Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Discrete-to-Continuum Metrics from Scalar Fields

Lars Rönnbäck

lars@uptochange.com

Stockholm University https://orcid.org/0000-0002-8810-7587

Research Article

Keywords: seam-driven geometry, discrete conformal metric, Gromov-Hausdorff convergence, cotangent Laplacian, inverse metric design

Posted Date: March 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9017238/v1

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

Additional Declarations: The authors declare no competing interests.

Highlights

Discrete-to-Continuum Metrics from Scalar Fields

Lars Rönnbäck

- A scalar-field (vertex-potential) rule reweights edge lengths locally and induces an intrinsic shortest-path metric.
- Quantitative $O(h)$ discrete-to-continuum convergence of the induced graph metric under standard distance-approximation assumptions.
- Inverse edge-weight fitting becomes a strictly convex QP in positive variables $X = e^s$; first-order curvature sensitivity is governed by the cotangent Laplacian at $s = 0$.

Discrete-to-Continuum Metrics from Scalar Fields

Lars Rönnbäck[a,*]

[a]Department of Computer and Systems Sciences, Stockholm University, Stockholm, Sweden

ARTICLE INFO

Keywords:
seam-driven geometry
discrete conformal metric
Gromov–Hausdorff convergence
cotangent Laplacian
inverse metric design

ABSTRACT

We present a scalar-first framework for geometry processing in which geometric structure is generated from a scalar field $s : U \to \mathbb{R}$ through an explicit local rule $\mathcal{R}$. We call this scalar field a *seam* (a potential/log-scale field; unrelated to UV/cut seams in parameterization). In the discrete setting, we study a conformal graph rule that assigns edge lengths using an endpoint quadrature of e^s, yielding an intrinsic shortest-path metric that is easy to optimize and differentiate. Our main results are (i) a quantitative discrete-to-continuum guarantee for the induced shortest-path metrics on graph families whose shortest-path distances approximate background geodesic distance up to an additive $O(h)$ error (e.g., sufficiently dense neighborhood graphs); (ii) a curvature sensitivity identity showing that first-order curvature variations are governed by the cotangent Laplacian; and (iii) a strictly convex inverse-design formulation for fitting target edge weights via a quadratic program in positive variables $X_u = e^{s(u)}$. These results position scalar-field parameterizations as a stable interface between differential-geometric objectives and practical optimization pipelines in mesh and graph processing.

1. Introduction

Discrete conformal geometry and mesh parameterization often represent intrinsic geometry through a *scalar* unknown (a vertex-based log-scale / conformal factor), which is then converted into edge lengths, angles, and curvature by an explicit local construction [13, 12, 24]. We adopt the same scalar-first ethos, but with a different computational emphasis: rather than building geometry primarily through angle-based formulations, we study *shortest-path* metric spaces induced by seam-reweighted edge lengths on graphs and meshes. In our terminology, the scalar field $s : U \to \mathbb{R}$ is a *seam* (a scalar potential / log-scale field), and a *Rule* $\mathcal{R}$ is the local map that turns (U, s) into geometric data.

This viewpoint unifies a range of classical constructions (e.g., conformal scaling, Hessian metrics, and optimal-transport potentials), but its main utility here is algorithmic. On discrete meshes, manipulating full tensor fields is cumbersome; by parameterizing geometry through a scalar seam, metric generation and inverse design reduce to sparse linear algebra and scalable optimization.

In this paper, we define the Seam-Rule framework (Section 2; a compact organizational language for seam-driven constructions) and catalog standard continuous and discrete Rules (Section 3). Section 4 illustrates how familiar continuous identities (e.g., Gauss–Bonnet) are cleanly expressed in the seam-rule viewpoint. Finally, Section 5 provides the computational core: proving that seam-generated shortest-path metrics on suitable graph families approximate smooth conformal metrics, identifying a cotangent-Laplacian curvature sensitivity, and enabling strictly convex inverse metric design.

Contributions (geometry processing view). Our contributions are geared toward mesh and graph processing, where one seeks parameterizations that are easy to optimize, differentiate, and analyze.

1. **A scalar-first metric parameterization on graphs/meshes.** We study a conformal graph rule that assigns edge lengths using an endpoint quadrature of e^s, producing a shortest-path metric with local control by a scalar seam.
2. **Quantitative discrete-to-continuum guarantees.** Under an asymptotically geodesic spanner hypothesis for the background shortest-path metric (as is standard for sufficiently dense neighborhood graphs), we prove an $O(h)$ Gromov–Hausdorff convergence rate between the seam-generated discrete metric and the smooth conformal metric (Theorem 3) and derive a corresponding uniform $O(h)$ metric error bound (Theorem 4).

*Corresponding author
lars@uptochange.com (L. Rönnbäck)
ORCID(s): 0000-0002-8810-7587 (L. Rönnbäck)

3. **Curvature sensitivity via cotangent weights.** We show that the *first-order* (Jacobian) sensitivity of angle-defect curvature with respect to the seam exactly matches the cotangent Laplacian (with the standard scaling used in discrete differential geometry) at $s = 0$ (Theorem 7), providing a differentiable link between scalar parameters and discrete curvature.
4. **Convex inverse metric design.** We reduce edge-weight fitting under the conformal graph rule to a strictly convex quadratic program in variables $X_u = e^{s(u)}$ on non-bipartite graphs (Theorem 5), leveraging a classical signless-Laplacian positive-definiteness criterion [9, 10]. We further provide gauge-fixing and conditioning guarantees (Theorem 6), and show that adding a graph-Laplacian smoothness prior preserves strict convexity while dramatically improving noise robustness (Corollary 2; on the Stanford Bunny, 7.3× error reduction at $\sigma = 0.10$).
5. **Genus-0 universality up to reparameterization.** We show that on the sphere (genus 0), seam-driven conformal scaling becomes *universal up to diffeomorphism*: any smooth metric can be written as a pullback of a fixed background metric followed by a conformal seam rule (Proposition 1).

1.1. Related Work

Scalar fields on meshes are a classical and widely used representation in geometry processing, in particular as vertex-based conformal factors and potentials. Discrete uniformization and discrete conformal parameterizations seek a scalar conformal scaling that induces a metric of prescribed curvature or that supports flattening/parameterization; see, e.g., Gu–Luo–Sun–Wu [12] and the variational formulations of Springborn–Schröder–Pinkall [13]. In this literature, a scalar field (often a log-scale factor) is the primary unknown, and the induced metric is obtained by a local rule that rescales discrete lengths/angles.

Our discrete conformal graph rule is closely aligned in spirit—a scalar log-scale drives intrinsic geometry—but differs in the induced object: we assign seam-dependent *edge lengths* (via an endpoint quadrature) and then take the induced *shortest-path* distance, whereas much of the classical discrete conformal literature emphasizes angle structures and curvature energies.

Circle packing provides another scalar-first route to discrete conformal geometry, in which circle radii (or log-radii) act as scalar variables whose induced intersection patterns and edge lengths encode a discrete conformal structure; see Stephenson [17] for an overview and references. Closely related, combinatorial Ricci flow evolves a vertex-based scalar parameterization (e.g., circle radii or log-scale factors) to drive discrete curvature toward a target, with Chow–Luo [16] as a foundational reference.

More broadly, discrete-to-continuum limits are a recurring theme in discrete differential geometry: as triangulations refine, one seeks conditions under which discrete conformal structures and their induced metrics converge (in appropriate senses) to their smooth counterparts; see, e.g., the discrete conformal/uniformization viewpoints in [11, 12] and references therein.

Our intent is not to claim novelty of using scalar fields per se, but to formalize a general "seam → rule → geometry" interface that accommodates multiple continuous and discrete constructions under one axiomatic umbrella, and to highlight algorithmic consequences of a particular discrete edge-length rule: a quantitative discrete-to-continuum guarantee for the resulting shortest-path metric, a curvature sensitivity identity, and a convex inverse-design formulation for fitting target edge weights.

Figure 1 provides a representative visualization of a seam field on a mesh and the intrinsic curvature induced by the seam-generated edge lengths.

2. Framework: Seams, Rules, and Composition

Let U be a paracompact Hausdorff space (or discrete set) equipped with a local background structure $\mathcal{T}_U$ (e.g., a differentiable atlas, a background metric h, or a graph adjacency). A *seam* is a scalar field $s : U \to \mathbb{R}$ in an admissible class $\mathcal{S}(U)$ (e.g., C^∞, Morse, convex, or discrete vertex array); it plays the role of a potential or (log) conformal factor and serves as the single generative control variable for local geometry. A *Rule* $\mathcal{R}$ is a local, natural assignment that sends admissible seam data $(U, s, \mathcal{T}_U)$ to a geometric output (e.g., a metric tensor, a pseudo-metric, or a weighted graph). We require three standard sheaf-theoretic properties—*locality* (output on $V \subseteq U$ depends only on $s|_V$ and $\mathcal{T}_U|_V$), *gluing* (compatible local outputs extend uniquely to a global one), and *functoriality* ($\mathcal{R}(U, \varphi^* s') = \varphi^* \mathcal{R}(U', s')$ for structure-preserving maps φ)—all of which are automatically satisfied by the concrete

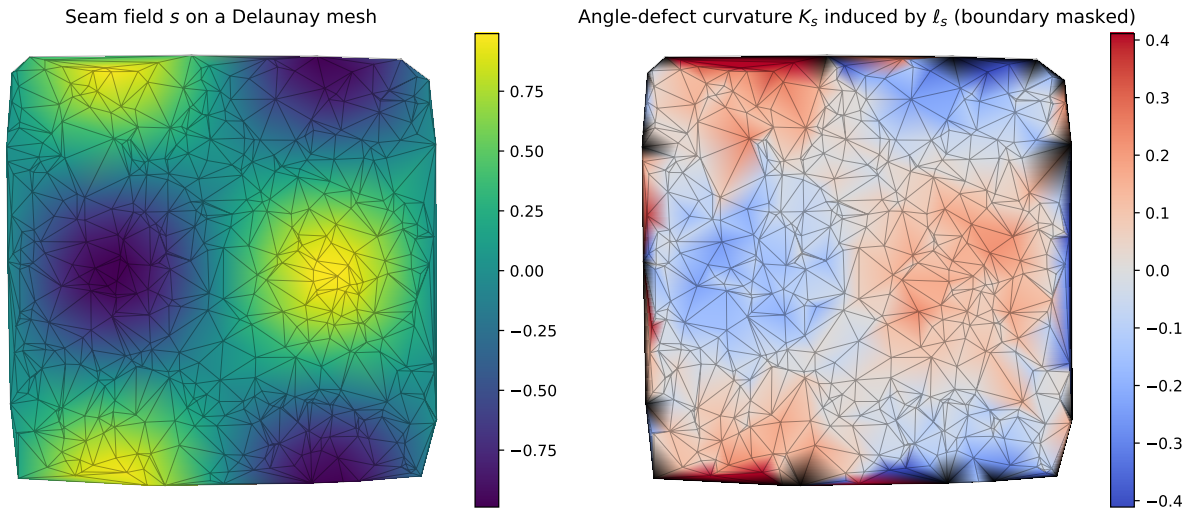


Figure 1: A small Delaunay mesh colored by the seam field s (left) and the induced discrete Gaussian curvature K_s (angle defect) computed from the seam-generated intrinsic edge lengths $\ell_s(u,v) = \ell_0(u,v)\frac{e^{s(u)}+e^{s(v)}}{2}$ (right; boundary vertices masked to avoid boundary angle-defect dominance).

rules we use. Rules compose: if $\mathcal{R}_1$ produces an intermediate background and $\mathcal{R}_2$ consumes it, the composed map $\mathcal{R}_{2\circ1}(U,(s_1,s_2)) := \mathcal{R}_2(U,s_2;\mathcal{R}_1(U,s_1))$ inherits all three properties (the verification is routine from the definitions).

For concreteness, keep in mind the conformal prototype throughout: the smooth rule $\mathcal{R}_{\text{Conf}}$ maps a seam s to the metric tensor $g = e^{2s}h$ on a background surface (U,h), while our discrete conformal graph rule maps s to seam-dependent edge lengths and hence a shortest-path distance.

3. The Repertoire of Rules

The generative power of the framework relies on specific definitions of $\mathcal{R}$.

3.1. Continuous Rules

For a smooth manifold U, we highlight three primary rules generating metric tensors g:

1. **The Hessian Rule ($\mathcal{R}_{\text{Hessian}}$):** $g = \nabla^2 s$ for a chosen torsion-free connection ∇ (in local affine coordinates: $g_{ij} = \partial_i\partial_j s$). Where g is positive definite (e.g., for strictly convex s), it defines a Riemannian metric locally (as proven later in Theorem 2). In general, $\nabla^2 s$ may be indefinite; allowing such outputs naturally leads to pseudo-Riemannian geometries, which we do not develop here. Hessian geometry is central to Information Geometry [1, 2].
2. **The Conformal Rule ($\mathcal{R}_{\text{Conf}}$):** $g_{ij} = e^{2s}h_{ij}$. Generates geometries conformally equivalent to a background metric h [3].
3. **The Gradient Rule ($\mathcal{R}_{\text{Grad}}$):** $g_{ij} = |\nabla s|_h^2 h_{ij}$. A subset of conformal geometries controlled by the eikonal magnitude of the seam [6].

As quick orientation, the Euclidean metric is recovered via $\mathcal{R}_{\text{Hessian}}$ with seam $s = \frac{1}{2}\sum(x^i)^2$, the Poincaré half-plane via $\mathcal{R}_{\text{Conf}}$ with $s = \ln R - \ln y$ on the flat background, and the stereographic sphere via $\mathcal{R}_{\text{Conf}}$ with $s = \ln(2R^2) - \ln(R^2 + r^2)$.

3.2. Discrete Rules ($\mathcal{R}_{\text{Graph}}$)

On a graph $G = (V, E)$ with background edge lengths $\ell_0(u,v)$, the seam $s : V \to \mathbb{R}$ generates a discrete shortest-path metric D via edge weights $w(u,v)$. Two vital implementations are:

1. **Conformal Graph Rule ($\mathcal{R}_{\text{Graph-Exp }s}$):** $w(u,v) = \ell_0(u,v)\frac{e^{s(u)}+e^{s(v)}}{2}$

2. **Gradient Graph Rule ($\mathcal{R}_{\text{Graph-}|\nabla s|}$):** $w(u,v) = \ell_0(u,v)\frac{|\nabla s|(u)+|\nabla s|(v)}{2}$

Remark 1 (Arithmetic mean vs. the classical geometric mean). In much of the discrete conformal literature on triangle meshes, a common vertex-based conformal scaling takes the *geometric mean* of endpoint factors, i.e.

$$\ell_s^{\text{geom}}(u,v) = \ell_0(u,v)\,\exp\Big(\tfrac{1}{2}\big(s(u)+s(v)\big)\Big) = \ell_0(u,v)\,\sqrt{e^{s(u)}e^{s(v)}}.$$

By contrast, our Conformal Graph Rule uses the *arithmetic mean*

$$\ell_s^{\text{arith}}(u,v) = \ell_0(u,v)\,\tfrac{1}{2}\big(e^{s(u)}+e^{s(v)}\big).$$

We deviate from the geometric-mean convention intentionally. The arithmetic mean is exactly the endpoint trapezoidal quadrature for the conformal line element $\int e^s \, \mathrm{d}\ell_0$ (Lemma 2), and it is the choice that makes inverse edge-weight fitting *quadratic* in the variables $X_u := e^{s(u)}$ (Theorem 5), yielding a strictly convex QP on non-bipartite graphs. In other words, the arithmetic rule is selected for its discrete-to-continuum consistency and its algorithmic convexity properties, not because it is the only possible discrete conformal scaling.

The two means also behave differently when adjacent seam values differ greatly. Writing $a = e^{s(u)}$ and $b = e^{s(v)}$, we have

$$\tfrac{1}{2}(a+b) = \sqrt{ab}\,\cosh\Big(\tfrac{1}{2}(s(u)-s(v))\Big),$$

so the arithmetic rule equals the geometric rule times a factor $\cosh(\frac{1}{2}\Delta s)$ that grows like $\frac{1}{2}e^{|\Delta s|/2}$. Thus for large seam gradients the arithmetic mean is dominated by the larger endpoint and can amplify edge-length contrast more aggressively than the geometric mean. Practically, this is precisely the regime where one typically uses damping/line-search, imposes smoothness priors, or enforces bounds on neighbor differences $|s(u)-s(v)|$ to avoid highly skewed triangles or early degeneracies.

These constructions are closely related to discrete conformal geometry and conformal parameterizations of triangle meshes; see, e.g., [13, 12]. They also connect to circle-packing/combinatorial Ricci flow viewpoints on discrete conformal scaling [16] and to related conformal-type transformations such as spin transformations [14].

4. Continuous Geometry: Curvature Identities

Framing continuous geometry in terms of seams often reduces complex tensor algebra to elegant scalar identities.

Theorem 1 (Gauss–Bonnet via a seam). *Let (M,g) be a closed oriented Riemannian surface. By the uniformization theorem, there exists a metric h of constant Gaussian curvature and a seam s such that $g = \mathcal{R}_{Conf}(s;h) = e^{2s}h$. Then $\int_M K_g \, dA_g = 2\pi\chi(M)$.*

Proof. Under the conformal rule, curvature transforms as $K_g = e^{-2s}(K_h - \Delta_h s)$, and the area form as $dA_g = e^{2s}\,dA_h$. Multiplying these yields $K_g\,dA_g = (K_h - \Delta_h s)\,dA_h$. Integrating over M, the Laplacian term $\int_M \Delta_h s\,dA_h$ vanishes identically by the divergence theorem. The seam's exact contribution perfectly cancels out globally, leaving $\int K_g\,dA_g = \int K_h\,dA_h = 2\pi\chi(M)$. □

Theorem 2 (Local non-degeneracy of the Hessian rule). *Let M be a smooth manifold equipped with a torsion-free connection ∇. Let $s \in C^\infty(M)$ and define the symmetric $(0,2)$-tensor $g := \nabla^2 s$ (the covariant Hessian). If $p \in M$ is a point where g_p is positive definite, then g defines a Riemannian metric on some neighborhood of p.*

Proof. Positive definiteness is an open condition: since $q \mapsto g_q$ varies smoothly, there exists a neighborhood $U \ni p$ such that g_q remains positive definite for all $q \in U$. On U, g is therefore a smooth Riemannian metric. □

Universality: solved on genus 0, open beyond. Beyond curvature identities, it is natural to ask whether seam-driven rules can represent *arbitrary* Riemannian metrics. For general manifolds this interacts with global degrees of freedom (e.g., moduli of conformal structures in dimension 2), but on genus 0 one can already obtain a clean universality theorem.

Proposition 1 (Genus-0 universality up to reparameterization). *Let M be a smooth closed surface of genus 0 and fix a smooth background metric h_0 on M. Then for every smooth Riemannian metric g on M there exist a diffeomorphism $\varphi : M \to M$ and a seam $s \in C^\infty(M)$ such that*

$$g = e^{2s}\, \varphi^* h_0.$$

Equivalently, on genus 0 surfaces the seam-driven conformal rule is universal up to pullback by diffeomorphisms.

PROOF. By uniformization, every smooth metric g on a closed surface is conformal to a constant-curvature metric; in particular there exists a metric $\widehat{h}$ on M and a seam $\widehat{s} \in C^\infty(M)$ such that $g = e^{2\widehat{s}}\widehat{h}$, where $\widehat{h}$ has constant Gaussian curvature (see, e.g., [3]).

On a genus 0 surface, the conformal structure is unique up to diffeomorphism (equivalently, every genus 0 Riemann surface is biholomorphic to the Riemann sphere; see, e.g., [4]). Thus, fixing the reference metric h_0, there exists a diffeomorphism $\varphi : M \to M$ and a constant $c \in \mathbb{R}$ such that $\widehat{h} = e^{2c}\, \varphi^* h_0$. Setting $s := \widehat{s} + c$ gives $g = e^{2s}\, \varphi^* h_0$. □

The content is a direct consequence of the uniformization theorem together with conformal rigidity in genus 0; we state it explicitly because it anchors the open problem for higher genus (below and in the conclusion).

Remark 2 (What remains open). Proposition 1 shows that on genus 0 surfaces, seams already capture *all* metrics once one allows the natural reparameterization operation $g \mapsto \varphi^* g$. For higher genus surfaces, conformal structures carry nontrivial moduli, so one must decide what additional admissible operations (additional rule types, restricted classes of diffeomorphisms, or auxiliary scalar fields) are needed to obtain universality. One natural augmentation is to separate *within-class* scaling from *between-class* variation: keep a seam s to represent the conformal factor inside a fixed conformal class, but also allow a small collection of global parameters ("moduli seams") that move the reference metric through conformal structures. Concretely, one could aim for a representation of the form

$$g = e^{2s}\, \varphi^* h(\tau),$$

where τ ranges over Teichmüller space (dimension 2 on genus 1 and $6g-6$ for genus $g \geq 2$) and $h(\tau)$ is a chosen smooth family of unit-area constant-curvature metrics realizing each conformal class. From an algorithmic/discrete viewpoint, τ could be parameterized by standard Teichmüller coordinates (e.g., Fenchel–Nielsen length–twist data [5]) or, more linearly, by periods of a basis of harmonic 1-forms (a "period matrix" description [4]), while s continues to provide the local scalar degree of freedom. We record this as an open problem in the conclusion (Open Problem 6).

5. Discrete Geometry and Computation

The true advantage of the seam framework is algorithmic. In this climax section, we prove that discrete graph rules strictly approximate continuous geometries, preserve curvature, and enable uniquely solvable inverse-design problems.

5.1. Quantitative Discrete-to-Continuum Limits

Lemma 1 (Vertex-path approximation of geodesics). *Let (M, g_0) be compact and let $\{G_n = (V_n, E_n)\}$ be a sequence of weighted graphs with vertex sets $V_n \subset M$ and background edge lengths ℓ_0 induced by g_0. Assume V_n is an h_n-net in (M, d_{g_0}) with $h_n \to 0$, and assume the edges are uniformly local in the sense that $\ell_0(u, v) = O(h_n)$ for all $(u, v) \in E_n$. Fix points $x, y \in M$ and choose vertices $u_n, v_n \in V_n$ with $d_{g_0}(u_n, x) \leq h_n$ and $d_{g_0}(v_n, y) \leq h_n$. Assume in addition that the 1-skeleton is an* asymptotically geodesic spanner *for (M, g_0) in the sense that there exists a constant C_{span} (independent of n) such that for all sufficiently large n and all vertices $a, b \in V_n$ there exists a vertex path P from a to b with*

$$L_0(P) \leq (1 + C_{\mathrm{span}}\, h_n)\, d_{g_0}(a, b).$$

Then there exists a vertex path P_n in G_n from u_n to v_n such that its background length satisfies

$$L_0(P_n) \le d_{g_0}(x, y) + Ch_n,$$

where C depends only on (M, g_0) and the constants implicit in the assumptions above.

PROOF. By the triangle inequality,

$$d_{g_0}(u_n, v_n) \le d_{g_0}(u_n, x) + d_{g_0}(x, y) + d_{g_0}(y, v_n) \le d_{g_0}(x, y) + 2h_n.$$

By the spanner assumption applied to $(a, b) = (u_n, v_n)$, there exists a vertex path P_n from u_n to v_n such that

$$L_0(P_n) \le (1 + C_{\text{span}}\, h_n)\, d_{g_0}(u_n, v_n) \le (1 + C_{\text{span}}\, h_n)\, (d_{g_0}(x, y) + 2h_n).$$

Since M is compact, $d_{g_0}(x, y) \le \operatorname{diam}(M, g_0)$, so the right-hand side is bounded by $d_{g_0}(x, y) + Ch_n$ for a constant C independent of n. □

Remark 3 (On the spanner assumption). The spanner hypothesis in Lemma 1 is the precise technical condition needed to convert continuous g_0-geodesic distances into comparable *edge-path* distances on the graph with an $O(h_n)$ error. Sampling/coverage alone ensures V_n is dense in M, but without a spanner-type assumption one may only obtain a mesh-dependent *multiplicative* distortion between graph shortest-path length and d_{g_0}. It is important to distinguish two related but different notions.

(i) In the planar Euclidean setting, the Delaunay triangulation is a *geometric t-spanner* (for an explicit universal constant $t > 1$); see, e.g., [18]. This is a constant-factor guarantee.

(ii) The hypothesis used in Lemma 1 is stronger: it asks for an *asymptotically geodesic* spanner with stretch $1+O(h_n)$, which is the ingredient that yields an *additive* $O(h_n)$ distance error for fixed $x, y \in M$. A constant-factor spanner bound alone does *not* imply such an asymptotically tight estimate.

For smooth manifolds, $(1 + o(1))$-type geodesic-graph approximation results are typically proved for sufficiently dense *neighborhood graphs* (e.g., ε-graphs or kNN graphs with neighborhood size tending to zero but containing many sample points), rather than for bounded-degree triangulation 1-skeleta; see, e.g., [19] for a representative uniform convergence statement in the Isomap setting. In our proofs, the convergence rate is driven by whatever stretch-to-1 estimate the chosen graph family provides (i.e., the rate at which the graph shortest-path geometry approximates d_{g_0}). Lemma 2 then shows the trapezoidal conformal edge rule contributes only a higher-order correction along any fixed vertex path. Equivalently, one can view the final $O(h_n)$ bound as the sum of two conceptually separate ingredients: (a) a graph-family approximation statement for (M, d_{g_0}) (encoded here by the spanner hypothesis), and (b) a local quadrature consistency statement for the seam rule (Lemma 2).

For background on Delaunay triangulations and mesh generation (independent of spanner/stretch guarantees), see [20, 21].

Lemma 2 (Trapezoidal consistency of the conformal edge rule). *Let $s \in C^2(M)$ and let $\gamma : [0, L] \to M$ be a unit-speed (g_0) geodesic segment. For a subsegment of length $\ell \le h$ with endpoints $p = \gamma(t_0)$ and $q = \gamma(t_0 + \ell)$, define the trapezoidal approximation*

$$T(p, q) := \ell\, \frac{e^{s(p)} + e^{s(q)}}{2}.$$

Then the conformal length satisfies

$$\left| \int_{t_0}^{t_0+\ell} e^{s(\gamma(t))}\, \mathrm{d}t - T(p, q) \right| \le C_s\, \ell^3,$$

where C_s depends on $\sup_M |\nabla^2(e^s)|$ (equivalently on $\sup_M(|\nabla s|, |\nabla^2 s|)$).

PROOF. Since γ is unit-speed in the background metric g_0, the parameter t is exactly g_0-arc length along γ, i.e. $\mathrm{d}t = \mathrm{d}\ell_0$ on the segment. The claim is then the standard trapezoidal error estimate for C^2 functions: if $f(t) := e^{s(\gamma(t))}$, then $f \in C^2$ and the error on an interval of length ℓ is bounded by $\frac{\ell^3}{12} \sup |f''|$. Boundedness of f'' follows from $s \in C^2$ and compactness of M. □

Theorem 3 (Discrete-to-continuum limit via correspondences). *Let (M, g_0) be a compact Riemannian manifold and $s \in C^2(M)$. Let $\{G_n = (V_n, E_n)\}$ be a sequence of weighted graphs with $V_n \subset M$ and mesh size $h_n \to 0$ such that the assumptions of Lemma 1 hold (in particular, V_n is an h_n-net and $\ell_0(u,v) = O(h_n)$ on edges). Let the discrete conformal graph metric d_n be generated by $\mathcal{R}_{\text{Graph-Exp } s}$ evaluated on V_n. Assume the asymptotically geodesic spanner property stated in Lemma 1 holds for G_n with respect to d_{g_0}. Then the Gromov–Hausdorff distance between (V_n, d_n) and (M, d_g) satisfies $d_{GH}\big((V_n, d_n), (M, d_g)\big) = O(h_n)$ (see, e.g., [7]), where $g := e^{2s} g_0$ and d_g is the induced geodesic distance on M. The $O(h_n)$ rate is driven entirely by the spanner hypothesis; the conformal edge rule's own trapezoidal quadrature contributes only $O(h_n^2)$ (Lemma 2), so the seam rule is never the convergence bottleneck.*

Proof. Define a correspondence $\mathcal{C}_n \subset V_n \times M$ by pairing each vertex $u \in V_n$ with itself viewed as a point in M, i.e. $\mathcal{C}_n := \{(u,u) : u \in V_n\}$. Since V_n is an h_n-net in M (by assumption), every $x \in M$ lies within g_0-distance $\le h_n$ of some $u \in V_n$, hence within g-distance $\le e^{\|s\|_\infty} h_n$ of some vertex as well. Thus $\mathcal{C}_n$ is an $O(h_n)$-surjective correspondence.

To bound the distortion, fix $u, v \in V_n$. By Lemma 1 (with $x = u$ and $y = v$), there exists a vertex path in the 1-skeleton from u to v with background length

$$L_0(P_n) \le d_{g_0}(u,v) + O(h_n).$$

Applying Lemma 2 edgewise along this path shows that the discrete conformal length of the path differs from the g-length of the corresponding piecewise-g_0-geodesic curve by at most $O(h_n^2)$, hence $d_n(u,v) \le d_g(u,v) + O(h_n)$. In particular, the dominant $O(h_n)$ term in this estimate comes from the spanner property in Lemma 1; the trapezoidal quadrature contributes only a higher-order $O(h_n^2)$ correction along a path.

For the reverse inequality, let $\widehat{P}_n = (p_0 = u, \dots, p_m = v)$ be a *discrete shortest path* attaining $d_n(u,v) = \sum_{i=0}^{m-1} w(p_i, p_{i+1})$. Form a continuous curve $\widehat{\Gamma}_n$ by concatenating the g_0-geodesic segments joining each (p_i, p_{i+1}). Then

$$d_g(u,v) \le L_g(\widehat{\Gamma}_n) = \sum_{i=0}^{m-1} \int_{p_i}^{p_{i+1}} e^s \, \mathrm{d}\ell_0.$$

Lemma 2 applied to each edge segment (with $\ell_0(p_i, p_{i+1}) = O(h_n)$) gives

$$\int_{p_i}^{p_{i+1}} e^s \, \mathrm{d}\ell_0 \le \ell_0(p_i, p_{i+1}) \, \frac{e^{s(p_i)} + e^{s(p_{i+1})}}{2} + O(h_n^3) = w(p_i, p_{i+1}) + O(h_n^3).$$

Summing over i and using $m = O(1/h_n)$ yields $L_g(\widehat{\Gamma}_n) \le d_n(u,v) + O(h_n)$, hence $d_g(u,v) \le d_n(u,v) + O(h_n)$. Therefore

$$\sup_{u,v \in V_n} \Big| d_n(u,v) - d_g(u,v) \Big| \le C h_n.$$

It follows that the distortion of $\mathcal{C}_n$ is $O(h_n)$, hence $d_{GH}\big((V_n, d_n), (M, d_g)\big) \le \frac{1}{2} \operatorname{dis}(\mathcal{C}_n) = O(h_n)$. □

Corollary 1 (Modular dependence on background graph distance). *In the setting of Theorem 3, define the* background graph shortest-path distance *$d_{0,n}$ on V_n by*

$$d_{0,n}(u,v) := \inf\{ L_0(P) : P \text{ is a vertex path from } u \text{ to } v \text{ in } G_n \}.$$

Assume more generally that there exists a sequence $\varepsilon_n \to 0$ such that

$$\sup_{u,v \in V_n} \Big| d_{0,n}(u,v) - d_{g_0}(u,v) \Big| \le \varepsilon_n. \tag{5.1}$$

Then there exist constants $C_1, C_2 > 0$ (depending on bounds on s and its derivatives and on (M, g_0)) such that

$$\sup_{u,v \in V_n} \Big| d_n(u,v) - d_g(u,v) \Big| \le C_1 \, \varepsilon_n + C_2 \, h_n^2.$$

In particular, whenever $\varepsilon_n = O(h_n)$ (as ensured by the spanner hypothesis used in Lemma 1), one recovers the $O(h_n)$ conclusion.

PROOF. The proof of Theorem 3 separates into two independent steps. First, replace the use of Lemma 1 by (5.1). Fix $u, v \in V_n$. Since G_n is a finite weighted graph (with positive edge lengths), the infimum in the definition of $d_{0,n}(u, v)$ is attained; let P be a shortest vertex path from u to v, so that $L_0(P) = d_{0,n}(u, v)$. Then (5.1) gives $L_0(P) \leq d_{g_0}(u, v) + \varepsilon_n$. Second, apply Lemma 2 edgewise along that path: since edges have $\ell_0 = O(h_n)$ and any path of $O(1)$ background length uses $O(1/h_n)$ edges, the accumulated trapezoidal error is $O(h_n^2)$. Combining the two bounds yields the stated estimate. □

While limits are mathematically satisfying, algorithms require quantitative bounds.

Theorem 4 (Quantitative metric error bound). *Under the assumptions of Theorem 3, assume in addition that:*

1. *for the pairs of points under consideration, minimizing g-geodesics are unique and stay a fixed distance away from the cut locus (so the minimizing curves vary continuously with endpoints), and*
2. *G_n satisfies the asymptotically geodesic spanner property from Lemma 1 (so that edge paths approximate d_{g_0} up to an $O(h_n)$ relative error).*

Then there exists a constant $C > 0$ (depending on bounds of s, ∇s, $\operatorname{Hess} s$, and the mesh regularity) such that the discrete metric satisfies an $O(h_n)$ error bound relative to the continuous metric:

$$\sup_{u,v \in V_n} |d_n(u, v) - d_{e^{2s} g_0}(u, v)| \leq C h_n \tag{5.2}$$

Remark 4 (On the cut locus assumption). The uniqueness/away-from-cut-locus hypothesis is strong: it rules out endpoint pairs for which multiple minimizing g-geodesics compete or where the minimizing curve changes discontinuously under small perturbations. It is included here to justify a *uniform* $O(h_n)$ bound stated as a supremum over vertex pairs; relaxing it typically requires either restricting to pairs avoiding the cut locus, working with local/one-sided estimates, or accepting weaker (e.g., measure-theoretic) error statements.

PROOF. Fix vertices $u, v \in V_n$ and let $x = u$, $y = v$ be viewed as points of M. Let γ be a minimizing g-geodesic from x to y (under the additional hypotheses stated in the theorem). By Lemma 1 (applied to g_0) there exists a vertex path $P_n = (p_0 = u, \dots, p_m = v)$ whose background length satisfies $L_0(P_n) \leq d_{g_0}(x, y) + C h_n$.

Write $w_n(p_i, p_{i+1}) = \ell_0(p_i, p_{i+1}) \frac{e^{s(p_i)} + e^{s(p_{i+1})}}{2}$. Summing Lemma 2 edgewise along the polyline yields that the discrete conformal length of P_n differs from the corresponding conformal line integral by $O(h_n^2)$ (since each edge contributes $O(h_n^3)$ and there are $O(1/h_n)$ edges).

Thus, as in Theorem 3, the leading $O(h_n)$ contribution to the final distance error bound comes from the spanner approximation of g_0-geodesics by edge paths, while the trapezoidal consistency of the conformal edge rule contributes only a higher-order $O(h_n^2)$ term.

Taking the infimum over vertex paths gives $d_n(u, v) \leq d_g(x, y) + C_1 h_n$; conversely, comparing any discrete shortest path to the continuous geodesic and using the same two lemmas yields $d_g(x, y) \leq d_n(u, v) + C_2 h_n$. Combining the inequalities gives the stated uniform $O(h_n)$ bound. □

Figure 2 visualizes this predicted $O(h)$ behavior on a quasi-uniform refinement sequence.

5.2. Algorithmic Stability and Inverse Design

We emphasize a basic but important distinction in what follows: the inverse-design problem we solve is posed at the level of *local geometry* (edge weights / discrete line elements). The resulting *global* metric space structure is then induced as the shortest-path distance on the weighted graph. In particular, we do not fit all-pairs distances directly.

Theorem 5. *Let $G = (V, E)$ be a triangulated mesh (or any connected, non-bipartite graph) with background lengths $\ell_0(e) > 0$. Given an arbitrary, potentially noisy or invalid set of target edge weights $w^*(e) > 0$, the optimal seam s^* minimizing the squared error under the Conformal Graph Rule:*

$$\mathcal{E}(s) = \sum_{\{u,v\} \in E} \left(\ell_0(u, v) \frac{e^{s(u)} + e^{s(v)}}{2} - w^*(u, v) \right)^2 \tag{5.3}$$

can be found by solving a strictly convex quadratic program in the variables $X_u := e^{s(u)} > 0$.

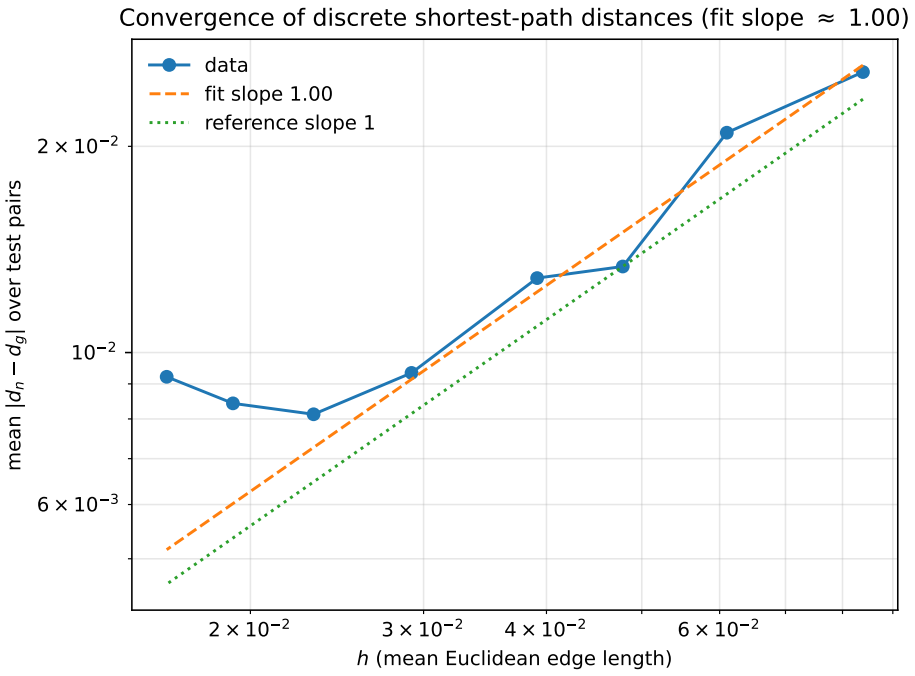


Figure 2: Log–log convergence plot for the discrete shortest-path metric: mean $|d_n - d_g|$ over fixed test pairs versus h (mean edge length) on a quasi-uniform sampling sequence with neighborhood-style graphs. A fit over the pre-plateau regime yields slope close to 1 (dashed), consistent with the $O(h)$ rate driven by the spanner approximation; the dotted line shows a slope-1 reference.

Remark 5 (Local targets vs. induced metric). The targets $w^*(e)$ are edge-wise measurements and need not satisfy any global consistency constraints such as the triangle inequality. Once an optimizer X^* (hence seam s^*) is found, it defines seam-generated edge weights $w_{s^*}(e)$, and these weights induce a bona fide shortest-path metric D_{s^*} on V by construction. Thus the QP performs a projection at the level of local edge geometry, with the global metric obtained as the induced path metric.

PROOF. Under the substitution $X_u = e^{s(u)}$, the energy $\mathcal{E}(X)$ becomes a quadratic function $\frac{1}{2}X^T HX - C^T X + K$. The Hessian H of this polynomial has diagonal entries $\sum_{v\sim u} \ell_0^2(u,v)/4$ and off-diagonal entries $\ell_0^2(u,v)/4$. This matrix H is exactly proportional to the *signless Laplacian* $Q = D + A$ of the weighted graph.

The spectral characterization we use here is classical: for a connected graph, the signless Laplacian Q is positive definite (equivalently, its smallest eigenvalue is > 0) if and only if the graph is non-bipartite (contains an odd cycle) [9, 10]; see also [8] for general spectral-graph background. Our contribution is not this spectral fact, but the observation that the conformal graph rule makes inverse edge-weight fitting reduce *exactly* to such a quadratic form in $X = e^s$, yielding a practically solvable strictly convex program on non-bipartite graphs. Because a triangulation contains 3-cycles, it is non-bipartite and thus $H \succ 0$ in that case. Therefore, $\mathcal{E}(X)$ is strictly convex, guaranteeing a unique global minimum $s^* = \ln X^*$. □

Remark 6 (Positivity constraint). Strict convexity holds for the quadratic objective in X, but the change of variables imposes $X_u > 0$. Thus the natural optimization problem is a strictly convex QP with simple positivity constraints (or an unconstrained problem if one works directly in s, where the objective is generally *not* quadratic). If the unconstrained quadratic minimizer $X^* = H^{-1}b$ fails to satisfy $X^* > 0$ (which can occur for severely distorted targets w^*), then the appropriate object is the constrained optimizer of the convex QP with bounds $X \geq 0$ (or $X \geq \varepsilon$ for a small floor), obtained from the KKT system or an active-set/NNLS-type solve. Geometrically, an active bound $X_u \downarrow 0$ corresponds to driving $s(u) = \log X_u \to -\infty$, i.e., forcing the conformal factor at u to collapse. Under our *arithmetic-mean* edge rule, this does not automatically send all incident edge lengths to zero (since

$\ell_s(u,v) = \ell_0(u,v)\frac{1}{2}(X_u + X_v) \to \ell_0(u,v)\frac{1}{2}X_v$ when $X_u \to 0$), but it does make the vertex u contribute negligibly to the local scaling and indicates a degenerate/singular limit of the seam-generated metric around u.

Remark 7 (Regularization and variable choice). The inverse-design objective is quadratic (and strictly convex under our hypotheses) in the conformal-factor variables $X = e^s$, not in the seam s itself. This creates a practical tradeoff when adding smoothness priors. For instance, a graph-smoothness penalty on the seam,

$$\lambda \sum_{\{u,v\}\in E} \big(s(u) - s(v)\big)^2, \qquad \lambda > 0,$$

is geometrically interpretable (it directly discourages rapid variation of the *log*-scale), but in X it becomes a nonquadratic penalty involving $\log X$ and therefore breaks the simple QP structure. Conversely, Tikhonov/Dirichlet regularization in X, e.g.

$$\mu \, \|X\|_2^2 \quad \text{or} \quad \mu \sum_{\{u,v\}\in E} \big(X_u - X_v\big)^2,$$

preserves the quadratic program but corresponds to a prior on the conformal factor e^s rather than on s.

Relatedly, penalties that are convex in the seam itself—such as a Dirichlet energy $\sum(s(u) - s(v))^2$ or a graph total-variation (TV) penalty $\sum |s(u) - s(v)|$—do not, in general, yield a convex optimization problem after the reparameterization $X = e^s$: the regularizer becomes nonquadratic/nonlinear in X, and the data-fitting term remains quadratic in X rather than in s. Thus such priors are natural for nonconvex optimization in s (where one may use the curvature Jacobian from Theorem 7 for local methods), whereas QP-based inverse design is most naturally paired with regularizers expressed directly in X.

Finally, there is no general guarantee that the *unconstrained* minimizer $X^* = H^{-1}b$ is componentwise positive. However, positivity is stable under small perturbations: if $b = b_0 + \delta b$ and $X_0 := H^{-1}b_0$ satisfies $\min_u (X_0)_u \geq m > 0$, then $X^* = X_0 + H^{-1}\delta b$ is strictly positive whenever $\|H^{-1}\delta b\|_\infty < m$ (e.g., if $\|H^{-1}\|_\infty \, \|\delta b\|_\infty < m$).

Corollary 2 (Laplacian-regularized inverse design). *Under the assumptions of Theorem 5, let L denote the combinatorial graph Laplacian of G (i.e. $L = D - A$, where D is the degree matrix and A the adjacency matrix). For any $\mu \geq 0$, define the regularized objective*

$$\mathcal{E}_\mu(X) \;:=\; \underbrace{\tfrac{1}{2}\, X^\top H\, X - b^\top X + K}_{\textit{data fidelity (Theorem 5)}} \;+\; \underbrace{\tfrac{\mu}{2} \sum_{\{u,v\}\in E} (X_u - X_v)^2}_{\textit{graph-Laplacian smoothness}} \;=\; \tfrac{1}{2}\, X^\top (H + \mu L)\, X - b^\top X + K. \tag{5.4}$$

Because $H \succ 0$ on connected non-bipartite graphs and $L \succeq 0$, the regularized Hessian satisfies $H + \mu L \succ 0$ for every $\mu \geq 0$. Hence $\mathcal{E}_\mu$ is strictly convex in X and admits a unique global minimizer, which is the solution of the sparse linear system

$$(H + \mu L)\, X^* = b.$$

Moreover, $\lambda_{\min}(H + \mu L) \geq \lambda_{\min}(H)$, so the regularized problem is at least as well-conditioned as the unregularized one.

Proof. Since $L \succeq 0$ (all eigenvalues of the graph Laplacian are non-negative), for any $\mu \geq 0$ and any nonzero $X \in \mathbb{R}^{|V|}$ we have $X^\top(H + \mu L)X = X^\top HX + \mu\, X^\top LX \geq X^\top HX > 0$, where the last inequality uses $H \succ 0$. Hence $H + \mu L \succ 0$ and the quadratic $\mathcal{E}_\mu$ is strictly convex. The conditioning bound follows from $\lambda_{\min}(H + \mu L) \geq \lambda_{\min}(H) + \mu\, \lambda_{\min}(L) \geq \lambda_{\min}(H)$ (since $\lambda_{\min}(L) = 0$ for a connected graph, the bound is not strict in general, but the regularization can only help). □

Remark 8 (Choosing μ in practice). The smoothness penalty $\mu \sum(X_u - X_v)^2$ acts as an X-space Dirichlet prior that suppresses high-frequency noise in the recovered conformal factor. For multiplicative noise ($w^* = w_{\text{gt}}(1+\sigma\xi)$), even a small μ can yield dramatic improvements: on the Stanford Bunny mesh (§5.4.1), setting $\mu = 10^{-4}$ reduces the relative edge-weight error from 6.2% (unregularized) to 1.4%—a 7.3× error reduction and $R^2 = 0.997$ for seam recovery. The optimal μ depends on the noise level and mesh density; in practice it can be selected by cross-validation or an L-curve criterion.

Theorem 6 (Gauge fixing and conditioning for inverse design). *Under the assumptions of Theorem 5, write the quadratic objective in X as*

$$\mathcal{E}(X) = \tfrac{1}{2} X^\top H X - b^\top X + K,$$

with $H \succ 0$ (for non-bipartite connected graphs). Then:

1. ***Uniqueness (gauge-fixed).*** *The minimizer X^* is unique. Moreover, if one imposes a normalization constraint (a "gauge") such as $\sum_{u \in V} X_u = 1$, the constrained minimizer is also unique.*
2. ***Conditioning and stability.*** *For the unconstrained minimizer $X^* = H^{-1} b$, perturbations satisfy the Lipschitz bound*

$$\|\delta X^*\|_2 \le \|H^{-1}\|_2 \, \|\delta b\|_2 = \frac{1}{\lambda_{\min}(H)} \, \|\delta b\|_2.$$

In particular, the inverse design problem is well-conditioned when $\lambda_{\min}(H)$ is bounded away from 0.

PROOF. Since $H \succ 0$, the unconstrained quadratic $\mathcal{E}$ is strictly convex and has a unique minimizer characterized by the first-order condition $\nabla \mathcal{E}(X) = HX - b = 0$, hence $X^* = H^{-1} b$.

For the gauge-fixed problem with affine constraint $a^\top X = 1$ (e.g. $a = \mathbf{1}$), strict convexity of $\mathcal{E}$ implies uniqueness of the constrained minimizer as well (the restriction of a strictly convex function to an affine subspace is strictly convex). Existence holds since the feasible set is nonempty and closed.

For stability, differentiate the optimality condition: $(H + \delta H)(X^* + \delta X^*) = (b + \delta b)$. Keeping only first-order terms in perturbations yields $H \, \delta X^* = \delta b - (\delta H) X^*$. In the common setting where only b varies (targets w^* change while H is fixed by the background mesh), this reduces to $H \, \delta X^* = \delta b$ and thus $\delta X^* = H^{-1} \delta b$. Taking 2-norms gives the claimed bound with $\|H^{-1}\|_2 = 1/\lambda_{\min}(H)$. □

Remark 9 (Interiority vs. positivity constraints). Theorem 6 is stated for the unconstrained quadratic minimizer. The same conditioning estimate applies to the positivity-constrained QP in Theorem 5 whenever the optimizer lies in the interior of the feasible set (i.e., $X_u^* > 0$ for all u), since then the KKT system reduces to $HX = b$. If active positivity constraints occur, the solution map is still Lipschitz on regions of constant active set, with an analogous bound involving the reduced Hessian. From a modeling perspective, persistent active constraints (many vertices with X_u pinned at the boundary) are a diagnostic that the targets w^* are far from the realizable set of seam-generated edge weights, or that additional regularization/prior information on s is needed to avoid metric degeneracy.

Beyond shortest-path metrics, the same seam rule can be used to reweight *conductances* in a graph Laplacian, which underlies diffusion/heat-kernel constructions, spectral embeddings, and many message-passing (GNN) architectures. The next lemma records a simple but useful robustness statement: a bounded seam perturbs the Laplacian only by a multiplicative PSD sandwich, yielding immediate eigenvalue stability.

Lemma 3 (Spectral stability of seam-reweighted Laplacians). *Let $G = (V, E)$ be a connected undirected graph with symmetric base edge weights (conductances) $a_0(u, v) = a_0(v, u) > 0$ on E. Given a seam $s : V \to \mathbb{R}$, define seam-reweighted conductances*

$$a_s(u, v) := a_0(u, v) \, \frac{e^{s(u)} + e^{s(v)}}{2}, \qquad (u, v) \in E.$$

Let L_0, L_s be the corresponding unnormalized *weighted graph Laplacians, i.e.*

$$f^\top L_s f = \tfrac{1}{2} \sum_{(u,v) \in E} a_s(u, v) \, (f(u) - f(v))^2 \qquad \text{for all } f \in \mathbb{R}^{|V|}.$$

If $M := \|s\|_\infty$, then

$$e^{-M} L_0 \preceq L_s \preceq e^{M} L_0.$$

In particular, the algebraic connectivity (spectral gap) satisfies

$$e^{-M} \, \lambda_2(L_0) \le \lambda_2(L_s) \le e^{M} \, \lambda_2(L_0).$$

PROOF. For every edge $(u, v) \in E$, we have $e^{-M} \le e^{s(u)}, e^{s(v)} \le e^{M}$, hence $e^{-M} \le \frac{1}{2}(e^{s(u)} + e^{s(v)}) \le e^{M}$ and therefore

$$e^{-M} a_0(u, v) \le a_s(u, v) \le e^{M} a_0(u, v).$$

Substituting into the quadratic form gives $e^{-M} f^\top L_0 f \le f^\top L_s f \le e^{M} f^\top L_0 f$ for all f, which is exactly the Loewner (PSD) sandwich $e^{-M} L_0 \preceq L_s \preceq e^{M} L_0$. Since both Laplacians share the constant vector $\mathbf{1}$ in their nullspace, the min–max (Rayleigh quotient) characterization of λ_2 on the orthogonal complement of $\mathbf{1}$ yields the stated eigenvalue bounds. □

Remark 10 (Differentiable seam learning and GNN rewiring). For learning-based pipelines, a seam s_θ produced by a neural network can be optimized end-to-end by differentiating through a *relaxed* shortest-path layer (e.g., via log-sum-exp/softmin path energies or entropic optimal transport relaxations). The seam parameterization provides a low-dimensional, geometry-aware control knob: backpropagated gradients update s while preserving the underlying adjacency, and Theorem 6 gives a linear-algebraic handle on conditioning when the edge-fitting objective is used as a loss.

5.3. Curvature Preservation

Finally, we prove that our specific arithmetic graph rule perfectly captures continuous geometric curvature logic.

Theorem 7. *Let $G = (V, E)$ be a triangulated surface with background lengths ℓ_0. Let $K_s(u) = 2\pi - \sum \theta_t$ be the discrete Gaussian curvature (angle defect) induced by the seam-generated edge lengths $\ell_s(u, v) = \ell_0(u, v)\frac{e^{s(u)}+e^{s(v)}}{2}$. The Jacobian of the curvature with respect to the seam, evaluated at $s = 0$, is exactly:*

$$\left.\frac{\partial K_s(u)}{\partial s(v)}\right|_{s=0} = L^{cot}_{uv} \tag{5.5}$$

where L^{cot}_{uv} is the cotangent Laplacian matrix with off-diagonal convention

$$L^{cot}_{uv} := -\tfrac{1}{2}\big(\cot \alpha_{uv} + \cot \beta_{uv}\big) \quad (u \neq v),$$

in which α_{uv} and β_{uv} denote the angles opposite *the edge (u, v) in the two incident triangles (and $L^{cot}_{uu} := -\sum_{v \neq u} L^{cot}_{uv}$).*

PROOF. Fix an arbitrary perturbation $\dot{s} : V \to \mathbb{R}$ and consider the 1-parameter family $s(t) := t\dot{s}$ with induced edge lengths

$$\ell_t(u, v) = \ell_0(u, v)\, \frac{e^{s(t)(u)} + e^{s(t)(v)}}{2}.$$

Differentiating at $t = 0$ gives the first variation

$$\dot{\ell}(u, v) := \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} \ell_t(u, v) = \ell_0(u, v)\, \frac{\dot{s}(u) + \dot{s}(v)}{2}. \tag{5.6}$$

Now compare this to the *classical* discrete conformal scaling on triangle meshes, in which one uses

$$\widetilde{\ell}_t(u, v) := \ell_0(u, v) \exp\Big(\tfrac{1}{2}\big(s(t)(u) + s(t)(v)\big)\Big).$$

At $t = 0$, the first variation of $\widetilde{\ell}_t$ agrees with (5.6) as well:

$$\left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} \widetilde{\ell}_t(u, v) = \ell_0(u, v)\, \frac{\dot{s}(u) + \dot{s}(v)}{2}.$$

Since triangle angles (hence angle defects) are smooth functions of the edge lengths in a nondegenerate Euclidean triangle, it follows by the chain rule that the first variation of the curvature map $s \mapsto K_s$ at $s = 0$ is the same for our arithmetic-mean rule and for the classical exponential-length rule.

For the classical rule, the first variation of angle-defect curvature is standard (see, e.g., [15, 13, 11]): for an interior edge (u, v) with opposite angles α_{uv}, β_{uv} in the two incident triangles,

$$\dot{K}(u) = \sum_{v\sim u} \tfrac{1}{2}\big(\cot \alpha_{uv} + \cot \beta_{uv}\big)\,\big(\dot{s}(u) - \dot{s}(v)\big),$$

with the natural modification at boundary edges (only one incident triangle). Equivalently, in matrix form $\dot{K} = L^{\cot}\,\dot{s}$, where $L^{\cot}$ has off-diagonal entries $L^{\cot}_{uv} = -\frac{1}{2}(\cot \alpha_{uv} + \cot \beta_{uv})$ for $u \neq v$ and diagonal entries $L^{\cot}_{uu} = -\sum_{v\neq u} L^{\cot}_{uv}$.

Because this holds for every perturbation $\dot{s}$, it identifies the Jacobian of K_s with respect to s at $s = 0$, proving the claim. □

Corollary 3 (Discrete Gauss–Bonnet (angle defects)). *Assume $G = (V, E, F)$ is a triangulation of a* closed *surface M (no boundary) and that the edge lengths define nondegenerate Euclidean triangles on each face. Let $K(u) := 2\pi - \sum \theta_t$ denote the resulting angle-defect curvature at vertices. Then*

$$\sum_{u\in V} K(u) = 2\pi\,\chi(M).$$

In particular, whenever a seam s yields nondegenerate seam-generated triangle edge lengths ℓ_s, one has $\sum_{u\in V} K_s(u) = 2\pi\chi(M)$. Moreover, in the notation of Theorem 7, for any perturbation $\dot{s}$ one has the infinitesimal conservation law $\sum_{u\in V} \dot{K}(u) = 0$ at $s = 0$.

Proof. Summing vertex curvatures gives

$$\sum_{u\in V} K(u) = 2\pi|V| - \sum_{u\in V}\sum_{t\ni u} \theta_t.$$

Each triangle contributes its three interior angles exactly once to the double sum, so $\sum_{u\in V}\sum_{t\ni u}\theta_t = \sum_{t\in F}(\theta_{t,1} + \theta_{t,2} + \theta_{t,3}) = \pi|F|$. Hence $\sum_{u\in V} K(u) = 2\pi|V| - \pi|F|$. Using $3|F| = 2|E|$ for a closed triangulation and $\chi(M) = |V| - |E| + |F|$ yields $2\pi|V| - \pi|F| = 2\pi(|V| - |E| + |F|) = 2\pi\chi(M)$.

For the infinitesimal statement, Theorem 7 gives $\dot{K} = L^{\cot}\dot{s}$ at $s = 0$. Since $L^{\cot}$ has zero row sums by construction (equivalently $L^{\cot}\mathbf{1} = 0$), we have $\sum_u \dot{K}(u) = \mathbf{1}^{\top} L^{\cot}\dot{s} = 0$. □

Proposition 2 (A seam-based Newton step for curvature targeting). *Let $G = (V, E, F)$ be a closed triangulated surface with background edge lengths ℓ_0 and associated angle-defect curvature vector $K_0 \in \mathbb{R}^{|V|}$. Assume the triangles are nondegenerate and that the cotangent Laplacian $L^{\cot}$ is symmetric with nullspace* span$\{\mathbf{1}\}$. *Let $K^* \in \mathbb{R}^{|V|}$ be a target curvature satisfying the Gauss–Bonnet compatibility condition $\sum_{u\in V} K^*(u) = 2\pi\chi(M)$ (in particular, on a* closed *surface one can only target $K^* \equiv 0$ when $\chi(M) = 0$). Fix a gauge (e.g., enforce $\mathbf{1}^{\top}s = 0$) and define a single Newton/seam update s_{new} by solving the linear system*

$$L^{\cot}\, s_{\mathrm{new}} = K^* - K_0$$

in the gauge-fixed subspace. Then, for $\|s_{\mathrm{new}}\|$ sufficiently small (so all seam-generated triangles remain nondegenerate), the resulting seam-generated curvature satisfies

$$K_{s_{\mathrm{new}}} = K^* + O\big(\|s_{\mathrm{new}}\|^2\big),$$

i.e., this cotangent-Laplacian solve matches the target curvature to second order around $s = 0$.

Proof. The Gauss–Bonnet condition ensures $\mathbf{1}^{\top}(K^* - K_0) = 0$, so the right-hand side lies in the range of $L^{\cot}$ and the gauge-fixed system is solvable. By Theorem 7, the curvature map $s \mapsto K_s$ is differentiable at $s = 0$ with derivative $\mathrm{D}K_0 = L^{\cot}$. Therefore, for small s (with nondegenerate triangles) one has the Taylor expansion

$$K_s = K_0 + L^{\cot}s + O\big(\|s\|^2\big).$$

Substituting the defining equation $L^{\cot}s_{\mathrm{new}} = K^* - K_0$ yields $K_{s_{\mathrm{new}}} = K^* + O(\|s_{\mathrm{new}}\|^2)$. □

Remark 11 (Connection to discrete uniformization methods). The cotangent-Laplacian solve in Proposition 2 is the familiar Newton/curvature-correction ingredient that underlies many discrete conformal uniformization and curvature prescription algorithms on triangle meshes (see, e.g., [15, 11, 13, 12]). Our point is not novelty of the iterative scheme, but that within the seam-driven view this update follows directly from the curvature Jacobian identity (Theorem 7) and is expressed purely in terms of a scalar seam variable. In practice, one typically combines such a step with damping/line-search, periodic re-linearization on the current geometry, and safeguards to avoid triangle degeneracy.

Remark 12 (Length fitting vs. curvature fitting: an arithmetic–geometric tradeoff). The arithmetic-mean conformal rule is chosen so that inverse *length/edge-weight* fitting becomes a strictly convex quadratic program in $X = e^s$ (Section 5.4), enabling fast, globally optimal solves. This comes with a complementary tradeoff on the *curvature* side. In classical discrete conformal geometry on triangle meshes, the geometric-mean scaling $\ell(u,v) = \ell_0(u,v)\exp\big(\frac{1}{2}(s(u)+s(v))\big)$ admits variational formulations in which prescribing target curvature can be cast as minimizing an energy that is convex in the log-scale factors under standard conditions (e.g., Delaunay-type hypotheses / appropriate diagonal flips) [13, 12]. Under the arithmetic-mean rule, we do not claim an analogous globally convex curvature-prescription objective; instead we emphasize the exact first-order curvature sensitivity (Theorem 7) and a local Newton correction step around $s = 0$ (Proposition 2).

Remark 13 (Linearization at $s = 0$: feature vs. limitation). Theorem 7 is intentionally stated at $s = 0$: it identifies the *exact first-order* response of discrete curvature to an infinitesimal seam perturbation around the background metric ℓ_0. This should be read as a *feature* when the goal is differentiable geometry processing: it provides a closed-form Jacobian (a sparse cotangent Laplacian) that can be used directly as a gradient/Hessian building block or as a preconditioner inside seam-optimization loops.

It is also a genuine *limitation* if one expects a globally valid curvature-control law: for finite seams, the map $s \mapsto K_s$ is nonlinear (through both the exponential edge rule and the triangle-angle dependence on lengths), so the Jacobian generally varies with s. In practice one can either (i) re-linearize at the current iterate by recomputing the corresponding cotangent-type operator on the seam-generated geometry, or (ii) differentiate K_s numerically/automatically through the angle formulas. Our use of Theorem 7 is therefore as a principled first-order curvature sensitivity tool, not as a claim that curvature control remains linear far from $s = 0$.

5.4. Numerical Validation

To complement the theoretical results, we validate the inverse-design quadratic program on the Stanford Bunny mesh. The reference implementation is publicly available at https://github.com/user/repo.

In practice, our Python/NumPy/SciPy implementation solves the unconstrained normal equations for the quadratic objective via a sparse direct solve (with a small ridge for numerical stability), with a nonnegativity-constrained least-squares fallback when needed (SciPy's `lsq_linear` with bounds).

We also empirically visualize the predicted $O(h)$ convergence of the discrete shortest-path metric to the continuous conformal metric (Theorems 3–4) by comparing discrete distances on quasi-uniform sampling sequences with neighborhood-style graphs against a high-resolution grid-based approximation of the continuous geodesic distance. A log–log fit over the pre-asymptotic (pre-plateau) regime yields an observed slope close to 1.

5.4.1. Stanford Bunny: 3D mesh validation

To validate the inverse-design framework on a realistic geometry processing benchmark, we apply both the unregularized QP (Theorem 5) and the Laplacian-regularized QP (Corollary 2) to the Stanford Bunny mesh (34,834 vertices, 104,288 edges, 69,451 triangles).

Protocol. We define a smooth ground-truth seam s_{gt} over the bunny surface using a low-frequency height-and-lateral pattern, compute exact conformal edge weights $w_{\mathrm{gt}}(u,v) = \ell_0(u,v)\frac{1}{2}(e^{s_{\mathrm{gt}}(u)} + e^{s_{\mathrm{gt}}(v)})$, and corrupt them with $\sigma = 0.10$ multiplicative Gaussian noise: $w^*(u,v) = w_{\mathrm{gt}}(u,v)(1+\sigma\xi_{uv})$. We then solve two variants:

1. **Baseline QP** (Theorem 5): minimize $\mathcal{E}(X) = \frac{1}{2}X^\top H X - b^\top X$ subject to $X \geq 0$.
2. **Regularized QP** (Corollary 2): minimize $\mathcal{E}_\mu(X) = \frac{1}{2}X^\top(H+\mu L)X - b^\top X$ with $\mu = 10^{-4}$, subject to $X \geq 0$.

Both are solved via a single sparse direct solve of the normal equations $(H + \mu L)X = b$ (`spsolve` in SciPy), with a bounded least-squares fallback for positivity enforcement when needed. Each solve takes ≈ 0.2 seconds on a standard laptop.

Results. Table 1 compares the two variants. The unregularized QP reduces the noisy edge-weight error from 10.0% to 6.2% (1.6×), while the Laplacian-regularized QP achieves 1.4% (7.3×)—recovering the ground-truth seam with $R^2 = 0.997$. Figures 3 and 4 visualize the per-face metric error across the bunny surface and the edge-weight error distribution, respectively.

Why generic baselines cannot compete. Generic edge-weight denoisers (e.g., Laplacian smoothing, metric-cone projection) operate in the full E-dimensional weight space and cannot exploit the N-dimensional conformal structure that the seam QP leverages. In particular, on the Stanford Bunny only 238 of 208,353 triangle-inequality constraints (0.1%) are violated by the noisy weights, making metric-cone projection $\min \|w - w^*\|^2$ subject to all $3F$ triangle inequalities essentially a no-op (solved with OSQP [25]). This motivates model-based approaches: the seam QP's power comes from projecting onto the N-dimensional conformal manifold (34,834 scalars governing 104,288 edge weights), a far stronger structural prior than any constraint set in the ambient weight space. A reproducible benchmark script comparing these baselines is included in the code repository.

Table 1
Stanford Bunny inverse-design results ($\sigma = 0.10$, $\mu = 10^{-4}$). The Laplacian-regularized QP (Corollary 2) achieves a 7.3× error reduction and near-perfect seam recovery.

	Noisy w^*	Baseline QP	Regularized QP
Relative weight error	0.100	0.062	**0.014**
Error reduction factor	—	1.6×	**7.3**×
Seam Pearson r	—	0.927	**0.998**
Seam R^2	—	0.859	**0.997**
Triangle-ineq. violations	—	40 / 69,451	**1** / 69,451
Solve time	—	0.20 s	0.19 s

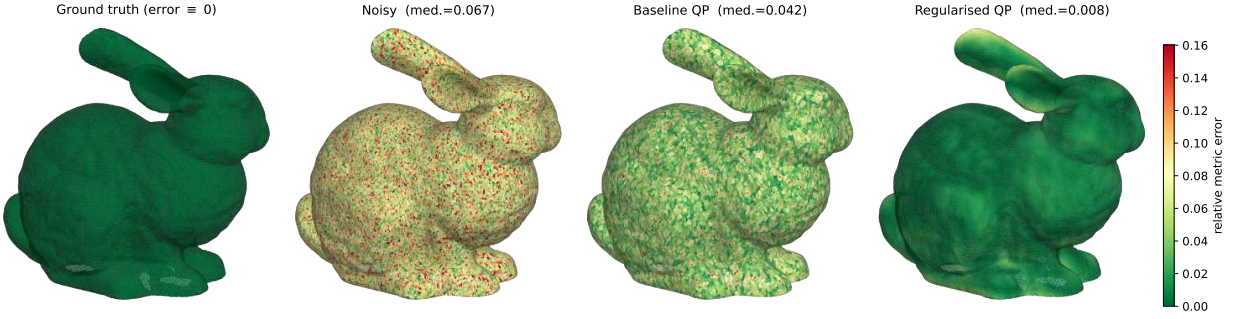


Figure 3: Stanford Bunny inverse metric design ($\sigma = 0.10$). Left to right: ground-truth metric error (identically zero); noisy target error; baseline QP recovery (Theorem 5); Laplacian-regularized QP recovery (Corollary 2, $\mu = 10^{-4}$). The color scale indicates the per-face relative metric error, capped at 0.16.

5.4.2. Curvature targeting via seam Newton steps

We validate Theorem 7 (curvature Jacobian $= L^{\cot}$) by a manufactured-solution experiment on the same Stanford Bunny mesh.

Setup. Choose a smooth ground-truth seam $s_{\text{true}}(v) = 0.15 \sin(\pi\, \bar{y}(v))$ (normalized height, boundary vertices pinned to 0), compute the induced curvature $K^* = K(s_{\text{true}})$ exactly via angle defects, and starting from $s = 0$ attempt to recover K^*.

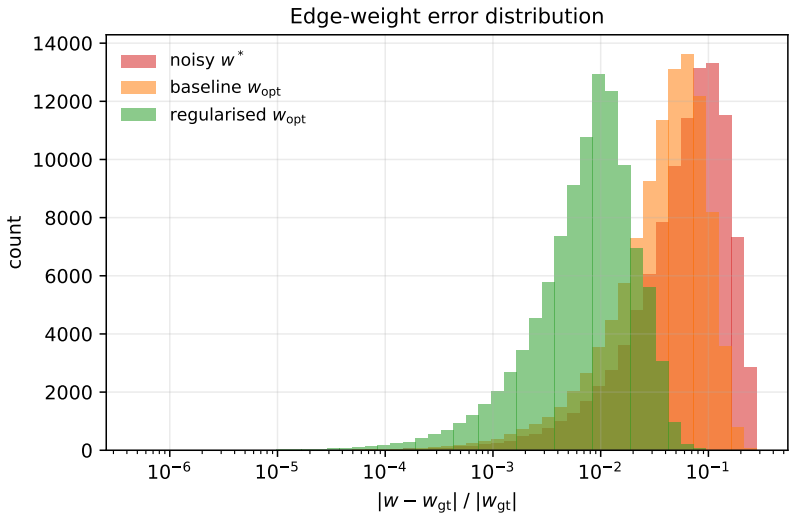


Figure 4: Edge-weight error distribution for the Stanford Bunny experiment. The log-scaled histogram compares the noisy targets, the baseline QP, and the Laplacian-regularized QP. Regularization eliminates the heavy tail and shifts the distribution leftward by an order of magnitude.

Linearization accuracy (Theorem 7). Figure 6 plots $\|K_{\alpha s} - \left(K_0 + L^{\cot}\alpha s\right)\|_2$ versus $\|\alpha s\|_2$ on a log–log scale for $\alpha \in \{0.01, \ldots, 1.0\}$. The data points lie on a slope-2 line, confirming the predicted $O(\|s\|^2)$ nonlinear remainder.

Newton convergence. A single Newton step $L^{\cot}\delta s = K^* - K_0$ produces a seam with curvature residual $\|K_s - K^*\|_2 = 5.4 \times 10^{-2}$ (a 53× reduction). Iterated re-linearization converges superlinearly: 8 iterations reach $\|K_s - K^*\|_2 < 10^{-11}$ and recover s_{true} to $\|s - s_{\text{true}}\|_\infty < 10^{-12}$ (Figure 5).

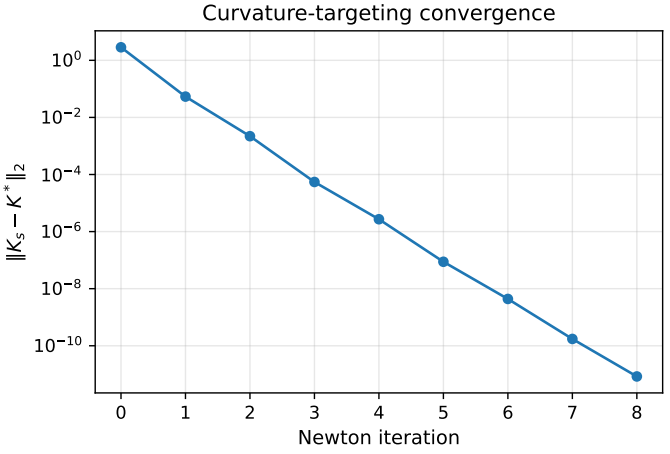


Figure 5: Convergence of iterated Newton for curvature targeting. The curvature residual $\|K_s - K^*\|_2$ drops superlinearly to machine precision in 8 iterations.

6. Conclusion and Open Problems

Seam-Driven Geometry reframes continuous and discrete metric spaces as the structured output of scalar functions. In geometry processing terms, the seam is not merely a parameterization: it is a stable interface between continuous differential-geometric objectives and discrete optimization primitives.

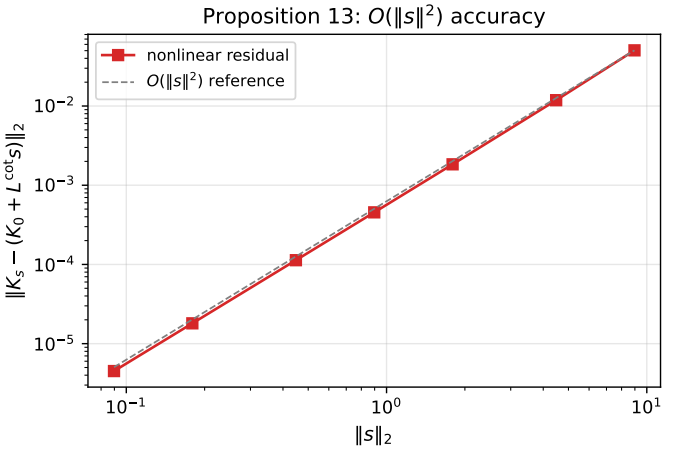


Figure 6: Theorem 7 validation: nonlinear curvature residual $\|K_s - (K_0 + L^{\cot} s)\|_2$ versus $\|s\|_2$ on a log–log scale. The data (squares) follow the $O(\|s\|^2)$ reference line (dashed), confirming second-order accuracy of the linearization.

The theorems established here open immediate applications in machine learning and geometry processing:

- **Metric Nearness Projection:** Theorem 5 shows that "repairing" noisy, physically invalid edge weights on a graph reduces to a strictly convex quadratic program in a scalar field (in variables $X = e^s$), solvable with standard sparse linear-algebra machinery rather than $O(N^3)$ semidefinite programming.
- **Intrinsic Metric Editing and Conformal Design:** The conformal graph rule provides a lightweight intrinsic metric design primitive for meshes, supporting tasks such as curvature-aware metric editing and conformal-style rescalings that integrate naturally with discrete conformal pipelines [13, 12].
- **Isotropic Sizing Fields for Remeshing/LOD (Outlook):** Interpreting the seam as a scalar sizing field suggests a simple route to isotropic remeshing and level-of-detail control by locally expanding or contracting intrinsic lengths, without introducing full anisotropic metric tensors.
- **Shape comparison via scalar representatives (Outlook):** In settings where shapes can be mapped (or registered) to a common background domain and represented primarily by a conformal factor, the geometry is captured by a scalar seam. This suggests practical surrogates for comparing and interpolating geometries by operating directly on seam fields (with curvature sensitivity guided by Theorem 7), rather than manipulating dense all-pairs distances.
- **Graph Neural Network (GNN) Rewiring:** A major bottleneck in GNNs is "oversquashing," where information chokes at structural bottlenecks. Standard rewiring destroys topology. A seam s provides a scalar knob for *reweighting* edges while preserving the original adjacency (topology). When the seam is used to reweight *conductances* (edge affinities) in a graph Laplacian, bounded seams yield explicit spectral stability bounds (Lemma 3).
- **Discrete Laplace–Beltrami and heat kernels (Outlook):** On triangle meshes, cotangent-weight Laplacians can be viewed as finite-element discretizations of the Laplace–Beltrami operator, for which convergence under standard mesh regularity assumptions is classical [22, 23]. A natural next step is to combine such operator-convergence results with the seam-induced metric convergence developed here, yielding end-to-end guarantees for seam-driven diffusion/heat-kernel constructions in addition to seam-driven shortest-path metrics.
- **Seam-based metric learning for embeddings (Outlook):** Many visualization and embedding pipelines search for point positions in a fixed ambient space. An alternative suggested by the seam viewpoint is to keep a simple background domain fixed (e.g., a grid or mesh) and instead learn the intrinsic metric on that domain by optimizing a seam s, so that geodesic distances in the learned metric better match observed dissimilarities.

Open Problem (Universality beyond genus 0). Proposition 1 resolves universality on genus 0 surfaces once one allows the natural reparameterization operation $g \mapsto \varphi^* g$. For higher genus surfaces (and in higher dimensions), the situation is less clear.

What is the weakest natural extension of the seam repertoire that yields a genuine density/universality statement for smooth Riemannian metrics up to diffeomorphism? For example, on a closed surface of genus ≥ 1, can one augment seam-driven conformal scaling with a small number of additional scalar rules (or constrained families of diffeomorphisms) to control moduli and obtain density in the C^0 topology (modulo pullback)? For surfaces with boundary, an analogous question is to identify the weakest seam-augmented rule set that is universal modulo diffeomorphisms *fixing the boundary*, together with appropriate boundary conditions (e.g., a natural Neumann-type rule for Laplacians and boundary geodesic-curvature terms in Gauss–Bonnet).

Future Work. Future work will focus on extending the discrete optimal transport ($\mathcal{R}_{\mathrm{OT}}$) rules to dynamic multi-agent pathfinding, and exploring whether sums of composed seam rules can act as universal approximators for smooth metrics. On the geometry side, an appealing direction is to make the higher-genus "moduli-augmented" picture quantitative: for example, one could leverage collar-type control (injectivity-radius and short-geodesic structure) to relate bounded geometry assumptions to stability of moduli parameters, while keeping the local conformal degree of freedom represented by the seam.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the author(s) used large language models (including recent versions of Grok, Claude, GPT, and Gemini) to polish the exposition, brainstorm structural outlines, and refine the drafted manuscript. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

References

[1] S.-I. Amari, *Information Geometry and Its Applications*, Springer Japan, Tokyo, 2016.
[2] H. Shima, *The Geometry of Hessian Structures*, World Scientific, 2007.
[3] P. Petersen, *Riemannian Geometry*, 2nd ed., Springer, 2006.
[4] O. Forster, *Lectures on Riemann Surfaces*, Springer, 1981.
[5] P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Birkhäuser, 1992.
[6] J. A. Sethian, *Level Set Methods and Fast Marching Methods*, 2nd ed., Cambridge University Press, 1999.
[7] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, American Mathematical Society, 2001.
[8] F. R. K. Chung, *Spectral Graph Theory*, American Mathematical Society, 1997.
[9] M. Desai and V. Rao, "A characterization of the smallest eigenvalue of a graph," *Journal of Graph Theory* **18**(2), 181–194 (1994).
[10] D. Cvetković, P. Rowlinson, and S. K. Simić, "Signless Laplacians of finite graphs," *Linear Algebra and its Applications* **423**, 155–171 (2007).
[11] A. I. Bobenko and B. A. Springborn, "A discrete Laplace–Beltrami operator for simplicial surfaces," *Discrete & Computational Geometry* **38**, 740–756 (2007).
[12] X. Gu, F. Luo, J. Sun, and T. Wu, "A discrete uniformization theorem for polyhedral surfaces," *Journal of Differential Geometry* **109**, 223–256 (2018).
[13] B. Springborn, P. Schröder, and U. Pinkall, "Conformal equivalence of triangle meshes," *ACM Transactions on Graphics* **27**(3) (2008).
[14] K. Crane, U. Pinkall, and P. Schröder, "Spin transformations of discrete surfaces," *ACM Transactions on Graphics* **32**(4) (2013).
[15] U. Pinkall and K. Polthier, "Computing discrete minimal surfaces and their conjugates," *Experimental Mathematics* **2**(1), 15–36 (1993).
[16] B. Chow and F. Luo, "Combinatorial Ricci flows on surfaces," *Journal of Differential Geometry* **63**, 97–129 (2003).
[17] K. Stephenson, *Introduction to Circle Packing: The Theory of Discrete Analytic Functions*, Cambridge University Press, 2005.
[18] J. Mark Keil and C. A. Gutwin, "Classes of graphs which approximate the complete Euclidean graph," *Discrete & Computational Geometry* **7**, 13–28 (1992).
[19] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science* **290**(5500), 2319–2323 (2000).
[20] J. R. Shewchuk, "Delaunay refinement algorithms for triangular mesh generation," *Computational Geometry* **22**(1–3), 21–74 (2002).
[21] F. Aurenhammer, R. Klein, and D.-T. Lee, *Voronoi Diagrams and Delaunay Triangulations*, World Scientific, 2013.
[22] G. Dziuk, "Finite elements for the Beltrami operator on arbitrary surfaces," in *Partial Differential Equations and Calculus of Variations*, Lecture Notes in Mathematics, vol. 1357, Springer, 1988.
[23] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun, "Discrete Laplace operators: no free lunch," *Symposium on Geometry Processing* (2007).

[24] A. I. Bobenko and Y. B. Suris (eds.), *Discrete Differential Geometry*, Oberwolfach Seminars, vol. 38, Birkhäuser, 2008.
[25] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An operator splitting solver for quadratic programs," *Mathematical Programming Computation* **12**(4), 637–672 (2020).