

Lifting Simplicies to Find Injectivity

XINGYI DU, Washington University in St. Louis, USA
NOAM AIGERMAN and QINGNAN ZHOU, Adobe Research, USA
SHAHAR Z. KOVALSKY, Duke University, USA
YAJIE YAN, Facebook, USA
DANNY M. KAUFMAN, Adobe Research, USA
TAO JU, Washington University in St. Louis, USA

Mapping a source mesh into a target domain while preserving local injectivity is an important but highly non-trivial task. Existing methods either require an already-injective starting configuration, which is often not available, or rely on sophisticated solving schemes. We propose a novel energy form, called Total Lifted Content (TLC), that is equipped with theoretical properties desirable for injectivity optimization. By lifting the simplices of the mesh into a higher dimension and measuring their contents (2D area or 3D volume) there, TLC is smooth over the entire embedding space and its global minima are always injective. The energy is simple to minimize using standard gradient-based solvers. Our method achieved 100% success rate on an extensive benchmark of embedding problems for triangular and tetrahedral meshes, on which existing methods only have varied success.

CCS Concepts: • **Computing methodologies** → **Mesh models**.

Additional Key Words and Phrases: Parameterization, injective embedding

ACM Reference Format:

Xingyi Du, Noam Aigerman, Qingnan Zhou, Shahar Z. Kovalsky, Yajie Yan, Danny M. Kaufman, and Tao Ju. 2020. Lifting Simplicies to Find Injectivity. *ACM Trans. Graph.* 39, 4, Article 1 (July 2020), 17 pages. <https://doi.org/10.1145/3386569.3392484>

1 INTRODUCTION

Computing constrained mappings between domains is a fundamental task, performed across a wide range of geometric and physical applications ranging from parameterization and UV-mapping, to deformation modeling and the simulation of elastica. In all of these applications, it is in most cases critical to generate a one-to-one, *injective* mapping. This ensures that the inverse map exists and that the correspondence between domains is well-defined. Injectivity is critical for various applications, such as painting textures in UV space, co-analyzing shapes based on correspondences, obtaining good-looking deformations, and generating physically correct simulations of materials, to name just a few.

Most of the time, (local) injectivity is formulated computationally as preservation of all mesh simplices' orientation, *i.e.*, no triangle

Authors' addresses: Xingyi Du, Washington University in St. Louis, USA, du.xingyi@wustl.edu; Noam Aigerman; Qingnan Zhou, Adobe Research, USA; Shahar Z. Kovalsky, Duke University, USA; Yajie Yan, Facebook, USA; Danny M. Kaufman, Adobe Research, USA; Tao Ju, Washington University in St. Louis, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

0730-0301/2020/7-ART1 \$15.00

<https://doi.org/10.1145/3386569.3392484>




Fig. 1. Injectively mapping a complex surface mesh (Lucy, 48K vertices) to a non-convex boundary (letter “G”, with zoom-ins), at the top, and mapping a tetrahedral mesh (Armadillo, 6K vertices) to a highly deformed target surface, at the bottom, as a result of minimizing our novel energy. These two examples are part of our new benchmark data set.

or tetrahedron is flipped. Unfortunately, the injectivity constraint is not only highly non-convex, but also an *open* set, making any optimization involving it non-trivial. As a result, many mapping and deformation algorithms focus on *preserving* triangle's orientation while improving the map's quality, *i.e.*, by minimizing distortion measures [Rabinovich et al. 2017; Zhu et al. 2018] that also act as a *barrier* that pushes them away from degenerating triangles on the closure of the locally-injective set. This in turn entails that they require a *feasible* embedding - one that is locally injective and satisfies all given constraints - as initialization to begin the minimization process. Indeed, if the input is non-injective, most distortion metrics that act as a barrier to prevent triangles from inverting also fight against *un-inverting* the initially-inverted triangles.

To find an un-inverted initializer for the above methods while exactly satisfying the given boundary constraints, one can opt to use one of two approaches: either (1) limit themselves to embeddings from the only known method with injectivity guarantee – Tutte’s embedding [Tutte 1963], whose guarantee is restricted to only 2D convex domains, or (2) use one of a number of recent methods [Aigerman and Lipman 2013; Fu and Liu 2016; Kovalsky et al. 2015; Su et al. 2019; Weber et al. 2012] that have been developed to produce injective and low-distortion mappings for given constraints, *without* requiring an injective initialization.

Unfortunately, the latter class of methods is in general not guaranteed to succeed in finding an injective mapping, and in practice they often fail on examples where injective mappings do exist (see Section 6). The main reason for that is that they focus on a much larger task, of computing *low-distortion* maps, which entails they are not tailor-made for injectivity. Indeed, for many methods, the low-distortion paradigm is engrained in their approach. Counter-intuitively, attempting to require a *less* strict distortion bound so as to optimize only injectivity often leads to deterioration of success rates in such methods instead of increasing them.

To address these issues, we focus solely on local injectivity, *i.e.*, the correct orientation of elements, without considering their geometric distortion. We revisit the idea of computing an injective mapping via a variational principal of minimizing some energy. Instead of a barrier energy, we devise a new energy tailor-made for recovering injectivity from a given non-injective embedding while satisfying positional constraints. We refer to this injectivity energy as *Total Lifted Content* (TLC). Intuitively, TLC measures the total content (2D area or 3D volume) of a mesh after *lifting* the simplices of the mesh to a higher dimension. We demonstrate two properties of our energy that shed light on our energy’s efficacy in enforcing injectivity:

- (1) TLC is a well-defined and smooth function over the entire embedding space, regardless of injectivity. This contrasts barrier energies (e.g., MIPS or Symmetric Dirichlet), which become undefined upon non-injectivity.
- (2) More importantly, the *global minimum* of TLC is *only* achieved by an injective embedding, if such an embedding exists.

We know of no existing energy in either 2D or 3D that possesses these properties. Additionally, we show the connection between minimizers of TLC and several known distortion-minimizing maps, including MIPS and harmonic maps. Unlike existing injectivity-recovery methods that rely on sophisticated and custom-made solvers to impose injectivity constraints, our energy can be easily minimized using standard solvers such as quasi-Newton and projected Newton.

To demonstrate the efficacy of our method, we introduce in Section 6 a benchmark set comprising of existing and many new challenging examples (including those in Figure 1) to extensively compare results of our method with state-of-the-art injective mapping methods. While our theoretical guarantee of injectivity only applies to global minima, it is not guaranteed we will achieve it in practice since our energy is *not* convex. However, we show empirically that converging to a global minimum is unnecessary for achieving

injectivity. In fact, simply terminating the optimization upon reaching an injective embedding obtains a success rate of 100% on all examples in the benchmark, whereas existing methods only have varied success. Importantly, we demonstrate that the output of our method then can be used to bootstrap standard distortion minimization methods to improve the distortion of the injective mapping. This enables injective distortion minimization for challenging examples that were previously not possible due to the unavailability of a starting, constraint satisfying, injective map.

2 PROBLEM STATEMENT AND OVERVIEW

We address the problem of injectively mapping an input simplicial (triangles or tetrahedra) mesh into a fixed boundary. We start with a 2D or 3D *rest* mesh M and a target mapping boundary B that is in one-to-one correspondence with the boundary of M . We then seek a (locally) injective embedding T of M into B , such that each simplex of T is positively oriented. Note that an injective embedding may not exist for certain choices of M and B . An example is when M has a single interior vertex and it is connected to all boundary vertices, and B is not a star-shape. As our goal is to find an injective mapping, we will assume that such a mapping exists for the given M and B .

We address this embedding task variationally by solving an energy minimization over the space of all possible embeddings. This *embedding space* has dimension $d \times n$ where $d = 2, 3$ is the dimension of the embedding and n is the number of interior vertices in the triangulation. The crux then is in forming an appropriate energy $E : T \rightarrow \mathbb{R}_+$ to minimize. A desirable energy should satisfy the following three criteria necessary for minimization to gain an injective embedding:

- (1) E is well-defined for all possible embedding T ;
- (2) E is at least C^2 over the embedding space; and
- (3) All global minima of E are injective embeddings.

Criteria (1) allows minimization to start from easily available, non-injective initial embeddings, e.g., as obtained by Tutte; (2) enables the effective application of efficient, gradient-descent based minimization; and (3) is necessary for minimizers of E to generate injective maps. Note that (3) alone may not be sufficient in practice; for example, the existence of non-injective, local minima can present significant challenges to gradient-based solvers. However, to our knowledge, no energy satisfying even these basic criteria has been previously proposed.

The rest of the paper is organized as follows. We start with a review of existing methods for injective mapping in Section 3. In Section 4 we introduce our new energy (Total Lifted Content), and detail not only how it meets the criteria above but also its connection with existing distortion energies. We discuss the minimization of the energy in Section 5, and present extensive experimental results on a new benchmark data set in Section 6. We conclude in Section 7 with discussions on venues of future research.

3 RELATED WORK

Injectivity is closely tied with surface parameterization, in which it is both the most needed as well as the most tractable, see [Floater and Hormann 2005; Hormann et al. 2007] for surveys. One of the

earliest known algorithms for injective mappings is Tutte’s embedding [Tutte 1963], which up to this day, is one of the only methods that *guarantee* an injective map without requiring an initialization. Although several works extended it [Aigerman and Lipman 2015; Floater 2003; Gortler et al. 2006] to other specific classes of mappings, its essential limitations remain: it can only map injectively to a prescribed convex boundary, without any interior constraints. Furthermore, its 3D extensions do not yield injective mappings, even in trivial cases.

Inversion-free optimization. A few works addressed the problem of inversion-free mappings in 2D and 3D, via the use of “barrier”-type energies, in which the objective function includes terms that grow asymptotically as an element becomes degenerate. All of these methods require an injective initializer, as the barrier term is unlikely (and in most cases can’t) recover from non-injective states. Hence they initialize from the identity map or from Tutte’s embedding. Locally Injective Mappings [Schüller et al. 2013] suggested incorporating a barrier term, using the log of the determinant. [Liu et al. 2016] followed a similar path by solving a sequence of convex programs. [Smith and Schaefer 2015] tailor a line search optimization that uses a maximal step size that avoids inversions. Instead of using an auxiliary injectivity barrier, several methods directly optimize distortion metrics that explode near degeneracies. One of the first works to suggest such an energy was MIPS [Hormann and Greiner 2000] which was used in inversion-free optimization in [Fu et al. 2015]. [Rabinovich et al. 2017] use ARAP as a pseudo-majorizer for other non-inverting energies. Other methods explored computational speedups of optimization of these energies via majorization [Shtengel et al. 2017], preconditioning [Claici et al. 2017], or modifying the quasi-Newton algorithm [Zhu et al. 2018]. [Liu et al. 2018] compute a parameterization by constructing a sequence of bounded distortion maps, without setting any positional constraints.

Inversion-free initialization. Another line of research directly focuses on computing mappings that adhere to given positional constraints, without requiring an initializer. As we demonstrate in our comparisons, since their methods are focused on low-distortion mappings (which are by construction inversion-free), they tackle a broader task, requiring a much more general approach, and hence fail on hard configurations in which our method succeeds. A common approach is to aim to project an existing mapping into a bounded-distortion one [Aigerman and Lipman 2013; Kovalsky et al. 2015]. [Fu and Liu 2016] compute a bounded-distortion embedding by minimizing gaps between the mesh simplices while keeping their distortion bounded. Similarly to us [Su et al. 2019] aim at finding an injective map, but use bounded-distortion type projection techniques that are not tailor-made for injectivity. We compare to these methods, and show we achieve a higher success rate at recovering injective mappings for given constraints.

A different approach is taken in [Weber et al. 2012], which is based on the observation that the extremal quasiconformal maps (maps minimizing the maximal conformal distortion), in the continuous setting, are guaranteed to be locally injective and can conform to given positional constraints. However, the injectivity guarantees do not translate well to a discrete approximation, and inversion of triangles often happen at concave corners [Weber and Zorin 2014].

Another work closely related to us is [Xu et al. 2011], who directly minimize the total unsigned area (TUA). However, as we will detail in Section 4.1, TUA in fact optimizes over the *closure* of the space of injective maps, whose boundary contains degenerate triangles with zero area. Furthermore, TUA suffers from derivative discontinuities and vanishing gradients, minimizing TUA often gets stuck in non-injective global minima. See comparisons in Figure 5.

The method of [Hefetz et al. 2019] uses a subspace for locally injective harmonic maps and can adhere to periodic cone conditions, but cannot map injectively into arbitrary boundary constraints. Several methods consider enabling the boundary vertices’ images to slide across the target boundary while ensuring the resulting assignment still induces a bijective mapping into the target domain. [Aigerman and Lipman 2013] map tetrahedral meshes into polycubes while allowing each vertex to slide on its assigned flat face of the polycube. [Aigerman and Lipman 2015] extend Tutte and show that for some target convex polygons, vertices may be allowed to slide on the boundary of the target polygon. Embeddings with sliding boundaries have been theoretically studied in [Lipman 2014], where local injectivity is shown to suffice for global injectivity even when the boundary is allowed to slide. While our method focuses on fixed boundary maps, it exhibits better success rates for achieving injectivity in comparison to [Aigerman and Lipman 2013], and can be applied to any target polygon, as opposed to the limited set of 4 convex polygons in [Aigerman and Lipman 2015].

Non-adhering injective mappings. A different line of approach computes maps that cannot be defined solely in terms of the input mesh (*i.e.*, the mapping changes the mesh structure). [Agarwal et al. 2008] untangle inversions introduced by 2-dimensional deformations of a triangular mesh by local remeshing. [Weber and Zorin 2014] compute an injective map based on a locally injective co-parameterization into an intermediate convex domain, requiring refinement of the input triangulation. [Campen et al. 2016] use piecewise-linear foliations to compute point-wise bijections of 2- and 3-dimensional objects onto canonical domains (*e.g.*, unit disk or ball); this bijection, however, does not readily conform with the input triangle or tetrahedral mesh and requires refinement to extract a piecewise-linear bijection. [Gu et al. 2018] show that injectivity can be achieved by a discrete Yamabe flow, but mesh surgeries are needed when singularities are developed. [Shen et al. 2019] propose a numerically robust method for computing a bijection onto a convex planar domain, but for a non-convex one, they have to modify the domain and add vertices. Furthermore, their method’s nature restricts it to 2D.

4 ENERGY

We turn to constructing an energy E that satisfies all three criteria enumerated in Section 2. Satisfying these criteria is not trivial, and we are unaware of any existing energy that meets all criteria. The construction of E stems from a central energy related to injectivity – the total unsigned area.

4.1 Total Unsigned Area

Xu et al. [2011] observed that when T is an injective triangular mesh, it minimizes the sum of the *unsigned* triangle areas among all




Fig. 2. Lifting a triangle t with vertices $\{x_i, y_i\}$ ($i = 1, 2, 3$) to a triangle \hat{t} in 4D via the $\sqrt{\alpha}$ -scaled auxiliary triangle \tilde{t} with vertices $\{u_i, v_i\}$.

embeddings into the target domain. This is due to the total unsigned area (TUA) being an upper bound to the sum of *signed* areas, which is constant for a fixed boundary B , and equal if all triangles have positive area.

However, TUA fails several criteria as mentioned above. First, it is not smooth, considered as a function of the embedding. In particular, TUA exhibits a C^1 discontinuity as a vertex moves across the supporting line of its opposite edge in a triangle. Second, while any injective embedding achieves the global minimum of TUA, the inverse is not true: a global minimum of TUA can also be achieved by a non-injective embedding where the only triangles that have non-positive areas are those having zero areas (*i.e.*, degenerate triangles). We call such an embedding, which lies on the closure of the space of injective embeddings, a *pseudo-injective* embedding.

In addition to these deficiencies, another problem of TUA is its vanishing gradient. Note that TUA has zero gradient with respect to any vertex surrounded by a ring of consistently oriented triangles. In practice, we have observed that minimizing TUA can easily get stuck in “plateaus”, or local minimizers of TUA with vanishing gradients, that are far from being injective (see Figure 5).

4.2 Total Lifted Content

To address the shortcomings of TUA, we propose to *lift* the triangles (or tetrahedra) to a higher-dimensional space, and consider their total area (or volume). The lifting is designed so that the total content (area or volume) of the lifted simplices is a smooth energy over the entire embedding space, and that *every* global minimum of the energy is achieved by an injective embedding. We call this energy *Total Lifted Content* (or TLC).

Specifically, to lift a d -dimensional ($d = 2, 3$) simplex t , we make use of another d -dimensional, non-degenerate *auxiliary* simplex \tilde{t} and a positive scalar α , both of which are fixed during embedding optimization. We construct a $2d$ -dimensional *lifted* simplex \hat{t} by concatenating the vertex coordinates of t with the corresponding coordinates of \tilde{t} scaled by $\sqrt{\alpha}$. An illustration for lifting a 2-dimensional triangle is shown in Figure 2. To lift a d -dimensional simplicial mesh T , we use a set of auxiliary simplices, one for each simplex of T . TLC is defined as the sum of the area (or volume) of the lifted simplices.

Before diving into more details, we give some intuition as to how TLC avoids the drawbacks of TUA. Recall that TUA of a mesh T is smooth except when a simplex becomes degenerate (having zero

content). Since the content of the lifted simplex has contributions from both the simplex of T and the auxiliary simplex (scaled by $\sqrt{\alpha}$), the lifted simplex is *never degenerate* if the auxiliary simplex is chosen to be non-degenerate and $\alpha > 0$. Hence TLC remains smooth even when some simplices of T become degenerate. Also, when $\alpha = 0$, TLC reduces to TUA and shares the same set of global minima as TUA (including the pseudo-injective embeddings). However, as we will show, there is a range of sufficiently small (but positive) α such that the global minimum of TLC remains a global minimum of TUA but has no degenerate simplices (*i.e.*, it is injective).

In the following, we first derive an explicit formula of TLC in terms of the geometric quantities of the input mesh and the auxiliary simplices. We then prove the two key properties of TLC, namely smoothness and injectivity at global minimum. We conclude this section by shedding light on the impact of the auxiliary simplices on the *shape* of the energy-minimizing embeddings.

4.2.1 Formula. We give a general formula for TLC in any dimension d . Consider a simplex t with auxiliary simplex \tilde{t} . Let X (respectively \tilde{X}) be a $d \times d$ matrix whose column vectors are the edge vectors from one vertex of the simplex t (respectively \tilde{t}) to the other d vertices of the simplex. The $2d \times d$ matrix of the edge vectors of the lifted simplex \hat{t} , denoted by \hat{X} , is therefore defined as

$$\hat{X} = \begin{pmatrix} X \\ \sqrt{\alpha} * \tilde{X} \end{pmatrix} \quad (1)$$

Consider the d -dimensional subspace of the $2d$ -dimensional lifted space that contains \hat{t} , and pick any orthonormal basis of this subspace. We can express each (column) edge vector of \hat{X} as a length- d vector in this basis, yielding another $d \times d$ matrix Y . Note that $Y^T Y = \hat{X}^T \hat{X} = X^T X + \alpha \tilde{X}^T \tilde{X}$. Using the volume formula of a d -dimensional simplex, the content of \hat{t} is $\|\text{Det}(Y)\|/d!$, where Det is the matrix determinant. By the multiplicativity of determinants,

$$\begin{aligned} E_{\tilde{t},\alpha}(t) &= \frac{1}{d!} \|\text{Det}(Y)\| \\ &= \frac{1}{d!} \sqrt{\text{Det}(Y^T Y)} \\ &= \frac{1}{d!} \sqrt{\text{Det}(X^T X + \alpha \tilde{X}^T \tilde{X})} \end{aligned} \quad (2)$$

Note that the lifted content reduces to the unsigned area (or volume) of t , $\|\text{Det}(X)\|/d!$, when $\alpha = 0$.

The TLC of a d -dimensional simplicial mesh T , for a given (and fixed) set of auxiliary simplices \tilde{T} and a scaling α , is

$$E_{\tilde{T},\alpha}(T) = \sum_{t \in T} E_{\tilde{t},\alpha}(t) \quad (3)$$

We ask that each auxiliary simplex in \tilde{T} has non-zero content, but they do not need to form a connected mesh. For example, \tilde{T} can be made up of equilateral triangles or tetrahedra of the same size. Again, note that TLC becomes TUA when $\alpha = 0$.

4.2.2 Smoothness. We first consider the lifted content of one simplex t . As a concrete example, consider the triangle t shown in the Figure 3 (top-left), and let \tilde{t} be an equilateral triangle (top-right). Observe from the plots in (a,b) that, as the red vertex of t moves along the vertical (blue) or horizontal (green) lines, the lifted content $E_{\tilde{t},\alpha}(t)$ of t is smooth and greater than the unsigned area of T (*i.e.*,

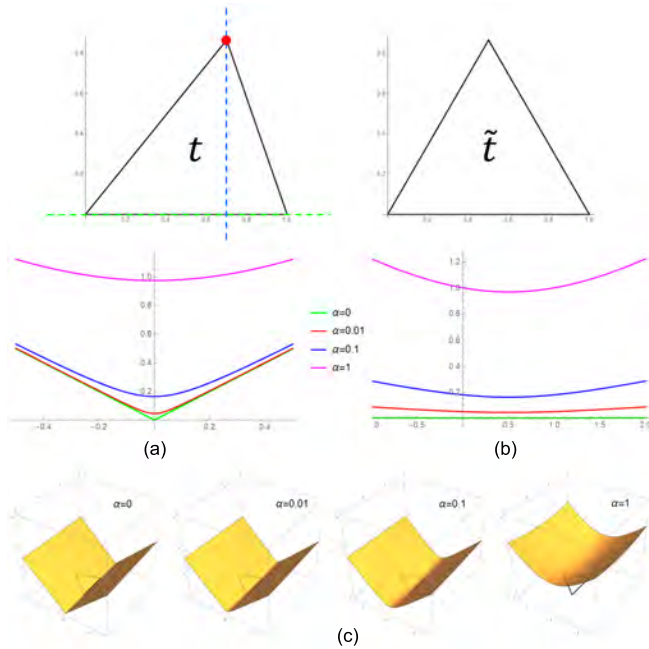


Fig. 3. Given a triangle t and an equilateral auxiliary triangle \tilde{t} (top), (a,b) plot the lifted content of t , $E_{\tilde{t},\alpha}(t)$, at different α values as the red vertex in t moves along the blue dotted line in (a) and the green dotted line in (b). (c) plots $E_{\tilde{t},\alpha}(t)$ as the red vertex moves over the plane.

$E_{\tilde{t},0}(t)$) for all $\alpha > 0$. We will confirm these observations below, for any dimension $d \geq 2$.

PROPOSITION 4.1. *The following holds for any simplex t , given a non-degenerate auxiliary simplex \tilde{t} and positive α :*

- (1) $E_{\tilde{t},\alpha}(t) > E_{\tilde{t},0}(t) \geq 0$.
- (2) $E_{\tilde{t},\alpha}(t)$ is differentiable to any order with respect to t 's vertices.


PROOF. (1) Since t is the orthogonal projection of the lifted simplex \hat{t} to the first d dimensions, the content of \hat{t} (i.e., $E_{\tilde{t},\alpha}(t)$) is no smaller than the unsigned content of t (i.e., $E_{\tilde{t},0}(t)$). The equality holds only when \hat{t} lies in a subspace parallel to the first d dimensions. This means that the projection of \hat{t} in the remaining d dimensions has no content, contradicting that \tilde{t} is non-degenerate.

- (2) The derivative of $E_{\tilde{t},\alpha}(t)$ with respect to vertices of t , to any order, is a summation of rational terms each with powers of $E_{\tilde{t},\alpha}(t)$ on the denominator and polynomials in t 's vertices on the numerator (see Appendix C). Since $E_{\tilde{t},\alpha}(t) > 0$, $E_{\tilde{t},\alpha}(t)$ has finite derivatives. \square

The above properties easily carry over to a simplicial mesh:

COROLLARY 4.2. *The Total Lifted Content, $E_{\tilde{T},\alpha}(T)$, is strictly positive, greater than the total unsigned areas or volumes of T , and differentiable to any order over the embedding space of T , given any set of non-degenerate auxiliary simplices \tilde{T} and positive α .*

4.2.3 Injectivity. We show that $E_{\tilde{T},\alpha}$ has only injective global minima for sufficiently small values of α . Recall that $E_{\tilde{T},0}$ reduces to the TUA energy, whose global minima are attained by injective and pseudo-injective embeddings. Note that an injective mapping lies in a region of the embedding space where TUA is constant, since any small perturbation to non-constrained vertices keeps the mapping injective and hence TUA is equal to the target domain's area. This means that the embedding lies in the flat plateau of TUA which is comprised of all injective embeddings, as well as some pseudo-injective embeddings that can be perturbed to become injective. On the other hand, a pseudo-injective embedding lies at a C^1 discontinuity of TUA (due to presence of degenerate triangles). However, for any $\alpha > 0$ $E_{\tilde{T},\alpha}$ is smooth everywhere and strictly above TUA (Corollary 4.2). To accomplish this transition from non-smoothness to smoothness as α increases, and for a small enough range of α , a pseudo-injective embedding must see a greater rise in energy than any injective embedding in order to “round off” the sharp bottom of the valley (see the illustration in the inset), and hence it is no longer a global minimizer of $E_{\tilde{T},\alpha}$.



We make a precise statement in the following proposition, which states that any injective embedding would have lower energy than all non-injective embeddings for some range of α , and the range depends solely on that injective embedding and the auxiliary simplices. The proof is given in Appendix B.

PROPOSITION 4.3. *Let T_0 be some d -dimensional ($d = 2, 3$) injective embedding into the target boundary and \tilde{T} be a set of non-degenerate auxiliary simplices. Then there exists some $\beta > 0$ such that $E_{\tilde{T},\alpha}(T) > E_{\tilde{T},\alpha}(T_0)$ for any non-injective embedding T and $\alpha < \beta$.*

4.2.4 Shape control. It is conceivable that there may exist other energies that also have the desired properties (being smooth and having only injective minima). A unique feature of TLC is that the choice of auxiliary simplices \tilde{T} offers additional control over the shape of the energy-minimizing embedding. As we shall see, when α takes on very small or large values, TLC converges to one of a few well-known energies.

As shown above, the global minimum of $E_{\tilde{T},\alpha}$ as α approaches 0 are attained only by injective embeddings. Since all injective embeddings have the same energy when $\alpha = 0$ (i.e., the content enclosed by the target boundary), the global minimum of TLC as α infinitesimally increases from zero is the injective embedding that sees the least instantaneous rise in energy. This instantaneous rise for a given simplex t , with auxiliary simplex \tilde{t} , is just the derivative of $E_{\tilde{t},\alpha}(t)$ with respect to α evaluated at $\alpha = 0$. This derivative has a surprisingly simple expression (see proof in Appendix A):

PROPOSITION 4.4. *Let t, \tilde{t} be two non-degenerate d -dimensional ($d = 2, 3$) simplices, and $D(t)$ the Dirichlet energy of linearly transforming t to \tilde{t} . Then*

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial \alpha} \Big|_{\alpha=0} = D(t) \quad (4)$$




Fig. 4. Left: Rest mesh (Bunny) and MIPS mapping into a target boundary (outline of letter “A”). Middle: Embeddings minimizing uniform-TLC (top) and rest-TLC (bottom) energy at various values of α . Right: Tutte and harmonic mapping of the rest mesh into the same boundary. Inverted triangles are colored red. Note the similarity between the two embeddings highlighted in boxes with the same color.

In other words, the injective embedding T that minimizes TLC as $\alpha \rightarrow 0$ converges to minimize the Dirichlet energy from T to the auxiliary simplices \tilde{T} . The latter energy, in two-dimensions, is in fact the MIPS energy [Hormann and Greiner 2000] from \tilde{T} (the “surface triangles”) to T (the “parameter triangles”).

At the other end of the spectrum, as α approaches ∞ , $E_{\tilde{T},\alpha}(t)$ converges to infinity as well for a given pair of t, \tilde{T} . However, its gradient $\nabla E_{\tilde{T},\alpha}(t)$ with respect to the simplex t has a well-defined limit up to a constant multiplier (see proof in Appendix A):

PROPOSITION 4.5. *Let t, \tilde{T} be two d -dimensional ($d = 2, 3$) simplices with \tilde{T} non-degenerate, and let $\tilde{D}(t)$ be the Dirichlet energy of the linear map from \tilde{T} to t . Then*

$$\lim_{\alpha \rightarrow \infty} \alpha^{1-\frac{d}{2}} \nabla E_{\tilde{T},\alpha}(t) = \nabla \tilde{D}(t) \quad (5)$$

As a result, the embedding T that minimizes TLC as $\alpha \rightarrow \infty$ tends to minimize the Dirichlet energy from the auxiliary simplices \tilde{T} to T . This minimizer is otherwise known as the Harmonic embedding from \tilde{T} to T .

In sum, the embedding that minimizes TLC straddles between the minimizer of the Dirichlet energy (for large α) and the MIPS energy (for small α). We visually depict this relation in two-dimensions in Figure 4 for two choices of \tilde{T} , either equilateral triangles of uniform size (top row) or triangles in the rest mesh (bottom row). We call the resulting energy with these two choices *uniform-TLC* and *rest-TLC*. Observe that as α increases, the minimizer of uniform-TLC approaches Tutte’s embedding of the rest mesh into the target boundary (blue boxes in Figure 4), which minimizes the Dirichlet energy of the map from the set of equilateral triangles. On the other hand, the embedding minimizing rest-TLC approaches the MIPS mapping of the rest mesh for small α (magenta boxes), and the harmonic mapping of rest mesh for large α (red boxes).

5 ALGORITHM

The TLC energy’s simplicity and smoothness enable efficient closed-form computation of its gradient and Hessian (see Appendix C). We

thus are able to explore minimization of TLC via a range of higher-order, nonlinear optimization methods. We begin by testing TLC optimization with both quasi-Newton and Newton-type strategies. For our quasi-Newton (QN) method we employed a standard, off-the-shelf limited-memory BFGS solver [Wright and Nocedal 1999]. Alternately, to exploit second-order information via Newton-type minimization more care is needed as the TLC Hessian can be indefinite. Rather than applying global offsets [Wright and Nocedal 1999] (as is standard in optimization packages) we follow recent developments in distortion optimization [Li et al. 2019; Teran et al. 2005] and ensure positive-definiteness of the global Hessian matrix by projecting per-simplex Hessians to positive-definite prior to assembly. The resulting projected-Newton (PN) method employs standard back-tracking line search and a direct solver [Chen et al. 2008] for each linear system solution.

Our theoretical guarantee applies only at global minima, and at potentially impractically small α . Given the nonlinearity and non-convexity of TLC, for a given choice of α optimizers may take a long time to reach convergence and could converge to non-injective local minima. Since our goal is solely to find an injective configuration, we choose to stop optimization as soon as the optimization encounters an injective solution along its search path. In turn, differing optimization methods can take dramatically different search paths that vary across examples so that their ability to reach injectivity differs per example.

In this regard we empirically observe that QN and PN are complementary. Due to its efficiency QN is faster at finding injectivity than PN on most examples, while PN then easily resolves the few examples where QN struggles to reach injectivity - albeit with increased computational cost. Motivated by this observation, we employ a two-stage solving method as follows (we use $N = 10000$):

- (1) Run QN solver for at most N iterations. If an injective mesh is found during the process, return this mesh as result. If the solver converges but the mesh is non-injective, report failure (i.e., a non-injective local minimum). If neither has happened after N iterations, then:

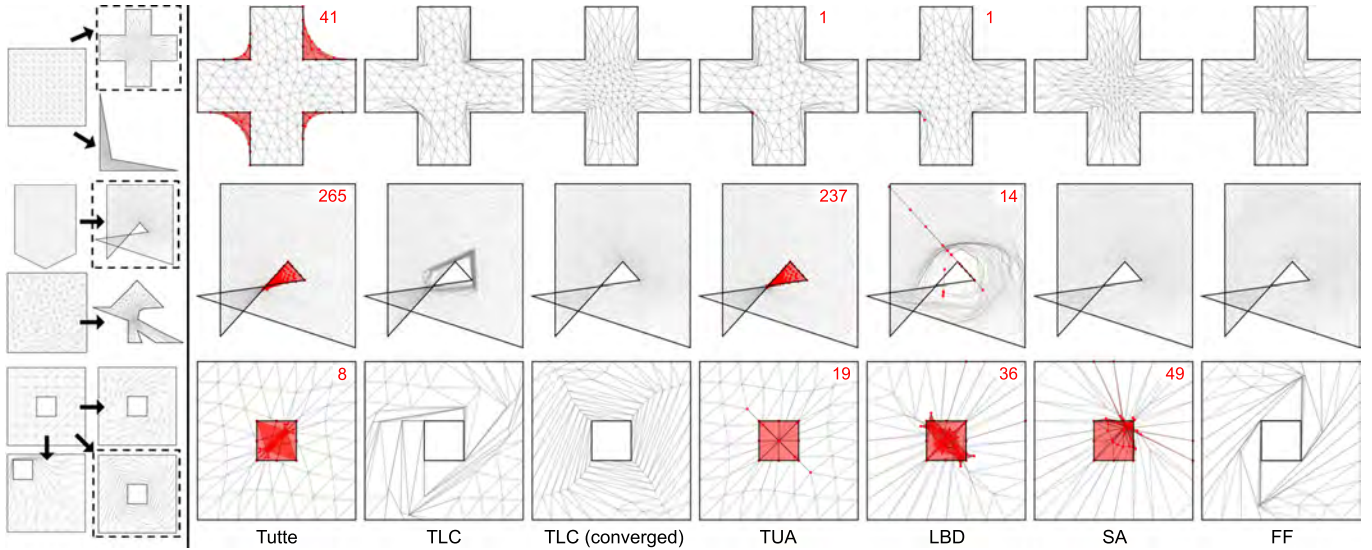


Fig. 5. Far-left: List of hand-made examples in the 2D parameterization category, each indicated by an arrow that goes from the rest mesh to the target domain. Middle: Three examples (top, middle, and bottom rows) from the list (highlighted in dotted lines), each showing Tutte embedding, result of minimizing TLC (stopped at injectivity), minimizing TLC (stopped at convergence), minimizing total unsigned area (TUA), and results of LBD, SA and FF. All inverted triangles and their vertices are colored red, and the total number of inverted triangles is marked in red.

- (2) Re-solve the starting problem by running PN solver for at most N iterations. As above, either return the first encountered injective mesh or report failure if the solver converges to a non-injective mesh. If neither happens after N iterations, report failure as well.

Initial configurations also have a significant impact on the search path taken by the solver, which again impact the ability of the optimization to reach injectivity. Ideally we would like to pick an initial configuration that is close to the global minimum when possible. A natural candidate is then Tutte embedding. While Tutte embedding cannot produce injective configurations for non-convex boundaries, we observe that the number of flipped elements are often small and concentrated near the boundary. In addition, as we discuss in the previous section, Tutte embedding is closely related to TLC when the auxiliary simplices are chosen to be uniformly sized regular elements. We evaluate the performance of our method under varying initializations in Section 6.5.

6 RESULTS

In this section, we evaluate our method on an extensive set of embedding problems for both triangular and tetrahedral meshes. We plan to distribute both our code and the data sets on the authors' websites.

6.1 Parameter choices

Our energy is controlled by two sets of parameters, the auxiliary simplices \tilde{T} and the scaling α . While the theoretical properties of our energy, smoothness and injective global minima, hold regardless

of the choice of \tilde{T} , different \tilde{T} result in not only different energy-minimizing embeddings (as studied in Section 4.2.4) but also different energy landscapes, and hence varied success rate for specific solvers. We have found that the QN solver in NLOpt is generally more successful in reaching injectivity when \tilde{T} consists of uniformly sized equilateral triangles or tetrahedra (*i.e.*, *uniform-TLC*) than when \tilde{T} is set to be the rest mesh M (*i.e.*, *rest-TLC*), for the same α values. In this work, we report the performance of uniform-TLC in both 2D and 3D, and leave a thorough investigation into the effect of \tilde{T} on the performance of different solvers for future work.

We observed that smaller α generally lead to higher chance of reaching injectivity by energy minimization (e.g., see Figure 4). This is in line with our theoretical guarantee of injectivity, which applies to only sufficiently small α . However, a too-small α results in a TLC energy that is very close to TUA and thus having close-to-vanishing gradients, which can significantly slow down the progress of gradient-based solvers. In our experiments, we found that choosing α such that the total content of \tilde{T} is 10^{-6} times of the content of the target domain strikes a good balance between success rate and efficiency, in both 2D and 3D.

6.2 A benchmark problem set

While existing injective embedding methods have all been tested on non-trivial examples, these examples are often different, making it challenging to perform direct comparisons between different methods. One of this work's contribution is developing an extensive problem set for fixed-boundary embedding in both 2D and 3D. The set includes as many examples as possible from existing works that we have access to, plus hundreds of new examples that we created. We hope that our problem set offers a benchmark for future research in this area.

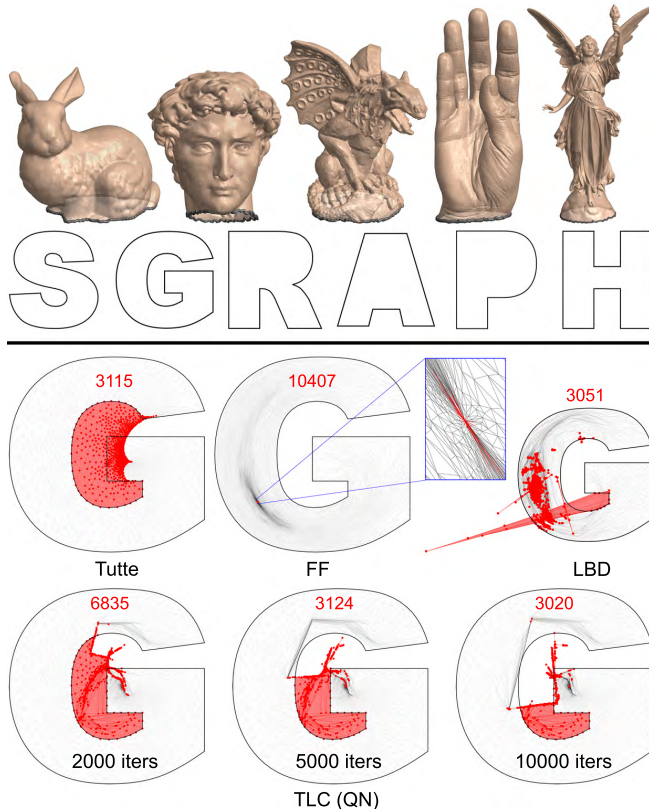


Fig. 6. Top: rest meshes and letter-like boundaries used to create the embedding problems in the 2D parameterization category. Bottom: results of various methods on embedding Lucy into letter “G”, including our QN solver at different iterations. The zoom-in takes a close look at the inverted triangles in the result of FF (all close to being collinear).

To provide fair and practically relevant evaluations, the data set is built with two criteria in mind. First, the problem should be feasible, meaning that an injective mapping satisfying the boundary constraints exists. If such mapping does not come with a problem, we ran all available methods (ours included) and include the problem in our data set if *any* method succeeded (that is, we did not exclude any example that our method failed while some other methods succeeded). Second, the problem should not be trivial. In particular, we discard all examples where Tutte’s embedding is already injective. Our data set consists of three broad categories that correspond to different use scenario of injective embedding:

- 2D parameterization: This category contains a few hand-made toy examples (Figure 5 far left), including two (in the middle) from previous works [Weber and Zorin 2014], and 30 examples that embed five surfaces (between 20K-50K vertices each) into the outline of six letters S,G,R,A,P,H (Figure 6 top; also in teaser). The bulk of this category comes from the impressive dataset put together by [Liu et al. 2018], which includes more than twenty thousand open surface meshes. For each mesh, we created an intersection-free target boundary using the

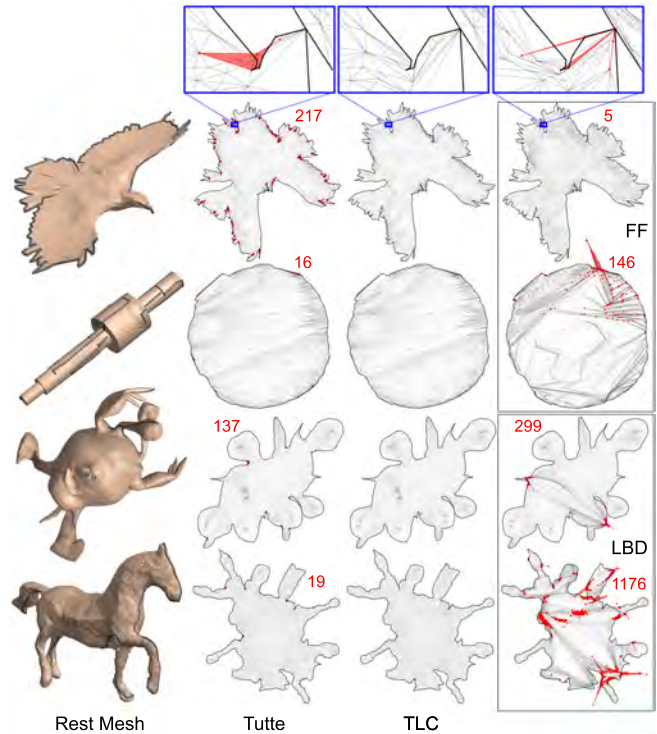


Fig. 7. Four examples in the 2D parameterization category derived from [Liu et al. 2018], where methods FF and LBD failed to find injective embeddings.

method of [Jiang et al. 2017], and included it in our data set if it is not solved by Tutte embedding (Figure 7).

- 3D parameterization: This category includes 116 polycube embedding problems from [Aigerman and Lipman 2013; Fu et al. 2016], 20 spherical mapping and 40 free-surface mapping problems from [Su et al. 2019] (Figure 8).
- 3D deformation: Simulating physical deformation often yields severely distorted meshes for which it is challenging to avoid inversion. Hence they provide an excellent test set for a method’s ability to recover from inverted elements. We capture frames from non-inverting simulations of twisting, generated by [Li et al. 2020], starting from a rest shape and then twisting it to introduce increasing levels of distortion. For each simulation we embed the starting mesh to the boundary surface of each successive frame, obtaining examples with generally increasing challenge. This category includes 728 examples from deforming a rod (Figure 9), a cube (Figure 10), and the armadillo (Figure 11; also in teaser).

6.3 Benchmark comparisons

We compare our method (TLC) to the three most, to our knowledge, competitive methods for injective embedding into a target domain without the need of an injective initialization: Large-scale Bounded Distortion Mappings (LBD) [Kovalsky et al. 2015], Simplex Assembly (SA) [Fu and Liu 2016] and Foldover-Free Volumetric Mapping (FF) [Su et al. 2019]. LBD needs to set the upper bound K on the




Fig. 8. Three examples from the 3D parameterization category, each mapping a rest tetrahedral mesh into a sphere (top), smooth surface (middle), and a polycube (bottom). Each example is a failure case for one of the three methods, FF, LBD and SA. Inverted tetrahedra are colored in red, and the numbers of inversion are marked in red.




Fig. 9. One example in the 3D deformation category (a twisting rod). All methods succeeded on this sequence.

distortion. While in theory an extremely large K is more permissive from a small one, in practice large values prevent convergence of the method [Kovalsky et al. 2015]. To choose K , we compute the maximal distortion between the rest mesh and the known injective solution (either coming with the example or produced by one of other methods), and set K to be double the maximal distortion. This guarantees the existence of an embedding with lower distortion than the upper bound. For SA and FF, we use the default parameters suggested by the authors. SA is not tested on examples that map a surface mesh to a planar boundary, since the corresponding code is not available (after confirmation with the authors of SA). For consistency, we feed all methods with the (non-injective) Tutte embedding as the initial map.

We are primarily interested in how often a method succeeds. Here, success is defined as producing an embedding without any inverted elements. As shown in Table 1, while every other method exhibits varied success across different categories, our method consistently solved all examples in each category. For our method, we also report (in parenthesis) the number of examples (if there are any) that

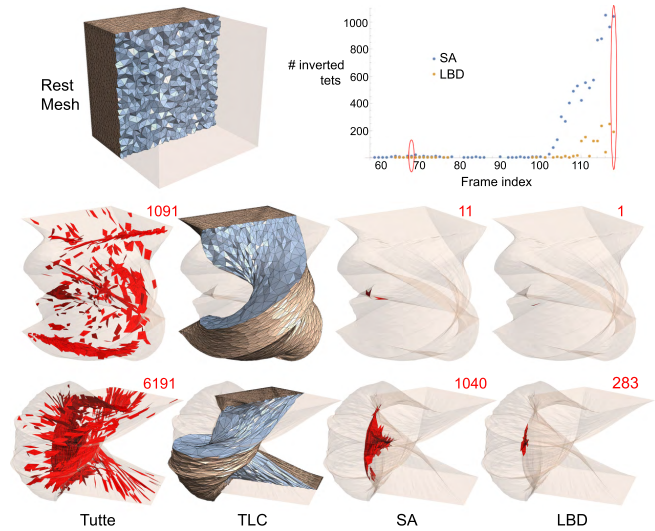


Fig. 10. Two example in the 3D deformation category (a twisting cube) where both SA and LBD failed to reach injectivity. The graph in the top-right shows the number of inverted tetrahedra for each of the 118 frames of the deformation sequence (ellipses indicate the frames from which the two examples were taken).

Table 1. Number of failed examples for each compared method in each category of the benchmark problem set.

		Total	TLC	FF	SA	LBD
2D (Param.)	Simple	7	0	0	1	6
	Letters	30	0 (17)	7	N/A	26
	[Liu et al. 2018]	10706	0 (19)	1292	N/A	267
3D (Param.)	Polycube	116	0	3	29	2
	Sphere	20	0	1	5	13
	Surface	40	0	0	0	6
3D (Deform)	Rod	8	0	0	0	0
	Cube	119	0	0	54	32
	Armadillo	601	0	75	177	378

require the second stage of the algorithm, which uses the PN solver. In the following, we present and discuss specific instances of each category of the benchmark.

2D parameterization. The hand-made examples shown in Figure 5 are designed to test an embedding algorithm’s ability to deal with large boundary deformations (e.g., convex-to-concave deformations, as in top and middle) and transformation of inner boundaries (e.g., the bottom example, where the inner square is rotated by 180 degrees). We found that both SA and (particularly) LBD tend to fail in these scenarios. Note that since our method stops as soon as (local) injectivity is obtained, the resulting embedding may contain many triangles with small angles. While reducing distortion is not our objective, for illustrative purposes, we show that the shape of such triangles significantly improves if optimization continues (see “TLC (converged)” column). This echoes our observations in Section 4.2.4 about the connection between the global minima of TLC and distortion-minimization maps. For these examples, we also show

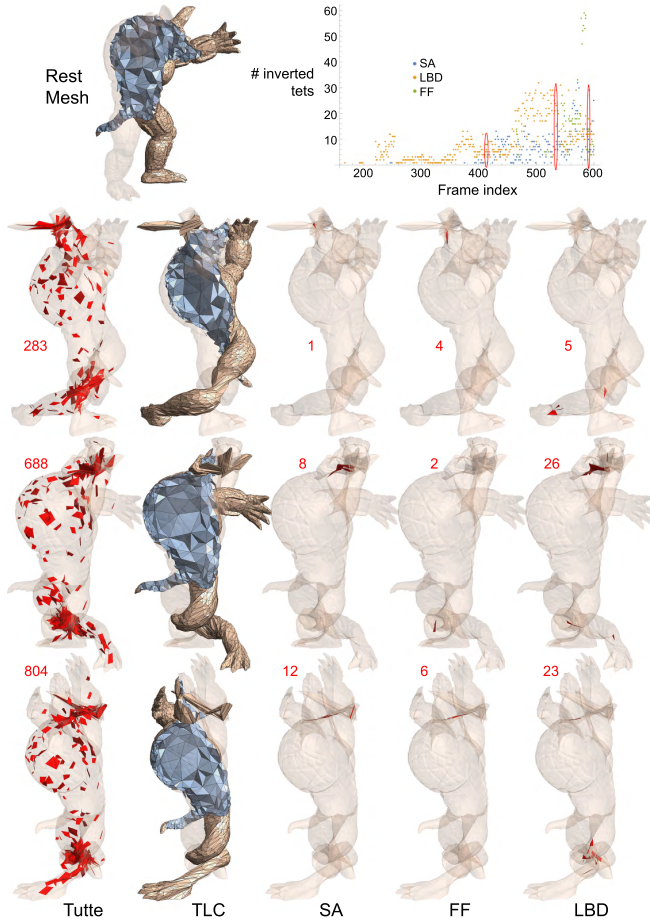


Fig. 11. Three example in the 3D deformation category (a twisting armadillo) where FF, SA and LBD all failed to reach injectivity. The graph in the top-right shows the number of inverted tetrahedra for each of the 600+ frames of the deformation sequence (ellipses indicate the frames from which the three examples were taken).

the results of minimizing the total unsigned area (TUA) energy, by setting $\alpha = 0$ in our method. Observe that the optimization easily gets stuck in local minima, where the gradient of the TUA vanishes.

The second group of examples, which map a surface mesh to a letter outline (Figure 6 top), prove challenging for all methods. Due to the highly detailed surfaces and the non-convex boundaries, achieving injectivity in these examples necessarily comes at the cost of significant triangle distortions. As shown in Figure 6 (bottom) for the example of mapping Lucy into the letter “G” (same as in Figure 1), Tutte embedding contains a large region of inverted triangles. Both LBD and FF failed, and the QN solver in the first stage of our algorithm failed to achieve injectivity within the maximum number of iterations. While it appears that QN might eventually reach injectivity if given more iterations to run, the PN solver in the second stage of our algorithm successfully finds an injective mesh with only 66 iterations (see result in Figure 1).

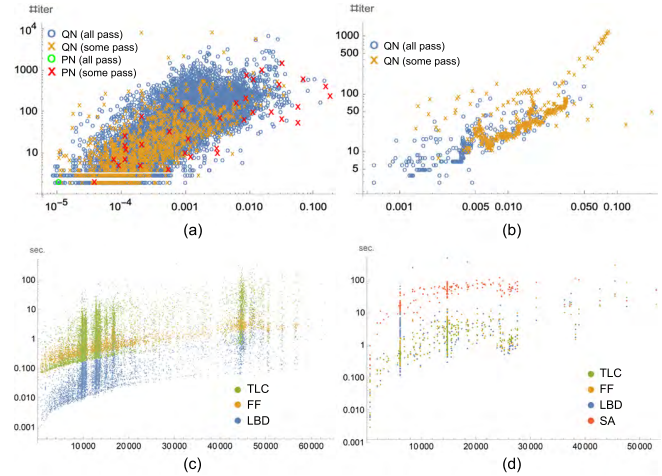


Fig. 12. Top: number of iterations taken by our PN or QN solvers over all benchmark examples in 2D (a) and 3D (b), where the horizontal axis is the ratio of the number of inverted simplices in the initial (Tutte) embedding over the total number of simplices (both axes are in log-scale). In the legend, “all pass” means that FF, LBD, and SA (3D) all succeed on that example; “some pass” means that at least one of these methods fails (our method succeeds on all examples). Bottom: running time (in seconds) of all methods in 2D (c) and 3D (d), where the horizontal axis is the number of vertices in the rest mesh (running time is in log scale).

In contrast to the surface-to-letter problems, the examples from [Liu et al. 2018] map a complex surface to a domain that is already adapted to the shape of the surface (created by a distortion-minimizing parameterization algorithm [Jiang et al. 2017]). As a result, Tutte embedding tends to create only small inverted triangles close to the target boundary (Figure 7). However, we observed that both FF and LBD often make the initial embedding worse by creating many more inverted triangles in the interior of the domain. In contrast, our method removes inversions without significantly impacting the interior triangulation.

3D parameterization. The embedding problem becomes significantly harder in 3D due to the added dimension. Tutte embedding, which is no longer guaranteed to be injective for convex 3D domains, generates inverted elements much more often than in 2D and in greater quantity. We found that existing methods are fairly effective in removing the majority of the inverted tetrahedra in this category (see Table 1). However, no method was able to resolve all inversions in all examples, except for ours. A few examples that compare our method with LBD, FF or SA in their failure cases are shown in Figure 8.

3D deformation. For examples arising from deforming sequences, existing methods struggle with many frames, particularly on more complex meshes and towards the end of the deforming sequence where there is a large amount of twisting (see Figures 10 and 11). On the other hand, our method successfully found injective mappings for all frames in the three deformation sequences.

Performance. We visualize the number of iterations taken by our QN solver or PN solver (for examples that QN reaches $N = 10000$




Fig. 13. Mapping a triangulated Hele-Shaw polygon (from [Shen et al. 2019]) to a square. Tutte embedding contains 46 inverted or co-linear triangles due to numerical errors (they are all along the boundary). Our method (QN solver) successfully restores the injectivity in 65 iterations.

iterations) to converge for every example in our benchmark in Figure 12 (a,b). Observe that our method generally requires more iterations when the initial (Tutte) embedding contains more inverted elements. On the other hand, examples where at least one other method (FF, SA, or LBD) fails generally do not pose particular challenges for our method, although they tend to require more iterations in 3D. We also show in Figure 12 (c,d) the wall-clock running time of all methods being compared (TLC, FF, SA, LBD) on the subset of the benchmark where all methods succeed. Observe that our method is comparable with FF but more expensive than LBD in 2D, while on-par with both FF and LBD but faster than SA in 3D. All experiments are performed on a Windows PC with Intel core-i7-4770 CPU at 3.40 GHz and 32GB of memory.

6.4 Comparison with other methods

We further compare our method with a few other injective mapping methods in 2D. Unlike methods compared above (FF, SA, LBD, and ours), these methods either may modify the mesh structure or are limited to convex target domains. We first compare with the method of [Weber and Zorin 2014], which uses an intermediate domain for co-parameterization and occasionally requires refinement of the input mesh. Our method (with QN) successfully found injective maps for all examples in that paper (provided by the authors), except those that do not admit feasible solutions (e.g., Figures 1 and 9 in their paper). Two examples were included in the “2D parameterization” category of our benchmark (middle row of Figure 5). For these two examples, [Weber and Zorin 2014] needs to add 2 and 9 new vertices to achieve injectivity. We next consider the recent method of [Shen et al. 2019], which restores injectivity to Tutte embeddings without mesh refinement but only for convex target domains. Our method can serve the same purpose but is not limited to convex domains. Figure 13 shows our result on a complex example from [Shen et al. 2019] (mapping a Hele-Shaw polygon to a square).

6.5 Impact of initial embedding

Since our energy is non-convex, the choice of the initial mesh could significantly affect the solver’s behavior. The effect can manifest in two ways, (1) time to reach injectivity, and (2) the shape of the embedding when injectivity is reached. We demonstrate these effects in Figures 14 and 15. In Figure 14, we ran our method on the same set of input but starting from three different types of initial embeddings, which result in rather different injective maps. Interestingly,

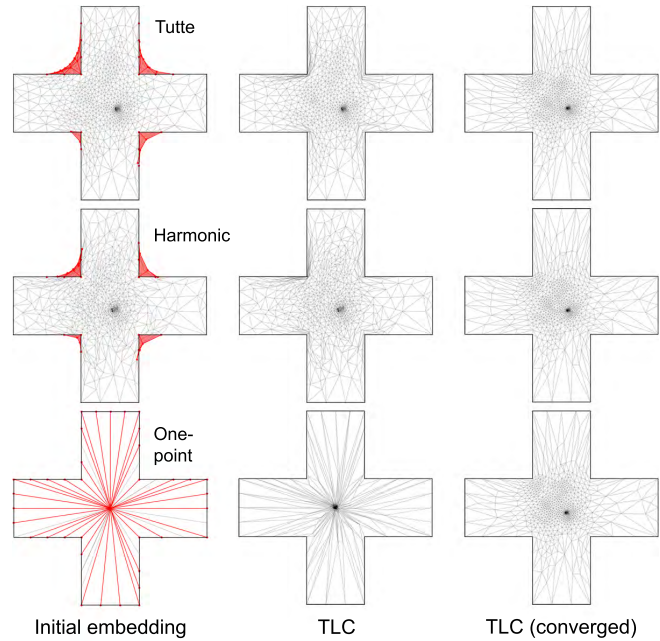


Fig. 14. Mapping a Bunny shape (same as Figure 4) into a cross with three different initial embeddings, Tutte embedding (top), harmonic map (middle), and one in which all interior vertices are collapsed onto a single point (bottom). Showing results of our method after terminating upon injectivity (middle) and after the energy has converged (right).




Fig. 15. Mapping a multi-boundary mesh using an initialization that has all vertices collapsed to a single point. The QN solver fails to reach injectivity within allowed number of iterations, while PN succeeds.

the solver produces very similar embeddings if we let it run until convergence of the energy (instead of terminating at injectivity). Figure 15 shows another example with a multi-connected target domain and a pathological initial embedding (note that the Tutte embedding is already injective). The QN solver failed to find an injective embedding within the maximum number of iterations, and while PN solver succeeded in reaching injectivity, it took nearly 5000 iterations.

6.6 Initializing distortion minimization

Many distortion-minimization algorithms can preserve injectivity if given a feasible injective initial map. Combining such an algorithm with our method allows to find injective *and* distortion-minimizing embeddings for challenging inputs that are challenging for existing approaches. We demonstrate such an example in Figure 16, where we




Fig. 16. Minimizing isometric distortion (right) starting from an injective mapping of Lucy into the letter “P” provided by our method (left).

compute an embedding of the Lucy model into the outline of letter “P” by first finding an injective embedding using our method (left), followed by minimizing an isometric energy (Symmetric Dirichlet) using a standard method [Zhu et al. 2018] (right). Observe from the zoom-ins that the isometric distortion is significantly reduced from the initial mapping produced by our method.

7 DISCUSSION

There are a number of promising avenues of future research. First, while we demonstrated success using only standard off-the-shelf solvers (quasi-Newton and projected Newton) for energy minimization, developing customized solvers has the potential to significantly improve convergence rate of the energy and hence reducing the time needed to reach injectivity. A second and related direction is exploring different types of auxiliary simplices, which can have a significant impact on the energy landscape and hence the search paths of the solver.

Last but not least, so far our theoretical analysis has only considered the injectivity of *global* minima of TLC. An even more interesting and practically relevant subject is characterizing the *local* minima and their injectivity. Interestingly, we did not observe any non-injective local minima in the benchmark data set (since our solver never converged to a non-injective embedding). Further theoretical investigation might lead to a better understanding of the energy landscape and in turn more effective solving techniques.

ACKNOWLEDGMENTS

This work is supported in part by NSF grant RI-1618685, NIH grant U2C CA233303-1, and Simons Math+X Investigators Award 400837. We would like to thank authors of several papers for providing code, data, and help with comparisons, and especially Hanxiao Shen, Ofir Weber, Alon Bright, Zohar Levi, and Xiao-Ming Fu.

REFERENCES

Pankaj K Agarwal, Bardia Sadri, and Hai Yu. 2008. Untangling triangulations through local explorations. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*. ACM, 288–297.

Noam Aigerman and Yaron Lipman. 2013. Injective and bounded distortion mappings in 3D. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 106.

Noam Aigerman and Yaron Lipman. 2015. Orbifold Tutte embeddings. *ACM Trans. Graph.* 34, 6 (2015), 190–1.

Marcel Campen, Cláudio T Silva, and Denis Zorin. 2016. Bijective maps from simplicial foliations. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 74.

Yanqing Chen, Timothy A Davis, William W Hager, and Sivasankaran Rajamanickam. 2008. Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate. *ACM Transactions on Mathematical Software (TOMS)* 35, 3 (2008), 22.

Sebastian Clatici, Mikhail Bessmeltsev, Scott Schaefer, and Justin Solomon. 2017. Isometry-Aware Preconditioning for Mesh Parameterization. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 37–47.

Michael Floater. 2003. One-to-one piecewise linear mappings over triangulations. *Math. Comp.* 72, 242 (2003), 685–696.

Michael S Floater and Kai Hormann. 2005. Surface parameterization: a tutorial and survey. In *Advances in multiresolution for geometric modelling*. Springer, 157–186.

Xiao-Ming Fu, Chong-Yang Bai, and Yang Liu. 2016. Efficient Volumetric PolyCube-Map Construction. *Comput. Graph. Forum* 35, 7 (2016), 97–106.

Xiao-Ming Fu and Yang Liu. 2016. Computing inversion-free mappings by simplex assembly. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 216.

Xiao-Ming Fu, Yang Liu, and Baining Guo. 2015. Computing locally injective mappings by advanced MIPS. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 71.

Steven Gortler, Craig Gotsman, and Dylan Thurston. 2006. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design* (2006).

Xianfeng Gu, Ren Guo, Feng Luo, Jian Sun, Tianqi Wu, et al. 2018. A discrete uniformization theorem for polyhedral surfaces II. *Journal of differential geometry* 109, 3 (2018), 431–466.

Eden Fedida Hefetz, Edward Chien, and Ofir Weber. 2019. A Subspace Method for Fast Locally Injective Harmonic Mapping. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 105–119.

Kai Hormann and Günther Greiner. 2000. MIPS: An Efficient Global Parameterization Method. *France on 1-7 July 1999. Proceedings, Volume 1. Curve and Surface Design. F61775-99-WF068* (2000), 153.

Kai Hormann, Bruno Lévy, and Alla Sheffer. 2007. Mesh parameterization: Theory and practice. (2007).

Zhongshi Jiang, Scott Schaefer, and Daniele Panozzo. 2017. Simplicial complex augmentation framework for bijective maps. *ACM Transactions on Graphics* 36, 6 (2017).

Shahar Z Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. 2015. Large-scale bounded distortion mappings. *ACM Trans. Graph.* 34, 6 (2015), 191–1.

Minchen Li, Zachary Ferguson, Teso Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental potential contact: intersection- and inversion-free, large-deformation dynamics. *ACM Transactions on Graphics (TOG)* 39, 4 (2020).

Minchen Li, Ming Gao, Timothy Langlois, Chenfanfu Jiang, and Danny M. Kaufman. 2019. Decomposed Optimization Time Integrator for Large-Step Elastodynamics. *ACM Transactions on Graphics* 38, 4 (2019).

Yaron Lipman. 2014. Bijective Mappings of Meshes with Boundary and the Degree in Mesh Processing. *SIAM Journal on Imaging Sciences [electronic only]* 7 (04 2014). <https://doi.org/10.1137/130939754>

Ligang Liu, Chunyang Ye, Ruiqi Ni, and Xiao-Ming Fu. 2018. Progressive parameterizations. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 41.

Tiantian Liu, Ming Gao, Lifeng Zhu, Eftychios Sifakis, and Ladislav Kavan. 2016. Fast and Robust Inversion-Free Shape Manipulation. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 1–11.

Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 2, 1 (1993), 15–36.

Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Scalable locally injective mappings. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 37a.

Christian Schüller, Ladislav Kavan, Daniele Panozzo, and Olga Sorkine-Hornung. 2013. Locally injective mappings. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing*. Eurographics Association, 125–135.

Hanxiao Shen, Zhongshi Jiang, Denis Zorin, and Daniele Panozzo. 2019. Progressive embedding. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 32.

Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahar Z Kovalsky, and Yaron Lipman. 2017. Geometric optimization via composite majorization. *ACM Trans. Graph.* 36, 4 (2017), 38–1.

Jason Smith and Scott Schaefer. 2015. Bijective parameterization with free boundaries. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 70.

Jian-Ping Su, Xiao-Ming Fu, and Ligang Liu. 2019. Practical Foldover-Free Volumetric Mapping Construction. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 287–297.

Joseph Teran, Eftychios Sifakis, Geoffrey Irving, and Ronald Fedkiw. 2005. Robust quasistatic finite elements and flesh simulation. In *Proceedings of the 2005 ACM*

- SIGGRAPH/Eurographics symposium on Computer animation. ACM, 181–190.
- William Thomas Tutte. 1963. How to draw a graph. *Proceedings of the London Mathematical Society* 3, 1 (1963), 743–767.
- Ofir Weber, Ashish Myles, and Denis Zorin. 2012. Computing extremal quasiconformal maps. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 1679–1689.
- Ofir Weber and Denis Zorin. 2014. Locally injective parametrization with arbitrary fixed boundaries. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 75.
- Stephen J Wright and Jorge Nocedal. 1999. *Numerical optimization*. Vol. 2. Springer New York.
- Yin Xu, Renjie Chen, Craig Gotsman, and Ligang Liu. 2011. Embedding a triangular graph within a given boundary. *Computer Aided Geometric Design* 28, 6 (2011), 349–356.
- Yufeng Zhu, Robert Bridson, and Danny M Kaufman. 2018. Blended cured quasi-newton for distortion optimization. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 40.

A LIFTED CONTENT AND DIRICHLET ENERGY

We shall prove Propositions 4.4 and 4.5 separately in 2D and 3D. In the process, we will derive alternative expressions of the lifted content of a simplex that reveal the connections to Dirichlet energy. These expressions lead to derivative formula that are used in the proofs of both this section and the next (on injectivity).

A.1 The 2D case

We first re-write the lifted content for a single triangle t given an auxiliary triangle \tilde{t} . Let X, \tilde{X}, \tilde{X} be the edge vectors of t, \tilde{t}, \tilde{t} as defined in Section 4.2, and L be the linear transformation from t to \tilde{t} (i.e., $\tilde{X} = LX$). We obtain

$$\begin{aligned} \text{Det}(X^T X + \alpha \tilde{X}^T \tilde{X}) &= \text{Det}(X^T X + \alpha X^T L^T L X) \\ &= \text{Det}(X^T (I + \alpha L^T L) X) \\ &= \text{Det}(X)^2 \text{Det}(I + \alpha L^T L) \\ &= \text{Det}(X)^2 (1 + \alpha \text{Tr}(L^T L) + \alpha^2 \text{Det}(L^T L)) \\ &= \text{Det}(X)^2 + \alpha \text{Det}(X)^2 \text{Tr}(L^T L) + \alpha^2 \text{Det}(\tilde{X})^2 \end{aligned} \quad (6)$$

where I is the identify matrix and Tr denotes the matrix trace. The third and fifth identities are due to the multiplicativity of determinants (i.e., $\text{Det}(AB) = \text{Det}(A)\text{Det}(B)$), and the fourth identity can be verified by hand (note that $L^T L$ is a 2×2 matrix).

Let $A(t)$ be the unsigned area of t , $A_{\tilde{t}}$ be the unsigned area of \tilde{t} , and $D(t)$ be the Dirichlet energy from t to \tilde{t} . Note that $D(t) = \text{Det}(X)\text{Tr}(L^T L)/4$ [Pinkall and Polthier 1993]. Substituting equation (6) into (2) gives the following equivalent expression for $E_{\tilde{t},\alpha}$,

$$E_{\tilde{t},\alpha}(t) = \sqrt{A(t)^2 + 2\alpha A(t)D(t) + \alpha^2 A_{\tilde{t}}^2} \quad (7)$$

By swapping t with \tilde{t} and repeating the same derivation, we get a symmetric expression using the Dirichlet energy $\tilde{D}(t)$ from \tilde{t} to t :

$$E_{\tilde{t},\alpha}(t) = \sqrt{A(t)^2 + 2\alpha A_{\tilde{t}} \tilde{D}(t) + \alpha^2 A_{\tilde{t}}^2} \quad (8)$$

We now prove both Propositions 4.4 and 4.5 using the two expressions above:

PROOF. Using expression (7), the derivative of $E_{\tilde{t},\alpha}(t)$ with respect to α has the form:

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial \alpha} = \frac{2A(t)D(t) + 2\alpha A_{\tilde{t}}^2}{2\sqrt{A(t)^2 + 2\alpha A(t)D(t) + \alpha^2 A_{\tilde{t}}^2}} \quad (9)$$

At $\alpha = 0$, and since $A(t) > 0$ by assumption of Proposition 4.4, we have:

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial \alpha} \Big|_{\alpha=0} = \frac{A(t)D(t)}{\sqrt{A(t)^2}} = D(t) \quad (10)$$

which proves Proposition 4.4.

Next, using expression (8), we find the gradient $\nabla E_{\tilde{t},\alpha}(t)$ with respect to t as

$$\begin{aligned} \nabla E_{\tilde{t},\alpha}(t) &= \frac{2A(t)\nabla A(t) + 2\alpha A_{\tilde{t}} \nabla \tilde{D}(t)}{2\sqrt{A(t)^2 + 2\alpha A_{\tilde{t}} \tilde{D}(t) + \alpha^2 A_{\tilde{t}}^2}} \\ &= \frac{\frac{2}{\alpha} A(t)\nabla A(t) + 2A_{\tilde{t}} \nabla \tilde{D}(t)}{2\sqrt{\frac{1}{\alpha^2} A(t)^2 + \frac{2}{\alpha} A_{\tilde{t}} \tilde{D}(t) + A_{\tilde{t}}^2}} \end{aligned}$$

As $\alpha \rightarrow \infty$, we evaluate the limit of the gradient as

$$\lim_{\alpha \rightarrow \infty} \nabla E_{\tilde{t},\alpha}(t) = \frac{A_{\tilde{t}} \nabla \tilde{D}(t)}{\sqrt{A_{\tilde{t}}^2}} = \nabla \tilde{D}(t) \quad (11)$$

which proves Proposition 4.5 for $d = 2$. \square

A.2 The 3D case

Let X, \tilde{X}, \tilde{X} be the edge vectors of tetrahedra t, \tilde{t}, \tilde{t} as defined in Section 4.2, and L, \tilde{L} be the linear transformations from t to \tilde{t} and from \tilde{t} to t (i.e., $\tilde{L} = L^{-1}$). Using a derivation similar to that in Equation 6, we obtain the identity:

$$\begin{aligned} \text{Det}(X^T X + \alpha \tilde{X}^T \tilde{X}) &= \text{Det}(X)^2 + \alpha \text{Det}(X)^2 \text{Tr}(L^T L) + \alpha^2 \text{Det}(\tilde{X})^2 \text{Tr}(\tilde{L}^T \tilde{L}) + \alpha^3 \text{Det}(\tilde{X})^2 \end{aligned} \quad (12)$$

Let $V(t)$ be the unsigned volume of t , $V_{\tilde{t}}$ be the unsigned volume of \tilde{t} , $D(t)$ be the Dirichlet energy from t to \tilde{t} , and $\tilde{D}(t)$ be the Dirichlet energy transforming \tilde{t} to t . Note that $D(t) = \text{Det}(X)\text{Tr}(L^T L)/12$ and $\tilde{D}(t) = \text{Det}(\tilde{X})\text{Tr}(\tilde{L}^T \tilde{L})/12$ [Pinkall and Polthier 1993]. Substituting these identifies into equation (12) and then into (2) gives the following equivalent expression for $E_{\tilde{t},\alpha}(t)$,

$$E_{\tilde{t},\alpha}(t) = \sqrt{V(t)^2 + 2\alpha V(t)D(t) + 2\alpha^2 V_{\tilde{t}} \tilde{D}(t) + \alpha^3 V_{\tilde{t}}^2} \quad (13)$$

Note that there are two equivalent expressions in 2D (7,8), each using the Dirichlet energy in one direction (D or \tilde{D}), whereas the 3D expression above contains the Dirichlet energy in both directions.

We now prove both Propositions 4.4 and 4.5 using the two expressions above:

PROOF. We first obtain the derivative of $E_{\tilde{t},\alpha}(t)$ with respect to α as

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial \alpha} = \frac{2V(t)D(t) + 4\alpha V_{\tilde{t}} \tilde{D}(t) + 3\alpha^2 V_{\tilde{t}}^2}{2\sqrt{V(t)^2 + 2\alpha V(t)D(t) + 2\alpha^2 V_{\tilde{t}} \tilde{D}(t) + \alpha^3 V_{\tilde{t}}^2}} \quad (14)$$

At $\alpha = 0$, and since $V(t) > 0$ by assumption of Proposition 4.4, we obtain

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial \alpha} \Big|_{\alpha=0} = \frac{V(t)D(t)}{\sqrt{V(t)^2}} = D(t) \quad (15)$$

which proves Proposition 4.4.

Next, we obtain the gradient $\nabla E_{\tilde{t},\alpha}^-(t)$ with respect to t as

$$\begin{aligned}\nabla E_{\tilde{t},\alpha}^-(t) &= \frac{2V(t)\nabla V(t)+2\alpha\nabla V(t)D(t)+2\alpha V(t)\nabla D(t)+2\alpha^2 V_{\tilde{t}}\nabla\tilde{D}(t)}{2\sqrt{V(t)^2+2\alpha V(t)D(t)+2\alpha^2 V_{\tilde{t}}\tilde{D}(t)+\alpha^3 V_{\tilde{t}}^2}} \\ &= \sqrt{\alpha} \frac{\frac{2}{\alpha^2} V(t)\nabla V(t)+\frac{2}{\alpha}\nabla V(t)D(t)+\frac{2}{\alpha} V(t)\nabla D(t)+2V_{\tilde{t}}\nabla\tilde{D}(t)}{2\sqrt{\frac{1}{\alpha^3} V(t)^2+\frac{2}{\alpha^2} V(t)D(t)+\frac{2}{\alpha} V_{\tilde{t}}\tilde{D}(t)+V_{\tilde{t}}^2}}\end{aligned}$$

As $\alpha \rightarrow \infty$, the limit of the gradient after scaling by $\alpha^{-1/2}$ is

$$\lim_{\alpha \rightarrow \infty} \alpha^{-1/2} \nabla E_{\tilde{t},\alpha}^-(t) = \frac{V_{\tilde{t}}\nabla\tilde{D}(t)}{\sqrt{V_{\tilde{t}}^2}} = \nabla\tilde{D}(t) \quad (16)$$

which proves Proposition 4.5 for $d = 3$. \square

B INJECTIVITY OF GLOBAL MINIMA

We shall prove the main result, Proposition 4.3, separately for dimensions $d = 2, 3$. In the discussions below, we assume that the auxiliary simplicies \tilde{T} are fixed and non-degenerate (i.e., having positive content). For notational simplicity, and since we are concerned with a specific embedding T , we use subscripts in place of the functional forms been used so far. That is, we use A_t (or V_t) to denote the unsigned area (or volume) of a triangle (or tetrahedron) $t \in T$, D_t the Dirichlet energy from t to its auxiliary triangle (or tetrahedron) \tilde{t} , and \tilde{D}_t the Dirichlet energy from \tilde{t} to t .

B.1 The 2D case

We will introduce several lemmas before proving the proposition. We start with some simple, but useful properties of the Dirichlet energy (see illustration in Figure 17):

LEMMA B.1. *Let e be one edge of t , h be the distance to e from the vertex opposite to e , and \tilde{h} be the corresponding distance in \tilde{t} . Then,*

- (1) $D_t \geq A_t(\tilde{h}/h)^2/2$
- (2) $h/|e| \geq \tilde{h}^2/4D_t$




Fig. 17. Illustration for Lemma B.1.

PROOF. Since the Dirichlet energy is invariant under rotation and translation, we first rotate and translate both t, \tilde{t} so that e and its corresponding edge \tilde{e} of \tilde{t} are both aligned with the X axis of the plane (see Figure 17). Note that $D_t = A_t \text{Tr}(L^T L)/2$, where L is the linear transformation matrix from t to \tilde{t} (as defined in Appendix A.1). With the aligned t, \tilde{t} , the bottom row of L is $\{0, \tilde{h}/h\}$. Hence $\text{Tr}(L^T L) \geq (\tilde{h}/h)^2$, proving the inequality in (1). The inequality in (2) can be derived from (1) by noting that $A_t = |e|h/2$. \square

The following lemma shows that if an embedding contains a sufficiently small triangle, it will contain some triangle t whose Dirichlet energy D_t (from t to \tilde{t}) is sufficiently large:

LEMMA B.2. *For any $\delta > 0$, there exists some $\epsilon > 0$ such that if an embedding T contains a triangle whose unsigned area is smaller than ϵ , then T must contain a triangle t such that $D_t > \delta$ and $A_t D_t > \epsilon\delta$.*

PROOF. We first show that if $D_t \leq \delta$ for all $t \in T$, then all edge lengths in T are lower bounded by some positive constant (invariant to the choice of T). For any edge e_0 in T , we find a sequence of triangles $\{t_1, \dots, t_k\} \subseteq T$ such that e_0 is incident to t_1, t_i and t_{i+1} share a common edge (denoted by e_i) for $i \in [1, k-1]$, and t_k contains an edge (denoted by e_k) on the target boundary B . See Figure 18 for an illustration. Consider triangle t_k . Let h be the height with respect




Fig. 18. Illustration for proof of Lemma B.2.

to base edge e_k , and \tilde{h} be the corresponding height in the auxiliary triangle \tilde{t}_k . Since $|e_{k-1}| \geq h$, $D_{t_k} \leq \delta$, and due to Lemma B.1(2),

$$\frac{|e_{k-1}|}{|e_k|} \geq \frac{h}{|e_k|} \geq \frac{\tilde{h}^2}{4D_{t_k}} \geq \frac{\tilde{h}^2}{4\delta}$$

Since e_k (being on the fixed boundary) has a fixed length invariant to T , and \tilde{h} is fixed, $|e_{k-1}|$ is lower-bounded by some constant. Iteratively applying this inequality to $|e_{k-2}|/|e_{k-1}|, \dots, |e_0|/|e_1|$ shows that $|e_0|$ is lower-bounded by some constant that depends only on the target boundary B , the auxiliary triangles \tilde{T} , and the combinatorial sequence of triangles $\{t_1, \dots, t_k\}$. Call a sequence *simple* if any triangle appears at most once in the sequence. Let σ_{e_0} be the minimal lower bound of $|e_0|$ over all possible simple triangle sequences, and let σ be any value smaller than the minimum of σ_{e_0} over all edges e_0 . Then σ lower-bounds the length of any edge in any embedding T where $D_t \leq \delta$ for all $t \in T$.

Next, consider an embedding T where $D_t \leq \delta$ for all $t \in T$, and consider any triangle $t \in T$. Let e be an edge of t , h be the height with respect to e , and \tilde{h} be the corresponding height in \tilde{t} . We can bound t 's unsigned area as follows (using Lemma B.1(2), $|e| > \sigma$, and $D_t \leq \delta$):

$$A_t = \frac{|e|^2 h}{2|e|} \geq \frac{|e|^2 \tilde{h}^2}{8D_t} > \frac{\sigma^2 \tilde{h}^2}{8\delta}$$

Define $\epsilon = \sigma^2 \lambda^2 / 8\delta$, where λ is the minimum height in all auxiliary triangles \tilde{T} . Due to the inequality above, if an embedding contains a triangle whose unsigned area is less than ϵ , then the embedding must contain a triangle t such that $D_t > \delta$.

Finally, we will show that if an embedding T contains some triangle t such that $D_t > \delta$, then it must contain some triangle s that satisfies both $D_s > \delta$ and $A_s D_s > \epsilon\delta$. If one of t 's edges, e , has length greater than σ , and let h be the height with respect to e , then we obtain (using Lemma B.1(1)):

$$A_t D_t \geq \frac{A_t^2 \tilde{h}^2}{2h^2} = \frac{|e|^2 \tilde{h}^2}{8} > \frac{\sigma^2 \lambda^2}{8} = \epsilon\delta$$

Hence t is the desired triangle s . Otherwise, all edges of t are no longer than σ . Consider the triangle sequence $\{t_1, \dots, t_k\} \subseteq T$ defined above that links one edge e_0 of t to the boundary of T . Since $|e_0| \leq \sigma$, there must exist some triangle t_i in the sequence such that $D_{t_i} > \delta$. Let t_i be the last such triangle in the sequence (i.e., with the highest index i). Since all triangles t_j for $j \in [i+1, k]$ satisfy $D_{t_j} \leq \delta$, the length of edge e_i must be greater than σ . Note that e_i is incident to triangle t_i , and so applying the inequality above for $t = t_i$ and $e = e_i$ shows that t_i is the desired triangle s . \square

Building on the previous lemma, the next lemma shows that if an embedding contains a sufficiently small triangle, then the rate of increase in TLC as α increases from 0 is sufficiently large for a non-trivial range of α :

LEMMA B.3. *For any $\delta > 0$, there exists some $\epsilon > 0$ and $\beta > 0$ such that if an embedding T contains a triangle whose unsigned area is smaller than ϵ , then for any $\alpha < \beta$, $\partial E_{\tilde{T},\alpha}(T)/\partial\alpha > \delta$.*

PROOF. By Lemma B.2, there exists some $\epsilon > 0$ such that if T contains some triangle whose unsigned area is smaller than ϵ , then T also contains a triangle t such that $D_t > 2\delta$ and $A_t D_t > 2\epsilon\delta$. We will focus on this specific triangle t , and show that there is some constant $\beta > 0$ (invariant to the choice of T) such that $\partial E_{\tilde{T},\alpha}(t)/\partial\alpha > \delta$ for any $\alpha < \beta$.

Recall Equation (9),

$$\frac{\partial E_{\tilde{T},\alpha}(t)}{\partial\alpha} = \frac{2A_t D_t + 2\alpha A_t^2}{2\sqrt{A_t^2 + 2\alpha A_t D_t + \alpha^2 A_t^2}}$$

Let $C = 2A_t D_t + 2\alpha A_t^2$. Note that A_t, D_t are non-negative, we have:

$$\frac{\partial E_{\tilde{T},\alpha}(t)}{\partial\alpha} > \frac{C}{2\sqrt{A_t^2 + \alpha C}} = \frac{\sqrt{\frac{C}{2}}}{\sqrt{\frac{2A_t^2}{C} + 2\alpha}} > \frac{\sqrt{A_t D_t}}{\sqrt{\frac{A_t}{D_t} + 2\alpha}}$$

The first inequality is due to $C > 2A_t D_t + \alpha A_t^2$, and the last inequality is due to $C > 2A_t D_t$. Let $\beta = 3\epsilon/4\delta$, we will next show that $\partial E_{\tilde{T},\alpha}(t)/\partial\alpha > \delta$ for any $\alpha < \beta$. We will separately examine the cases when $A_t \geq \epsilon$ and when $A_t < \epsilon$. Suppose $A_t \geq \epsilon$, and noting that $D_t > 2\delta$, we obtain for all $\alpha < \beta$:

$$\frac{\sqrt{A_t D_t}}{\sqrt{\frac{A_t}{D_t} + 2\alpha}} = \frac{\sqrt{D_t}}{\sqrt{\frac{1}{D_t} + \frac{2\alpha}{A_t}}} > \frac{\sqrt{2\delta}}{\sqrt{\frac{1}{2\delta} + \frac{3}{2\delta}}} = \delta$$

Now suppose $A_t < \epsilon$, and noting that $D_t > 2\delta$ and $A_t D_t > 2\epsilon\delta$, we again obtain for all $\alpha < \beta$:

$$\frac{\sqrt{A_t D_t}}{\sqrt{\frac{A_t}{D_t} + 2\alpha}} > \frac{\sqrt{2\epsilon\delta}}{\sqrt{\frac{\epsilon}{2\delta} + \frac{3\epsilon}{2\delta}}} = \delta$$

Hence we conclude that $\partial E_{\tilde{T},\alpha}(t)/\partial\alpha > \delta$ for any $\alpha < \beta$. The proof is completed by noting that

$$\frac{\partial E_{\tilde{T},\alpha}(T)}{\partial\alpha} = \sum_{s \in T} \frac{\partial E_{\tilde{T},\alpha}(s)}{\partial\alpha} \geq \frac{\partial E_{\tilde{T},\alpha}(t)}{\partial\alpha} > \delta$$

\square

Finally, we prove Proposition 4.3 for $d = 2$ using the previous lemma:

PROPOSITION B.4 (4.3 IN 2D). *Given any injective 2D embedding T_0 , there exists some $\beta > 0$ such that $E_{\tilde{T},\alpha}(T) > E_{\tilde{T},\alpha}(T_0)$ for any non-injective embedding T and $\alpha < \beta$.*

PROOF. Since T_0 is injective, the derivative $\partial E_{\tilde{T},\alpha}(T_0)/\partial\alpha$ is bounded. Pick an arbitrary but small positive value τ , and define δ as the maximum derivative for all $\alpha < \tau$.

Now suppose T has a triangle whose unsigned area is smaller than ϵ , which is found by Lemma B.3 for δ . By that Lemma, there exists some $\gamma > 0$ such that for any $\alpha < \gamma$, $\partial E_{\tilde{T},\alpha}(T)/\partial\alpha > \delta$. As a result, the following holds for all $\alpha < \min(\tau, \gamma)$:

$$\frac{\partial(E_{\tilde{T},\alpha}(T) - E_{\tilde{T},\alpha}(T_0))}{\partial\alpha} = \frac{\partial E_{\tilde{T},\alpha}(T)}{\partial\alpha} - \frac{\partial E_{\tilde{T},\alpha}(T_0)}{\partial\alpha} > \delta - \delta = 0$$

Since $E_{\tilde{T},0}(T) \geq E_{\tilde{T},0}(T_0)$ (as T_0 achieves the least total unsigned areas), we conclude $E_{\tilde{T},\alpha}(T) > E_{\tilde{T},\alpha}(T_0)$ for any $\alpha \in (0, \min(\tau, \gamma))$.

Otherwise, suppose T has no triangle whose unsigned area is smaller than ϵ . Since T is non-injective, it must contain at least one inverted triangle whose unsigned area is no smaller than ϵ . Due to Corollary 4.2, for any $\alpha > 0$,

$$E_{\tilde{T},\alpha}(T) > E_{\tilde{T},0}(T) \geq E_{\tilde{T},0}(T_0) + 2 * \epsilon$$

Since the derivative $\partial E_{\tilde{T},\alpha}(T_0)/\partial\alpha$ is bounded, there exists some $\kappa > 0$ such that, for all $\alpha < \kappa$,

$$E_{\tilde{T},\alpha}(T_0) < E_{\tilde{T},0}(T_0) + 2 * \epsilon \leq E_{\tilde{T},\alpha}(T)$$

The proof is completed by letting $\beta = \min(\tau, \gamma, \kappa)$. \square

B.2 3D

The proof in 3D proceeds similarly as in 2D, with an essentially same set of lemmas and similar proofs. Most differences appear in extending Lemmas B.1 and B.2 to 3D.

We start by extending the properties of the 2D Dirichlet energy in Lemma B.1 to 3D, while adding a new property (3):

LEMMA B.5. *Let f be one face of t , h be the distance to f from the vertex opposite to f , \tilde{h} be the corresponding distance in \tilde{t} , and D_f be the Dirichlet energy from f to its corresponding face \tilde{f} in \tilde{t} . Then,*

- (1) $D_t \geq V_t(\tilde{h}/h)^2/2$
- (2) $h/A_f \geq \tilde{h}^2/6D_t$
- (3) $D_t \geq D_f h/3$

PROOF. We first rotate and translate both t, \tilde{t} so that f and \tilde{f} are both on the XY plane, with one corner at the origin. Note that $D_t = V_t \text{Tr}(L^T L)/2$, where L is the linear transformation matrix from t to \tilde{t} (as defined in Section A.2). After the rotation and translation, L has the form:

$$L = \begin{pmatrix} L_f & v \\ \mathbf{0} & \frac{\tilde{h}}{h} \end{pmatrix}$$

where L_f is the linear transformation matrix from f to \tilde{f} . As a result, $\text{Tr}(L^T L) \geq (\tilde{h}/h)^2$, which leads to properties (1) and (2) as shown in the proof of Lemma B.1, and $\text{Tr}(L^T L) \geq \text{Tr}(L_f^T L_f)$, which yields property (3) due to $D_f = A_f \text{Tr}(L_f^T L_f)/2$ and $V_t = A_f h/3$. \square

Extending Lemma B.2 to 3D, we argue that if an embedding contains a sufficiently small tetrahedron, it will contain some tetrahedron t whose Dirichlet energy to \tilde{t} , D_t , is sufficiently large:

LEMMA B.6. *For any $\delta > 0$, there exists some $\epsilon > 0$ such that if an embedding T contains a tetrahedron whose unsigned volume is smaller than ϵ , then T must contain a tetrahedron t such that $D_t > \delta$ and $V_t D_t > \epsilon \delta$.*

PROOF. We first show that if $D_t \leq \delta$ for all tetrahedra $t \in T$, then all triangle areas in T are lower bounded by some positive constant (invariant to the choice of T). For any triangle f_0 in T , we find a sequence of tetrahedra $\{t_1, \dots, t_k\} \subseteq T$ such that f_0 is incident to t_1 , t_i and t_{i+1} share a common face (denoted by f_i) for $i \in [1, k-1]$, and t_k contains a face (denoted by f_k) on the target boundary B .

Consider tetrahedron t_k . Let h be the height with respect to base face f_k , \tilde{h} be the corresponding height in the auxiliary tetrahedron \tilde{t}_k , e be the common edge shared by f_k and f_{k-1} , and h' be the height of f_{k-1} with respect to base e . See Figure 19 for an illustration. Since $h' \geq h$, $D_{t_k} \leq \delta$, and due to Lemma B.5(2),

$$\frac{A_{f_{k-1}}}{A_{f_k}} = \frac{|e|h'}{2A_{f_k}} \geq \frac{|e|h}{2A_{f_k}} \geq \frac{|e|\tilde{h}^2}{12D_{t_k}} \geq \frac{|e|\tilde{h}^2}{12\delta}$$

Since f_k is on the fixed boundary and so is e , both A_{f_k} and $|e|$ are invariant to T . Since \tilde{h} is also fixed, $A_{f_{k-1}}$ is lower-bounded by some constant.




Fig. 19. Illustration for proof of Lemma B.6.

We hope to lower-bound A_{f_0} by iteratively applying the above inequality to the triangle sequence $A_{f_{k-2}}, \dots, A_{f_0}$, as we have done in 2D. However, the rightmost term of the inequality also involves the length of an edge $|e|$ of f_k . To bound $A_{f_{k-2}}$, for example, we need to additionally show that the length of any other edge of f_{k-1} (which might be incident to f_{k-2}) is also lower-bounded by some constant invariant to T . Consider some other edge e' of f_{k-1} (see Figure 19). Let h' , h'' be the corresponding heights of h' , h'' in \tilde{t}_k , and $D_{f_{k-1}}$ be the Dirichlet energy from f_{k-1} to its corresponding face \tilde{f}_{k-1} in \tilde{t} . Due to Lemmas B.1 (2) and B.5 (2,3),

$$\frac{|e'|}{|e|} \geq \frac{h'}{|e|} \geq \frac{\tilde{h}'^2}{4D_{f_{k-1}}} \geq \frac{\tilde{h}'^2 h''}{12D_{t_k}} \geq \frac{A_{f_{k-1}} \tilde{h}'^2 h''^2}{72D_{t_k}^2} \geq \frac{A_{f_{k-1}} \tilde{h}'^2 h''^2}{72\delta^2}$$

Since $|e|$, \tilde{h}' , \tilde{h}'' are fixed, and $A_{f_{k-1}}$ is lower-bounded by some constant (as shown above), $|e'|$ is lower-bounded by some constant

as well. Now we can iteratively apply the two inequalities above to $A_{f_{k-2}}/A_{f_{k-1}}, \dots, A_{f_0}/A_{f_1}$, and to the shared edges between successive triangles f_i, f_{i-1} for $i \in [k-1, 1]$, to lower-bound A_{f_0} by some constant that depends only the target boundary B , the auxiliary tetrahedra \tilde{T} , and the combinatorial sequence of tetrahedra $\{t_1, \dots, t_k\}$. Call a sequence *simple* if any tetrahedron appears at most once in the sequence. Let σ_{f_0} be the minimal lower bound of A_{f_0} over all possible simple tetrahedra sequences, and let σ be any value smaller than the minimum of σ_{f_0} over all triangles f_0 . Then σ lower-bounds the area of any triangle in any embedding T where $D_t \leq \delta$ for all $t \in T$.

The rest of the proof proceeds like the proof of Lemma B.2. Consider an embedding T where $D_t \leq \delta$ for all $t \in T$, and consider any tetrahedron $t \in T$. Let f be a face of t , h be the height with respect to f , and \tilde{h} be the corresponding height in \tilde{t} . We can bound t 's unsigned volume as follows (using Lemma B.5(2), $A_f > \sigma$, and $D_t \leq \delta$):

$$V_t = \frac{A_f^2 h}{3A_f} \geq \frac{A_f^2 \tilde{h}^2}{18D_t} > \frac{\sigma^2 \tilde{h}^2}{18\delta}$$

Define $\epsilon = \sigma^2 \lambda^2 / 18\delta$, where λ is the minimum height in all auxiliary tetrahedra \tilde{T} . Due the inequality above, if an embedding contains a tetrahedron whose unsigned volume is less than ϵ , then the embedding must contain a tetrahedron t such that $D_t > \delta$.

Finally, we will show that if an embedding T contains some tetrahedron t such that $D_t > \delta$, then it must contain some tetrahedron s that satisfies both $D_s > \delta$ and $V_s D_s > \epsilon \delta$. If one of t 's face, f , has unsigned area greater than σ , and let h be the height with respect to f , then we obtain (using Lemma B.5(1)):

$$V_t D_t \geq \frac{V_t^2 \tilde{h}^2}{2h^2} = \frac{A_f^2 \tilde{h}^2}{18} > \frac{\sigma^2 \lambda^2}{18} = \epsilon \delta$$

Hence t is the desired tetrahedron s . Otherwise, all faces of t have unsigned areas no greater than σ . Consider the tetrahedra sequence $\{t_1, \dots, t_k\} \subseteq T$ defined above that links one face f_0 of t to the boundary of T . Since $A_{f_0} \leq \sigma$, there must exist some tetrahedron t_i in the sequence such that $D_{t_i} > \delta$. Let t_i be the last such tetrahedron in the sequence (i.e., with the highest index i). Since all tetrahedra t_j for $j \in [i+1, k]$ satisfy $D_{t_j} \leq \delta$, the unsigned area of triangle f_i must be greater than σ . Note that f_i is a face of tetrahedron t_i , and so applying the inequality above for $t = t_i$ and $f = f_i$ shows that t_i is the desired tetrahedron s . \square

Building on the previous lemma, and extending Lemma B.3 to 3D, the next lemma shows that if an embedding contains a sufficiently small tetrahedron, then the rate of increase in TLC as α increases from 0 is sufficiently large for a non-trivial range of α :

LEMMA B.7. *For any $\delta > 0$, there exists some $\epsilon > 0$ and $\beta > 0$ such that if an embedding T contains a tetrahedron whose unsigned volume is smaller than ϵ , then for any $\alpha < \beta$, $\partial E_{T, \alpha}(T) / \partial \alpha > \delta$.*

PROOF. By Lemma B.6, there exists some $\epsilon > 0$ such that if T contains some tetrahedron whose unsigned volume is smaller than ϵ , then T also contains a tetrahedron t such that $D_t > 2\delta$ and $A_t D_t > 2\epsilon \delta$. We will focus on this specific tetrahedron t , and show

that there is some constant $\beta > 0$ (invariant to the choice of T) such that $\partial E_{\tilde{t},\alpha}(t)/\partial\alpha > \delta$ for any $\alpha < \beta$.

Recall Equation (14),

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial\alpha} = \frac{2V_t D_t + 4\alpha V_{\tilde{t}} \tilde{D}_t + 3\alpha^2 V_{\tilde{t}}^2}{2\sqrt{V_t^2 + 2\alpha V_t D_t + 2\alpha^2 V_{\tilde{t}} \tilde{D}_t + \alpha^3 V_{\tilde{t}}^2}}$$

Let $C = 2V_t D_t + 4\alpha V_{\tilde{t}} \tilde{D}_t + 3\alpha^2 V_{\tilde{t}}^2$. Note that $V_t, V_{\tilde{t}}, D_t, \tilde{D}_t$ are non-negative, we have:

$$\frac{\partial E_{\tilde{t},\alpha}(t)}{\partial\alpha} \geq \frac{C}{2\sqrt{V_t^2 + \alpha C}} = \frac{\sqrt{C}}{2\sqrt{\frac{V_t^2}{C} + \alpha}} \geq \frac{\sqrt{V_t D_t}}{\sqrt{\frac{V_t}{D_t} + 2\alpha}}$$

Proceeding exactly as the rest of the proof of Lemma B.3 (replacing A_t with V_t), we can show that $\partial E_{\tilde{t},\alpha}(t)/\partial\alpha > \delta$ for any $\alpha < \beta$ where $\beta = 3\epsilon/4\delta$. \square

Finally, we prove Proposition 4.3 for $d = 3$ using the previous lemma:

PROPOSITION B.8 (4.3 IN 3D). *Given any injective 3D embedding T_0 , there exists some $\beta > 0$ such that $E_{\tilde{T},\alpha}(T) > E_{\tilde{T},\alpha}(T_0)$ for any non-injective embedding T and $\alpha < \beta$.*

PROOF. The proof proceeds exactly as the proof of Proposition B.4, after making the dimensionality-specific replacements: “triangle” by “tetrahedron”, “area” by “volume”, and “Lemma B.3” by “Lemma B.7”. \square

C GRADIENT AND HESSIAN OF TLC

We provide explicit expressions for the gradient and hessian of TLC. As TLC is accumulative, it suffices to consider a single simplex t with an auxiliary simplex \tilde{t} , whose TLC is $E_{\tilde{t},\alpha}(t)$ (which we shall shorthand as E). Denote the vertices of the d -dimensional simplex t as v_1, \dots, v_{d+1} , and similarly \tilde{v}_i, \tilde{v}_i for simplices \tilde{t}, \tilde{t} . Here, we only consider the cases of $d = 2, 3$.

To simplify the expressions, we use an alternative formula for the volume of a simplex known as the Cayley-Menger determinant. The formula calculates the volume E of the lifted simplex, \hat{t} , as $E = \sqrt{D}/c$, where c is a dimension-dependent constant and D is the determinant of a matrix involving squared edge lengths of \hat{t} . For a triangle t , $c = 4$ and

$$D = \text{Det} \begin{pmatrix} 2\hat{d}_{12} & \hat{d}_{12} + \hat{d}_{13} - \hat{d}_{23} \\ \hat{d}_{12} + \hat{d}_{13} - \hat{d}_{23} & 2\hat{d}_{13} \end{pmatrix}$$

where \hat{d}_{ij} is the squared edge length between vertices \hat{v}_i, \hat{v}_j (note that $\hat{d}_{ij} = d_{ij} + \tilde{d}_{ij}$, where d_{ij}, \tilde{d}_{ij} are squared lengths of the corresponding edges in t, \tilde{t}). For a tetrahedron t , $c = 12\sqrt{2}$ and

$$D = \text{Det} \begin{pmatrix} 2\hat{d}_{12} & \hat{d}_{12} + \hat{d}_{13} - \hat{d}_{23} & \hat{d}_{12} + \hat{d}_{14} - \hat{d}_{24} \\ \hat{d}_{12} + \hat{d}_{13} - \hat{d}_{23} & 2\hat{d}_{13} & \hat{d}_{13} + \hat{d}_{14} - \hat{d}_{34} \\ \hat{d}_{12} + \hat{d}_{14} - \hat{d}_{24} & \hat{d}_{13} + \hat{d}_{14} - \hat{d}_{34} & 2\hat{d}_{14} \end{pmatrix}$$

C.1 Gradient

It suffices to derive the gradient of E with respect to one vertex v_i , which is a length- d vector. Differentiating $E = \sqrt{D}/c$ yields:

$$\frac{\partial E}{\partial v_i} = \frac{\partial D}{\partial v_i} / 2c^2 E. \quad (17)$$

Applying the chain rule to $\partial D/\partial v_i$ yields:

$$\frac{\partial D}{\partial v_i} = \sum_{j \neq i} \frac{\partial D}{\partial \hat{d}_{ij}} \frac{\partial \hat{d}_{ij}}{\partial v_i}.$$

We provide expressions for the terms on the right-hand side. First,

$$\frac{\partial \hat{d}_{ij}}{\partial v_i} = \frac{\partial d_{ij}}{\partial v_i} = 2(v_i - v_j). \quad (18)$$

Next, if t is a triangle, and k is the vertex index other than i, j , we get

$$\frac{\partial D}{\partial \hat{d}_{ij}} = 2(\hat{d}_{ik} + \hat{d}_{jk} - \hat{d}_{ij}). \quad (19)$$

Finally, if t is a tetrahedron, and k, l are the vertex indices other than i, j , we get

$$\begin{aligned} \frac{\partial D}{\partial \hat{d}_{ij}} &= 2(\hat{d}_{il} - \hat{d}_{ik})(\hat{d}_{jk} - \hat{d}_{jl}) + \\ &2\hat{d}_{kl}(\hat{d}_{ik} + \hat{d}_{il} + \hat{d}_{jk} + \hat{d}_{jl} - \hat{d}_{kl} - 2\hat{d}_{ij}). \end{aligned} \quad (20)$$

C.2 Hessian

It suffices to derive the hessian of E with respect to two vertices v_i, v_j , which is a $d \times d$ matrix. Differentiating equation (17) yields:

$$\frac{\partial^2 E}{\partial v_i \partial v_j} = \left(\frac{\partial^2 D}{\partial v_i \partial v_j} - 2c^2 \frac{\partial E}{\partial v_j} \left(\frac{\partial E}{\partial v_i} \right)^T \right) / 2c^2 E.$$

As we already have expression for the gradient $\partial E/\partial v_i$, we only need to provide expression for $\partial^2 D/\partial v_i \partial v_j$. Application of the chain rule yields:

$$\frac{\partial^2 D}{\partial v_i \partial v_j} = \sum_{k \neq i} \left(\left(\sum_{l \neq j} \frac{\partial^2 D}{\partial \hat{d}_{ik} \partial \hat{d}_{jl}} \frac{\partial \hat{d}_{jl}}{\partial v_j} \right) \left(\frac{\partial \hat{d}_{ik}}{\partial v_i} \right)^T + \frac{\partial D}{\partial \hat{d}_{ik}} \frac{\partial^2 \hat{d}_{ik}}{\partial v_i \partial v_j} \right).$$

Expressions for first-order terms on the right-hand side are given in equations (18,19,20). We next provide expressions for the remaining second-order terms. Let I denote the identity matrix, and define δ_{pq} as 1 if $p = q$ and 0 otherwise. We have:

$$\frac{\partial^2 \hat{d}_{ik}}{\partial v_i \partial v_j} = 2(\delta_{ij} - \delta_{jk})I.$$

Next, if t is a triangle, we have:

$$\frac{\partial^2 D}{\partial \hat{d}_{ik} \partial \hat{d}_{jl}} = \begin{cases} -2, & i = j, k = l \\ 2, & \text{otherwise.} \end{cases}$$

Finally, if t is a tetrahedron, and let p, q be vertex indices other than i, j, k, l , we can derive:

$$\frac{\partial^2 D}{\partial \hat{d}_{ik} \partial \hat{d}_{jl}} = \begin{cases} -4\hat{d}_{pq}, & i = j, k = l \\ 2(\hat{d}_{pk} + \hat{d}_{pl} - \hat{d}_{lk}), & i = j, k \neq l \\ 2(\hat{d}_{ij} + \hat{d}_{il} + \hat{d}_{jk} + \hat{d}_{kl} - 2\hat{d}_{ik} - 2\hat{d}_{jl}), & \text{otherwise.} \end{cases}$$