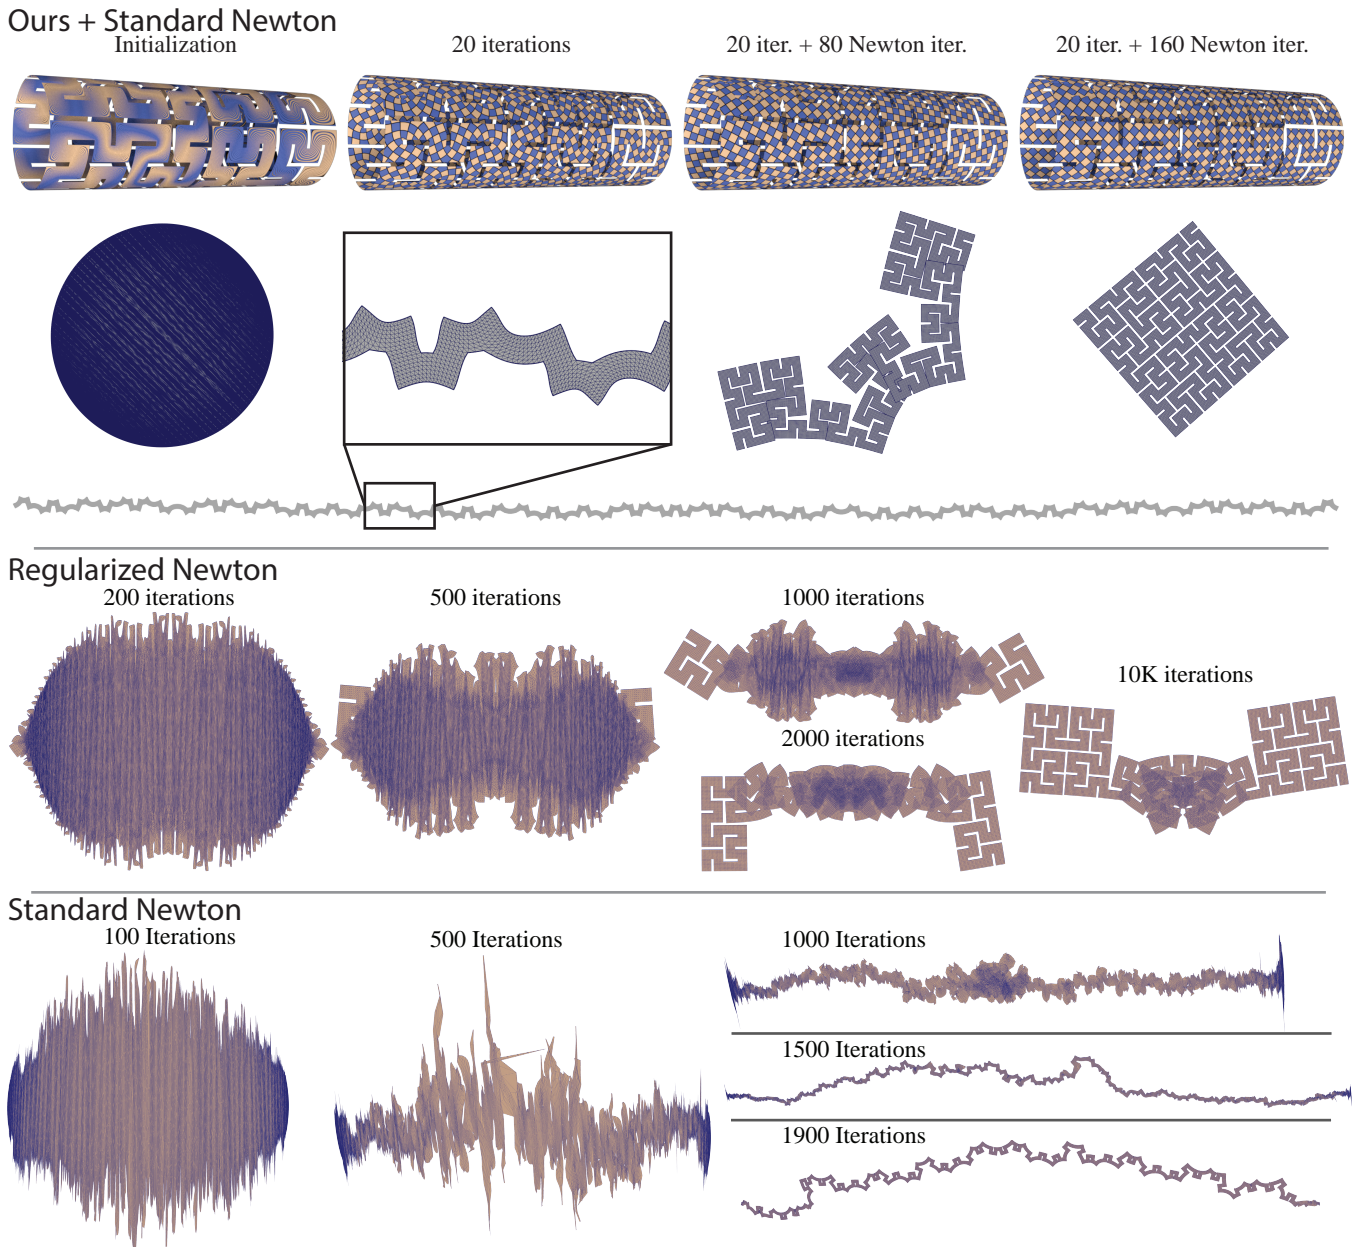


Fig. 16. Example of running our algorithm on a stress test. The highly distorted Tutte’s embedding is quickly unrolled by our algorithm in about 20 iterations, producing a map close to being isometric. However, from this point on the progress of our algorithm slows down considerably, although the minimum is still reached after ca. 500K iterations. The regularized Newton method is much slower (middle row) and the system is severely ill-conditioned due to the extremely distorted triangles. The standard, non-regularized Newton fails to find a descent direction after 1900 iterations. By combining our method with the standard Newton’s method, we are able to reach the global minimum in a total of 180 iterations. Note that we used Newton iterations with our method as an initialization *only* in this figure. All the other results in the paper are produced using solely our algorithm.

LEVI, Z. AND ZORIN, D. 2014. Strict minimizers for geometric optimization. *ACM Trans. Graph.* 33, 6, 185:1–185:14.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM Trans. Graph.* 21, 3, 362–371.

LIPMAN, Y. 2012. Bounded distortion mapping spaces for triangular meshes.

ACM Trans. Graph. 31, 4, 108:1–108:13.

LIU, L., ZHANG, L., XU, Y., GOTSMAN, C., AND GORTLER, S. J. 2008. A local/global approach to mesh parameterization. In *Proc. Symposium on Geometry Processing*. 1495–1504.

LIU, T., GAO, M., ZHU, L., SIFAKIS, E., AND KAVAN, L. 2016. Fast and robust inversion-free shape manipulation. *Comput. Graph. Forum* 35, 2.

MARTIN, S., THOMASZEWSKI, B., GRINSPUN, E., AND GROSS, M. 2011. Example-based elastic materials. *ACM Trans. Graph.* 30, 4, 72:1–72:8.

MARTINEZ ESTURO, J., RÖSSL, C., AND THEISEL, H. 2014. Smoothed quadratic energies on meshes. *ACM Trans. Graph.* 34, 1, 2:1–2:12.

MULLEN, P., TONG, Y., ALLIEZ, P., AND DESBRUN, M. 2008. Spectral conformal parameterization. In *Proc. Symposium on Geometry Processing*. 1487–1494.

MYLES, A., PIETRONI, N., AND ZORIN, D. 2014. Robust field-aligned global parametrization. *ACM Trans. Graph.* 33, 4.

MYLES, A. AND ZORIN, D. 2012. Global parametrization by incremental flattening. *ACM Trans. Graph.* 31, 4, 109:1–109:11.

NOCEDAL, J. AND WRIGHT, S. J. 2006. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer, Berlin. NEOS guide <http://www-fp.mcs.anl.gov/otc/Guide/>.

PIGHIN, F. AND LEWIS, J. 2007. Practical least-squares for computer graphics. In *ACM SIGGRAPH 2007 Courses*.

PORANNE, R., OVREIU, E., AND GOTSMAN, C. 2013. Interactive planarization and optimization of 3D meshes. *Comput. Graph. Forum* 32, 1, 152–163.

RABINOVICH, M. 2016. Scalable locally injective mappings. <https://github.com/MichaelRabinovich/Scalable-Locally-Injective-Mappings>.

SANAN, P. D. 2014. Geometric elasticity for graphics, simulation, and computation. Ph.D. thesis. California Institute of Technology.

SCHENK, O., BOLLHÖFER, M., AND RÖMER, R. A. 2008. On large-scale diagonalization techniques for the Anderson model of localization. *SIAM Rev.* 50, 1, 91–112.

SCHENK, O., WCHTER, A., AND HAGEMANN, M. 2007. Matching-based preprocessing algorithms to the solution of saddle-point problems in large-scale nonconvex interior-point optimization. *Computational Optimization and Applications* 36, 2-3, 321–341.

SCHREINER, J., ASIRVATHAM, A., PRAUN, E., AND HOPPE, H. 2004. Inter-surface mapping. *ACM Trans. Graph.* 23, 3.

SCHÜLLER, C., KAVAN, L., PANOZZO, D., AND SORKINE-HORNUNG, O. 2013. Locally injective mappings. *Computer Graphics Forum* 32, 5, 125–135.

SHEFFER, A. AND DE STURLER, E. 2001. Parameterization of faceted surfaces for meshing using angle-based flattening. *Engineering with Computers* 17, 3, 326–337.

SHEFFER, A., LÉVY, B., MOGILNITSKY, M., AND BOGOMYAKOV, A. 2005. ABF++: Fast and robust angle based flattening. *ACM Trans. Graph.* 24, 2, 311–330.

SHEFFER, A., PRAUN, E., AND ROSE, K. 2006. Mesh parameterization methods and their applications. *Found. Trends Comput. Graph. Vis.* 2, 2, 105–171.

SMITH, J. AND SCHAEFER, S. 2015. Bijective parameterization with free boundaries. *ACM Trans. Graph.* 34, 4, 70:1–70:9.

SORKINE, O. AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *Proc. Symp. Geometry Processing*. 109–116.

TERAN, J., SIFAKIS, E., IRVING, G., AND FEDKIW, R. 2005. Robust quasistatic finite elements and flesh simulation. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA '05. ACM, New York, NY, USA, 181–190.

TUTTE, W. T. 1963. How to draw a graph. *Proceedings of the London Mathematical Society* 13, 743–767.

WEBER, O., MYLES, A., AND ZORIN, D. 2012. Computing extremal quasiconformal maps. *Comput. Graph. Forum* 31, 5.

YOSHIZAWA, S., BELYAEV, A., AND SEIDEL, H.-P. 2004. A fast and simple stretch-minimizing mesh parameterization. In *Proceedings of the*

Shape Modeling International 2004. SMI '04. IEEE Computer Society, Washington, DC, USA, 200–208.

ZAYER, R., LÉVY, B., AND SEIDEL, H.-P. 2007. Linear angle based parameterization. In *Proc. SGP*. 135–141.

APPENDIX

A. SOLVING EQ. (13)

In order to solve Eq. (13), we write it in matrix form, and in terms of the coordinates \mathbf{x} . Then Eq. (13) is transformed into

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \quad (34)$$

where the structure of \mathbf{A} and \mathbf{b} in the 2D case is as follows. Assuming a set of orthogonal frames per element are prescribed, we let D_x, D_y be the FE gradient matrices of the mesh w.r.t. the frames. Additionally, we define four diagonal matrices, \mathbf{W}_{ij} for $i, j = 1, 2$, where the diagonal of \mathbf{W}_{ij} holds the (i, j) entries of all of the weights \mathbf{W}_f . In other words, $\mathbf{W}_{ij} = \text{diag}(\{\mathbf{W}_f(i, j)\}_f)$. Similarly, we define \mathbf{R}_{ij} to be the column vector holding the (i, j) entries of all of the \mathbf{R}_f

$$\mathbf{A} = \mathbf{W}\mathbf{D} = \begin{pmatrix} \mathbf{W}_{11} & \mathbf{W}_{12} & 0 & 0 \\ \mathbf{W}_{21} & \mathbf{W}_{22} & 0 & 0 \\ 0 & 0 & \mathbf{W}_{11} & \mathbf{W}_{12} \\ 0 & 0 & \mathbf{W}_{21} & \mathbf{W}_{22} \end{pmatrix} \begin{pmatrix} D_x & 0 \\ 0 & D_y \\ D_x & 0 \\ 0 & D_y \end{pmatrix} \quad (35)$$

$$\mathbf{b} = \begin{pmatrix} \mathbf{R}_{11} \\ \mathbf{R}_{21} \\ \mathbf{R}_{12} \\ \mathbf{R}_{22} \end{pmatrix} \quad (36)$$

This can be readily solved by any least squares minimization algorithm.

B. SECTION 3 PROOFS

B.1 Lemma 3.1

Lemma 3.1 *Let $\mathbf{J} = \mathbf{U}\mathbf{S}_J\mathbf{V}^\top$ be the Singular Value Decomposition of \mathbf{J} . Then,*

$$\mathcal{D}(\mathbf{J}) = \mathcal{D}(\mathbf{S}_J) \quad (37)$$

$$\nabla_{\mathbf{J}}\mathcal{D}(\mathbf{J}) = \mathbf{U}\nabla_{\mathbf{S}_J}\mathcal{D}(\mathbf{S}_J)\mathbf{V}^\top \quad (38)$$

PROOF. Eq. (37) is immediate from the definition in Eq. (1). As for Eq. (38), we use the formula for the derivative of the singular values (see [Giles 2008])

$$\nabla_{\mathbf{J}}\mathcal{D}(\mathbf{S}_J) = \mathbf{U}\nabla_{\mathbf{S}_J}\mathcal{D}(\mathbf{S}_J)\mathbf{V}^\top \quad (39)$$

□

B.2 Local step, general construction: proof of Eq. (32)

For a rotation invariant $\mathcal{D}(\mathbf{J})$ that is separably strictly convex in singular values, Eq. (31) can be satisfied by setting $(\mathbf{S}_A)_i$ such that:

$$\frac{\partial}{\partial \sigma_i} \mathcal{D}(\sigma_1, \dots, \sigma_{i-1}, \sigma_i, \dots, \sigma_d)(\sigma_i) = 0 \quad (40)$$

In particular, in the case of a true isometric distortion measure, (31) is satisfied by setting the local step as the closest rotation $\Lambda = \mathbf{U}\mathbf{V}^\top$.

This can be seen from the fact that every partial function of $\mathcal{D}(\sigma)$ is strictly convex on $\mathbb{R}_{>0}$, and therefore has a single minimum, $(\mathbf{S}_\Lambda)_i$. Hence, for every $\sigma_i < (\mathbf{S}_\Lambda)_i$, $\frac{\partial}{\partial \sigma_i} \mathcal{D}(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_d)(\sigma_i) < 0$, and for every $\sigma_i > (\mathbf{S}_\Lambda)_i$, $\frac{\partial}{\partial \sigma_i} \mathcal{D}(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_d)(\sigma_i) > 0$. This is also true for $(\mathbf{S}_J - \mathbf{S}_\Lambda)_i$ and so Eq. (32) is satisfied.

B.3 Conformal energy local step derivation

Let $\mathcal{D}(\mathbf{J}) = \frac{\text{tr}(\mathbf{J}^\top \mathbf{J})}{\det(\mathbf{J})^{2/d}}$. $\mathcal{D}(\mathbf{J})$ is rotation invariant and can be written as $\mathcal{D}(\sigma) = \frac{\sum_{i=1}^d \sigma_i^2}{\sigma_1 \dots \sigma_d}$. By differentiating the distortion measure w.r.t. the singular values in the 2D case, we find that

$$(\nabla_{\mathbf{S}_J} \mathcal{D}(\mathbf{S}_J))_1 = \frac{1}{\sigma_2} - \frac{\sigma_2}{\sigma_1^2},$$

and similarly, $(\nabla_{\mathbf{S}_J} \mathcal{D}(\mathbf{S}_J))_2 = \frac{1}{\sigma_1} - \frac{\sigma_1}{\sigma_2^2}$. Assuming \mathbf{J} is *not* a similarity already, then, since $\sigma_1 > \sigma_2 > 0$, the first entry is negative, while the second is positive. By choosing $\sigma_1 > (\mathbf{S}_\Lambda)_i > (\sigma_2)$, this holds true for $(\mathbf{S}_J - \mathbf{S}_\Lambda)_i$ and so Eq. (31) is satisfied.

For the 3D case,

$$(\nabla_{\mathbf{S}_J} \mathcal{D}(\mathbf{S}_J))_i = \frac{-2\sigma_{i+1}\sigma_{i+2}(\sigma_{i+1}^2 + \sigma_{i+2}^2 - 2\sigma_i^2)}{(3\sigma_i\sigma_{i+1}\sigma_{i+2})^{5/3}},$$

where the index i cycles from 1 to 3 (i.e., $\sigma_4 = \sigma_1$). This is zero only for $\sigma_i = \sqrt{\frac{\sigma_{i+1}^2 + \sigma_{i+2}^2}{2}}$, and so $(\nabla_{\mathbf{S}_J} \mathcal{D}(\mathbf{S}_J))_1 < 0$,

$(\nabla_{\mathbf{S}_J} \mathcal{D}(\mathbf{S}_J))_3 > 0$. We note that $\bar{\sigma} = \sqrt{\frac{\sigma_1^2 + \sigma_3^2}{2}}$ satisfies $\sigma_3 < \bar{\sigma} < \sigma_1$. Therefore, by choosing $\mathbf{S}_\Lambda = \bar{\sigma} \mathbf{U} \mathbf{V}^\top$, we get $(\mathbf{S}_J - \mathbf{S}_\Lambda)_1 < 0$, $(\mathbf{S}_J - \mathbf{S}_\Lambda)_3 > 0$, and by construction, same as the proof for Eq. (32), we get that the sign of $(\mathbf{S}_J - \mathbf{S}_\Lambda)_2$ is equal to the sign of $(\nabla_{\mathbf{S}_J} \mathcal{D}(\mathbf{S}_J))_2$.