

Constrained Willmore Surfaces

YOUSUF SOLIMAN, Caltech

ALBERT CHERN, UCSD

OLGA DIAMANTI, TU Graz, Austria

FELIX KNÖPPEL, TU Berlin

ULRICH PINKALL, TU Berlin, Germany

PETER SCHRÖDER, Caltech

Smooth curves and surfaces can be characterized as minimizers of squared curvature bending energies subject to constraints. In the univariate case with an isometry (length) constraint this leads to classic non-linear splines. For surfaces, isometry is too rigid a constraint and instead one asks for minimizers of the Willmore (squared mean curvature) energy subject to a conformality constraint. We present an efficient algorithm for (conformally) constrained Willmore surfaces using triangle meshes of arbitrary topology with or without boundary. Our conformal class constraint is based on the discrete notion of conformal equivalence of triangle meshes. The resulting non-linear constrained optimization problem can be solved efficiently using the competitive gradient descent method together with appropriate Sobolev metrics. The surfaces can be represented either through point positions or differential coordinates. The latter enable the realization of abstract metric surfaces without an initial immersion. A versatile toolkit for extrinsic conformal geometry processing, suitable for the construction and manipulation of smooth surfaces, results through the inclusion of additional point, area, and volume constraints.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; **Mesh geometry models**; • **Mathematics of computing** → **Discretization**; **Mathematical optimization**.

Additional Key Words and Phrases: discrete differential geometry, geometric flows, Willmore surfaces, conformal geometry, geometric modeling, differential coordinates

ACM Reference Format:

Yousuf Soliman, Albert Chern, Olga Diamanti, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2021. Constrained Willmore Surfaces. *ACM Trans. Graph.* 40, 4, Article 112 (August 2021), 17 pages. <https://doi.org/10.1145/3450626.3459759>

1 INTRODUCTION

A draft person’s physical spline, *i.e.*, a thin elastic rod, has long been a model for smooth curves. Subject to point constraints, the spline minimizes bending, *i.e.*, the squared curvature integrated along the curve, while maintaining its length. This captures both physical (least bending) and aesthetic (“as-round-as-possible”) aspects.

Authors’ addresses: Yousuf Soliman, Caltech; Albert Chern, UCSD, 9500 Gilman Dr, La Jolla, CA, 92093; Olga Diamanti, TU Graz, Rechbauerstraße 12, 8010, Graz, Austria; Felix Knöppel, TU Berlin; Ulrich Pinkall, TU Berlin, Straße des 17. Juni 135, 10623, Berlin, Germany; Peter Schröder, Caltech, 1200 E California Blvd, Pasadena, CA, 91125.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).
0730-0301/2021/8-ART112
<https://doi.org/10.1145/3450626.3459759>

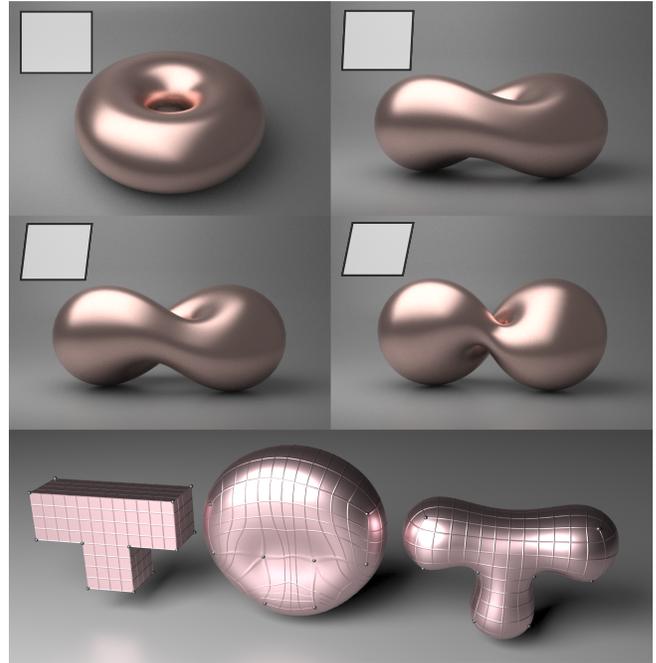


Fig. 1. Top: Tori which minimize the Willmore energy subject to a conformality constraint. The inset parallelepiped visualizes the fundamental domain of each torus. Its “tilt” controls the twist of the torus. Without constraining the conformal class the minimum in all four cases would be achieved at the standard Clifford torus. Bottom: A more complex shape with point constraints showing the minimizer of the Willmore energy without (left) and with (right) the conformal class constraint.

In the case of surfaces $f: M \rightarrow \mathbb{R}^3$ the bending energy is given by the Willmore energy

$$\mathcal{W}(f) = \int_M H^2 dA,$$

which measures the squared L^2 -norm of the mean curvature H of the surface. As a bending energy it is of great interest in physical modeling ranging from bio-physics [Canham 1970; Helfrich 1973] to computer animation [Bridson et al. 2003; Grinspun et al. 2003]. Based on its “as-round-as-possible” aspect, the Willmore energy also finds application in surface modeling and form finding [Joshi and Séquin 2007; Vaxman et al. 2015, 2018; Gruber and Aulisa 2020].

Unfortunately without any additional constraints the surface “material” underlying the Willmore energy is too flexible. For example, all surfaces of genus 1 have their Willmore minimum at the Clifford torus [Marques and Neves 2014]. Or consider the case of surfaces which are topological spheres, constraining any 4 points always yields the round sphere through those 4 points as minimizer. What is missing is the analog of the isometry constraint from the univariate setting. Requiring the deformation of a surface to be isometric though leaves it too rigid. Instead the natural choice is a conformality constraint [Garsia 1961; Rüedy 1971], making the study of minimizers of the conformally constrained Willmore functional an active research area [Bohle et al. 2008; Rivière 2008; Schätzle 2013; Kuwert and Schätzle 2013; Heller and Ndiaye 2019]. Instead of an arbitrarily flexible material the resulting surfaces are now made of an isotropic auxetic (negative Poisson ratio) material [Konaković et al. 2016]. Limiting the flexibility of the surface in this manner enforces intrinsic (conformal class) constraints on the extrinsic appearance. In the case of the torus this fixes the aspect ratio and skewness of the fundamental domain and the minimizer reflects this (Figure 1, top as well as Figure 2). Similarly, for the spherical shape with 4 point constraints the surface deforms in a more intuitive way following the control points (Figure 19). Figure 1 (bottom) shows a more complex example comparing the minimizer without and with conformal constraint.

More control over the shape is not the only practical benefit of the conformal class constraint. Conformality also ensures that the deformation is locally a similarity. In the case of triangle meshes this preserves mesh quality. Degeneration of meshes during optimization is otherwise a common problem [Bobenko and Schröder 2005] (see also Figure 13) and has motivated methods based on conformal deformations [Crane et al. 2011, 2013] or the addition of conformal penalty forces [Barrett et al. 2016; Gruber and Aulisa 2020]. Neither of these approaches though are entirely satisfactory since they exhibit numerical drift (Figure 12) in the conformal class resp. modify the energy being minimized.

In this paper we take a different approach: we employ an *exact* discrete conformality constraint for extrinsic variational conformal geometry processing of oriented discrete surfaces (simplicial 2-manifolds of arbitrary genus with or without boundary). This constraint can maintain the conformal class of a surface exactly, or be used to achieve a desired target conformal class. It is based on the discrete notion of *conformal equivalence of triangle meshes* (CETM) [Springborn et al. 2008; Pinkall and Springborn 2021; Gillespie et al. 2021; Campen et al. 2021] and is independent of the precise form of the energy.

For efficient numerical minimization we use gradient flow with an appropriate Sobolev metric [Renka and Neuberger 1995; Eckstein et al. 2007; Schumacher 2017]. Constraints are incorporated via Lagrange multipliers and the associated min/max problem is solved with the recently proposed competitive gradient descent (CGD) method [Schäfer and Anandkumar 2019]. Of interest beyond the conformality constraint are constraints on point positions, surface area, and enclosed volume (see for example Gruber and Aulisa [2020]) and we demonstrate examples of these in various combinations (see Figures 1, 3 and 19).

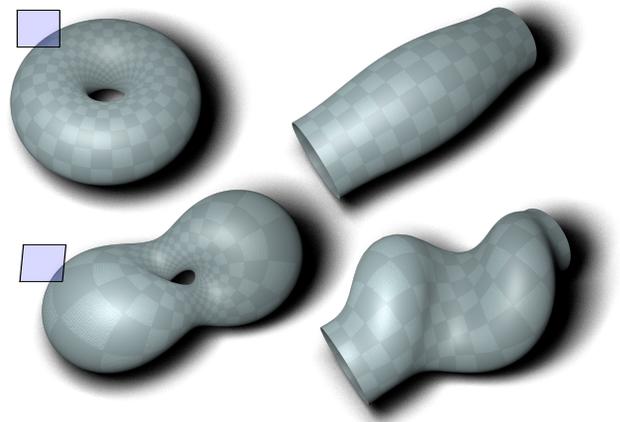


Fig. 2. The conformal class (visualized by inset rectangles denoting the fundamental domain) controls the aspect ratio and twist of the surface. Here we see how the same twists are realized by cylinders with boundary constraints (positions and tangents) and tori (see also Figure 1, top). The conformality of the deformation is visualized here with a checkerboard texture.

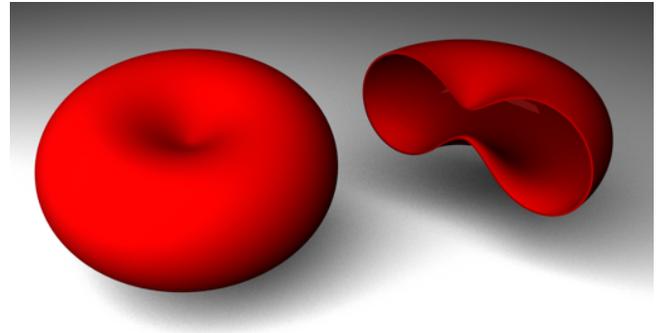


Fig. 3. A sphere minimizing the Willmore energy subject to constraints on surface area, enclosed volume, and conformal class. Physically such shapes model vesicles, for example red blood cells, which are enclosed by a thin-shell of high membrane and low bending stiffness [Canham 1970; Helfrich 1973; Gruber and Aulisa 2020]. Here, the conformality constraint only serves to preserve mesh quality.

While triangle mesh optimization can be parameterized by vertex positions, in particular when starting with a given immersion, we also give a formulation in terms of differential coordinates [Sorkine 2006; Xu and Zhou 2009]. Aside from making our approach compatible with the many existing geometry processing algorithms which use differential coordinates, it also enables us to find conformal immersions for abstract metric surfaces, valuable in mathematical visualization. For triangle meshes such differential coordinates are piecewise maps, constant per triangle, and we allow them to vary independently [Custers and Vaxman 2020]. Surprisingly, optimization over such *triangle fields* is well defined even if triangles do not “glue” together and provides a richer search space for the optimization,

often allowing more rapid progress towards convergence. Recovery of an actual surface is then ensured through the inclusion of a (linear) integrability constraint. Differentials as variables furthermore allow us to interpret the Lagrange multipliers associated with the conformal class constraint as *quadratic differentials* [Weber et al. 2012] and, in a mechanical analogy, as stress tensors with corresponding forces acting along edges to counteract anisotropic distortion.

2 DISCRETE FORMULATION

In this section, we fix our notation and define the basic geometric objects needed to formulate variational problems for surfaces.

2.1 Discrete Surface

Topology. Throughout we describe the underlying surface by an abstract triangulation $M = (V, E, F)$. The vertices, edges, and faces are denoted by indices $i \in V$, pairs $ij \in E$, and triples $ijk \in F$, respectively. The dual mesh is denoted by $M^* = (V^*, E^*, F^*)$ and we refer to dual cells by their corresponding primal cell, that is, we identify $V^* \cong F$, $E^* \cong E$, and $F^* \cong V$ with no further comment. We assume that M is a closed oriented surface of arbitrary topology. Treating surfaces with boundary is relegated to Appendix D.

Differential forms. We use discrete differential forms from Discrete Exterior Calculus [Desbrun et al. 2008]. 0-, 1-, and 2-forms are assigned to vertices, edges, and facets respectively and represent integral quantities over their respective subdomain. Recall that forms on M^* can be paired with forms on M and integrated: for $\alpha \in \Omega^{2-k}(M^*; \mathbb{R}^m)$ and $\beta \in \Omega^k(M; \mathbb{R}^m)$ the primal-dual pairing is defined as

$$\langle\langle \alpha | \beta \rangle\rangle := \sum_{\sigma} \langle \alpha_{\sigma}, \beta_{\sigma} \rangle,$$

where the sum is taken over all oriented k -cells σ in M and the angle brackets indicate that the product used in the definition is the inner product on \mathbb{R}^m . This pairing identifies the dual space of $\Omega^k(M; \mathbb{R}^m)$ with $\Omega^{2-k}(M^*; \mathbb{R}^m)$ and the pairing $\langle\langle \alpha | \beta \rangle\rangle$ is interpreted as the integral $\int \langle \alpha \wedge \beta \rangle$ over M .

Geometry. The intrinsic geometry of the surface is specified by positive edge lengths $\ell : E \rightarrow \mathbb{R}_{>0}$ satisfying the triangle inequality in each face. We call such an ℓ a *discrete metric* on M ; this data endows each face with the geometry of a Euclidean triangle. Vertex positions $f : V \rightarrow \mathbb{R}^3$, interpolated linearly in the faces, provide a geometric realization of M in \mathbb{R}^3 . We say that f is an *immersion* if every vertex star is embedded. An immersion induces a discrete metric given by the edge lengths of the realization $\ell_{ij} = |f_j - f_i|$.

2.2 Geometric Energies

Consider a discrete surface energy $\mathcal{E}(f)$. A variation of vertex positions $\dot{f} : V \rightarrow \mathbb{R}^3$ induces an infinitesimal change of energy

$$\dot{\mathcal{E}} = \left. \frac{d}{dt} \right|_{t=0} \mathcal{E}(f + t\dot{f}).$$

We assume that the energy is differentiable in the sense that all these partial derivatives exist. In this case, we define the *gradient 2-form* of \mathcal{E} , denoted as $\text{grad } \mathcal{E}(f) \in \Omega^2(M^*; \mathbb{R}^3)$, by the requirement

$$\dot{\mathcal{E}} = \langle\langle \text{grad } \mathcal{E}(f) | \dot{f} \rangle\rangle = \sum_{i \in V} \langle \text{grad } \mathcal{E}(f)_i, \dot{f}_i \rangle$$

for all surface variations \dot{f} . The components of the gradient 2-form are the partial derivatives with respect to the variables. These must be derived for an implementation. Critical points of \mathcal{E} satisfy the Euler-Lagrange equations

$$\text{grad } \mathcal{E}(f) = 0. \quad (1)$$

From now on, we write $\text{grad } \mathcal{E}$ for the gradient 2-form if the surface f is clear from context.

Area. The area of a discrete surface is defined by

$$\mathcal{A}(f) := \sum_{ijk \in F} A_{ijk} = \frac{1}{2} \langle\langle Lf | f \rangle\rangle,$$

where A_{ijk} is the area of the triangle ijk and $L := -d*d : \Omega^0(M; \mathbb{R}^3) \rightarrow \Omega^2(M^*; \mathbb{R}^3)$ is the positive-definite cotan Laplacian. The gradient 2-form is $\text{grad } \mathcal{A} = Lf$. The quantity $\frac{1}{2}Lf$ is the *integrated* discrete mean curvature vector, which we informally denote by $\text{HN } dA$.

Willmore energy. Measuring the L^2 -norm of the mean curvature gives rise to the Willmore energy, which can be discretized most simply as

$$\mathcal{W}(f) := \langle\langle \text{HN } dA | * \text{HN } dA \rangle\rangle = \frac{1}{4} \langle\langle Lf | * Lf \rangle\rangle,$$

where $*$ is the Hodge-star on dual 2-forms, defined as the inverse of the diagonal matrix of vertex dual areas (Appendix D).

3 CONFORMAL CONSTRAINTS

In this section, we derive the Euler-Lagrange equation for variational problems with a conformal constraint. We begin with preliminaries, recalling the notion of discrete conformal equivalence and the role of length cross ratios and their logarithms in parameterizing the set of discretely conformally equivalent metrics, before examining the overall structure of conformally constrained critical points of a geometric energy. Finally we give concrete expressions for the discrete Euler-Lagrange equations and see that the associated Lagrange multipliers are quadratic differentials which can be related to earlier work and interpreted as conformal stress tensors.

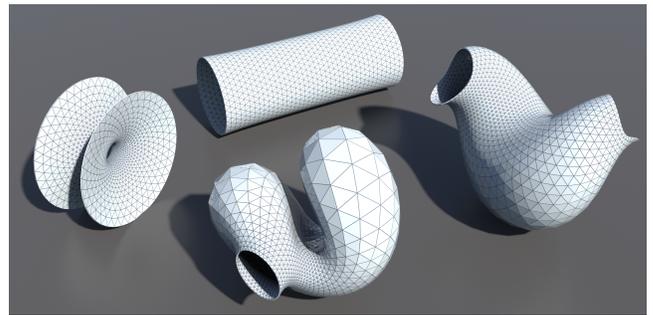


Fig. 4. Different representatives of the same conformal class (identical length cross ratios), each minimizing the Willmore energy, but differing in shape due to different boundary conditions.

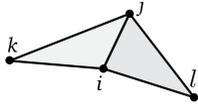
Discrete Conformal Equivalence. In the smooth setting two metrics g and \tilde{g} are conformally equivalent if there exists a function $u: M \rightarrow \mathbb{R}$ such that $\tilde{g} = e^{2u}g$. In the discrete setting we follow [Luo 2004; Springborn et al. 2008; Bobenko et al. 2015; Pinkall and Springborn 2021], calling two discrete metrics $\ell, \tilde{\ell}$ *conformally equivalent* if there exist vertex scale factors $u: V \rightarrow \mathbb{R}$ such that for all $ij \in E$

$$\tilde{\ell}_{ij} = e^{\frac{u_i+u_j}{2}} \ell_{ij} \quad (2)$$

This definition has been successful in solving a number of conformal variational problems involving edge lengths [Chow and Luo 2003; Jin et al. 2007; Glickenstein 2011; Springborn 2017; Ge 2018]. The solutions reproduce structures found in the smooth setting, and exhibit convergence under refinement [Luo et al. 2020] as shown by theoretical results, as well as by numerical experiments. Importantly, every infinitesimal conformal deformation can be realized by infinitesimal extrinsic deformations of vertex positions [Lam and Pinkall 2017], except in the case of isothermic surfaces.

Cross ratios. To each metric ℓ we can assign length cross ratios $\mathfrak{c}: E \rightarrow \mathbb{R}_{>0}$,

$$\mathfrak{c}_{ij} = \frac{\ell_{il}\ell_{jk}}{\ell_{ij}\ell_{ki}}.$$



Clearly, conformally equivalent metrics have equal length cross ratios. Thus \mathfrak{c} can be considered as a map defined on the set $\mathcal{C}(M)$ of conformal equivalence classes of discrete metrics on M . Moreover, the length cross ratios completely characterize the conformal class:

PROPOSITION 3.1 ([SPRINGBORN ET AL. 2008]). *Two discrete metrics are discretely conformally equivalent if and only if their length cross ratios are the same.*

Logarithmic coordinates. To each metric ℓ on M we can assign its logarithmic edge lengths $\lambda = \log \ell$. Taking the logarithm of Equation 2 reveals a linear relationship for the logarithmic edge lengths of conformally equivalent metrics: If $\lambda = \log \ell$ and $\tilde{\lambda} = \log \tilde{\ell}$, then $[\ell] = [\tilde{\ell}]$ holds if and only if $\tilde{\lambda} = \lambda + Au$ for some $u: V \rightarrow \mathbb{R}$, where $A \in \mathbb{R}^{E \times V}$ is the averaging map defined as

$$(Au)_{ij} = \frac{1}{2}(u_i + u_j).$$

Thus the conformal classes $[\ell] \in \mathcal{C}(M)$ are identified with elements $[\lambda] \in \mathbb{R}^E / \text{im } A$. Moreover,

$$\log \mathfrak{c}_{ij} = \log \left(\frac{\ell_{il}\ell_{jk}}{\ell_{ij}\ell_{ki}} \right) = \lambda_{il} - \lambda_{ij} + \lambda_{jk} - \lambda_{ki}. \quad (3)$$

Thus $\log \mathfrak{c} = C\lambda$ for a linear map $C \in \mathbb{R}^{E \times E}$ and in terms of A and C the equivalence in Proposition 3.1 is expressed as

$$\text{im } A = \ker C. \quad (4)$$

Consequently, the map $\mathcal{C}(M) \ni [\ell] \mapsto \log \mathfrak{c} \in \text{im } C$ defines a global chart for $\mathcal{C}(M)$.

3.1 Constrained Euler-Lagrange Equation

Here we characterize the critical surfaces of conformally constrained variational problems. We begin with a high level view before giving detailed discrete expressions.

An infinitesimal conformal deformation is a variation \dot{f} such that $(\log \mathfrak{c})^\circ = C \dot{\lambda} = 0$. By definition, f is then critical for a surface energy



Fig. 5. Constrained Willmore surfaces of higher genus.

\mathcal{E} under all infinitesimal conformal deformations if and only if

$$\langle \text{grad } \mathcal{E} | \dot{f} \rangle = 0 \text{ for all } \dot{f} \in \ker(CD\lambda)$$

where $D\lambda(\dot{f}) = \dot{\lambda}$ denotes the differential of λ in terms of surface variations \dot{f} . This can also be expressed equivalently as

$$\ker(CD\lambda) \subset \text{grad } \mathcal{E}^\perp \Leftrightarrow \text{grad } \mathcal{E} \in (\ker CD\lambda)^\perp.$$

We derive the Euler-Lagrange equations by describing the orthogonal complement of $\ker(CD\lambda)$.

Using the fact that the kernel of a linear operator is the orthogonal complement of the image of its adjoint, we have

$$(\ker CD\lambda)^\perp = \text{im}(CD\lambda)^* = \text{im}(D\lambda^*C^*) = D\lambda^*(\text{im } C^*).$$

By taking the orthogonal complement of Equation 4 we also have $\text{im } C^* = \ker A^*$. Furthermore, if we identify $(\mathbb{R}^E)^*$ with \mathbb{R}^E using the standard inner product then

$$\ker A^* \cong \left\{ q: E \rightarrow \mathbb{R} : \sum_{ij} q_{ij} = 0, \forall i \in V \right\}.$$

So far, we have shown that f is a conformally constrained critical point if and only if there exists $q \in \ker A^*$ satisfying $\text{grad } \mathcal{E} = D\lambda^*q$. It remains to compute $D\lambda^*q$: for a surface variation $\dot{f}: V \rightarrow \mathbb{R}^3$ we have

$$\langle D\lambda^*q | \dot{f} \rangle = \langle q | \dot{\lambda} \rangle = \sum_{ij \in E} q_{ij} \frac{\langle df_{ij}, \dot{f}_j - \dot{f}_i \rangle}{|df_{ij}|^2} = - \sum_{i \in V} \left(\sum_{ij} q_{ij} \frac{\langle df_{ij}, \dot{f}_i \rangle}{|df_{ij}|^2} \right),$$

where we used

$$\dot{\lambda}_{ij} = \frac{\dot{\ell}_{ij}}{\ell_{ij}} = \frac{\langle df_{ij}, \dot{df}_{ij} \rangle}{|df_{ij}|^2} = \frac{\langle df_{ij}, \dot{f}_j - \dot{f}_i \rangle}{|df_{ij}|^2}, \quad (5)$$

Consequently,

$$(D\lambda^*q)_i = - \sum_{ij} q_{ij} \frac{df_{ij}}{|df_{ij}|^2}, \quad (i \in V)$$

and so we have the Euler-Lagrange equations below.

THEOREM 3.2. *A discrete immersion of a closed surface $f: V \rightarrow \mathbb{R}^3$ is critical for a surface energy \mathcal{E} under all infinitesimal conformal variations if and only if there exists $q: E \rightarrow \mathbb{R}$ satisfying*

$$\sum_{ij} q_{ij} = 0 \quad (6)$$

$$(\text{grad } \mathcal{E})_i + \sum_{ij} q_{ij} \frac{df_{ij}}{|df_{ij}|^2} = 0 \quad (7)$$

for all vertices $i \in V$.

Equation 7 can be interpreted as the condition that $\text{grad } \mathcal{E}$ is orthogonal to the space of discrete conformal immersions, and so we call q the Lagrange multiplier of the conformally constrained critical surface.

3.2 Discrete Quadratic Differentials

We call the conformal Lagrange multipliers *discrete quadratic differentials*, and we write $\mathcal{Q}(M) := \ker A^*$. This notion has appeared in the literature [Lam and Pinkall 2016, 2017] in the context of discrete minimal surfaces, circle packings, and discrete harmonic functions. In the smooth theory, holomorphic quadratic differentials also appear as Lagrange multipliers for conformally constrained variational problems. This discrete notion plays an important role in our numerical optimization—below, we interpret q as a surface stress and then discretize a metric on the space of quadratic differentials.

A Lagrange multiplier q can be interpreted as tension along the edges balancing the energy force at the vertices. Define $\tau \in \Omega^1(M^*; \mathbb{R}^3)$ by

$$\tau_{ij} := q_{ij} \frac{df_{ij}}{|df_{ij}|^2}. \quad (8)$$

Notice that τ is a *dual* 1-form, so $d\tau$ is the dual 2-form obtained by totaling the contributions τ_{ij} around a vertex i . Pairing the Euler-Lagrange equations with an arbitrary surface variation \dot{f} gives

$$\dot{\mathcal{E}} = - \sum_{ij \in E} \langle \tau_{ij}, d\dot{f}_{ij} \rangle = \sum_{i \in V} \langle (d\tau)_i, \dot{f}_i \rangle = \langle d\tau | \dot{f} \rangle.$$

Each summand $\langle \tau_{ij}, d\dot{f}_{ij} \rangle$ measures the (virtual) work done in response to the edge strain $d\dot{f}_{ij}$, and so we can interpret τ_{ij} as an edge-stress. We also see that Euler-Lagrange equations can be expressed as the force balancing $d\tau = \text{grad } \mathcal{E}$. This shows that q_{ij} generates the force τ_{ij} at the vertex i , and so if $q_{ij} > 0$ then there is a compressive stress on the edge; conversely, if $q_{ij} < 0$ then there is a tensile stress. In this context we think of q as a force per unit length squared, which differs from the traditional notion of stress as a force per cross-sectional area. Figure 6 visualizes the forces that the Lagrange multiplier induces on the vertices at a conformally constrained Willmore resp. area minimizing surface.

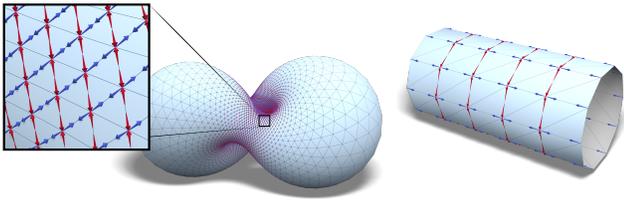


Fig. 6. We visualize the conformal tension force at a constrained Willmore surface (Left) and at a constrained minimal surface (Right). Notice that the cylinder is constrained minimal since the induced force counteracts the mean curvature normal, which tries to squeeze the cylinder.

Measuring Infinitesimal Conformal Distortion. The interpretation of the Lagrange multipliers as a conformal tension force provides a natural metric on $\mathcal{Q}(M)$, which is needed for our numerical optimization. Since q induces the conformal tension force τ (Equation 8)

we can define the L^2 -norm of q by the L^2 -norm of τ :

$$\|q\|_{L^2\Omega}^2 := \|\tau\|_{L^2\Omega^1}^2 = \sum_{ij \in E} \frac{w_{ij}^*}{\ell_{ij}^2} q_{ij}^2 = \sum_{ij \in E} \frac{q_{ij}^2}{\ell_{ij} \ell_{ij}^*}, \quad (9)$$

where $w_{ij}^* = \ell_{ij}/\ell_{ij}^*$ are dual edge weights, ℓ_{ij} are primal edge lengths, and ℓ_{ij}^* are dual edge lengths. This gives us a diagonal mass matrix that we call $I_{\mathcal{Q}} \in \mathbb{R}^{E \times E}$. One can define other L^2 -metrics on $\mathcal{Q}(M)$ by choosing a metric different from the one induced by the immersion. Two natural choices are given by the uniformized metric of constant curvature, which discretizes the Weil-Petersson metric, or using the Bergman metric [Crane et al. 2011].

4 OPTIMIZATION

In this section we develop flows to minimize a geometric energy $\mathcal{E}(f)$, the principal example of which is *gradient descent*. To describe a variety of different descent methods in a common framework we derive the gradient descent update step as the minimization of a quadratically regularized, local linear approximation of the given \mathcal{E} :

$$\min_{\dot{f}} \underbrace{\mathcal{E}(f) + \langle \text{grad } \mathcal{E} | \dot{f} \rangle}_{\text{linearization}} + \underbrace{\frac{1}{2t} \langle \dot{f}, \dot{f} \rangle}_{\text{regularizer}}. \quad (10)$$

Here the first two terms are from the Taylor series expansion of \mathcal{E} at f in the direction \dot{f} and the minimization is over \dot{f} . The step size $t > 0$ is encoded in the quadratic regularization, and describes our trust in the local linear approximation. If the inner product is described by the mass matrix I , i.e., if $\langle \dot{f}, \dot{f} \rangle = \dot{f}^T I \dot{f}$, then the minimizer of the local approximation is the solution to the linear system

$$I \dot{f} = -t \text{grad } \mathcal{E}. \quad (\text{GD})$$

So we see that the approximation in Equation 10 results in the usual definition of gradient descent. Of course, this gradient descent will not preserve any of the desired constraints, in particular, this flow will change the discrete conformal class.

Sobolev gradients. Choosing a different metric (inner product) in the regularizer above defines a different notion of gradient. We can then design efficient gradient flows by carefully choosing the metric so that Equation 10 closely approximates $\mathcal{E}(f + \dot{f})$. For example, replacing the metric by the Hessian of \mathcal{E} results in Newton's method; the Hessian, however will generally not be positive-definite and so does not actually define an inner product. For many geometric energies an appropriate metric is given by the *Sobolev metric* matching the energy space of \mathcal{E} [Schumacher 2017; Yu et al. 2020].

Since the bi-Laplacian energy is the linearization of the Willmore energy, the appropriate inner product is the Sobolev H^2 -metric:

$$\langle \dot{f}, \dot{f} \rangle_{H^2} := \langle L\dot{f}, L\dot{f} \rangle_{L^2} + \langle \dot{f}, \dot{f} \rangle_{L^2}. \quad (11)$$

The existence and analysis of this Willmore gradient flow in the smooth setting is found in [Schumacher 2017], and the numerical results show that using Sobolev gradients makes Willmore minimization remarkably easy (see the comparison in Figure 7).

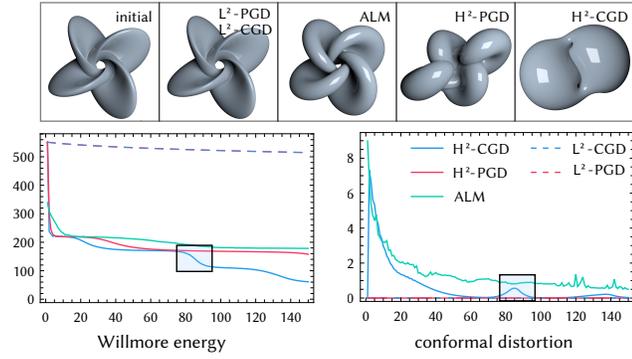


Fig. 7. Starting with the initial surface on the top left different optimization methods make varying amounts of progress for the same number of iterations with the Sobolev CGD method performing best. Here we can clearly see the importance of using Sobolev rather than L^2 gradients. The highlighted box illustrates how CGD can improve over PGD by relaxing constraint satisfaction throughout the flow. For comparison we also included an augmented Lagrangian method (ALM).

4.1 Projected gradient flow

To preserve the conformal class we consider the projected gradient descent defined by restricting our local approximation to the set of infinitesimally conformal variations. Recall that $\mathring{f} : \mathbb{V} \rightarrow \mathbb{R}^3$ is infinitesimally conformal if $C D\lambda(\mathring{f}) = 0$. Minimizing Equation 10 subject to this linear constraint results in the saddle point system

$$\begin{pmatrix} \frac{1}{t}I & D\lambda^* C^* \\ -C D\lambda & 0 \end{pmatrix} \begin{pmatrix} \mathring{f} \\ \mathring{\mu} \end{pmatrix} = - \begin{pmatrix} \text{grad } \mathcal{E} \\ 0 \end{pmatrix}. \quad (\text{PGD})$$

Note that we use $\mathring{\mu}$ as a variable to parameterize $\mathring{q} = C^* \mathring{\mu}$ to avoid the constraint on \mathring{q} (Equation 6).

This is *projected gradient descent*, and the variables $\mathring{\mu}$ are the linearized Lagrange multipliers which ensure that \mathring{f} is an infinitesimal conformal deformation. Since the conformal constraint is nonlinear, taking an explicit step in the direction \mathring{f} will not exactly preserve the conformal class and so one still needs to project the surface back to the space of conformal immersions [Schumacher 2017].

4.2 Competitive gradient flow

As we just saw, evolving the surface by a projected gradient step and a constraint projection is inefficient, due to the nonlinear dependence of the discrete conformal class on the vertex positions. Put another way, restricting the evolution to the space of conformal immersions means that the projected gradient descent takes very small steps along the nonlinear subspace in \mathbb{R}^{3V} of conformal immersions.

Rather than choosing the Lagrange multipliers $\mathring{\mu}$ to make \mathring{f} a conformal variation, as in (PGD), we will update both the surface and the Lagrange multipliers according to the *competitive gradient descent* (CGD) of [Schäfer and Anandkumar 2019]. This approach aims to find saddle points of the minimax problem on the Lagrangian

$$\min_f \max_{\mu} \mathcal{L}(f, \mu) := \mathcal{E}(f) + \langle C^* \mu | \lambda(f) \rangle,$$

which correspond to constrained critical points of \mathcal{E} . We extend this evolution to the Riemannian setting and we will see how it naturally generalizes the projected gradient flow even when not on the prescribed constraint manifold.

Analogous to the derivation of gradient descent through a linear approximation of \mathcal{E} , we consider a local *bilinear* approximation of \mathcal{L} in terms of the surface (\mathring{f}) and the Lagrange multiplier ($\mathring{\mu}$) variations:

$$\begin{aligned} \min_{\mathring{f}} \mathcal{L}(f, \mu) + \langle \text{grad } \mathcal{E} + D\lambda^* C^* \mu | \mathring{f} \rangle + \langle D\lambda^* C^* \mathring{\mu} | \mathring{f} \rangle + \frac{1}{2t} \langle \mathring{f}, \mathring{f} \rangle \\ \max_{\mathring{\mu}} \mathcal{L}(f, \mu) + \langle \mathring{\mu} | C\lambda(f) \rangle + \langle \mathring{\mu} | C D\lambda(\mathring{f}) \rangle - \frac{1}{2t} \langle \mathring{\mu}, \mathring{\mu} \rangle. \end{aligned}$$

Notice that the bilinear (third) term is present in both equations and couples them. This particularly simple form allows us to explicitly obtain the updates \mathring{f} and $\mathring{\mu}$ by solving the generalized saddle point system

$$\begin{pmatrix} \frac{1}{t}I_f & X^T \\ -X & \frac{1}{t}I_{\mu} \end{pmatrix} \begin{pmatrix} \mathring{f} \\ \mathring{\mu} \end{pmatrix} = - \begin{pmatrix} \text{grad } \mathcal{E} + X^T \mu \\ C\lambda(f) \end{pmatrix}, \quad (\text{CGD})$$

where $X := C D\lambda$ is the linearization of the constraints, and where I_f and I_{μ} are the mass matrices in the space of surfaces and Lagrange multipliers, respectively. For the Lagrange multipliers, we use the L^2 -inner product of the induced conformal tension, defined in Equation 9. Notice that the right hand side is the same as the one that occurs in (PGD) when f lies on the constraint manifold and μ is orthogonal to the constraint manifold—this shows that the competitive gradient descent provides a minimal modification of projected gradient descent, which can however also explore the space of non-conformal immersions (see Figure 7 for a comparison). In the same way that the Sobolev gradient descent defines a quasi-Newton method for $\text{grad } \mathcal{E}(f) = 0$, we see that the Sobolev CGD also defines a quasi-Newton method for the constrained Euler-Lagrange equation from Theorem 3.2. Since this simple modification to PGD makes sure the flow points back to the constraint manifold, it is no longer necessary to explicitly project the surface back to a conformal immersion. This generalization has an additional practical benefit over projected methods: one can now prescribe the conformal class to be different than that of the initial surface (see Figure 2).

5 RELAXING INTEGRABILITY

So far we have considered vertex positions as the variables of our constrained optimization. Since differential coordinates are used extensively in geometry processing [Sorkine 2006; Xu and Zhou 2009], it is worthwhile to see how our conformality constraint and the attendant constrained minimization problems can be parameterized with differentials. As we will see in this section, doing so reveals the underlying connection between the discrete Lagrange multipliers and their smooth counterparts. It also provides additional computational benefits on challenging examples. To make this concrete we introduce *triangle fields*, which are natural extensions of differential coordinates beyond just parameterizing surface deformations, and study discrete conformal variational problems at the level of differential coordinates (see Figure 8).

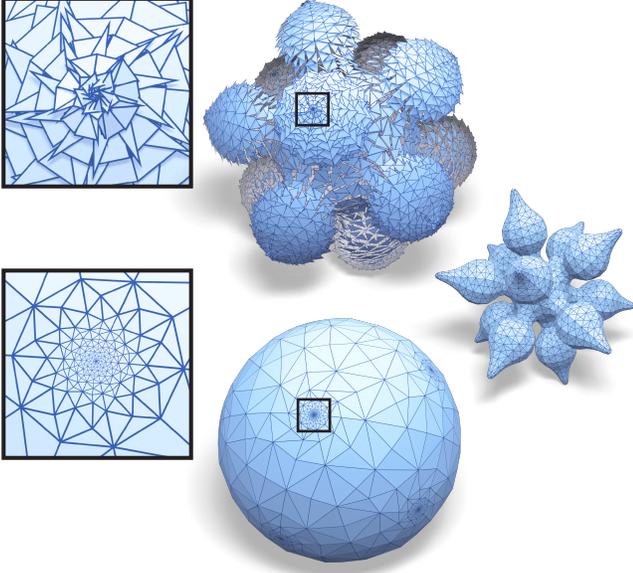


Fig. 8. Numerically difficult optimizations such as this “spiky sphere” Willmore minimization can be much improved with the use of triangle fields. These do not have to integrate into a surface (“glue together”) during the optimization. Enforcing integrability only near the end of the optimization then provides a richer set of paths.

5.1 Triangle Fields and the Lens Complex

The differential df of a piecewise linear immersion $f: M \rightarrow \mathbb{R}^3$ is a piecewise constant \mathbb{R}^3 -valued 1-form on M , i.e.

$$df \in \Omega_F^1(M; \mathbb{R}^3) := \{(\omega_{ijk})_{ijk \in F} \mid \omega_{ijk} : T_{ijk}M \rightarrow \mathbb{R}^3 \text{ linear}\}.$$

Clearly, the differential of an immersion is injective on all triangles ijk . If we relax integrability this leads us to consider the open subset

$$\mathcal{M} = \{\omega \in \Omega_F^1(M; \mathbb{R}^3) \mid \omega_{ijk} \text{ injective for all } ijk \in F\}$$

of $\Omega_F^1(M; \mathbb{R}^3)$ —a $6|F|$ -dimensional vector space. The elements of \mathcal{M} map the abstract triangles in M to triangles in \mathbb{R}^3 (up to translation) and we call them *triangle fields*. They induce geometric quantities like angles, lengths, and normals on the abstract triangles.

To this end we need to distinguish between the two halfedges comprising an edge in E . For any such halfedge $ij \in H$ let $X_{ij} \in T_{ijk}M$ be the corresponding tangent vector, and for $\omega \in \Omega_F^1(M; \mathbb{R}^3)$ we write

$$\omega_{ij} := \omega_{ijk}(X_{ij}) = \int_{ij} \omega_{ijk}$$

for the evaluation of ω along the halfedge ij .

Clearly, for an immersion f we have $df \in \Omega^1(M; \mathbb{R}^3)$. But for a generic triangle field ω this is no longer true—there are $ij \in H$ where

$$\omega_{ij} = \omega_{ijk}(X_{ij}) \neq -\omega_{jil}(X_{ji}) = -\omega_{ji}.$$

In particular, also the lengths of corresponding halfedges may differ. Therefore it is convenient to consider all objects as living on a modified cell complex, whose discrete differential forms naturally contain the discrete forms of M as well as the piecewise constant

1-forms $\Omega_F^1(M; \mathbb{R}^3)$. The idea is closely related to the halfedge forms introduced in [Custers and Vaxman 2020].

Lens Complex. The lens complex \tilde{M} of M is obtained by gluing digons in between opposite halfedges in M (Figure 10). Hence we have the following immediate correspondences,

$$\tilde{V} \cong V, \quad \tilde{E} \cong H, \quad \tilde{F} \cong E \cup F.$$

Whenever clear from context we label the cells of \tilde{M} by the corresponding elements in M —for example, when working over the lens complex we choose to write V to indicate the vertex set of \tilde{M} .

Under these identifications we immediately get the equalities:

$$\Omega^0(\tilde{M}; \mathbb{R}^3) = \Omega^0(M; \mathbb{R}^3), \quad \Omega^2(\tilde{M}; \mathbb{R}^3) = \Omega_E^2(\tilde{M}; \mathbb{R}^3) \oplus \Omega^2(M; \mathbb{R}^3),$$

where $\Omega_E^2(\tilde{M}; \mathbb{R}^3) \cong \mathbb{R}^{3E}$. Moreover, note that each edge in \tilde{E} comes with a canonical orientation given by the orientation of the corresponding half edge in H . Hence $\Omega^1(\tilde{M}; \mathbb{R}^3) \cong \mathbb{R}^{3H}$ —with no further conditions imposed. Among the elements of $\Omega^1(\tilde{M}; \mathbb{R}^3)$ the 1-forms ω on M are distinguished by the condition that $\omega_{ij} = -\omega_{ji}$ for all $ij \in H$, i.e. ω is closed on digons:

$$\Omega^1(M; \mathbb{R}^3) \cong \{\omega \in \Omega^1(\tilde{M}; \mathbb{R}^3) \mid d\omega \in \Omega^2(M; \mathbb{R}^3)\}.$$

A similar condition characterizes the piecewise constant 1-forms among the 1-forms on \tilde{M} : Since the edge vectors of a triangle sum up to zero, every $\omega \in \Omega_F^1(M; \mathbb{R}^3)$ defines a 1-form on \tilde{M} which is closed on triangles F . Conversely, every 1-form $\omega \in \Omega^1(\tilde{M}; \mathbb{R}^3)$ which is closed on triangles F defines a piecewise constant 1-form on M . Hence

$$\Omega_F^1(M; \mathbb{R}^3) \cong \{\omega \in \Omega^1(\tilde{M}; \mathbb{R}^3) \mid d\omega \in \Omega_E^2(\tilde{M}; \mathbb{R}^3)\}.$$

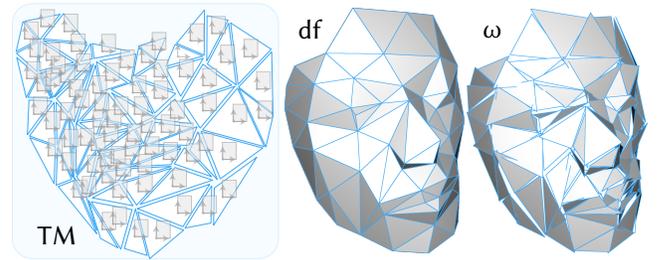


Fig. 9. \mathbb{R}^3 -valued discrete piecewise constant 1-forms identify intrinsic tangent spaces (shown as per-triangle copies of the 2D plane, left) with subspaces of \mathbb{R}^3 spanned by the embedded triangles (middle, right). Discrete 1-forms df that arise as differentials of \mathbb{R}^3 -valued zero forms f “stitch up” to discrete immersions (middle). However, using this tangent space identification we can define geometric quantities also for more general, not necessarily integrable candidate differentials ω (triangle fields, right).

L^2 inner product. To express orthogonality between forms as it appears, for example, in the Hodge decomposition we will use the L^2 inner product on $\Omega_F^1(M; \mathbb{R}^3)$

$$\langle\langle \dot{\omega}, \dot{\omega} \rangle\rangle = \int_M \langle \dot{\omega} \wedge * \dot{\omega} \rangle.$$

Here star denotes the Hodge star with respect to the triangle field metric and the integral of the piecewise constant 2-form $\langle \dot{\omega} \wedge * \dot{\omega} \rangle$ is

given as a sum of its integrals over triangles. Since $\hat{\omega} \in \Omega_{\mathbb{F}}^1(M; \mathbb{R}^3)$ is constant (and hence closed) on each triangle ijk , its integral over ijk can be computed by the usual cotangent formula:

$$\int_{ijk} \langle \hat{\omega}_{ijk} \wedge * \hat{\omega}_{ijk} \rangle = \frac{1}{2} (\cot \theta_{ij} |\hat{\omega}_{ij}|^2 + \cot \theta_{jk} |\hat{\omega}_{jk}|^2 + \cot \theta_{ki} |\hat{\omega}_{ki}|^2),$$

where $\theta_{ij}, \theta_{jk}, \theta_{ki} \in (0, \pi/2)$ denote the interior angles in ijk opposite the halfedges ij, jk, ki . In particular, $(*\alpha)_{ij} := \frac{1}{2} \cot \theta_{ij} \alpha_{ij}$ defines a map $\Omega^1(\tilde{M}; \mathbb{R}^3) \rightarrow \Omega^1(\tilde{M}^*; \mathbb{R}^3)$ such that

$$\langle \hat{\omega}, \hat{\omega} \rangle = \langle * \hat{\omega} | \hat{\omega} \rangle,$$

where on the right hand side $\hat{\omega}$ is regarded as a 1-form $ij \mapsto \hat{\omega}_{ij}$ on the lens complex.

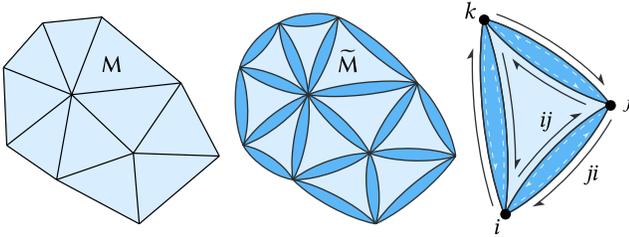


Fig. 10. Starting with a triangle mesh (left) we obtain the lens complex (middle) by blowing up every edge into a “lens”. A pair of oppositely oriented halfedges ij, ji in M bound a lens facet (2-cell) in \tilde{M} (right).

Integrability. A triangle field ω is locally the differential of a surface if it is closed. Since ω is already closed on the triangles, the only obstruction to local integrability is whether ω is closed on the lenses. For a lens $ij \in E \subset \tilde{F}$

$$(d\omega)_{ij} = -(\omega_{ij} + \omega_{ji}).$$

If ω is closed on the lenses then ω defines an \mathbb{R}^3 -valued discrete 1-form on M . The DEC Hodge decomposition shows that the remaining obstruction to global integrability is the harmonic component of ω . We conclude that ω is the differential of a surface if and only if

$$d\omega = 0 \in \Omega_{\mathbb{F}}^2(\tilde{M}; \mathbb{R}^3) \quad (12)$$

$$\langle \omega, \eta_i \rangle = 0 \in \mathbb{R}^3 \quad (13)$$

where $\{\eta_i\}_{i=1}^{2g}$ is a basis of harmonic 1-forms on M .

5.2 Conformal Lagrange Multipliers

Here we derive the conformal Lagrange multipliers for triangle field energies. We explain how they can be interpreted as conformal stress tensors (quadratic differentials) and how this discretization is related to other notions in geometry processing.

Length Multi-Ratios. A triangle field ω does not define a metric on M but induces edge lengths ℓ on the lens complex \tilde{M} which satisfy the triangle inequality in each face of M with no compatibility condition over the lenses—we call such a collection of halfedge lengths a *triangle field metric* on M .

The notion of conformal equivalence extends to triangle field metrics in the obvious way—two triangle field metrics ℓ and $\tilde{\ell}$

are called equivalent if there is a function $u: V \rightarrow \mathbb{R}$ such that $\tilde{\ell}_{ij} = e^{(u_i+u_j)/2} \ell_{ij}$ for all $ij \in H$. Following [Bobenko et al. 2016], where discrete conformal equivalence of polyhedral surfaces is studied, we characterize discrete conformal equivalence over the lens-complex by *logarithmic length multi-ratios* $\log m: \tilde{E} \rightarrow \mathbb{R}$

$$\log m(\ell)_{ij} = \log \left(\frac{\ell_{il} \ell_{jk} \ell_{ji}}{\ell_{lj} \ell_{ki} \ell_{ij}} \right) = (\lambda_{il} - \lambda_{lj} + \lambda_{jk} - \lambda_{ki}) + (\lambda_{ji} - \lambda_{ij})$$

where k, l are the vertices across the edge ij as before.

Now, if we define—analogously to Section 3—the averaging and multi-ratio matrices, $\tilde{A} \in \mathbb{R}^{\tilde{E} \times V}$ and $\tilde{C} \in \mathbb{R}^{\tilde{E} \times \tilde{E}}$, by

$$(\tilde{A}u)_{ij} = \frac{1}{2}(u_i + u_j), \quad \tilde{C}\lambda = \log m(\ell). \quad (14)$$

then the relationship $\text{im } \tilde{A} = \ker \tilde{C}$ (Lemma A.1) directly implies:

PROPOSITION 5.1. *Two triangle field metrics are discretely conformally equivalent if and only if their length multi-ratios are the same.*

Euler-Lagrange Equations. Given that $\text{im } \tilde{A} = \ker \tilde{C}$, the derivation of the Euler-Lagrange equations in the triangle field setting follows verbatim the arguments of Section 3. Carrying this out one ends up with the following characterization: A triangle field ω is a critical point of a triangle field energy $\tilde{\mathcal{E}}$ under infinitesimal conformal deformations if and only if $\text{grad } \tilde{\mathcal{E}} = D\lambda^* q$ for

$$q \in \ker \tilde{A}^* \cong \{q: \tilde{E} \rightarrow \mathbb{R}: \sum_{ij} q_{ij} + q_{ji} = 0, \forall i \in V\}.$$

A variation $\hat{\omega} \in \Omega_{\mathbb{F}}^1(M; \mathbb{R}^3)$ of ω yields a variation $\hat{\lambda}$ of its logarithmic lengths

$$\hat{\lambda}_{ij} = \frac{\langle \omega_{ij}, \hat{\omega}_{ij} \rangle}{|\omega_{ij}|^2}$$

(cf., Equation 5). In particular, expressing this with respect to the L^2 -metric, we obtain that the L^2 -gradient of λ_{ij} is given by

$$\text{grad } \lambda_{ij} = \pi_{ij} := \begin{cases} \frac{1}{A_{ijk}} \omega_{ijk} \circ P_{ij} & \text{on } ijk \\ 0 & \text{elsewhere} \end{cases}$$

where $P_{ij} := \frac{\langle \cdot, X_{ij} \rangle}{|X_{ij}|^2} X_{ij}$ is the projector onto the halfedge ij in the tangent space $T_{ijk}M$. Thus we obtain the following – note the difference between Equation 6 and Equation 15:

THEOREM 5.2. *A triangle field $\omega \in \mathcal{M}$ is critical for an energy $\tilde{\mathcal{E}}: \mathcal{M} \rightarrow \mathbb{R}$ under all infinitesimal conformal variations if and only if there exists $q: \tilde{E} \rightarrow \mathbb{R}$ satisfying*

$$\sum_{ij} q_{ij} + q_{ji} = 0 \quad \forall i \in V, \quad (15)$$

$$\text{grad } \tilde{\mathcal{E}} + \sum_{ij \in \tilde{E}} q_{ij} \pi_{ij} = 0, \quad (16)$$

where $\text{grad } \tilde{\mathcal{E}}$ is the L^2 -gradient of the energy.

To relate this to the surface case we now consider infinitesimal conformal triangle field variations that preserve integrability. By the DEC Hodge decomposition on the lens complex we know that a variation $\hat{\omega} \in \Omega_{\mathbb{F}}^1(M; \mathbb{R}^3)$ is exact if and only if it is orthogonal to all

co-closed 1-forms. In this case the left hand side of Equation 16 is not necessarily zero, but co-closed:

COROLLARY 5.3. *A triangle field $\omega \in \mathcal{M}$ is critical for an energy $\tilde{\mathcal{E}} : \mathcal{M} \rightarrow \mathbb{R}$ under all infinitesimal exact conformal variations if and only if there exists $q : E \rightarrow \mathbb{R}$ satisfying*

$$\sum_{ij} q_{ij} = 0 \quad \forall i \in V, \quad (17)$$

$$d * \left(\text{grad } \tilde{\mathcal{E}} + \sum_{ij \in \tilde{E}} q_{ij} \pi_{ij} \right) = 0, \quad (18)$$

where $\text{grad } \tilde{\mathcal{E}}$ is the L^2 -gradient of the energy.

Discrete Conformal Stress. These Euler-Lagrange equations show that conformal Lagrange multipliers $q \in \ker \tilde{A}^*$ result in a force

$$\tilde{\tau}_{ijk} := q_{ij} \pi_{ij} + q_{jk} \pi_{jk} + q_{ki} \pi_{ki}. \quad (19)$$

that acts on deformations of the extrinsic triangle field. Noticing that this is a tangential \mathbb{R}^3 -valued 1-form, we can actually write $\tilde{\tau}_{ijk} = \frac{1}{A_{ijk}} \omega_{ijk} \circ S_{ijk}$ where $S_{ijk} : T_{ijk} \mathcal{M} \rightarrow T_{ijk} \mathcal{M}$ is the *conformal stress tensor*

$$S_{ijk} := q_{ij} P_{ij} + q_{jk} P_{jk} + q_{ki} P_{ki}. \quad (20)$$

An infinitesimal deformation of the intrinsic geometry of the faces is described by a self-adjoint linear map $D_{ijk} : T_{ijk} \mathcal{M} \rightarrow T_{ijk} \mathcal{M}$ per face $ijk \in F$, and as usual in mechanics the stress tensor gives rise to the energy change $\langle S_{ijk}, D_{ijk} \rangle$ when the strain D_{ijk} is attempted.

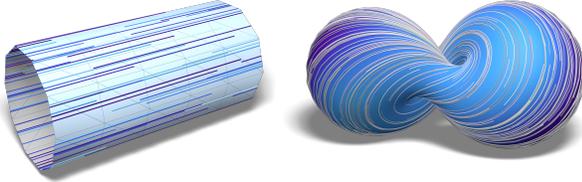


Fig. 11. A conformal stress tensor can be visualized through its principal stress directions. In each face ijk we compute the eigenspace of S_{ijk} with the largest eigenvalue which defines a line field (aka the *horizontal foliation* of the quadratic differential q) and we visualize its integral curves.

Since $\text{tr } P_{ij} = 1$, the condition that q is a conformal Lagrange multiplier (Equations 15 and 17) states that the conformal stress tensor is trace-free when averaged around a vertex. To better understand what this means one can compare this to the case when S_{ijk} is trace-free inside every face: in this case, S_{ijk} does not resist scaling deformations and only resists deformations that change the corner angles. The trace-free condition averaged to the vertices means that conformal stress tensor does not measure resistance to isotropic scaling, and only resistance to anisotropic distortion. In this context, the notion of anisotropic distortion is described precisely through discrete conformal equivalence.

Holomorphic Quadratic Differentials. A similar description of quadratic differentials has appeared in related work on extremal quasi-conformal mappings [Weber et al. 2012, Eq. 8]. There a holomorphic

quadratic differential is discretized in a piecewise linear way per face as

$$S_{ijk} = \frac{1}{6A_{ijk}} (h_{ij} P_{ij} + h_{jk} P_{jk} + h_{ki} P_{ki})$$

where $h \in \Omega^1(M; \mathbb{R})$ is harmonic. Other quasi-conformal parameterization settings also give rise to holomorphic quadratic differentials (see, e.g., [Lui et al. 2014; Lei et al. 2017]).

6 RESULTS

We begin this section with a discussion of our implementation and its numerical properties before turning to examples of our method in action.

Our algorithm was implemented in *C++* and we performed numerical experiments on a MacBookPro class computer (2.8GHz Intel Core i7-7700HQ, 8GB of RAM). To solve the saddle point systems from Section 4 we precomputed symbolic factorizations via PAR-DISO [Alappat et al. 2020; Bollhöfer et al. 2020, 2019] and then updated the numerical factorizations and applied backsubstitution for each subsequent problem. We also precomputed the numerical factorization of the cotan Laplacian.

The iterations of our optimization are dominated by the solution of a sparse linear system of size $3|V| + |E|$. For an overall sense of the runtime we note that for a mesh of 125k triangles, preprocessing takes about 6 seconds; the following iterations take approximately 3 seconds—computing geometric quantities and derivatives takes about 600ms, and constructing and solving the saddle point problem takes an additional 3 seconds.

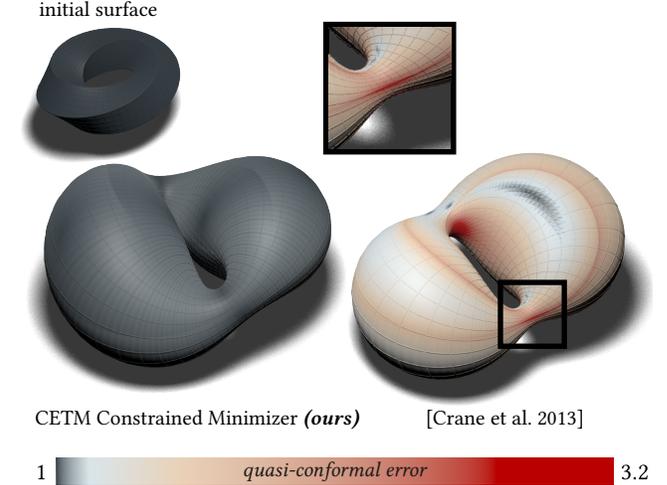


Fig. 12. Starting with a twisted torus (upper left) our method finds a conformally equivalent Willmore minimizer with very low quasi-conformal distortion (lower left). Using the method of [Crane et al. 2013], based on conformal deformations, numerical drift leads to accumulating quasi-conformal error (lower right). Note in particular the “neck” region where parametric squares are distorted into high aspect ratio rectangles.

A distinction of our approach relative to earlier works is the use of discrete conformal equivalence as the basis for the conformality

constraint. This ensures that the conformal class is preserved up to double precision accuracy. Practically this implies that our approach preserves mesh quality at convergence (and in practice throughout the flow). In Figure 13 we visualize the quasiconformal error Q of Willmore minimizers comparing presence resp. absence of the conformality constraint. There is no quasiconformal error when $Q = 1$. Without any conformality constraint the element quality can be severely degraded since the reparameterization degree of freedom is wholly uncontrolled. This deterioration of the mesh is a well known issue [Bobenko and Schröder 2005; Barrett et al. 2016; Gruber and Aulisa 2020]. Simply discretizing a smooth notion of conformal deformations [Crane et al. 2011, 2013] is still insufficient due to numerical discretization errors accumulating along the flow leading to significant drift away from the initial conformal class (Figure 12). Such discretization errors depend heavily on the aspect ratios of triangles which our method is quite insensitive to. Figure 14 illustrates this by showing that two quite different triangulations of a given shape yield conformally constrained minimizers of the same shape.

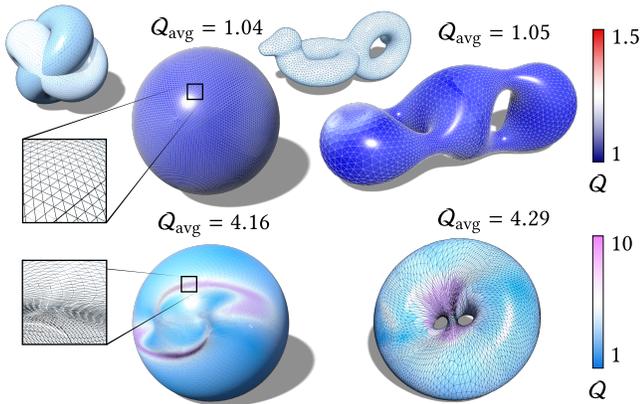


Fig. 13. Although Willmore minimization with no conformal class constraint (bottom row) yields smooth surfaces, there can be severe distortion in the mesh quality. Fixing the discrete conformal class (top row) yields low quasiconformal error as expected. Recall that a conformal deformation is locally a similarity hence tends to preserve element quality.

Our choice of competitive gradient descent as a method for constrained optimization was motivated by comparing it with standard constrained optimization methods (Figure 7). In particular, we considered the projected gradient descent (Section 4.1) and an augmented Lagrangian method [Bertsekas 1996; Nocedal and Wright 2006]—the projected and competitive gradients were defined with respect to both the L^2 and H^2 inner products, and updates were computed via explicit forward steps. Overall, our strategy was both more efficient and more robust to the quality of the initial surface (or triangle field). As our approach can be seen as a relaxation of the projected gradient descent, we found that the projected gradient descent performed quite well in general. However, it was comparatively slow since it enforced constraint satisfaction every iteration. Not only did it increase the per-iteration computation time, but the evolution would get “stuck” in regions where the projected Willmore gradient is

nearly zero—Figure 7 shows an example where our method used non-conformal deformations (highlighted box in figure) to resolve these challenging situations. The Augmented Lagrangian method was the least consistent and it would often get stuck along the way or in undesirable local minima (inset); this method had slow convergence and its performance was very sensitive to the choice of penalty parameters.

Since we also consider a non-convex optimization problem, there are situations where our approach has difficulty converging. In particular, initializing the optimization from a nearly degenerate and non-conformal triangle field results in seemingly random evolution. In practice, we only ran into these problems when trying to dramatically change the prescribed discrete conformal class of a non-integrable triangle field.

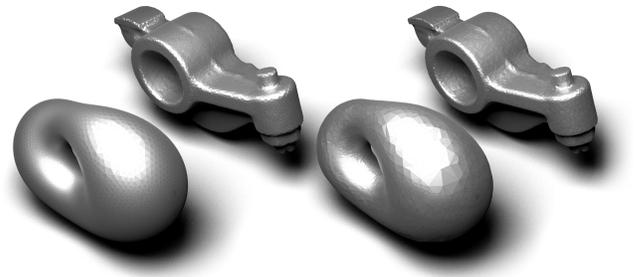


Fig. 14. Length cross ratios are a discrete notion of conformal equivalence and remain invariant under conformal change. Consequently our method is insensitive to the quality of the underlying triangulation. We demonstrate this here by starting with two different triangulations of a given input shape, resulting in conformally constrained Willmore minimizers of the same shape.

6.1 Constrained Willmore Surfaces

Finding optimal shapes has an extensive history in mathematics, with motivation that goes beyond aesthetics. This search often reveals an understanding of deep mathematical concepts and the physical phenomena governed by these ideas. Willmore surfaces arose from the desire to combine minimal surface theory with (Möbius) conformal geometry of the ambient space [Blaschke 1929]. Constrained Willmore surfaces then arise naturally by additionally constraining the intrinsic conformal geometry of the surface—they are the subject of an active mathematical research field [Babich and Bobenko 1993; Bohle 2010; Kuwert and Schätzle 2013; Heller 2015; Heller and Ndiaye 2019].

Besides general results about existence and uniqueness of constrained minimizers, the existing theory and examples almost exclusively concern tori. Neither explicit constructions nor numerical visualizations of constrained Willmore surfaces have been available in the general case. By minimizing the Willmore energy in a fixed conformal class, we provide a new numerical framework to experimentally study constrained Willmore surfaces of any genus.

6.1.1 Genus 0 Surfaces (Spheres). According to the uniformization theorem, all genus zero surfaces are conformally equivalent to the

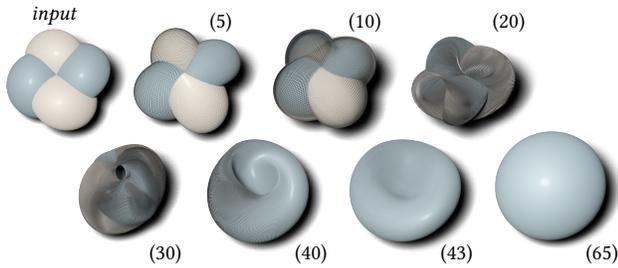


Fig. 15. Conformal Willmore-minimizing flow starting from a slightly perturbed Morin surface, the fourfold rotationally symmetric surface that appears as the half-way model of the sphere eversion [Francis 1987]. Here we visualize the evolving triangle fields at different iterations.

sphere. We then expect our Willmore minimizing algorithm to produce round spheres for any genus zero input. In Figure 15 we start with a *Morin surface*, a fourfold rotationally symmetric surface that appears as the half-way model of a sphere eversion [Francis 1987]. The surface was slightly perturbed to break the symmetry, and given as input to our algorithm, which converges to a round sphere. Figure 3 gives an example of a sphere minimizing the Willmore energy under conformal class, area, and volume constraints.

6.1.2 Genus 1 Surfaces (Tori). Several classes of explicit examples of constrained Willmore tori have been constructed [Pinkall 1985; Heller 2013], but there is still little known about minimizers in conformal classes far away from the square torus (Clifford torus). Using triangle fields, we find constrained Willmore tori that are *local* minimizers for a given conformal class—this also yields a visualization of every abstract conformal type as a surface in space.

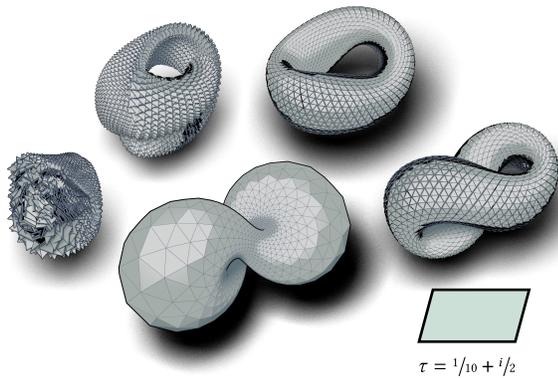


Fig. 16. To construct a constrained Willmore torus with a given conformal class we triangulate the lattice representing the flat torus in the same conformal class and start with a random triangle field on this mesh. Here we visualize iterations 1, 10, 20, 30, and 47 of the optimization. (See also the top of Figure 1.)

To find these constrained minimizers via surface flows it is necessary to begin with an initial immersion of the abstract conformal class, which is just part of the original problem we’re trying to solve.

On the other hand, triangle fields of every conformal class are easily constructed by simply taking a triangulation of the fundamental domain. Competitive gradient descent of the triangle field gives a simple algorithm to compute constrained Willmore tori directly from a description of the conformal class. We build on the spinorial framework of [Chern et al. 2018] to ensure that our minimizers do not have pinch points. Technical and implementation details are given in Appendices C and D. Convergence to a critical point is measured by the residual norm of the Euler-Lagrange equations—we stop our optimization once this error is below 10^{-6} .

When starting from an initial immersion that exhibits certain symmetries (Figure 17), our algorithm may first flow towards a critical point that exhibits those same symmetries. We can use this feature to produce symmetric critical points of the Willmore energy that are not minimizers.

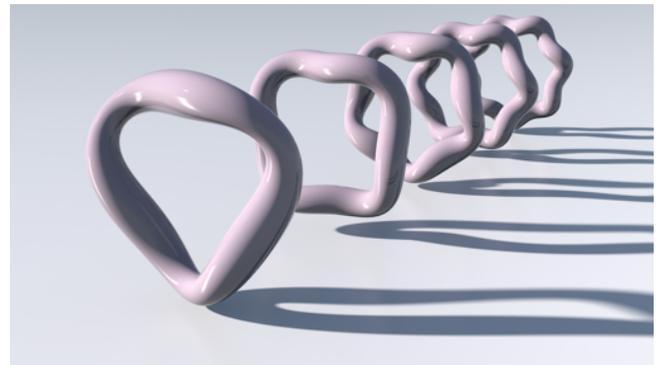


Fig. 17. Constrained Willmore tori with n -fold rotational symmetry.

6.1.3 Higher Genus Surfaces. Since we consider general conformal variational problems we are able to produce numerical realizations of general constrained Willmore minimizers of higher genus for the very first time. A gallery of such results, produced by giving various high genus surfaces as input, is shown in Figure 5.

6.2 Conformal Surface Modeling

Conformal variational problems provide an extensible surface modeling framework where many important geometric features can be easily controlled. For example, since conformal maps preserve textures our conformal Willmore minimization provides an excellent replacement for standard fairing methods which can distort textures even with a very small amount of smoothing. Below, we consider several illustrative examples of the conformal constraint applied to problems in geometry processing.

6.2.1 Point Constraints. In the presence of conformality, fixing positions of some points provides control of surface geometry. In Figure 1 (bottom), we see that fixing vertex positions while minimizing the Willmore energy without the conformal constraint results in a surface that does not resemble the initial geometry at all. On the other hand requiring the deformation to be conformal is enough for few point constraints to encode the overall geometry of the initial surface. Figure 19 helps explain this behavior in the simple case of a

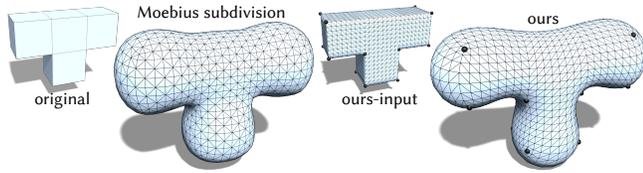


Fig. 18. Our Willmore-minimizing flows (right) produce interpolatory surfaces that are locally as sphere-like as possible. On the left, a result of canonical Möbius subdivision [Vaxman et al. 2018], which specifically targets sphere-reproducing interpolation, albeit using a subdivision approach.

sphere. Changing the cross-ratio of any four point constraints leads to constrained Willmore surfaces that are not Möbius equivalent. Instead conformally constrained Willmore surfaces with additional point constraints provide as-spherical-as-possible interpolants of the points. This property has been recognized as desirable in architectural geometry processing. Subdivision schemes based on Möbius geometry were developed to provide approximations of such sphere preserving interpolants [Vaxman et al. 2018]. Figure 18 compares our conformally constrained Willmore surface with Möbius subdivision.

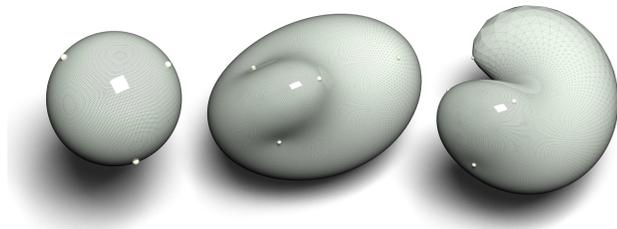


Fig. 19. Four points selected on a round sphere fix the edge length cross-ratio on the two triangles formed by them. Moving one of the points asymmetrically changes the induced cross-ratio means, so no Möbius transformation can map between the old and the new point positions. This causes the shape to deviate from the original sphere.

6.2.2 Gradient Domain Surface Processing. Another way to encode the input shape into our framework is to modify the energy functional so that the minimizing surfaces balance smoothness and closeness to the input shape. We used the L^2 -norm as a measure of similarity, resulting in the modified energy

$$\mathcal{E}_\varepsilon(f) := \frac{1}{2} \|f - f_0\|_{L^2}^2 + \varepsilon \mathcal{W}(f), \quad (21)$$

where ε is a smoothness parameter.

These energies are commonplace in gradient domain image/surface processing [Mumford and Shah 1989] and computer graphics [Chuang et al. 2016] in the context of image and surface denoising. In contrast to using point constraints, minimizing Equation 21 in a fixed conformal class gives only a “soft” constraint on the vertex positions (Figure 20).

6.2.3 Constrained Area Minimal Surfaces. The discrete conformal class constraint can be applied to other than curvature energies. A classical example is the area energy, whose minimizers are minimal surfaces modeling the behavior of soap films. Applying a conformal

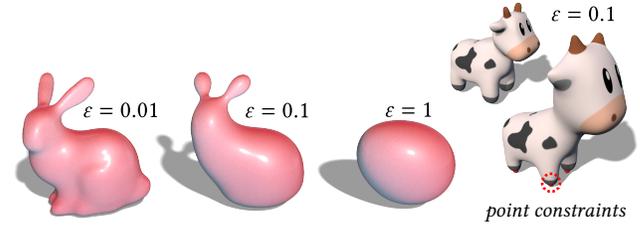


Fig. 20. Willmore-regularized surface fairing which jointly optimizes for fidelity to the input, and minimal Willmore energy with the parameter ε controlling the regularizing effect of the Willmore energy.

class constraint in this setting yields a material which can shrink but only isotropically, in effect a type of isotropic auxetic soap film. Recall that auxetic materials are characterized by a negative Poisson ratio and in the isotropic case best modeled by conformal maps [Konaković et al. 2016]. Figure 6 (cylinder) shows an example of such a conformally constrained minimal surface. The usual shrinkage towards the middle of the cylinder (with fixed boundaries) is not present here since it would require anisotropic deformation. In fact the conformal forces counterbalance the mean curvature normal force. Figure 21 shows more complex examples of a cylinder with twist and a sphere with six caps cut off.

We leave exploration of such materials in the context of geometric and physical modeling to future work.



Fig. 21. Minimal surfaces which minimize the area energy often collapse due to their propensity for shrinkage. Once the conformality constraint is added, as done here, any area scaling must be isotropic. The resulting shapes, effectively made of an isotropic auxetic soap film material, are distinctly different.

7 CONCLUSION

We presented conformally constrained variational problems as a new framework for solving problems in mathematical visualization and geometry processing. The conformal constraint is discretized using the notion of discrete conformal equivalence, and we show that this constraint on the intrinsic parameterization has profound effects on the extrinsic shape of a surface. We used this framework to numerically compute constrained Willmore for surfaces of any genus and conformal type for the very first time. The numerical optimization is based on a manifold reformulation of competitive

gradient descent, which we found to be more robust than standard methods. The conformal constraint, in the presence of Willmore optimality, produces interesting surfaces with rich geometric properties that may be desirable for form finding purposes. Our algorithm is sufficiently general and can be adapted to other geometric energies or user-defined modeling constraints. Investigating the effect of the conformal constraint on other variational problems is an interesting avenue for future work. On the theory side, we look forward to deepen our understanding of the relationship between our results and isothermic triangulated surfaces, discrete holomorphic quadratic differentials, Dirac operators, and Teichmüller spaces, which would nicely couple smooth and discrete theory.

ACKNOWLEDGMENTS

This work was supported in part by a National Science Foundation graduate research fellowship, the Kortschak Scholars program at Caltech, the Einstein Foundation Berlin, the Deutsche Forschungsgemeinschaft (SFB/TRR 109 “Discretization in Geometry and Physics”), and the Taiwanese National Center for Theoretical Sciences (Mathematics Division). Additional support was provided by the Information Science and Technology Initiative at Caltech and SideFX Software.

REFERENCES

- Christie Alappat, Achim Basermann, Alan R. Bishop, Holger Fehske, Georg Hager, Olaf Schenk, Jonas Thies, and Gerhard Wellein. 2020. A Recursive Algebraic Coloring Technique for Hardware-Efficient Symmetric Sparse Matrix-Vector Multiplication. *ACM Trans. Par. Comput.* 7, 3 (2020), 19:1–37.
- Mikhail V. Babich and Alexander I. Bobenko. 1993. Willmore Tori with Umbilic Lines and Minimal Surfaces in Hyperbolic Space. *Duke Math. J.* 72, 1 (1993), 151–185.
- John W. Barrett, Harald Garcke, and Robert Nürnberg. 2016. Computational Parametric Willmore Flow with Spontaneous Curvature and Area Difference Elasticity Effects. *SIAM J. Numer. Anal.* 54, 3 (2016), 1732–1762.
- Dimitri P. Bertsekas. 1996. *Constrained Optimization and Lagrange Multiplier Methods*. Athena Scientific.
- Wilhelm Blaschke. 1929. *Vorlesungen über Differentialgeometrie und geometrische Grundlagen von Einsteins Relativitätstheorie III*. Springer.
- Alexander I. Bobenko, Ulrich Pinkall, and Boris Springborn. 2015. Discrete Conformal Maps and Ideal Hyperbolic Polyhedra. *Geom. & Top.* 19, 4 (2015), 2155–2215.
- Alexander I. Bobenko and Peter Schröder. 2005. Discrete Willmore Flow. In *Proc. Symp. Geom. Proc.* Eurographics, 101–110.
- Alexander I. Bobenko, Stefan Sechelmann, and Boris Springborn. 2016. Discrete Conformal Maps: Boundary Value Problems, Circle Domains, Fuchsian and Schottky Uniformization. In *Advances in Discrete Differential Geometry*, Alexander I. Bobenko (Ed.). Springer, 1–56.
- Christoph Bohle. 2010. Constrained Willmore Tori in the 4-Sphere. *J. Diff. Geom.* 86, 1 (2010), 71–132.
- Christoph Bohle, G. Paul Peters, and Ulrich Pinkall. 2008. Constrained Willmore Surfaces. *Calc. Var.* 32 (2008), 263–277.
- Matthias Bollhöfer, Aryan Eftekhar, Simon Scheidegger, and Olaf Schenk. 2019. Large-scale Sparse Inverse Covariance Matrix Estimation. *SIAM J. Sci. Comp.* 41, 1 (2019), A380–A401.
- Matthias Bollhöfer, Olaf Schenk, Radim Janalik, Steve Hamm, and Kiran Gullapalli. 2020. State-of-the-Art Sparse Direct Solvers. (2020), 3–33.
- Robert Bridson, S. Marino, and Ronald Fedkiw. 2003. Simulation of Clothing with Folds and Wrinkles. In *Proc. Symp. Comp. Anim.* 28–36.
- Marcel Campen, Ryan Capouellez, Hanxiao Shen, Leyi Zhu, Daniele Panozzo, and Denis Zorin. 2021. Efficient and Robust Discrete Conformal Equivalence with Boundary. (2021). arXiv:2104.04614.
- P. B. Canham. 1970. The Minimum Energy of Bending as a Possible Explanation of the Biconcave Shape of the Human Red Blood Cell. *J. Theor. Biology* 26, 1 (1970), 61–81.
- Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2018. Shape from Metric. *ACM Trans. Graph.* 37, 4 (2018), 63:1–17.
- Bennet Chow and Feng Luo. 2003. Combinatorial Ricci Flows on Surfaces. *J. Diff. Geom.* 63, 1 (2003), 97–129.
- Ming Chuang, Szymon Rusinkiewicz, and Misha Kazhdan. 2016. Gradient-Domain Processing of Meshes. *J. Comp. Graph. Tech.* 5, 4 (2016), 44–55.
- Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2011. Spin Transformations of Discrete Surfaces. *ACM Trans. Graph.* 30, 4 (2011), 104:1–10.
- Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Robust Fairing via Conformal Curvature Flow. *ACM Trans. Graph.* 32, 4 (2013), 61:1–10.
- Bram Custers and Amir Vaxman. 2020. Subdivision Directional Fields. *ACM Trans. Graph.* 39, 2 (2020), 11:1–23.
- Mathieu Desbrun, Eva Kanso, and Yiyi Tong. 2008. Discrete Differential Forms for Computational Modeling. In *Discrete Differential Geometry*, Alexander I. Bobenko, Peter Schröder, John M. Sullivan, and Günther M. Ziegler (Eds.). Oberwolfach Seminars, Vol. 38. Birkhäuser Verlag.
- Ilya Eckstein, Jean-Philippe Pons, Yiyi Tong, C.-C. Jay Kuo, and Mathieu Desbrun. 2007. Generalized Surface Flows for Mesh Processing. In *Proc. Symp. Geom. Proc.* Eurographics, 183–192.
- George K. Francis. 1987. *A Topological Picturebook*. Springer.
- Adriano M. Garsia. 1961. An Imbedding of Closed Riemann Surfaces in Euclidean Space. *Comm. Math. Helv.* 35 (1961), 93–110.
- Huabin Ge. 2018. Combinatorial Calabi Flows on Surfaces. *Trans. Amer. Math. Soc.* 370, 2 (2018), 1377–1391.
- Mark Gillespie, Boris Springborn, and Keenan Crane. 2021. Discrete Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 40, 2 (2021).
- David Glickenstein. 2011. Discrete Conformal Variations and Scalar Curvature on Piecewise Flat Two-and Three-Dimensional Manifolds. *J. Diff. Geom.* 87, 2 (2011), 201–238.
- Eitan Grinspun, Anil Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete Shells. In *Proc. Symp. Comp. Anim.* 62–67.
- Anthony Gruber and Eugenio Aulisa. 2020. Computational P-Willmore Flow with Conformal Penalty. *ACM Trans. Graph.* 39, 5 (2020), 161:1–16.
- W. Helfrich. 1973. Elastic Properties of Lipid Bilayers: Theory and Possible Experiments. *Z. Naturforsch. C* 28, 11–12 (1973), 693–703.
- Lynn Heller. 2013. Constrained Willmore tori and elastic curves in 2-dimensional space forms. (2013). arXiv:1303.1445.
- Lynn Heller. 2015. Constrained Willmore and CMC Tori in the 3-Sphere. *Diff. Geom. Appl.* 40 (2015), 232–242.
- Lynn Heller and Cheikh Brahim Ndiaye. 2019. First Explicit Constrained Willmore Minimizers of Non-Rectangular Conformal Class. (2019). arXiv:1710.00533.
- Miao Jin, Junho Kim, and Xianfeng David Gu. 2007. Discrete Surface Ricci Flow: Theory and Applications. In *IMA International Conference on Mathematics of Surfaces*. Springer, 209–232.
- Pushkar Joshi and Carlo Séquin. 2007. Energy Minimizers for Curvature-Based Surface Functionals. *Comp. Aid. Des. Appl.* 4, 5 (2007), 607–617.
- Mina Konaković, Keenan Crane, Bailin Deng, Sofien Bouaziz, Daniel Piker, and Mark Pauly. 2016. Beyond Developable: Computational Design and Fabrication with Auxetic Materials. *ACM Trans. Graph.* 35, 4 (2016), 89:1–11.
- Ernst Kuwert and Reiner Schätzle. 2013. Minimizers of the Willmore Functional under Fixed Conformal Class. *J. Diff. Geom.* 93, 3 (2013), 471–530.
- Wai Yeung Lam and Ulrich Pinkall. 2016. Holomorphic Vector Fields and Quadratic Differentials on Planar Triangular Meshes. In *Advances in Discrete Differential Geometry*, Alexander I. Bobenko (Ed.). Springer, 241–265.
- Wai Yeung Lam and Ulrich Pinkall. 2017. Isothermic Triangulated Surfaces. *Math. Ann.* 368, 1–2 (2017), 165–195.
- Na Lei, Xiaopeng Zheng, Jian Jiang, Yu-Yao Lin, and David Xianfeng Gu. 2017. Quadrilateral and Hexahedral Mesh Generation Based on Surface Foliation Theory. *Comp. Meth. Appl. Mech. & Eng.* 316 (2017), 758–781.
- Lok Ming Lui, Ka Chun Lam, Shing-Tung Yau, and Xianfeng Gu. 2014. Teichmüller Mapping (T-Map) and its Applications in Landmark Matching Registration. *SIAM J. Img. Sci.* 7, 1 (2014), 391–426.
- Feng Luo. 2004. Combinatorial Yamabe Flow on Surfaces. *Comm. Contemp. Math.* 6, 5 (2004), 765–780.
- Feng Luo, Jian Sun, and Tianqi Wu. 2020. Discrete Conformal Geometry of Polyhedral Surfaces and its Convergence. (2020). arXiv:2009.12706.
- Fernando C. Marques and André Neves. 2014. Min-Max Theory and the Willmore Conjecture. *Ann. Math.* 179, 2 (2014), 683–782.
- David Mumford and Jayant Shah. 1989. Optimal Approximations by Piecewise Smooth Functions and associated Variational Problems. *Comm. Pure Appl. Math.* 42, 5 (1989), 577–685.
- Jorge Nocedal and Stephen J. Wright. 2006. *Numerical Optimization* (2 ed.). Springer.
- Ulrich Pinkall. 1985. Hopf Tori in S^3 . *Invent. Math.* 81, 2 (1985), 379–386.
- Ulrich Pinkall and Boris Springborn. 2021. A Discrete Version of Liouville’s Theorem on Conformal Maps. *Geom. Dedicata* (2021).
- R. J. Renka and J. W. Neuberger. 1995. Minimal Surfaces and Sobolev Gradients. *SIAM J. Sci. Comp.* 16, 6 (1995), 1412–1427.
- Tristan Riviere. 2008. Analysis Aspects of Willmore Surfaces. *Invent. Math.* 174, 1 (2008), 1–45.
- Reto A. Ruedy. 1971. Embeddings of Open Riemann Surfaces. *Comm. Math. Helv.* 46 (1971), 214–225.

- Florian Schäfer and Anima Anandkumar. 2019. Competitive Gradient Descent. (2019). arXiv:1905.12103v2.
- Reiner Michael Schätzle. 2013. Conformally Constrained Willmore Immersions. *Adv. Calc. Var.* 6, 4 (2013), 375–390.
- Henrik Schumacher. 2017. On H^2 -Gradient Flow for the Willmore Energy. (2017). arXiv:1703.06469v1.
- Olga Sorkine. 2006. Differential Representations for Mesh Processing. *Comp. Graph. Forum* 25, 4 (2006), 789–807.
- Boris Springborn. 2017. Hyperbolic Polyhedra and Discrete Uniformization. (2017). arXiv:1707.06848.
- Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 27, 3 (2008), 77:1–11.
- Amir Vaxman, Christian Müller, and Ofir Weber. 2015. Conformal Mesh Deformations with Möbius Transformations. *ACM Trans. Graph.* 34, 4 (2015), 55:1–11.
- Amir Vaxman, Christian Müller, and Ofir Weber. 2018. Canonical Möbius Subdivision. *ACM Trans. Graph.* 37, 6 (2018), 227:1–15.
- Ofir Weber, Ashish Myles, and Denis Zorin. 2012. Computing Extremal Quasiconformal Maps. *Comp. Graph. Forum* 31, 5 (2012), 1679–1689.
- Wei-Wei Xu and Kun Zhou. 2009. Gradient Domain Mesh Deformation—A Survey. *J. Comp. Sci. Tech.* 24, 1 (2009), 6–18.
- Christopher Yu, Henrik Schumacher, and Keenan Crane. 2020. Repulsive Curves. (2020). arXiv:2006.07859.

A THE LENS COMPLEX

This appendix details the discrete conformal equivalence over the lens complex \tilde{M} .

A.1 Discrete Conformal Equivalence

Below, we prove that discrete conformal equivalence over the lens complex is characterized by logarithmic multi-ratios.

LEMMA A.1. *The averaging and multi-ratio maps defined in Equation 14 satisfy $\text{im } \tilde{A} = \ker \tilde{C}$.*

PROOF. Consider $\lambda \in \text{im } \tilde{A}$. Notice that $\lambda_{ij} = \lambda_{ji}$, and so λ descends to a well-defined function over the edges of M .

$$(\tilde{C}\lambda)_{ij} = (C\lambda)_{ij} + (\lambda_{ij} - \lambda_{ji}) = (C\lambda)_{ij} = 0,$$

where the last equality follows from Proposition 3.1. This shows that $\text{im } \tilde{A} \subseteq \ker \tilde{C}$.

Consider $\lambda \in \ker \tilde{C}$ and fix a vertex $i \in V$. For every face $ijk \in F$ define candidate scale factors $u_i^{ijk} = \lambda_{ij} - \lambda_{jk} + \lambda_{ki}$. Since

$$u_i^{jil} - u_i^{ijk} = (\tilde{C}\lambda)_{ij} = 0$$

we see that $u_i^{ijk} = u_i^{jil}$ for any neighboring faces. Since M is manifold all of the faces containing i are connected by a path of neighboring faces. Thus, the candidate scale factors u_i^{ijk} only depend on the vertex i and we get a well defined function $u : V \rightarrow \mathbb{R}$. Direct substitution yields

$$(\tilde{A}u)_{ij} = \frac{1}{2}(u_i + u_j) = \lambda_{ij} - \frac{1}{2}(\tilde{C}\lambda)_{ij} = \lambda_{ij},$$

and so $\lambda \in \text{im } \tilde{A}$. This shows that $\ker \tilde{C} \subseteq \text{im } \tilde{A}$ and we conclude that $\text{im } \tilde{A} = \ker \tilde{C}$, as desired. \square

B TRIANGLE FIELD EXTENSIONS

To extend geometric variational problems to the space of triangle fields we need to reformulate functions that are defined in terms of vertex positions in terms of the halfedge vectors. More precisely, we say that the triangle field function $\tilde{\mathcal{F}}(\omega)$ is an extension of a function $\mathcal{F}(f)$ that depends on vertex positions f if the relationship

$$\tilde{\mathcal{F}}(df) = \mathcal{F}(f)$$

is satisfied for all discrete surfaces $f : V \rightarrow \mathbb{R}^3$. Since df is translationally invariant, an extension does not exist unless $\mathcal{F}(f)$ is translationally invariant. In this section we obtain triangle field extensions via the Poisson reconstruction;

Recall that the Poisson reconstruction, $f(\omega)$, is the projection onto the space of integrable (exact) 1-forms, and it solves the Poisson equation

$$Lf = d * \omega.$$

By applying the function to this projection we get an extension:

$$\tilde{\mathcal{F}}(\omega) := \mathcal{F}(f(\omega)).$$

Now we compute the gradient in the space of triangle fields.

An infinitesimal deformation $\dot{\omega}$ induces a deformation \dot{f} of the Poisson reconstruction, and so the variation can be computed from the gradient of \mathcal{F} :

$$\langle \text{grad } \tilde{\mathcal{F}}, \dot{\omega} \rangle = \dot{\tilde{\mathcal{F}}} = \langle \text{grad } \mathcal{F}, \dot{f} \rangle.$$

Since Poisson equation is linear, \dot{f} is also the Poisson reconstruction of $\dot{\omega}$. This implies that a co-closed triangle field variation causes no change in the reconstruction $\dot{f} \equiv 0$.

By considering co-closed $\dot{\omega}$ in the equation above we find that $\text{grad } \tilde{\mathcal{F}}$ is orthogonal to all co-closed 1-forms. By the Hodge decomposition, the gradient is exact and we can find $\varphi : V \rightarrow \mathbb{R}^3$ satisfying $\text{grad } \tilde{\mathcal{F}} = d\varphi$. The gradient relationship, in terms of φ , becomes

$$\begin{aligned} \langle \text{grad } \tilde{\mathcal{F}}, \dot{\omega} \rangle &= \langle d\varphi, \dot{\omega} \rangle = \langle d * \dot{\omega}, \varphi \rangle \\ &= \langle L\dot{f}, \varphi \rangle = \langle L\varphi, \dot{f} \rangle = \langle \text{grad } \mathcal{F}, \dot{f} \rangle, \end{aligned}$$

where we used the self-adjointness of L and the infinitesimal Poisson reconstruction $L\dot{f} = d * \dot{\omega}$. Since \dot{f} is arbitrary, we conclude that $L\varphi = \text{grad } \mathcal{F}$. Summarizing, the gradient of the extension can be computed by solving a Poisson equation with the gradient 2-form of $\mathcal{F}(f)$ on the right-hand side:

$$\begin{cases} \text{grad } \tilde{\mathcal{F}} = d\varphi, & \varphi \in \Omega^0(M; \mathbb{R}^3) \\ L\varphi = \text{grad } \mathcal{F}. \end{cases} \quad (22)$$

Recall that a Poisson equation has a solution if and only if the right hand side has zero mean. The gradient 2-form $\text{grad } \mathcal{F}$ integrates to zero precisely when \mathcal{F} is translationally invariant, and so Equation 22 can always be solved in our setup. We now give some examples of triangle field extensions.

B.1 Point constraints

Since an integrable triangle field only describes a surface up to translation, it does not make sense to talk about the vertex positions of the triangle field $df \in \mathcal{M}$. Nevertheless, it is still meaningful to consider the translationally invariant relative point positions,

$$\mathcal{P}(f) := f_u - f_v,$$

between distinct vertices $u, v \in V$. It is straightforward to check that

$$\text{grad } \mathcal{P} = \delta_u - \delta_v,$$

where δ_v is the unit impulse located at the vertex $v \in V$. We use Equation 22 to find that $\text{grad } \tilde{\mathcal{P}}$ is the harmonic 1-form with source and sink at u and v , respectively—that is, $\text{grad } \tilde{\mathcal{P}} = dp$ where $p :$

$V \rightarrow \mathbb{R}$ is the harmonic function with source and sink at u and v , respectively

$$\mathbb{L}p = \delta_u - \delta_v.$$

This shows that $\text{grad } \tilde{\mathcal{P}}(\omega) = dp$ does not depend on ω , and so $\tilde{\mathcal{P}}$ is linear. Explicitly, we have

$$\tilde{\mathcal{P}}(\omega) = \langle \omega, dp \rangle.$$

Notice the similarity between these point constraints and the integrability condition in Equation 13: by considering the inner product with a harmonic 1-form dual to a homology generator γ we can prescribe the periods $\int_\gamma \omega$. On the other hand, considering the inner product with the harmonic 1-form with sources and sinks we can prescribe the difference of point positions, which can also be expressed as the integral $f_u - f_v = \int_{\gamma_v^u} \omega$, where γ_v^u is a curve from v to u . In both cases, inner products with harmonic 1-forms encode the integrals of the triangle field over curves in the mesh.

B.2 Enclosed volume

The enclosed volume of a surface is the other notable constraint that we extend to triangle fields in this way:

$$\mathcal{V}(f) := \frac{1}{6} \sum_{ijk} f_i \cdot (f_j \times f_k).$$

Recall that the gradient 2-form is the discrete area vector:

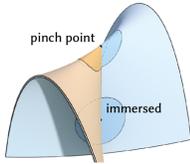
$$(\text{grad } \mathcal{V})_i = \frac{1}{3} \sum_{ijk} A_{ijk} N_{ijk}. \quad (23)$$

The gradient of our extension $\tilde{\mathcal{V}}$ is then $d\vartheta \in \Omega^1(M; \mathbb{R}^3)$ where ϑ is the solution of the Poisson equation

$$d * d\vartheta = \text{grad } \mathcal{V}.$$

C STARTING FROM RANDOM TRIANGLE FIELDS

The integrability conditions of a triangle field ensure that all faces can be glued together in a consistent way, but they do not ensure that the integrated surface comes from an immersion since there may be pinch points where the vertex star is not embedded (inset). Practically speaking, pinch points very rarely develop when we initialize the optimization with $\omega_0 = df$, where df is the differential of a smooth surface. On the other hand, when no initial immersion is available (e.g. in Figure 16 and Figure 1) energy minimization of an arbitrary triangle field ω_0 tends to develop many pinch points. In [Chern et al. 2018] the authors introduce discrete spin structures and show that by parameterizing rotations with spinors one can eliminate pinch points through energy minimization. Similarly, by parameterizing triangle fields with spinors we are able to find immersed minimizers of a much larger class of challenging variational problems—for example, computing constrained Willmore surfaces when only the discrete conformal class is given.



C.1 Discrete Spinors

We quickly review the necessary concepts from [Chern et al. 2018], relegating the technical details to their paper. Given a background metric ℓ , a *discrete spinor field* is a function $\psi : F \rightarrow S^3 \subset \mathbb{H}$ that we

interpret as defining a triangle field by $\omega_\psi = \bar{\psi} dz \psi$ inside each face, where z_{ijk} are complex local coordinates. A *discrete spin connection* is then a choice of square-root of parallel transport $\tau_{ij} \in \mathbb{C}$ across the edges $ij \in E$. The remarkable feature of this framework is that pinch points are encoded in the smoothness of ψ : an integrable triangle field $\omega_\psi = df$ has no pinch points if $|\psi_{jil} - \tau_{ij}\psi_{ijk}| < \sqrt{2}$ for all edges $ij \in E$. Consequently, the connection Dirichlet energy $\mathcal{E}_\tau(\psi)$ provides a bound on the number of pinch points.

C.2 Spinorial Triangle Fields

General triangle fields can be parameterized by a triangle field metric $\ell : \tilde{E} \rightarrow \mathbb{R}$ and face spinors $\psi : F \rightarrow S^3$:

$$\begin{aligned} \mathcal{S} : \mathbb{R}^{\tilde{E}} \times \mathbb{H}^F &\rightarrow \Omega_F^1(M; \mathbb{R}^3) \\ (\ell, \psi) &\mapsto (\bar{\psi} dz \psi)_{ijk \in F} \end{aligned}$$

where (z_{ijk}) are isometric local coordinates in each face with respect to the metric ℓ .

Conformal Triangle Fields. By fixing a reference metric ℓ_0 we can specialize the parameterization above to the space of conformal triangle fields using vertex scale factors $u : V \rightarrow \mathbb{R}$:

$$\begin{aligned} \mathcal{S}_c : \mathbb{R}^V \times \mathbb{H}^F &\rightarrow \Omega_F^1(M; \mathbb{R}^3) \\ (u, \psi) &\mapsto \mathcal{S}(e^u \ell_0, \psi). \end{aligned}$$

Since this parameterization only produces conformally equivalent triangle fields the discrete conformal constraint can safely be dropped from the optimization when using these coordinates.

C.3 Immersive Regularization

Finally, we get to the main point: when parameterizing triangle fields with spinors we can regularize our variational problem by adding a small multiple of the spin Dirichlet energy

$$\mathcal{E}_\tau(\psi) := \frac{1}{2} \sum_{ij \in E} w_{ij} |\psi_{jil} - \tau_{ij}\psi_{ijk}|^2.$$

This simple regularization almost completely eliminates the emergence of any pinch points, and constrains the search space to immersed surfaces.

D IMPLEMENTATION

In this section we give the details needed to implement the flow from Section 4. We give an overview of the algorithm before describing explicit matrix expressions:

The conformal constraint is computed through the logarithmic edge lengths, $\lambda_{ij} = |f_j - f_i|$, and the cross-ratio map $C \in \mathbb{R}^{E \times E}$ (Equation 3), which can be expressed as a sum of sparse matrices $C_{ijk} \in \mathbb{R}^{E \times E}$ defined per face, with non-zero elements as follows:

$$C_{ijk} = \begin{pmatrix} ij & jk & ki \\ 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \begin{matrix} ij \\ jk \\ ki \end{matrix}. \quad (24)$$

ALGORITHM 1: Constrained Willmore Evolution

Input: Triangle mesh $M = (V, E, F)$ with vertex positions $f : V \rightarrow \mathbb{R}^3$, prescribed logarithmic length cross ratios c_0 , and step size $t > 0$.

- 1: Build cross-ratio map $C \in \mathbb{R}^{E \times E}$ ▷Eqn. 24
- 2: Initialize Lagrange multipliers $\mu = 0 \in \mathbb{R}^E$
- 3: **do**
- 4: Compute log edge lengths λ Section 3
- 5: Build cotan Laplacian L and h ▷Eqn. 27
- 6: Build w ▷Eqn. 29
- 7: $g \leftarrow Lh + w$ ▷Willmore gradient
- 8: $c \leftarrow C\lambda - c_0$ ▷constraint violation
- 9: $X \leftarrow CD\lambda$ ▷constraint derivative
- 10: Build metrics I_{H^2} and I_Q ▷Eqns. 25 and 26
- 11: $(\Delta f, \Delta \mu) \leftarrow \text{solve} \begin{pmatrix} \frac{1}{t} I_{H^2} & X^T \\ X & -\frac{1}{t} I_Q \end{pmatrix} \begin{pmatrix} \Delta f \\ \Delta \mu \end{pmatrix} = \begin{pmatrix} -g - X^T \mu \\ -c \end{pmatrix}$
- 12: $f \leftarrow f + \Delta f$ ▷take forward step
- 13: $\mu \leftarrow \mu + \Delta \mu$
- 14: **while** not satisfied

Its derivative is computed through the derivative of the logarithmic edge lengths $D\lambda \in \mathbb{R}^{E \times 3V}$, which has nonzero entries

$$(D\lambda)_{ij,j} = -(D\lambda)_{ij,i} = f_j - f_i.$$

The inner product on the space of discrete quadratic differentials is given by the diagonal matrix $I_Q \in \mathbb{R}^{E \times E}$ with nonzero entries

$$(I_Q)_{ij,ij} = \frac{1}{w_{ij} \ell_{ij}^2}, \quad (25)$$

for each edge $ij \in E$, where w_{ij} are the usual cotan edge weights. The Sobolev H^2 -metric is discretized as

$$I_{H^2} := *_0 + L *_0^{-1} L, \quad (26)$$

and where $*_0 \in \mathbb{R}^{V \times V}$ is the diagonal matrix of vertex dual areas and $L \in \mathbb{R}^{V \times V}$ is the cotan Laplacian.

Willmore energy. To define the Willmore energy, we use the discrete mean curvature vector $h := \frac{1}{2} *_0^{-1} Lf$, which at a vertex $i \in V$ can be expressed as the sum over the neighboring edges

$$h_i = \frac{1}{2A_i} \sum_j w_{ij} df_{ij}. \quad (27)$$

The Willmore energy is

$$\mathcal{W}(f) = \langle *_0 h, h \rangle = \frac{1}{4} \text{tr}(f^T L *_0^{-1} Lf), \quad (28)$$

and its gradient 2-form is

$$\text{grad } \mathcal{W} = Lh + w.$$

Here w is the vector with \mathbb{R}^3 blocks

$$w_i = \sum_j \left(\frac{\partial w_{ij}}{\partial f_i} \langle h_j - h_i, df_{ij} \rangle - \frac{\partial A_j}{\partial f_i} |h_j|^2 \right), \quad (29)$$

where the gradient of the cotan weights and vertex areas are

$$\begin{aligned} \frac{\partial w_{ij}}{\partial f_i} &= \frac{1}{\sin^2 \theta_k^{ij}} N_{ijk} \times \frac{df_{ki}}{|df_{ki}|} - \frac{1}{\sin^2 \theta_l^{ji}} N_{jil} \times \frac{df_{il}}{|df_{il}|}, \\ \frac{\partial A_j}{\partial f_i} &= \frac{1}{2} (N_{ijk} \times df_{jk} + N_{jil} \times df_{lj}). \end{aligned}$$

Constraints. Additional constraints, $\xi(f) = 0$, can be incorporated by joining them to conformal constraint vector and differential. This corresponds to making the changes

$$c \leftarrow \begin{pmatrix} C\lambda - c_0 \\ \xi(f) \end{pmatrix}, \quad X \leftarrow \begin{pmatrix} CD\lambda \\ D\xi \end{pmatrix}$$

in the algorithmic description above.

Boundary conditions. For surfaces with boundary, discrete conformal equivalence is characterized by the logarithmic length-cross ratios of all the *interior* edges, and so we use the modified cross-ratio map $C_0 \in \mathbb{R}^{E_0 \times E}$ which is obtained by removing all of the rows from C corresponding to boundary edges.

Dirichlet boundary conditions can be incorporated by restricting the updates to the interior vertices. Practically, that means that we remove all of the rows and columns of the saddle point system that correspond to the boundary vertices.

First-order boundary conditions can be treated as additional constraints in the optimization. For example, one can constrain the normals of the boundary faces

$$N_{F_\partial}(f) = (N_{ijk})_{ijk \in F_\partial},$$

where $F_\partial \subset F$ is the set of boundary faces. The derivative of this constraint mapping $DN_{F_\partial} \in \mathbb{R}^{3F_\partial \times 3V}$ is the sparse matrix with $\mathbb{R}^{3 \times 3}$ -blocks

$$(DN_{F_\partial})_{ijk,i} = -\frac{1}{A_{ijk}} (N_{ijk} \times df_{jk}) N_{ijk}^T.$$

We remark that the simpler strategy of fixing all of the vertices along the boundary strip may not work in the presence of a conformality constraint—in particular, it may not be possible to achieve a prescribed conformal class when the vertices along the boundary strip are fixed.

D.1 Triangle Fields

There are few conceptual differences when working with triangle fields, but the objects and notations differ in some important ways. We now describe the implementation at the level of triangle fields:

The cross-ratio matrix $\tilde{C} \in \mathbb{R}^{H \times H}$ is defined per-face

$$\tilde{C}_{ijk} = \begin{pmatrix} ij & jk & ki \\ -1 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \begin{matrix} ij \\ jk \\ ki \end{matrix}. \quad (30)$$

The local integrability conditions are described via the matrix $d_1 \in \mathbb{R}^{E \times H}$ with

$$(d_1)_{ij_E, ij} = \pm 1, \quad (31)$$

where ij_E denotes the unoriented edge, and the sign is chosen according to the orientation of the halfedge. The Willmore energy is extended to triangle fields as the Dirichlet energy:

$$\tilde{\mathcal{W}}(\omega) := \frac{1}{4} \langle d * \omega, d * \omega \rangle = \frac{1}{4} \sum_{i \in V} \frac{1}{A_i} \left| \sum_{ij} w_{ij} \omega_{ij} \right|^2.$$

Notice that $\tilde{\mathcal{W}}(df) = \mathcal{W}(f)$. Since we reduce the order of integrability when working with ω the Sobolev metric we use is the

ALGORITHM 2: Constrained Willmore Evolution (Triangle Fields)

Input: Triangle mesh $M = (V, E, F)$ with initial triangle field ω , prescribed logarithmic length cross ratios c_0 , and step size $t > 0$.

- 1: Build cross-ratio map $\tilde{C} \in \mathbb{R}^{H \times H}$ ▷Eqn. 30
- 2: Build a basis of $\{Y_i\}_{i=1}^{2g}$ of harmonic dual 1-forms
- 3: Build $d_1 \in \mathbb{R}^{E \times H}$ ▷Eqn. 31
- 4: Initialize Lagrange multipliers $\mu = 0 \in \mathbb{R}^{|H|+3(|E|+2g)}$
- 5: **do**
- 6: Compute log edge lengths λ
- 7: $g \leftarrow \text{grad } \tilde{W}$ ▷Willmore gradient
- 8: $c \leftarrow \begin{pmatrix} \tilde{C}\lambda - c_0 \\ d_1\omega \\ \langle\langle Y_i | \omega \rangle\rangle \end{pmatrix}$ ▷constraint violation
- 9: $X \leftarrow \begin{pmatrix} \tilde{C}D\lambda \\ d_1 \\ Y_i \end{pmatrix}$ ▷constraint derivative
- 10: Build metrics I_{H^1} and I_Q ▷Eqns. 25 and 32
- 11: $(\Delta\omega, \Delta\mu) \leftarrow \text{solve} \begin{pmatrix} \frac{1}{t}I_{H^1} & X^\top \\ X & -\frac{1}{t}I_Q \end{pmatrix} \begin{pmatrix} \Delta\omega \\ \Delta\mu \end{pmatrix} = \begin{pmatrix} -g - X^\top\mu \\ -c \end{pmatrix}$
- 12: $\omega \leftarrow \omega + \Delta\omega$ ▷take forward step
- 13: $\mu \leftarrow \mu + \Delta\mu$
- 14: **while** not satisfied

 H^1 -metric defined by

$$I_{H^1} := *_1 + *_1 d_0 *_0^{-1} d_0^\top *_1 \quad (32)$$

Received January 2021