# Implicit Minimal Surfaces for Bijective Correspondences

ETIENNE CORMAN, Université de Lorraine, CNRS, LORIA, France
YOUSUF SOLIMAN, Side Effects Software, Canada
ROBIN MAGNET, Inria, France and Université Paris Cité, France
MARK GILLESPIE, Inria, France and University of Utah, USA

We introduce an *implicit* representation of continuous, bijective, orientation-preserving maps between genus zero surfaces with or without boundary. The distortion of these maps can easily be minimized by optimizing the Ginzburg-Landau functional—a ubiquitous model in physics and differential geometry—leading to a simple algorithm for computing bijective correspondences using only standard tools of the tangent vector field toolbox. The method avoids combinatorial mesh modifications and does not require barrier functions to enforce bijectivity making it more robust to noise and simpler to implement. Moreover, the algorithm does not assume a bijective initialization and can untangle non-bijective correspondences generated by computationally cheaper methods such as functional maps. It supports the use of both landmark points and landmark curves to guide the correspondence. The key idea is that a bijection between surfaces defines a two-dimensional mapping surface sitting inside the four-dimensional product space of the two inputs, and this mapping surface can be stored implicitly as the zero set of a *complex section*—essentially a complex function defined on the product space. Now the distortion of the map can be optimized by minimizing the area of this mapping surface, which amounts to minimizing the Ginzburg-Landau functional of the complex section. We demonstrate the practical benefits of our method by comparing to state-of-the-art correspondence algorithms and show that our implicit representation offers improved stability and naturally supports constraints that are difficult to enforce with explicit map representations.

CCS Concepts: • **Mathematics of computing → Mesh generation**.

Additional Key Words and Phrases: shape matching

**ACM Reference Format:**
Etienne Corman, Yousuf Soliman, Robin Magnet, and Mark Gillespie. 2026. Implicit Minimal Surfaces for Bijective Correspondences. *ACM Trans. Graph.* 1, 1 (March 2026), 20 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

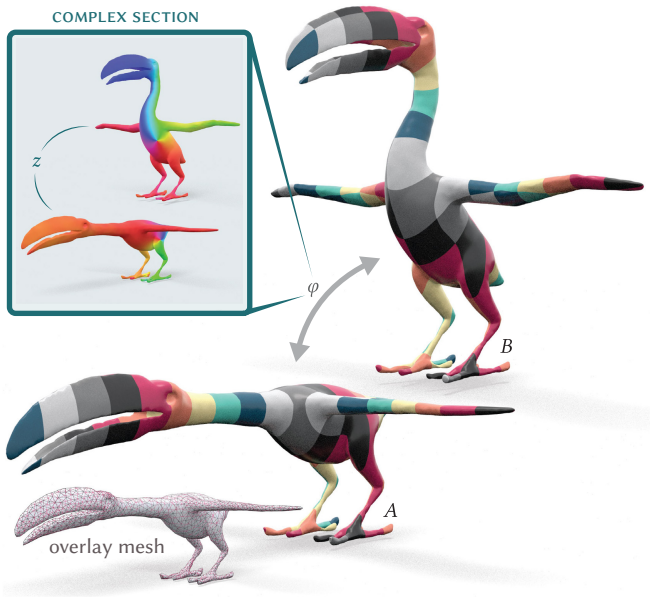


Fig. 1. We introduce a new implicit representation of maps $\varphi : A \to B$ between triangle meshes $A$ and $B$, encoding the maps as the zero sets of complex functions $z$ on the product space $A \times B$. In this representation we can compute high-quality orientation-preserving bijections between $A$ and $B$ by minimizing a simple Ginzburg-Landau energy, without requiring any combinatorial mesh modifications, barrier functions, or a bijective initialization. Once an implicit map has been computed, it encodes not only the vertex map, but also the entire overlay mesh under the correspondence.

## 1 Introduction

Computing a map between two surfaces $A$ and $B$ is a ubiquitous problem in geometry processing. It is a prerequisite for analyzing data across collections of shapes and is essential for transferring data such as textures, segmentations, or other semantic attributes, which are crucial to a wide range of applications from studying brain-fold variations in medical imaging to reconstructing morphological correspondences in paleontology, and many more tasks across computer graphics, computer vision, and scientific modeling.

When the surfaces share the same topological class, the goal is typically to find a map that is bijective, continuous, and aligns salient geometric or semantic features. This goal is often formalized by looking for as-isometric-as possible maps. Yet it remains highly challenging to compute such maps: the underlying distortion energies are non-convex, and ensuring continuity and bijectivity during optimization imposes strict constraints that are difficult to enforce.

In practice, even describing an exact map between two triangle meshes is cumbersome, yet it is necessary for evaluating a map's distortion. For every point on $A$, the map must specify its image on $B$, not only for the vertices of $A$ but also for all interior points of its faces. The most common approach is to describe the map by *remeshing* both surfaces so that they share identical connectivity, either by intersecting the meshes [Schmidt et al. 2020] or by replacing them altogether [Schmidt et al. 2023]. However, this representation is

Authors' Contact Information: Etienne Corman, etienne.corman@cnrs.fr, Université de Lorraine, CNRS, LORIA, Vandœuvre-lès-Nancy, France; Yousuf Soliman, yousufs@sidefx.com, Side Effects Software, Toronto, Canada; Robin Magnet, robin.magnet@inria.fr, Inria, Paris, France and Université Paris Cité, Paris, France; Mark Gillespie, mark.gillespie81@gmail.com, Inria, Palaiseau, France and University of Utah, Salt Lake City, USA.


ACM 1557-7368/2026/3-ART
https://doi.org/10.1145/nnnnnnn.nnnnnnn

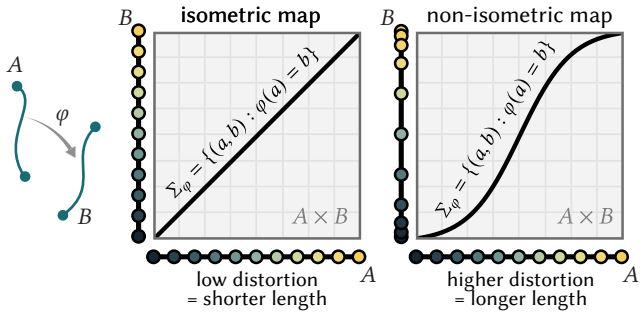


Fig. 2. To illustrate the idea behind our method, we move down a dimension and consider a map $\varphi : A \to B$ between two curves *(left)*. Such a map can be represented by its "graph", *i.e.* the set of identified points $\Sigma_\varphi = \{(a, b) : \varphi(a) = b\}$ inside the product space $\subset A \times B$ *(center, right)*. The more that $\varphi$ distorts the curves, the longer its graph becomes—thus, we can compute high-quality maps by minimizing the length of the graph.

deeply coupled to the mesh connectivity, whereas the map itself is a smooth, coordinate-invariant geometric object. Instead, we propose a new representation of correspondences that avoids combinatorial modifications entirely and relies only on well-established tools for processing tangent vector fields (Figure 1).

*Mappings as Surfaces in Product Space.* The representation we study in this paper is based on the observation that any mapping $\varphi : A \to B$ between two surfaces $A$ and $B$ induces a two-dimensional manifold $\Sigma_\varphi$ embedded in the product space $A \times B$:

$$\Sigma_\varphi \coloneqq \{(p, \varphi(p)) \ : \ p \in A\}. \tag{1}$$

Because $\varphi$ is continuous and bijective, each point $p \in A$ is mapped to a unique point $\varphi(p) \in B$. Consequently, the projection of $\Sigma_\varphi$ onto $B$ covers $B$ exactly once. Conversely, if a two-dimensional submanifold $\Sigma \subset A \times B$ has the property that its projection onto each input surface covers that surface exactly once, then $\Sigma$ represents a bijective continuous mapping between $A$ and $B$. Moreover, if the projections of $\Sigma$ onto $A$ and $B$ are orientation-preserving, then the mapping $\varphi$ will also preserve orientation (Figure 5).

The case of matching two curves provides an intuitive illustration (Figure 2). In this setting, the product space is a rectangle, and a map is represented by the graph of the function $a \mapsto \varphi(a)$. This representation enables a direct measure of the distortion of the map: the length of the graph. The identity map, represented by a straight line, clearly minimizes distortion, while any other mapping necessarily increases both the distortion and the length of the graph.

Crucially, these principles extend to higher dimensions. A submanifold $\Sigma \subset A \times B$ whose projections cover each input surface exactly once defines a continuous and bijective correspondence. And if $\Sigma$ has minimal area, it represents a distortion-minimizing map (Figure 3). (For the mathematical expression of this distortion, and its relation to other well known distortion measures, see Section 1.1.)

We thus shift the focus: rather than working with map $\varphi$ directly, we instead compute a surface in the product manifold $A \times B$ which implicitly encodes a distortion-minimizing correspondence.

*Minimal Surfaces and Complex Fields.* Now that we have framed the matching problem using minimal surfaces, a central question remains: *How can we effectively represent and compute a minimal surface in a four-dimensional manifold?* Fortunately a substantial body of literature addresses this problem using tools remarkably similar to those used in geometry processing. A key insight is that minimal surfaces can be characterized through the zero set of a vector-field like object known as a *complex section*. In geometry processing, we often use the property that the zeros of a smooth vector field on a surface lie at isolated points, and the zeros of complex functions in 3D space lie along curves. Similarly, the zeros of a complex section on a 4D space lie on a two-dimensional subset.

This leads to our core representation: a map $\varphi : A \to B$ is encoded as a complex-valued section $z$ on the product space $A \times B$ which is zero on the mapping surface: so $z(a, b) = 0$ if and only if $b = \varphi(a)$.

The bijectivity of $\varphi$ can also be expressed via the field $z$. For each fixed $p \in A$, the restriction $b \mapsto z(p, b)$ must vanish at exactly one point $b \in B$, and vice versa for each fixed $q \in B$. In practice, this condition is incorporated into a special choice of connection on $A$ and $B$, similar to the standard treatment of cross fields (see Section 3.1).

Using this representation, the computation of a minimal surface becomes accessible. Minimizers of the Ginzburg-Landau functional

$$\mathcal{GL}_\varepsilon(z) \coloneqq \int_{A\times B} \left(\tfrac{1}{2}|\nabla z|^2 + \frac{1}{4\varepsilon^2}(1 - |z|^2)^2\right) \mathrm{vol}_{A\times B}, \tag{2}$$

converge, as $\varepsilon \to 0$, to fields whose zero sets form minimal surfaces. Intuitively, if we were to normalize the field $z$, zeros of $z$ would become singularities where the Dirichlet energy blows up. The parameter $\varepsilon$ forces the field to approach unit length almost everywhere, leading to a large Dirichlet energy near the zeros. Thus, minimizing Dirichlet energy minimizes the area of the zero set. A more detailed discussion of this phenomenon is provided in Section 3.2.

The behavior of the Ginzburg-Landau functional and related energies on spaces with nontrivial topology is still an active area of mathematical research, so in this work we mostly restrict our attention to surfaces with *sphere-like topology*. However, we discuss the treatment of genus-zero surfaces with boundary in Section 4.7.

*Contributions.* In summary, we introduce a method to compute high quality continuous, bijective, orientation-preserving maps between genus zero surfaces with or without boundary, by minimizing the Ginzburg-Landau energy on a four-dimensional product space. This approach yields several advantages:

- Distortion minimization is achieved without barrier terms that prohibit non-bijective configurations, resulting in improved stability and simpler optimization;
- The method does not require a bijective initial map and can robustly "untangle" invalid or overlapping correspondences;
- In addition to landmark points, the algorithm naturally supports landmark curves, allowing points to slide along feature lines without explicit parameterization;
- The formulation is fully intrinsic, but preserves orientation;
- No combinatorial modifications to the input meshes are needed;

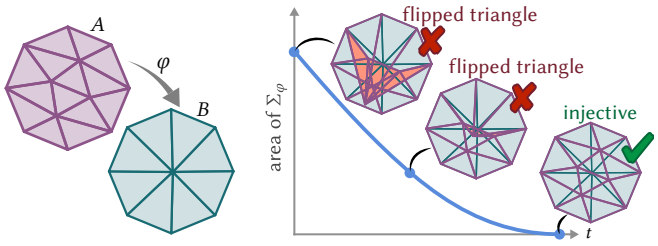


Fig. 3. As in the case of curves, a map $\varphi$ between surfaces $A$ and $B$ defines a surface $\Sigma_\varphi$ in the product space whose area encodes the distortion of $\varphi$. In particular, maps which flip triangles have larger surface area.

- The implementation relies solely on widely used operators from the tangent vector field processing toolbox. Leveraging the tensor product structure, we avoid constructing dense operators on the full product space.

### 1.1 Area Minimization and Metric Distortion

We now give a brief analysis of our distortion measure: the area of the "graph" $\Sigma_\varphi$ associated to a map $\varphi : A \to B$ (Equation 1). The area of $\Sigma_\varphi$ can be written in terms of the singular values $\sigma_1, \sigma_2$ of the Jacobian of $\varphi$, as shown in Appendix A:

$$\mathrm{Area}(\Sigma_\varphi) = \int_A \sqrt{(1+\sigma_1^2)(1+\sigma_2^2)}\,\mathrm{vol}_A\,. \tag{3}$$

This area based distortion measure is bounded by the Dirichlet energy[1] and the area distortion[2] of $\varphi$, with equality (up to the addition of a constant) if and only if $\varphi$ is conformal:

$$\int_A |\det d\varphi|\ \mathrm{vol}_A \le \mathrm{Area}(\Sigma_\varphi) - \mathrm{Area}(A) \le \int_A \tfrac{1}{2}|d\varphi|^2\ \mathrm{vol}_A\,. \tag{4}$$

*Definition 1.1.* A bijective correspondence $\varphi : A \to B$ is an *area minimizing correspondence* if $\Sigma_\varphi$ is a minimal surface.

The bound in Equation 4 also shows that maps with low total harmonic energy [Ezuz et al. 2019b, Eq. 5] have small area by noticing that $\Sigma_\varphi = \Sigma_{\varphi^{-1}}$, so a similar bound involving the area of $B$ and the Dirichlet energy of $\varphi^{-1}$ holds. Averaging the two inequalities shows that the area of $\Sigma_\varphi$ is controlled by the sum of the Dirichlet energy of $\varphi$ and its inverse. Typical distortion energies for mappings, that may also be related, are listed in [Abulnaga et al. 2023, Table 1] (see also [Poya et al. 2023]). It would also be interesting to understand relationships between area minimizing correspondences and Gromov-Wasserstein distances [Mandad et al. 2017].

### 1.2 Bijective Correspondences and Homology

We conclude with a brief comment on the constraints that we will later impose on the mapping surface $\Sigma_\varphi$. Topological properties of $\varphi$, like orientation preservation and bijectivity, are reflected in the topology of $\Sigma_\varphi$. For instance, if we intersect $\Sigma_\varphi$ with a "vertical" slice at a fixed $a \in A$, or a "horizontal" slice at a fixed $b \in B$,

$$\Sigma_\varphi \cap (\{a\} \times B) = \{(a, \varphi(a))\}, \quad \Sigma_\varphi \cap (A \times \{b\}) = \{(\varphi^{-1}(b), b)\},$$

we find a single intersection point. In particular, the projection maps $\pi_A : \Sigma_\varphi \to A$ and $\pi_B : \Sigma_\varphi \to B$ have degree $\pm 1$. The orientation preservation of $\varphi$ is encoded in the sign of the degree, which in turn is encoded in the homology class of $\Sigma_\varphi$. Thus we see that the homology class $[\Sigma_\varphi] \in H_2(A \times B)$ is not arbitrary—constant maps into $B$, for example, induce graphs that are in the same homology class of $A$ but do not satisfy the correct degree constraint for $\pi_B$.

The homology class of surfaces $\Sigma_\varphi$ arising from bijections $\varphi$ is described explicitly in Milnor and Stasheff [1974, Theorem 11.11, p.128] and is known as the *diagonal homology class* $[\Delta]$ in $A \times B$ since under the identification $A \cong B$ induced by $\varphi$, the graph $\Sigma_\varphi$ is the diagonal in $A \times A$. When $A$ and $B$ are simply-connected, the intersection of $\Sigma_\varphi$ with horizontal and vertical slices completely determines its homology class, and thus we will need to constrain the topology of our mapping surfaces to agree with $[\Delta]$ on slices.

## 2 Related Work

Computing mappings between surfaces is a long-standing challenge, and many algorithms have been proposed to address it. Our method is distinguished from existing alternatives by a combination of properties: it minimizes isometric distortion while guaranteeing orientation preservation, though it does not strictly enforce bijectivity on discrete domains. Below we contrast it to several alternative families of approaches, but a comprehensive review lies beyond the scope of this paper—see the surveys by van Kaick et al. [2011]; Tam et al. [2012] and Sahillioğlu [2020] for broader context.

### 2.1 Relaxed Map Representation

The quality of a map, including bijectivity, continuity, and orientation preservation, varies significantly with its representation.

*Registration.* One approach to surface mapping is to deform one surface to match the other. While these methods simplify distortion analysis [Beg et al. 2005; Sharf et al. 2006; Huang et al. 2008; Li et al. 2008b; Tam et al. 2012; Eisenberger et al. 2019, 2020], they offer limited guarantees of bijectivity. Moreover, obtaining a close fit between the surfaces can involve expensive optimization, and aligning shape features often requires user intervention.

Ezuz et al. [2019a,b] minimizes the harmonic energy using vertex coordinates as variables. Unlike our method, their mapping remains undefined within triangles. However, like ours, it supports arbitrary initialization and can untangle correspondences.

*Higher-Dimensional Relaxation.* Many works relax the bijectivity constraint by embedding the problem in higher-dimensional spaces, though this may introduce discontinuities, suboptimal distortion minimization, or orientation violations. Kantorovich's optimal transport formulation replaces bijections with probability measures, sacrificing bijectivity and continuity. Mandad et al. [2017] and Brifault et al. [2025] mitigate the lack of continuity by regularizing the transport plan. Functional maps [Ovsjanikov et al. 2012, 2017] represent mappings as linear operators acting on function spaces. However, converting these to continuous, bijective maps remains challenging [Melzi et al. 2019; Ren et al. 2021], and orientation preservation is non-trivial in the presence of intrinsic symmetries [Ren et al.

[1] $f^{\mathrm{Area}}(\sigma) \le 1 + \frac{1}{2}(\sigma_1^2+\sigma_2^2)$ follows from the elementary inequality $xy \le \frac{1}{2}(x^2+y^2)$
[2] $f^{\mathrm{Area}}(\sigma) \ge 1 + |\sigma_1\sigma_2|$ follows from the elementary inequality $(1+x^2)(1+y^2) \ge (1+|xy|)^2$

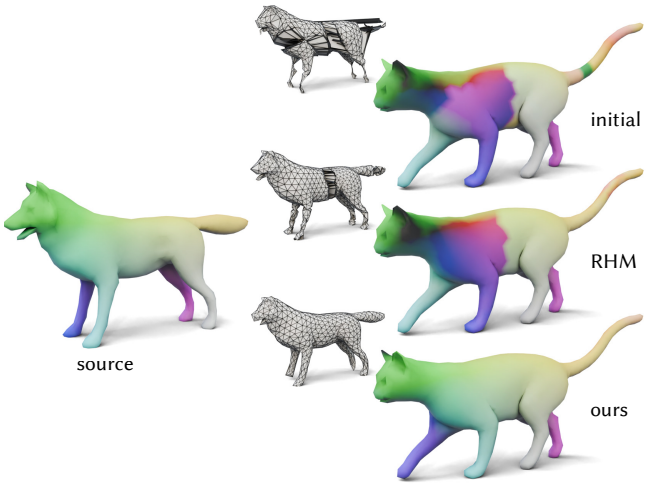


Fig. 4. The maps computed by our algorithm are constrained to preserve normal orientation of the base surface. Intrinsic algorithms (such as functional maps) often fail to distinguish between symmetric parts. Here, we initialize with a map computed via [Ren et al. 2018] that incorrectly swaps the left and right front legs. While RHM [Ezuz et al. 2019b] improves smoothness, it fails to resolve the orientation reversal. In contrast, our approach successfully untangles the map, visualized by both color and geometry transfer.

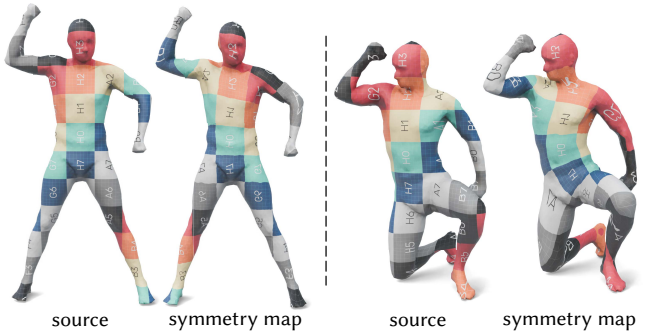


Fig. 5. Our method preserves mapping orientation by encoding surface orientation in the connection form. This enables the computation of a symmetric map from an object to itself by simply inverting the normals of the target domain.

2018, 2020; Donati et al. 2022]. Unlike exact bijective representations, functional maps enable partial mappings [Litany et al. 2016]. Solomon et al. [2012] frames correspondences as probabilistic measures, but it is subject to the same limitations. Vestner et al. [2017] share our premise of computing maps as surfaces in a product space, noting that minimal surfaces reduce distortion—though without formal characterization for surfaces. Their method approximates $\Sigma_\varphi$ via kernel density estimation, reducing the problem to linear assignment. Consequently, their mapping is vertex-defined—with no guarantee of orientation preservation—and relies on search over the space of all possible vertex matchings.

## 2.2 Strict Bijectivity

*Mapping via Common Domains.* A provably bijective map can be obtained by composing two bijective maps from surfaces into a common reference domain. Methods leveraging Tutte embeddings ensure bijectivity by mapping into simple reference domains: the plane [Kanai et al. 1997; Litke et al. 2005; Aigerman et al. 2014, 2015], periodic tilings [Aigerman and Lipman 2015], the sphere [Aigerman et al. 2017; Baden et al. 2018], or hyperbolic space [Tsui et al. 2013; Aigerman and Lipman 2016; Shi et al. 2016]. But distortion is typically controlled only for the surface-to-domain mapping, except for conformal maps [Li et al. 2008a; Baden et al. 2018] whose composition preserves conformality. Other work has attempted to address these limitations by mapping to geometrically closer domains [Schreiner et al. 2004], blending multiple conformal maps [Kim et al. 2011], or computing mesh intersections in the common domain to optimize distortion [Schmidt et al. 2019, 2020]. The latter achieves high accuracy but is computationally intensive and often converges to suboptimal solutions. Morreale et al. [2021] proposed using neural networks to encode surfaces, which alleviates some of the piecewise-linear limitation of existing mesh-based methods, but their method remains limited to pairs of surfaces which are both mapped into a common domain.

*Remeshing.* A straightforward approach to encode correspondences is to approximate both surfaces with a mesh of shared connectivity. Distortion minimization then reduces to a remeshing problem, often solved via coarse-to-fine refinement [Michikawa et al. 2001; Peng and Timalsena 2016]. Early methods prioritized mesh approximation error [Kraevoy and Sheffer 2004; Yang et al. 2020], while later works explicitly minimized distortion using volumetric mappings [Yang et al. 2018], intrinsic remeshing [Takayama 2022], or spherical parameterizations [Schmidt et al. 2023]. These methods guarantee bijectivity and orientation preservation but require bijective initialization, precluding untangling.

*Explicit Product Space.* Another line of work enforces geometric consistency by explicitly searching for a submanifold in the discrete product space. Windheuser et al. [2011a,b] initially formulated the search for an orientation-preserving diffeomorphism as an integer program on the product triangulation. Recent work have attempted to improve the scalability of this approach with novel formulations and solvers [Roetzer et al. 2022; Amrani et al. 2025; Roetzer and Bernard 2025]. However, these methods fundamentally need to define the connectivity of the product graph with massive constraint matrices, leading to prohibitive memory requirements. The resulting map is furthermore only combinatorial and doesn't define a continuous mapping within the faces.

In contrast, our method never explicitly meshes the product space, even though the implicit formulation encodes the geometric consistency constraints in the topology of the field's connection.

### 2.3 Surface Reconstruction in Higher Dimensions

Our algorithm builds a 2-dimensional surface embedded in a 4-dimensional space. Kohlbrenner et al. [2023] address the broader challenge of reconstructing a manifold from point samples in higher dimensions, but our setting is dramatically simplified because the surface of interest is defined directly as the zero set of an implicit function—not approximated from a sparse points set.

### 2.4 Minimal Surfaces

In the mathematical literature, the problem of computing minimal surfaces is often relaxed to computing minimal "currents" instead, which yields a convex relaxation of the problem in the space of generalized surfaces defined by Federer and Fleming [1960]. For a computational introduction to the theory, see Wang and Chern [2021]. This approach has been used in computer graphics in applications ranging from neural surface modeling [Palmer et al. 2022] to quad meshing [Palmer et al. 2024]. Unfortunately, when dealing with surfaces of codimension greater than 1, this relaxation is no longer tight. Indeed, even when considering two-dimensional surfaces bounded by a one-dimensional loop in $\mathbb{R}^4$, the relaxation to currents can already produce currents which have a lower area than any valid surface [Young 1963; De Lellis 2014].

In the past twenty years, zero sets of critical points of the Ginzburg-Landau energy have been shown to be related to codimension-2 minimal surfaces [Jerrard and Soner 2002; Alberti et al. 2005; De Philippis and Pigati 2022; Canevari et al. 2023], with very recent work showing that a generalization of the Ginzburg-Landau energy known as the *self-dual Yang-Mills-Higgs functional*[3] provides a phase-field approximation of the codimension-2 area functional [Pigati and Stern 2021; Parise et al. 2024]. These results are formulated in the language of complex line bundles since it provides the tools to generalize codimension-1 implicit surfaces represented as the level sets of real valued functions to codimension-2 implicit surfaces represented as the zero sets of complex sections with prescribed topology. Since we wish to compute two-dimensional minimal surfaces in a four-dimensional space of prescribed topology, we will formulate our algorithm as a discretization of aspects of this theory.

## 3 Background

In Sections 3.1 and 3.2 we explain more background about the underlying smooth mathematical theory, and we review the discretizations that we use in Sections 3.3 and 3.4. But the practically-minded reader can jump to Section 4 for a concrete description of our algorithm.

### 3.1 Implicit Representation via Complex Line Bundles

Our algorithm searches for a two-dimensional surface belonging to a specified homology class inside of a four-dimensional space. Here we describe how these surfaces can be represented using mathematical objects known as *complex line bundles*.

*Warmup.* As a simpler example, we start with codimension-two subsets of a surface $S$, *i.e.* sets of oriented points. We can try to represent these points as the shared zero level set of a pair of functions $f_1, f_2 : S \to \mathbb{R}$, or equivalently as the zero level set of a single complex function $f : S \to \mathbb{C}$. But not every collection of points can be represented as such a zero set. For instance, if $f$ is a projection onto the $xy$-plane then the zero level set is precisely the set of signed intersections between the surface and a ray shot up from the origin. But now, we can see that our point sets will always have even size, and if we add up the signs of each intersection point, we will always get a sum of 0. Indeed, the same is true for all complex functions on $S$, not just projections—for any complex function $f : S \to \mathbb{C}$, the summing the signs of the zeros yields 0.

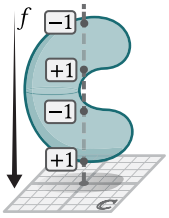


Alternatively, we could represent our points as the zeros of a vector field. Now the Poincaré-Hopf theorem guarantees that adding the indices of the zeros must give the Euler characteristic $\chi(S)$, providing a different constraint on our zero set. Although any vector field can *locally* be represented as a complex function, the space of vector fields and the space of complex functions have different global structures—and thus different constraints on their zero sets.

The space of vector fields and space of complex functions and are both examples of *complex line bundles*. Formally, a complex line bundle on a manifold $M$ is a space which associates a copy $\mathbb{C}_x$ of the complex plane to each point $x \in M$, so it locally looks like the product space $M \times \mathbb{C}$. However, it may have a different global topological structure. The analogue of a vector field on a general complex line bundles is a smooth mapping which assigns to each point $x \in M$ a value $z(x) \in \mathbb{C}_x$. Such mappings are known as *sections*. If we start with a $d$-dimensional manifold $M$, then the zero set of a section will be an an oriented submanifold of dimension $d - 2$. And just as the zeros of a complex function or vector field always sum up to a fixed constant, the zeros of a given section of a complex line bundle on $M$ always lie in the same *homology class* of $H_{d-2}(M; \mathbb{Z})$. Conversely, for any homology class in $H_{d-2}(M; \mathbb{Z})$, one can construct a complex line bundle whose the zero sets always lie in the given class.

*Connections and Curvature.* The geometry and topology of sections on a complex line bundle can be studied using *connections*, just as with vector fields. A connection $\nabla$ on a complex line bundle is a differential operator which provides a way of "parallel transporting" a value $z(x) \in \mathbb{C}_x$ along a path $\gamma$ from $x$ to $y$ to obtain a value $z(y) \in \mathbb{C}_y$. However, parallel transporting $z(x) \in \mathbb{C}_x$ along two different paths can result in different values in $\mathbb{C}_y$. Indeed, parallel transporting $z(x)$ along a closed loop may yield a different value $z$ upon returning to $x$. This failure to close up is quantified by the curvature 2-form $\Omega^\nabla$. See [Milnor and Stasheff 1974, Appendix C] for more details.

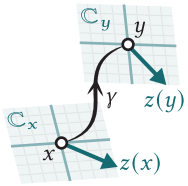


*Chern Classes.* The curvature 2-form $\Omega^\nabla$ is particularly important because it characterizes the topological class of the line bundle—and thus the topological class of our zero sets, via the *first Chern class*. The first Chern class of a complex line bundle is the cohomology class $c_1 := [\frac{1}{2\pi}\Omega^\nabla] \in H^2(M)$. It encodes the homology class of the zero sets of our sections in the following sense: for any closed

[3] it additionally includes a connection as a free variable of the optimization, penalizing the $L^2$-norm of its curvature in addition to the usual Ginzburg-Landau energy

2-dimensional subset $S \subseteq M$, the total number of oriented intersections with the $(d-2)$-dimensional zero set of a generic section is precisely $\frac{1}{2\pi}\int_S \Omega^\nabla$. Remarkably, these values are always integers, and do not depend on the specific section or connection that we started from. Taking the curvature of *any* connection on the complex line bundle always yields the same cohomology class.

*Connection Laplacians.* Just as the ordinary Laplacian measures the smoothness of scalar functions, connections can also be used to measure the smoothness of a section $z$ on a complex line bundle. The connection Dirichlet energy

$$\mathcal{E}^D(z) = \tfrac{1}{2}\int_M |\nabla z|^2 \ \mathrm{vol}_M \tag{5}$$

then gives rise to the *connection Laplacian* $L^\nabla$ as the gradient of this Dirichlet energy. When $M = A \times B$ is the product of two surfaces $A$ and $B$, the Dirichlet energy can be expressed in a factorized form

$$\mathcal{E}^D(z) = \int_A \left(\int_B |\nabla z|_B^2 \ \mathrm{vol}_B\right)\mathrm{vol}_A + \int_B \left(\int_A |\nabla z|_A^2 \ \mathrm{vol}_A\right)\mathrm{vol}_B, \tag{6}$$

where $|\nabla z|_A^2$ (*resp.*, $|\nabla z|_B^2$) denotes the norm of the covariant derivative of the section restricted to $A$ (*resp.*, $B$).

*Zero Set Distribution.* The placement of zeros of a smooth section in a complex line bundle is very closely related to the curvature of the connection. Experiments show that in low energy states of the connection Laplacian, zeros concentrate in regions of high curvature [Knöppel et al. 2013; Weißmann et al. 2014]. Recent theoretical work supports this intuition by showing that by starting with random smooth sections where the zeros are distributed according to the curvature, the corresponding heat flow will connect these random zeros to the ground state zero set [Nicolaescu and Savale 2017]. We will use this intuition to design line bundles where curvature direct the zero set, producing high quality and controllable correspondences.

## 3.2 Implicit Area Minimization

To use the implicit representation of codimension-2 surfaces, we need a way of minimizing the area of the surface associated to a section. Fortunately, we can use the Ginzburg-Landau functional to do exactly this—reducing the problem of bijective mapping to the computation of smooth complex sections on the product mesh.

*Allen-Cahn Energy.* Before describing the Ginzburg-Landau functional that our method uses, we review the main idea behind implicit area minimization in the codimension-1 case. In this case, we represent a surface as the zero set of a real-valued function $u : M \to \mathbb{R}$ and consider the Allen-Cahn energy

$$\mathcal{AC}_\varepsilon(u) := \int_M \tfrac{1}{2}|du|^2 + \frac{1}{4\varepsilon^2}(1-u^2)^2 \,\mathrm{vol}_M\,. \tag{7}$$

The main contribution of the energy comes from the double well potential which is minimized when $u \equiv \pm 1$. The Dirichlet energy, on the other hand, prevents discontinuous jumps from a region where $u \equiv +1$ to a region where $u \equiv -1$. However, in the limit as $\varepsilon \to 0$ the minimizers $u_\varepsilon$ converge to function that has a jump across a minimal surface [Modica and Mortola 1977].

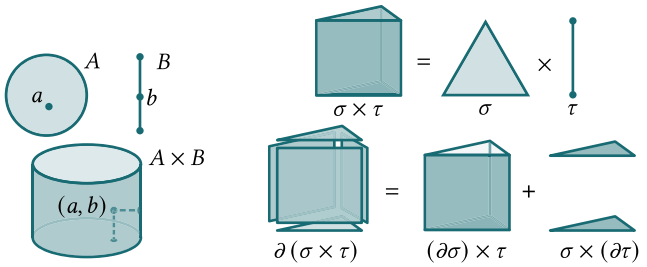


Fig. 6. *Left*: Geometrically, the product space $A \times B$ is the result of extruding $A$ along $B$. Algebraically, points of the product space correspond to ordered pairs $(a, b)$ for $a \in A$ and $b \in B$. *Right*: The boundary operator obeys a product rule on product cells.

*Ginzburg-Landau Energy.* The Ginzburg-Landau functional on a complex line bundle closely resembles the Allen-Cahn equation

$$\mathcal{GL}_\varepsilon(z) := \int_M \tfrac{1}{2}|\nabla z|^2 + \frac{1}{4\varepsilon^2}(1-|z|^2)^2 \,\mathrm{vol}_M, \tag{8}$$

The unit norm penalty is now minimized when $z \in S^1$, and so we call it the *circular well potential*. Now in the $\varepsilon \to 0^+$ limit minimizers produce $S^1$-valued harmonic maps away from a codimension-2 set of zeros. Compared to the scalar valued case, the relationship between the Ginzburg-Landau functional and the codimension-2 area of its zero set is more subtle [Parise et al. 2024]. Nevertheless, in certain situations minimizers of the Ginzburg-Landau functional still have zeros that concentrate on minimal surfaces [Lin and Riviere 1999; Canevari et al. 2023]. The intuition to keep in mind is (1) that $\varepsilon$ controls the interface width between the zero set and the set where $|z| \approx 1$ and (2) since the energy blows up as $\varepsilon \to 0^+$ the blow up should happen on a set that is area minimizing. In Section 4, we describe our algorithm that computes a correspondence as the zero set of a low energy state of the Ginzburg-Landau energy.

In our specific setting, the domain $M = A \times B$ is the product of two surfaces. The remainder of this section introduces the construction of this discrete product space (Section 3.3), and the discretization of complex line bundles (Section 3.4) required to evaluate the energy.

## 3.3 Product Meshes and Cell Complexes

Mathematically, a $d$-dimensional mesh is an object called a *cell complex*. Just as a two-dimensional mesh is made up of zero-dimensional vertices, one-dimensional edges, and two-dimensional faces, a cell complex $M = (M^0, M^1, \ldots, M^d)$ is made up of $k$-dimensional $k$-cells $\sigma_i^k \in M^k$ for all $k = 0, \ldots, d$. When the index $k$ is unnecessary we following the usual convention for meshes and call the vertex set $V_M := M^0$, the edge set $E_M := M^1$, and the face set $F_M := M^2$.

The boundary $\partial\sigma_i^k$ of a $k$-cell $\sigma_i^k$ is a collection of oriented $(k-1)$-cells, which we write as a formal sum $\partial\sigma_i^k = \sum_j \sigma_j^{k-1}$. The boundary operator can be written as a matrix $\partial_k \in \mathbb{Z}^{M^{k-1} \times M^k}$ where $(\partial_k)_{j,i}$ is 1 if $\sigma_j^{k-1}$ appears in $\partial\sigma_i^k$ with positive orientation, $-1$ if it appears with negative orientation, and zero otherwise. The transpose of these boundary operators $\mathrm{d}_{k-1} := \partial_k^\top$ are the discrete exterior derivatives in discrete exterior calculus [Desbrun et al. 2005], and we denote the space discrete differential $k$-forms by $\Omega^k(M; \mathbb{R})$.

*Product Spaces.* Given two topological spaces $A$ and $B$, the product $A \times B$ is the set of ordered pairs $(a, b)$ for $a \in B$ and $b \in B$ (Figure 6). Similarly, if we have a pair of triangle meshes $A = (V_A, E_A, F_A)$ and $B = (V_B, E_B, F_B)$, their product space is a four-dimensional cell complex whose cells are products of cells from $A$ and $B$. Its vertex set is precisely the set $V_{A\times B} = V_A \times V_B$ of pairs of vertices, while its edge set is $E_{A\times B} = E_A \times V_B \cup V_A \times E_B$, its face set is $F_{A\times B} = V_A \times F_B \cup E_A \times E_B \cup F_A \times V_B$. In general its set of $k$-cells is

$$(A \times B)^k = \bigcup_{i+j=k} A^i \times B^j. \tag{9}$$

The boundary operator on $A \times B$ obeys a product rule (Figure 6)

$$\partial(\sigma_i^k \times \sigma_j^l) = (\partial\sigma_i^k) \times \sigma_j^l + (-1)^k \sigma_i^k \times (\partial\sigma_j^l), \tag{10}$$

which we can write in matrix notation as

$$\partial_k^{A\times B} = \sum_{i+j=k} \partial_i^A \otimes \mathbb{I}_j^B + (-1)^i \mathbb{I}_i^A \otimes \partial_j^B. \tag{11}$$

## 3.4 Discrete Complex Line Bundles

We use the structure-preserving discretizations of complex line bundles with connections given by Knöppel and Pinkall [2016]. Here we summarize the construction on a cell complex $M = (V, E, F, \ldots)$. A discrete complex line bundle on $M$ consists of:

(1) a copy $\mathbb{C}_i$ of the complex plane for each vertex $i \in V$,
(2) a connection $r_{ij} \in \mathbb{C}$ of unit norm for each oriented edge $ij \in E$, satisfying $r_{ji} = r_{ij}^{-1}$. These complex numbers act as discrete parallel transport maps $r_{ij} : \mathbb{C}_i \to \mathbb{C}_j$,
(3) a curvature $\Omega_{ijk} \in \mathbb{R}$ for each face $ijk \in F$, that satisfies $d\Omega = 0$ (trivially true on a surface). It must also be compatible with the connection: for each face $ijk \in F$, we have $r_{ki} r_{jk} r_{ij} = \exp(\imath\, \Omega_{ijk})$.

Just as in the smooth setting, the topological class of a discrete complex line bundle (and the implicit surfaces it can represent) are determined by the curvature 2-form $[\frac{1}{2\pi}\Omega] \in H^2(M)$. Since the parallel transport maps only determine $\Omega$ modulo $2\pi$, we have the freedom to modify the topological class of the bundle without changing the connection—in Section 4.1, we use this ability to construct an appropriate bundle for our minimal surfaces.

Often, a discrete connection is represented a rotation angle $\rho_{ij}$ for each oriented edge, with the parallel transport maps $r_{ij} = e^{\imath\rho_{ij}}$.

*Sections and Zeros.* A discrete section $z$ is an assignment of a complex number $z_i \in \mathbb{C}_i$ to each vertex $i$. Given a section $z$, the rotation of $z$ along edge $ij$ is measured by the *angular one-form*

$$\omega_{ij}^z := \arg\left(\frac{z_j}{r_{ij} z_i}\right) \in [-\pi, \pi). \tag{12}$$

If $\omega_{ij}^z = -\pi$, then $z$ has a zero along edge $ij$. In the generic case where $\omega_{ij}^z \in (-\pi, \pi)$, Knöppel and Pinkall define the index 2-form

$$\mathrm{ind}^z := \tfrac{1}{2\pi}(d\omega^z + \Omega), \tag{13}$$

and prove that $\mathrm{ind}_{ijk}^z$ takes integer values, which can be interpreted as the sum of the indices of all zeros within face $ijk$.

Since $d\omega^z$ is a closed discrete 2-form, $[\mathrm{ind}^z] \equiv [\frac{1}{2\pi}\Omega]$ as cohomology classes in $H^2(M; \mathbb{R})$. Since $\mathrm{ind}^z$ encodes the zero set of $z$, we conclude that no matter which section $z$ we pick, the zero set always lies in the same homology class—just as in the smooth setting.

In cases where the section $z$ is clear from context, we will drop the superscript and refer to $\omega_{ij}$ and $\mathrm{ind}_{ijk}$.

*3.4.1 Discrete Levi-Civita Connection.* One important example of a discrete complex line bundle is given by the discrete Levi-Civita connection on the tangent bundle $TM$ of a triangle mesh $M$, which Knöppel et al. [2013] construct by "flattening out" the tangent space at each vertex. If we write the corner angles incident on vertex $i$ as $\theta_i^{jk}$ and let $\Theta_i := \sum_{ijk \succ i} \theta_i^{jk}$, then scaling the angles around $i$ by a factor of $\frac{2\pi}{\Theta_i}$ yields scaled corner angles $\tilde{\theta}_i^{jk} := \frac{2\pi}{\Theta_i}\theta_i^{jk}$ which sum to $2\pi$. Thus, the tangent space of vertex $i$ can naturally be associated with the complex plane $T_iM \cong \mathbb{C}$. Each halfedge $ij$ emanating from vertex $i$ can be associated with a unit complex number $\hat{h}_{ij} \in T_iM$ giving its direction in the tangent space. Now, after choosing any coordinates on $T_iM$ and $T_jM$, the discrete Levi-Civita connection is

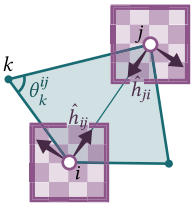


$$r_{ij}^{LC} := -\hat{h}_{ji}\hat{h}_{ij}^{-1}, \tag{14}$$

which maps $T_iM$ to $T_jM$ while preserving the direction of edge $ij$. To compute the complex line bundle, we assign curvature to each face $ijk$ based on its rescaled corner angles

$$\Omega_{ijk}^{LC} = \tilde{\theta}_i^{jk} + \tilde{\theta}_j^{ki} + \tilde{\theta}_k^{ij} - \pi. \tag{15}$$

These curvatures are compatible with $r^{LC}$ and satisfy the Gauss-Bonnet theorem $\sum \Omega_{ijk}^{LC} = 2\pi\chi(M)$.

*3.4.2 Trivial Connections.* Starting from any initial connection $r^0$ with curvature $\Omega^0$, we can modify the curvature by solving a Poisson equation. If we are given another discrete 2-form $\Omega \in \Omega^2(M; \mathbb{R})$ on a simply connected triangle mesh $M$ that has the same integral, *i.e.*,

$$\sum_{ijk\in F} \Omega_{ijk}^0 = \sum_{ijk\in F} \Omega_{ijk}, \tag{16}$$

we can find a discrete 1-form $\alpha \in \Omega^1(M; \mathbb{R})$ that satisfies

$$\mathrm{d}_1\alpha = \Omega - \Omega^0. \tag{17}$$

Using the Hodge decomposition, this amount to solving a Poisson equation on the dual mesh[4]. The discrete connection

$$r_{ij} = e^{\imath\alpha_{ij}} r_{ij}^0 \tag{18}$$

is then compatible with the prescribed curvature 2-form $\Omega$. If we concentrate all of the prescribed curvature on a finite number of points this is the trivial connections algorithm of Crane et al. [2010]. We use this procedure in Section 4.1 to construct a connection with the correct curvature on each input surface, and in Section 4.4 to compute an initialization from a noisy correspondence.

[4] letting $\mathsf{L}^- = \mathrm{d}_1 *_1^{-1} \mathrm{d}_1^\top$ we have that $\alpha = *_1^{-1}\mathrm{d}_1^\top\beta$ where $\beta$ solves $\mathsf{L}^-\beta = \Omega - \Omega^0$

*3.4.3 Finite Element Space.* To define continuous maps between triangle meshes, we need to be able to evaluate a section $z$ not only at mesh vertices, but also inside triangles. To do so, we use the finite element space for complex line bundles on surfaces constructed by Knöppel and Pinkall [2016] and Liu et al. [2016]. In our setting, the basis function for vertex $i$ within triangle $ijk$ can be expressed as:

$$\phi_i(p) = b_i(p) \exp\left[-\imath\left(\int_{p_i \to p} \rho_{ijk}\right)\right], \tag{19}$$

where $b_i(p)$ is the barycentric coordinate of point $p$ with respect to vertex $i$, $\rho_{ijk}$ is the connection 1-form interpolated using Whitney interpolation [Whitney 1957; Desbrun et al. 2005], and $p_i \to p$ is the straight line connecting $i$ to $p$.

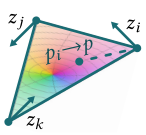


Using this finite element space on a triangle mesh $S$, Knöppel et al. [2013, §6.1.1] give discretizations of the connection Laplacian $L_S^\nabla \in \mathbb{C}^{V\times V}$ and the mass matrix $M_S^\nabla \in \mathbb{C}^{V\times V}$, which depend on the parallel transport maps $r_{ij}$ and face curvatures $\Omega_{ijk}$. We describe the full constructions in the supplementary material.

*3.4.4 Product Space Finite Elements.* Recall that the complex field $z : A \times B \to \mathbb{C}$ is defined at each vertex and stored as a complex matrix $Z$ of dimensions $|V_A| \times |V_B|$. In order to interpolate these values over the 4-cells which make up the product space $A \times B$, we define basis functions by taking tensor products of the basis functions $\phi^A$ and $\phi^B$ defined on surfaces $A$ and $B$ respectively.

Any 4-cell $\sigma^4_{ijk,abc} \in F_{A\times B}$ is the product of triangles $ijk \in F_A$ and $abc \in F_B$. The interpolated field at a point $(p, q) \in \sigma^4_{ijk,abc}$ is:

$$z(p,q) = \sum_{\substack{u\in\{i,j,k\}\\ v\in\{a,b,c\}}} Z_{u,v}\, \phi^A_u(p)\, \phi^B_v(q). \tag{20}$$

Note that if we fix $q$ and allow $p$ to vary, we recover the standard surface interpolant on $A$ up to a global rotation.

This interpolation scheme enables the definition of a connection Laplacian and the precise localization of zeros within each cell.

*3.4.5 Connection Laplacian.* We define a connection Laplacian on the product space by discretizing the Dirichlet energy (Equation 6). The discrete Dirichlet energy of a discrete section $Z \in \mathbb{C}^{V_A\times V_B}$ on $A \times B$ can be written in terms of left and right multiplication by the connection Laplacian and mass matrix defined on each manifold:

$$\mathcal{E}^D(Z) = \frac{1}{2}\langle Z, L_A^\nabla Z (M_B^\nabla)^\top\rangle_{\mathbb{C}} + \frac{1}{2}\langle Z, M_A^\nabla Z (L_B^\nabla)^\top\rangle_{\mathbb{C}}, \tag{21}$$

Where we write $\bullet^\top$ for the ordinary transpose, $\bullet^\dagger$ for the conjugate transpose, and use the inner product $\langle A, B\rangle_{\mathbb{C}} = \mathrm{Re}\,\mathrm{tr}[A^\dagger B]$ for complex matrices. For a full derivation, see Appendix B.

Similarly, the connection mass matrices $M_A^\nabla \in \mathbb{C}^{V_A\times V_A}, M_B^\nabla \in \mathbb{C}^{V_B\times V_B}$ determine the $L^2$ energy of product space sections via:

$$\int_{A\times B} |z|^2 \,\mathrm{vol}_{A\times B} = \langle Z, M_A^\nabla Z (M_B^\nabla)^\top\rangle_{\mathbb{C}}, \tag{22}$$

and the scalar mass matrices $M_A \in \mathbb{R}^{V_A\times V_A}, M_B \in \mathbb{R}^{V_B\times V_B}$ determine the $L^2$ energy of real functions $F \in \mathbb{R}^{V_A\times V_B}$ on $A \times B$:

$$\int_{A\times B} f(a,b)^2 \,\mathrm{vol}_{A\times B} = \langle F, M_A F (M_B)^\top\rangle_{\mathbb{R}}. \tag{23}$$

*3.4.6 Encoding the Correct Homology Class.* In order to compute bijections between $A$ and $B$, we must prescribe a specific curvature on the product space $A \times B$. Recall that the map $\varphi : A \to B$ is defined implicitly: for any vertex $i \in V_A$, the image $\varphi(i)$ is the zero of the section $z$ restricted to the slice $\{i\} \times B$. This defines a section $z^{(i)}$ on $B$ where $z^{(i)}(p) = z(i, p)$ for any point $p \in B$. Since vertex $i$ should map to a single point on $B$, this restricted section must contain exactly one zero. Recalling that the total number of signed zeros on a the slice is $\frac{1}{2\pi}\int \Omega$ (Section 3.1), this implies the curvature of the bundle restricted to any copy of $B$ must integrate to $2\pi$. By symmetry, the curvature of the bundle restricted to any slice on $A$ must also integrate to $2\pi$. We satisfy these constraints by starting with two base connections on $A$ and $B$ separately, with curvatures $\Omega^A$ and $\Omega^B$ respectively, each integrating to $2\pi$, and combining them to form a product connection on $A \times B$. This construction guarantees that the curvature associated to this resulting connection restricts to $\Omega^A$ on any horizontal slice and to $\Omega^B$ on any vertical slice, as desired. Formally, when $A$ and $B$ are topological spheres, this product connection and curvature precisely encode the diagonal class $[\Delta]$, which is the homology class in which all bijections lie.

## 4 Algorithm

Our algorithm for computing correspondences has four key steps:

(1) Designing a discrete connection on a complex line bundle such that each slice $A$ and $B$ of the product mesh contains exactly one zero (Section 4.1);
(2) Initializing the complex field based on input correspondences (Section 4.4);
(3) Minimizing the discrete Ginzburg–Landau energy for $z$ (Section 4.2);
(4) Evaluating the correspondences by identifying the zero set of $z$ after interpolation using a finite element basis (Section 4.3).

To extend the algorithm's applicability, we introduce several straightforward modifications: support for surfaces with boundaries (Section 4.7), support for point or curve landmarks (Section 4.6), and a simple coarse-to-fine acceleration strategy (Section 4.5).

Comprehensive pseudocode can be found in the supplementary material.

*Notation.* Throughout, $A = (V_A, E_A, F_A)$ and $B = (V_B, E_B, F_B)$ are closed genus zero triangle meshes, with edge lengths $\ell_A : E_A \to \mathbb{R}$ and $\ell_B : E_B \to \mathbb{R}$. We write the discrete section on $A \times B$ as a matrix $Z \in \mathbb{C}^{V_A\times V_B}$. We write $\bullet^\top$ for the ordinary transpose, $\bullet^\dagger$ for the conjugate transpose, and use the inner products $\langle A, B\rangle_{\mathbb{C}} = \mathrm{Re}\,\mathrm{tr}[A^\dagger B]$ for complex matrices, and $\mathrm{tr}[A^\top B]$ for real matrices. On each mesh $S$ we let $L_S^\nabla, M_S^\nabla \in \mathbb{C}^{V_S\times V_S}$ be the connection Laplacian and mass matrix for complex sections on $S$ (Section 3.4.5), and let $M_S \in \mathbb{R}^{V_S\times V_S}$ be the real diagonal mass matrix for scalar functions.

### 4.1 Constructing Surface Connections

We begin by finding connections on $A$ and $B$ with total curvature $2\pi$, which will allow us to find sections on $A$ or $B$ with a single zero—see Section 3.4.6 for more discussion of the curvature constraint.

We can construct such a connection on each mesh by taking advantage of the $2\pi$ ambiguity in the curvature of a discrete complex

line bundle. We start with the trivial bundle $\mathbb{C}_i := \mathbb{C}$ and trivial parallel transport maps $r^0_{ij} := 1$. We set the initial curvature $\Omega^0$ to zero everywhere except on an arbitrarily chosen face $f \in F$ where we set $\Omega^0_f := 2\pi$. This is a discretization of the skyscraper bundle [Knöppel 2020, §7.2], which has the correct total curvature, but contains no geometric information. To obtain a more meaningful connection, we solve a linear system as in Section 3.4.2 to obtain a curvature of half the Gaussian curvature $\Omega^{LC}_{ijk}$ on each face (Equation 15), *i.e.* we solve for a new connection $r_{ij}$ with curvature $\Omega_{ijk} = \frac{1}{2}\Omega^{LC}_{ijk}$.

Remark 1. *Any other discrete connection with the same curvature 2-form can, in principle, be used instead. In Appendix D we describe two alternative constructions that are directly related to tangent vectors and the Levi-Civita connection.*

## 4.2 Discrete Ginzburg-Landau Minimization

We discretize the Ginzburg-Landau energy, Equation 8, via a finite-element Dirichlet energy and a lumped discretization of the circular well penalty term. This energy can be minimized with L-BFGS using the expressions for the energy and gradient provided below.

*Dirichlet Energy.* As explained in Section 3.4.5, the Dirichlet energy on $A \times B$ can be assembled from the connection Laplacians and mass matrices of $A$ and $B$ as follows

$$\mathcal{E}^D(Z) = \tfrac{1}{2}\left\langle Z\overline{M}^{\nabla}_B, L^{\nabla}_A Z\right\rangle_{\mathbb{C}} + \tfrac{1}{2}\left\langle \overline{M}^{\nabla}_A Z, Z L^{\nabla}_B\right\rangle_{\mathbb{C}}. \tag{24}$$

*Circular Well Potential.* We discretize the penalty term at the vertices, weighted by their barycentric dual volumes in $A \times B$:

$$\mathcal{W}(Z) = \sum_{(i,j)\in V_A\times V_B} (1 - |Z_{i,j}|^2)^2 M_{i,i} M_{j,j}. \tag{25}$$

To write $\mathcal{W}$ in matrix notation, we define the norm deviation matrix $U \in \mathbb{R}^{V_A\times V_B}$ by $U_{i,j} = |Z_{i,j}|^2 - 1$, so that $\mathcal{W}(Z) = \langle U, M_A U M_B^{\top}\rangle_{\mathbb{R}}$. This vertex-based discretization minimizes the stencil size—and computational cost—of function and gradient evaluations.

*Ginzburg-Landau Energy.* The discrete energy of a section $Z$ is

$$\begin{aligned}\mathcal{GL}_\lambda(Z) &:= \mathcal{E}^D(Z) + \tfrac{\lambda}{4}\mathcal{W}(Z)\\ &= \tfrac{1}{2}\left\langle Z\overline{M}^{\nabla}_B, L^{\nabla}_A Z\right\rangle_{\mathbb{C}} + \tfrac{1}{2}\left\langle \overline{M}^{\nabla}_A Z, Z L^{\nabla}_B\right\rangle_{\mathbb{C}}\\ &\quad + \tfrac{\lambda}{4}\langle U M_B, M_A U\rangle_{\mathbb{R}},\end{aligned} \tag{26}$$

where the variable $\lambda$ is called the *Ginzburg-Landau parameter*. Its gradient can also be written as the following matrix in $\mathbb{C}^{V_A\times V_B}$:

$$\nabla_Z \mathcal{GL}_\lambda = L^{\nabla}_A Z \left(M^{\nabla}_B\right)^{\top} + M^{\nabla}_A Z \left(L^{\nabla}_B\right)^{\top} + \lambda\left(M_A U M_B^{\top}\right) \odot Z, \tag{27}$$

where $\odot$ denotes the element-wise product.

The choice of Ginzburg-Landau parameter significantly influences the quality of the correspondence. If $\lambda$ is below the minimal eigenvalue of $L^{\nabla}_{A\times B}$ then the only critical point of the energy is $z \equiv 0$ (Lemma C.1). But if $\lambda$ is too large, the circular well potential repels zeros away from the vertices. Since the Dirichlet term carries a comparatively low weight, the field $z$ tends to remain close to its initial value.

One simple strategy is to set $\lambda$ relative to the smallest eigenvalue $\lambda_{A\times B}$ of $L^{\nabla}_{A\times B}$. In the smooth setting, this eigenvalue decomposes as $\lambda_{A\times B} = \lambda_A + \lambda_B$, where $\lambda_A$ and $\lambda_B$ are the smallest eigenvalues of the connection Laplacians on $A$ and $B$. While this equality only holds approximately in the discrete setting due to differing mass matrix discretizations (see Appendix C for details), it remains a good estimate. Experimentally, we find $\lambda \approx 100\lambda_{A\times B}$ provides near-optimal results across a variety of meshes (Figure 7).

However, for challenging initializations, a smaller value of $\lambda$ is often necessary to facilitate the smoothing of correspondences and ensure correct matching between surface features. For example, in Figures 4 and 20, the optimization is performed in two stages: first, with $\lambda = 10\lambda_{A\times B}$ to align large-scale features, and second, with $\lambda = 100\lambda_{A\times B}$ to refine the correspondences and address high-frequency matching.

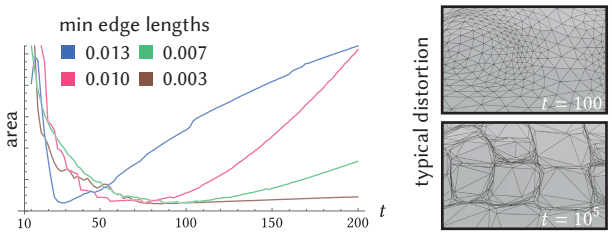


Fig. 7. An experimental search for the optimal choice of Ginzburg-Landau parameter $\lambda = t\,\lambda_0$ where $\lambda_0 = \lambda_A + \lambda_B$. Since the area depends on the distortion between the surfaces, to consistently analyze the optimal parameter choice we normalize the area by the maximal area measured for each example separately. Large values of $\lambda$ cannot produce smooth correspondences unless the mesh has a sufficiently high resolution, as indicated by having a small minimum edge length.

## 4.3 Evaluating the Correspondence

By construction (Section 4.1), for any vertex $v \in V_B$, the $v$-th column of $Z$ is a section $z^{(v)}$ on $A$— simply denoted $z$ below—with a single zero (see Figure 8). To locate this zero, we first compute the integer-valued index 2-form $\mathrm{ind}^z$ using Equation 13. This form takes a non-zero value on exactly one face $ijk \in F_A$, where the zero lies.

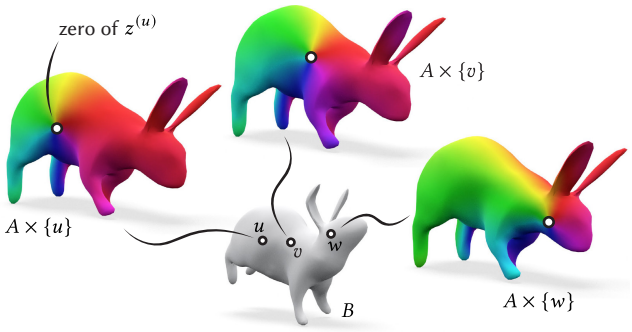


Fig. 8. To evaluate the map at a vertex $u \in V_B$, we take the $u$'th column of $Z$ as a section $z^{(u)}$ on $A$. We can visualize this section as a complex function in a chart that covers all but a single face of $A$. Now the image $\varphi(u)$ of vertex $u$ is given by the location of the zero on $A$.

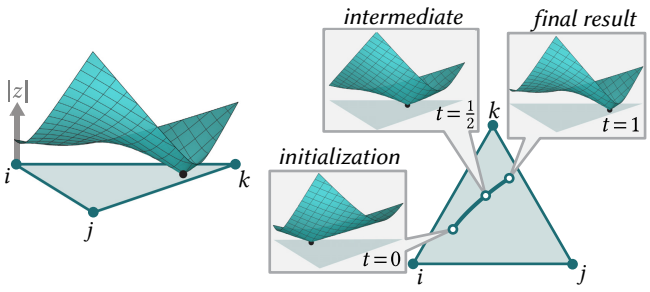


Fig. 9. *Left:* In order to evaluate the correspondence, we have to find a root of the interpolated section $z$ within some triangle $ijk$. Unfortunately the interpolant is nonlinear and non-convex, so solving directly with Newton's method will not necessarily find the desired root. *Right:* Instead, we interpolate the face curvature from flat at $t = 0$—in which case the interpolant becomes linear—to the full curvature $\Omega_{ijk}$ at $t = 1$, solving with Newton's method at each time step.

We then compute the barycentric coordinates $(b_j, b_k)$ of the zero of the interpolation within $ijk$ (Section 3.4.3) by solving the following system:

$$(1-b_j-b_k)|z_i|+b_j|z_j|e^{\imath\left(b_k\Omega+\omega_{ij}\right)}+b_k|z_k|e^{-\imath\left(b_j\Omega+\omega_{ki}\right)}=0. \quad (28)$$

Directly solving Equation 28 with Newton's method is unlikely to yield the solution within triangle $ijk$. Instead, we adopt a homotopy continuation approach, and deform the geometry from a flat triangle ($t = 0$) to the target curvature $\Omega_{ijk}$ ($t = 1$), tracking the zero as it moves. We use Newton's method to find the zero at each time step (see Figure 9). Pseudocode is provided in the supplementary material.

To prevent the zero from "escaping" the triangle, we ensure that the zero never hits an edge during the interpolation, *i.e.* that $\omega^z_{ij}(t), \omega^z_{jk}(t), \omega^z_{ki}(t) \in (-\pi, \pi)$ for $0 \le t < 1$. This is achieved using the following linear interpolation:

$$\omega^z_{ij}(t) = \omega^z_{ij} + (1-t)\tfrac{1}{3}\left(\Omega_{ijk} - 2\omega^z_{ij} + \omega^z_{jk} + \omega_{ki}\right) \quad (29)$$

$$\omega^z_{jk}(t) = \omega^z_{jk} + (1-t)\tfrac{1}{3}\left(\Omega_{ijk} + \omega^z_{ij} - 2\omega^z_{jk} + \omega_{ki}\right) \quad (30)$$

$$\omega^z_{ki}(t) = \omega^z_{ki} + (1-t)\tfrac{1}{3}\left(\Omega_{ijk} + \omega^z_{ij} + \omega^z_{jk} - 2\omega^z_{ki}\right) \quad (31)$$

$$\Omega_{ijk}(t) = t\Omega_{ijk} \quad (32)$$

One can easily check that $\omega^z_{ij}(t) \in (-\pi, \pi)$ for $0 \le t < 1$ and that for all time $t$ the index form is preserved, ensuring that the zero remains inside triangle $ijk$:

$$\forall t, \quad \omega^z_{ij}(t) + \omega^z_{jk}(t) + \omega^z_{ki}(t) + \Omega_{ijk}(t) = 2\pi\,\mathrm{ind}^z_{ijk}\,. \quad (33)$$

*Edge-Edge Intersections.* The section $Z$ encodes not only the locations of vertices, but also the entire embedding of the edges under the correspondence (Figures 1 and 10). The intersection of an edge of mesh $A$ and an edge of mesh $B$ may be calculated by locating a zero inside of an "edge-edge" face in the product space $A \times B$. Since these faces have zero curvature, this can be done by solving a single quadratic equation. The equation and its derivation are found in Appendix E.

*Evaluation Inside a Face.* We can also evaluate the correspondence for any point $p$ within triangle $ijk \in F_B$ using the basis functions in Section 3.4.3. Using the complex weights $\phi_\bullet(p)$ for each vertex $i, j, k$ of $B$ given by Equation 19, we define a section $z^{(p)}$ on $A$ as the linear combination of the columns of $Z$ corresponding to vertices $i, j, k$. Solving for its zero exactly as above yields the image of $p$ on $A$.

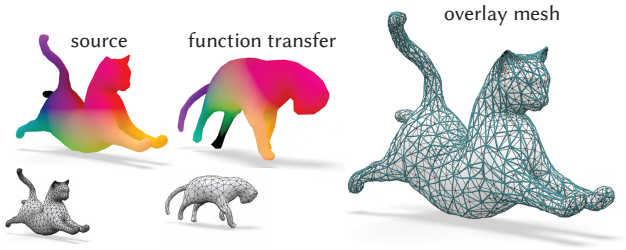


Fig. 10. Our algorithm produces bijective and continuous mappings even when the meshes exhibit substantially different connectivity and sampling. The zero set of the field may place multiple vertices of $A$ within the same face of mesh $B$, while preserving low distortion and bijectivity.

## 4.4 Initialization

Our method can start from a (potentially noisy) input function given as vertex-face maps $\varphi : V_A \to F_B$ and $\psi : V_B \to F_A$. To use these in our algorithm, we must convert them into an initial complex section $z : A \times B \to \mathbb{C}$ whose zero set approximates the graphs of $\varphi$ and $\psi$.

*Strategy.* We leverage the principle that low-energy states of a connection Laplacian concentrate their zeros in regions of high curvature (Section 3.1). To do so, we construct a specific connection $r^{\varphi,\psi}$ on $A \times B$ that concentrates all its curvature exactly along the graphs of $\varphi$ and $\psi$. We then define our initial section $z$ as the smallest eigenvector of the associated Laplacian. Because this bundle lies in the same homology class as the one constructed in Section 4.1, this section serves as a valid initialization for the following Ginzburg-Landau minimization.

*Construction.* We build $r^{\varphi,\psi}$ slice-by-slice using the trivial connections algorithm from Section 3.4.2. For each vertex $i_A \in V_A$, we compute a connection $r^{i_A}$ on $B$ that concentrates $2\pi$ curvature into the single face $\varphi(i_A)$. Then, for every edge of the form $i_A \times e_B$ in the product space, we assign $r^{\varphi,\psi}_{i_A \times e_B} = r^{i_A}_{e_B}$ Symmetrically, for each $i_B \in V_B$, we run trivial connection on $A$, concentrating curvature in face $\psi(i_B)$ and assign this connection to product space edges of the form $e_A \times i_B$. This results in a fully defined connection on the product space whose curvature is concentrated to the graph of the input maps. The zero set of the smallest eigenvector of the associated Laplacian acts as a good representation of the input correspondences (see Figure 17).

*Laplacian.* The connection constructed above does not factor as the product of independent connections on $A$ and $B$, which would normally require assembling the full four-dimensional finite element Laplacian. Instead, we consider a simpler discretization of the

Dirichlet energy using a lumped discretization for the outer integrals of Equation 6 while retaining a finite-element discretization for the inner integrals over the 2D slices. Explicitly, for every vertex $i_A \in V_A$, let $L_B^{\nabla,i_A}$ denote the finite element Laplacian of $r^{i_A}$ on $B$, and let $z^{(i_A)}$ be the restriction of section $Z$ to the slice. Using similar notation for the slices fixed at $i_B$, we can express the Dirichlet energy of the connection $r^{\varphi,\psi}$ as:

$$\mathcal{E}^D_{\varphi,\psi}(Z) = \sum_{v\in V_A} (M_A)_{v,v}\langle L_B^{\nabla,v} z^{(v)}, z^{(v)}\rangle_{\mathbb{C}} + \sum_{v\in V_B} (M_B)_{v,v}\langle L_A^{\nabla,v} z^{(v)}, z^{(v)}\rangle_{\mathbb{C}}, \tag{34}$$

and so the blocks in the product-space connection Laplacian are given by the surface connection Laplacians on the individual slices weighted by the dual area of the vertex on the complementary mesh.

The smallest eigenvalue can then be computed via an iterative matrix-free solver such as LOBPCG [Knyazev et al. 2007].

*Initialization from Distributions.* Rather than starting from a fixed map $\varphi : V_A \to F_B$, one can also start from a *distribution* on mesh $B$ associated to each vertex of $A$, provided *e.g.* by a functional map or an optimal transport plan. The procedure is almost exactly the same: the only change is that instead of concentrating all curvature in one face $\varphi(i_A)$, we spread curvature over $B$ proportional to the density associated to vertex $i_A$.

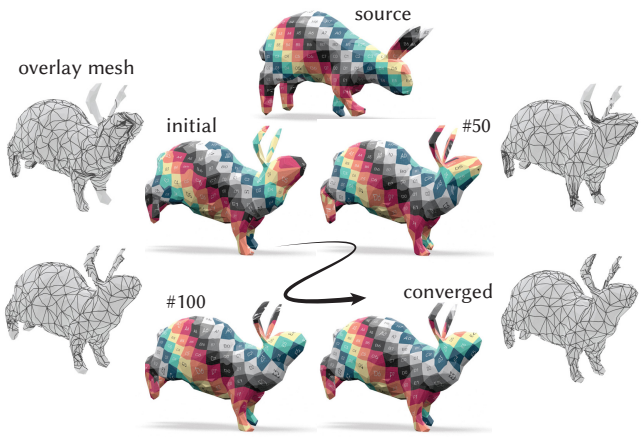


Fig. 11. High quality correspondences can be computed even on very coarse tessellations. On this example, with $|V_A| = 474$ and $|V_B| = 502$, the solution is computed in 21 seconds. The robustness to tessellation density is a consequence of the finite-element representation of the implicit surface: the complex section $z$ smoothly extends to the interior of the faces and does not require fine sampling to represent a continuous bijection.

## 4.5 Multiresolution Hierarchy

Our method can be trivially accelerated using a multiresolution hierarchy of meshes. We first compute an optimal complex section $z^{(0)}$ on coarsened meshes $A^{(0)}$ and $B^{(0)}$ by minimizing the Ginzburg-Landau energy. To transfer this solution to the fine meshes $A$ and $B$,

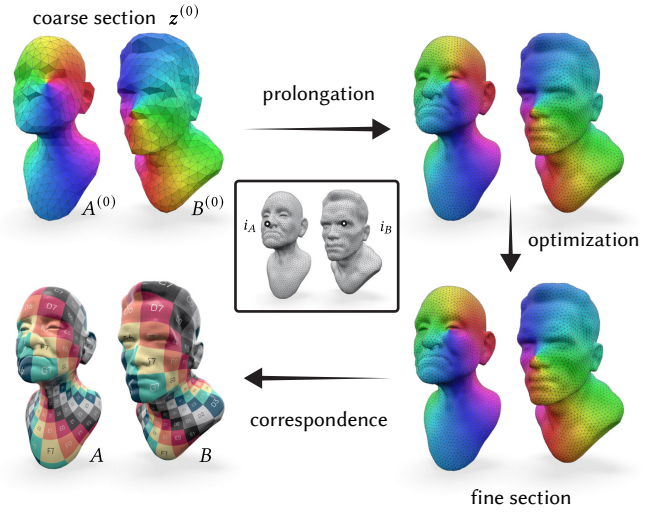


Fig. 12. Our implicit representation of the map can be prolongated to a finer mesh by closest point interpolation of the complex section. Therefore, we can initialize our optimization from a coarse solution, which already provides a very good approximation of the desired map. To visualize the sections, we plot the complex values on the slices $\{i_A\}\times B$ and $A\times\{i_B\}$, and to visualize the correspondence we transfer a conformal parameterization from $B$.

we distinguish between two cases depending on the bundle structure.

*Direct Upsampling.* For the surface connections constructed in Section 4.1, the fibers $\mathbb{C}_i$ are globally identified with the complex plane. This allows us to transfer the complex values of the section directly without explicit change of basis (see Figure 12). We upsample $z^{(0)}$ by closest point projection: for every vertex pair $(i_A, i_B)$ in the fine product mesh $A \times B$, we find the closest geometric points $p_A \in A^{(0)}$ and $p_B \in B^{(0)}$ and evaluate $z^{(0)}$ at $(p_A, p_B)$, using interpolation defined in Section 3.4.4. This upsampled section is then used as initialization for the fine-scale Ginzburg-Landau minimization, or for the iterative eigenvector solver. In practice, this can provide a 12× speedup in computation time (Section 6). Note that for this transfer to be as accurate as possible, the face selected in Section 4.1, which concentrates the curvature in the coarse mesh, should contain the corresponding face in the fine mesh.

*Geometric Initialization.* When working with more general connections, such as those based on tangent vector fields (Appendix D), direct transfer of complex values might not be possible. In these cases, we instead transfer the *geometric correspondence*, by first computing the coarse correspondences $\varphi^{(0)}$ and $\psi^{(0)}$ from the section $Z^{(0)}$, as described in Section 4.3. We then construct initial vertex-to-face maps between the fine meshes $A$ and $B$ by composing these coarse correspondences with closest point projections. The algorithm then follows the initialization and minimization procedures from Sections 4.2 and 4.4 to obtain refined correspondences. This strategy significantly reduces the number of Ginzburg-Landau iterations required on the fine meshes.

### 4.6 Landmarks

The Ginzburg-Landau energy can be modified to accommodate a user-specified set of landmarks $\mathcal{L} = \{(l_i^A, l_i^B)\}_{i=1}^p$ (Figure 13). Since the graph of the correspondence passes through a point $p \in A \times B$ if and only if $z(p) = 0$, we explicitly encourage this constraint by penalizing non-zero values at the landmarks. Hence, we consider the specifications as soft-constraints.

We replace the standard circular well potential with a spatially varying version, defined by a non-negative potential function $V : A \times B \to \mathbb{R}$ having isolated zeros at $\mathcal{L}$:

$$\mathcal{W}(Z) := \frac{1}{4} \int_{A\times B} (V(p) - |z(p)|^2)^2 \ \mathrm{vol}_{A\times B},$$

In our implementation, we construct $V$ using Gaussians centered at the landmarks:

$$V(p) := \min_{(p_A,p_B)\in\mathcal{L}} \left(1 - \exp\left(-\frac{1}{2\sigma_A^2} d_A(p,p_A)^2 - \frac{1}{2\sigma_B^2} d_B(p,p_B)^2\right)\right),$$

where $\sigma_A, \sigma_B > 0$ determine the kernel width, and $d_A(\cdot,\cdot)$ and $d_B(\cdot,\cdot)$ denote the geodesic distance on $A$ and $B$, respectively. Finer granularity over the pinning potential can be achieved by specifying a landmark dependence choice of $\sigma_A$ and $\sigma_B$.

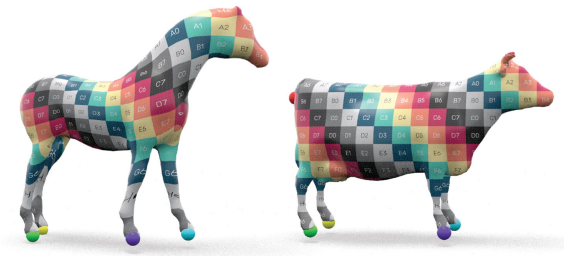

Fig. 13. Landmarks are specified using a pinning potential that has local minima at the specified points in the product mesh.

This modification of the Ginzburg-Landau functional is known as *singularity pinning* in the physics literature[5], and Aftalion et al. [2001] proved that in the small $\varepsilon$ limit that zeros of the Ginzburg-Landau minimizers (in a two-dimensional background) are "pinned" at the local minimizers of $V$. For this attraction to be effective in practice, the width of the Gaussian kernel must be sufficiently large. Since $V$ modifies the energy gradient only locally, a zero far outside the well, where $V \approx 1$, will not be affected by the potential.

The pinning potential can also be modified to support unparameterized curve to curve correspondences (Figure 14). For a pair of curves $\gamma_A \subset A$ and $\gamma_B \subset B$ we modify $V$, and replace the point-to-point distances with point-to-curve distances $d_A(p,\gamma_A)^2$ and $d_B(p,\gamma_B)^2$. The resulting contribution encourages the section to vanish along the two-dimensional surface $\gamma_A \times \gamma_B \subset A \times B$.

### 4.7 Surfaces with Boundary

In Figure 15 we compute a correspondence between two topological disks by filling in each boundary component with a disk to obtain topological spheres. To ensure that the correspondence identifies the boundaries together we also include a curve-to-curve singularity pinning potential for them. We use a triangle fan to fill in the boundary components for simplicity, though we remark that a more isotropic remeshing of the boundary disk provides a better finite element space that may be necessary if the aspect ratio of these triangles becomes too large.

[5]it has been observed experimentally that the relevant singularities configurations are related to material defects that can be encoded in $V$

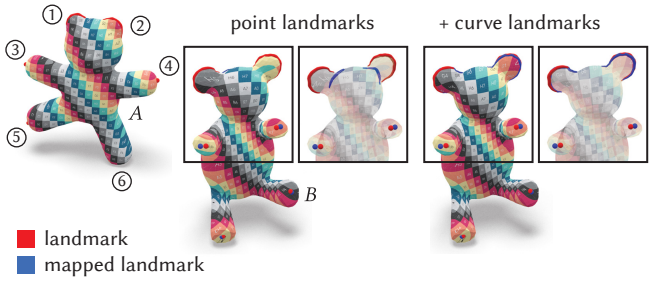


Fig. 14. Enforcing curve to curve correspondences where the points are allowed to slide along the specified curves is notoriously challenging. Our implicit representation can elegantly handle these as soft-constraints by modifying the circular well potential. The red points and curves are the specified landmarks, and the blue points and curves visualize their image in $B$ under the computed correspondence.

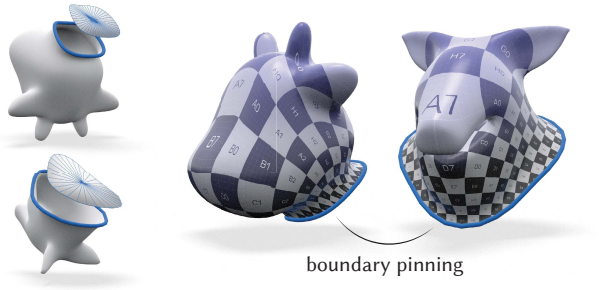


Fig. 15. Our method handles surfaces with boundary by filling in a disk and pinning the boundaries together.

### 4.8 Intrinsic Triangulations

Since the Ginzburg-Landau functional is defined solely in terms of intrinsic quantities, we can improve accuracy by using the intrinsic Delaunay triangulation (iDT) [Bobenko and Springborn 2007]. While our discretization is already robust to triangulation quality and sampling density (Figure 11), on severe examples, flipping to the intrinsic Delaunay triangulation produces lower distortion correspondences (Figure 16). Moreover, the intrinsic Delaunay triangulation is known to improve the quality of PDE based discretizations of geodesic distance that we use to constrain landmarks. We used the integer coordinates introduced in [Gillespie et al. 2021] to represent intrinsic triangulations, and applied this intrinsic preprocessing throughout. Further improvements to the triangulation quality can potentially be achieved by further considering intrinsic Delaunay refinement [Sharp et al. 2019; Gillespie et al. 2021], especially in situations where the input resolution is too coarse.

## 5 Results and Evaluation

In this section, we demonstrate the fundamental features and applications of our approach and compare with state-of-the-art techniques for computing bijective correspondences. As surface correspondences are general purpose tools, our method has further applications beyond the operations we consider here (mesh transfer, surface interpolation, *etc.*).

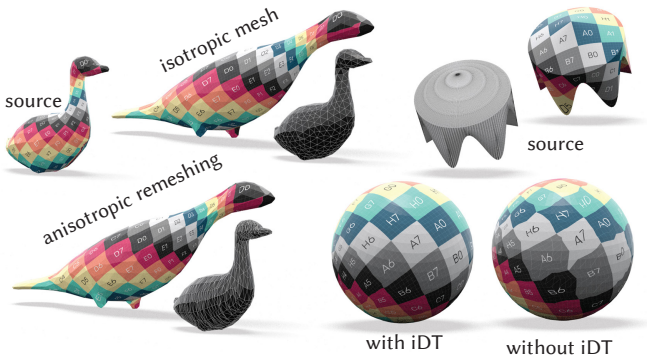


Fig. 16. Our method is robust to triangulation quality, computing the same correspondence when the source is triangulated with both an isotropic and highly non-isotropic mesh *(left)*. Using the intrinsic Delaunay triangulation, our method produces smooth correspondences even in the most extreme situations *(right)*.

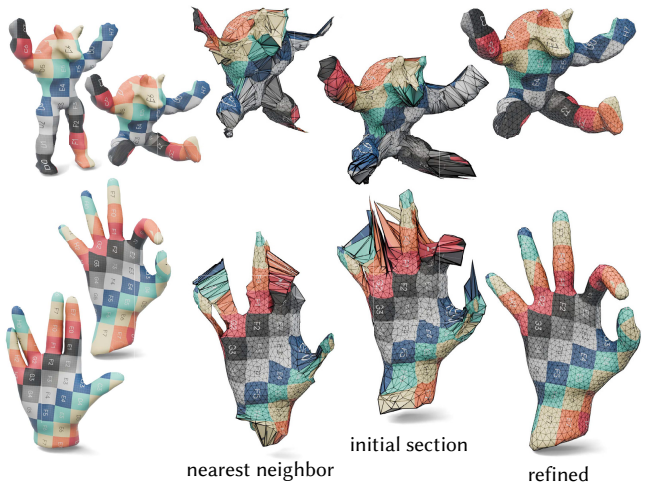


Fig. 17. Correspondences obtained from closest-point initialization, visualized with texture transfer *(left)*. Our initialization algorithm finds an initial complex section whose zero set approximates the graph of the nearest neighbor correspondences. Minimizing the Ginzburg-Landau energy refines this section to produce a high-quality correspondence *(right)*. Correspondences are visualized using both geometry and texture transfer.

### 5.1 Landmark-Free Correspondences

Many surface correspondence algorithms require landmarks or other user-specified constraints to guide the optimization to a meaningful solution. Our method can compute high-quality correspondences even in the absence of landmarks, relying only on the extrinsic geometry of the input surfaces. Given two surfaces $A$ and $B$, we first scale and rigidly align them. Then we consider the nearest-neighbor map which assigns to each vertex $v_A \in V_A$ the face of $B$ containing the closest point, and vice versa.

While this raw nearest-neighbor map is often noisy and discontinuous, our initialization procedure (Section 4.4) can effectively regularize it. By computing the initial section $z_0$ as the lowest eigenmode of the connection Laplacian, we naturally smooth the correspondences. As shown in Figure 17, this step alone can sometimes significantly correct the rough geometric input, producing a coherent initial map even before the energy minimization begins. Minimizing the Ginzburg-Landau energy eventually refines this robust initialization to produce a high-quality correspondence, as seen on the right of Figure 17.

Due to its simplicity and effectiveness, we use this closest-point initialization on all examples (including those with landmark constraints) unless otherwise stated. Figures 10 to 12, 16 and 17 show landmark-free correspondences computed between pairs of nearly isometric shapes differing by large deformations. Despite the lack of supervision, our method produces smooth, minimal-distortion maps that correctly untangle and align salient geometric features.

### 5.2 Untangling Surface Maps

As demonstrated in Section 5.1, our algorithm recovers meaningful correspondences even from noisy nearest neighbor maps. Because the Ginzburg-Landau functional remains well-defined regardless of whether the zero set represents a valid bijection, our method can also repair invalid correspondences produced by external algorithms.

In particular, thanks to its structural orientation preserving constraints, our method is able to correct correspondences exhibiting severe orientation reversals due to intrinsic symmetries, as shown in Figure 4. Since this constraint effectively excludes the identity map from the search space, we can efficiently compute symmetric self-maps (see Figure 5).

The method is also particularly well-suited for fixing local geometric collapses, which are common artifacts in spectral-based shape matching (*e.g.* functional maps), where thin structures shrink to points. Figure 18 shows a typical example of such a collapse in the output of the ZoomOut algorithm [Melzi et al. 2019], appearing despite the use of initial landmarks. In these cases, our minimization drives the zero set to expand and cover the target surface, recovering a precise bijective map from a degenerate input.

### 5.3 Curve-to-Curve Correspondences

Like many shape correspondence algorithms, our method supports using point landmarks to constrain specific locations (Figure 13). However, many applications require more flexible constraints, such as mapping a curve on surface $A$ to a corresponding curve on surface $B$ without prescribing the pointwise alignment. These constraints arise naturally when matching surfaces with sharp feature curves,

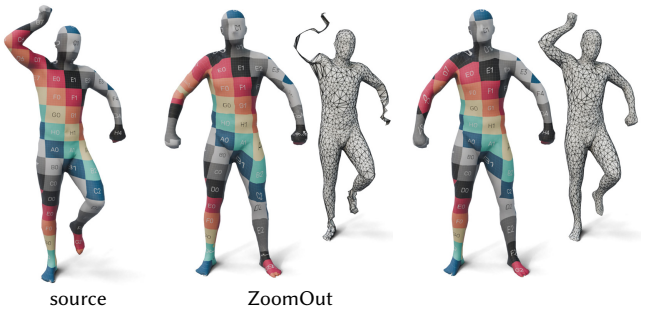


Fig. 18. Spectral methods, such as Functional Maps, usually fail preserve thin shape features. A ZoomOut algorithm [Melzi et al. 2019] initialized with landmarks produces a map that appears smooth (center). Visualizing coordinate transfer however reveals that the arms of the human shape have collapsed. Using this map as initialization, our algorithm (right) is able to correctly match the limbs thanks to its built-in topology and orientation preserving constraints.

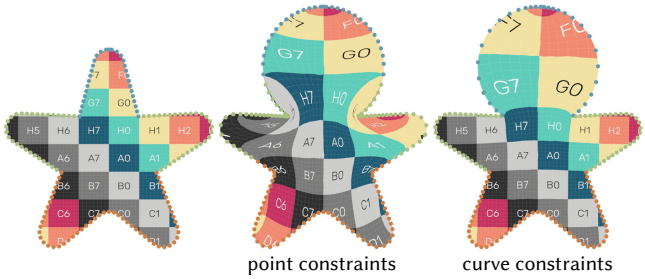


Fig. 19. Enforcing curve-to-curve correspondences can be challenging, and many algorithms rely on constraining the correspondences of each point on the curve. Matching curves based on curvilinear coordinates, however, may introduce significant distortion (middle). Our framework enables the easy enforcement of curve-to-curve correspondences, allowing points to slide along the specified curves (right). This approach substantially reduces overall distortion by accurately matching the concave corners of the boundary.

such as boundaries or creases (Figures 14 and 15). They also appear frequently in geometric morphometrics under the name of "semi-landmarks" [Gunz and Mitteroecker 2013] where curves are easier to identify than specific points.

A naïve, albeit common, approach is to sample points along each curve and enforce explicit point-to-point correspondences. However, this arbitrarily fixes the parameterization between the curves; even the "natural" choice of arc-length parameterization often twists the map, leading to poor interior correspondences (Figure 19). Our implicit framework, on the other hand, handles such curve-to-curve correspondences with ease—the singularity pinning potential attracts the zeros (and therefore the mapping surface) to the two-dimensional patch $\gamma_A \times \gamma_B \subset A \times B$ traced out by the pair of corresponding curves, without enforcing any preferential parameterization. The correspondence between the curves then emerges according to what is energetically favorable. Importantly, the two input curves do not need to have the same number of sample points. Finally, we note that extending this approach to mixed point-to-curve constraints is straightforward.

## 5.4 Comparisons with Prior Work

We evaluate our method against state-of-the-art approaches, checking robustness to initialization, and quality of the resulting correspondences. In particular, we compare to the constant curvature metric based inter-surface mapping (ISM) approach from [Schmidt et al. 2020], the adaptive triangulations (AT) method from [Schmidt et al. 2023], along with the reversible harmonic maps (RHM) approach of [Ezuz et al. 2019b].

*Robustness to Initialization.* A key advantage of our approach lies in its ability to recover high-quality maps even from poor initial correspondences (Sections 5.1 and 5.2). In Figure 20, we replicate the initialization experiment from [Schmidt et al. 2020, Fig. 13], computing correspondence from increasingly distorted initial states. RHM, which directly minimizes the Dirichlet energy, consistently becomes trapped in local minima and fails to realign thin features or correct large distortions. While ISM performs significantly better,

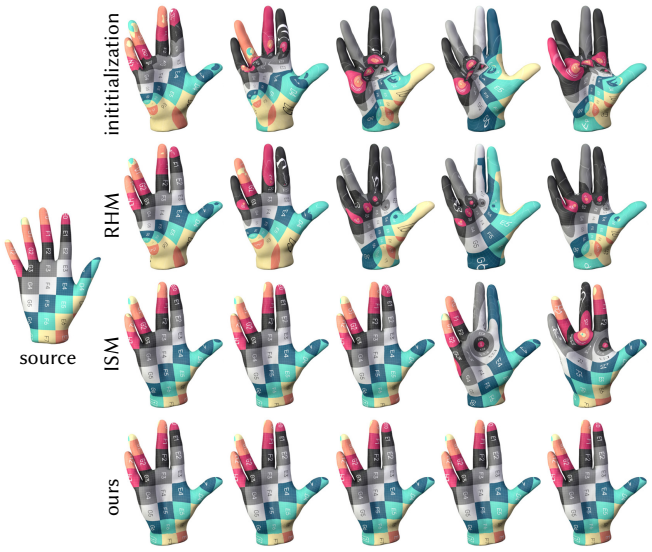


Fig. 20. We replicate the experiment from Schmidt et al. [2020] to evaluate robustness against poor initializations (top row). The RHM method [Ezuz et al. 2019b] becomes trapped in local minima of distortion for each initialization. While Schmidt et al. [2020] (ISM) performs better and achieves consistent results for the first three initializations, our method produces nearly identical low-distortion mappings across all initializations.

it achieves consistent results only for the first three initializations before diverging. In contrast, our method converges to a nearly identical, low-distortion map across all initializations, verifying its untangling capability discussed in Section 5.2.

*Map Quality and Distortion.* We further compare the quality and distortion of the final maps against ISM, which can also represent the overlay mesh induced by the correspondences. As shown in Figure 21, while both algorithms converge to geometrically similar correspondences, our implicit Ginzburg-Landau minimization results in consistently lower distortion energies (measured by both

the area of the correspondence graph and symmetric Dirichlet energy). The difference is particularly noticeable in the smoothness of the distortion distribution and the regularity of the texture transfer. Our method avoids the localized spikes in distortion often seen in explicit remeshing-based approaches.

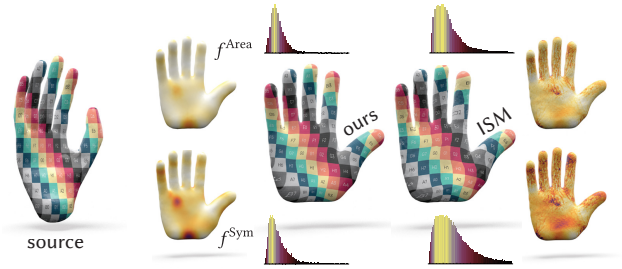


Fig. 21. Both our algorithm and the algorithm of Schmidt et al. [2020] (ISM) represent the common subdivision of the overlay mesh induced by the correspondence. It is needed to evaluate distortion energies of the correspondence intrinsically. While both algorithms find similar correspondence, implicit area minimization produces a less distorted map, even at the finer scales. This is reflected in the smoother distortion distributions, measured both by the area of the correspondence graph given by integrating $f^{\text{Area}}(\sigma_1, \sigma_2) = ((1+\sigma_1^2)(1+\sigma_2^2))^{1/2}$, and the symmetric Dirichlet energy, integrating $f^{\text{Sym}}(\sigma_1, \sigma_2) = \sigma_1^2 + \sigma_2^2 + \sigma_1^{-2} + \sigma_2^{-2}$. The histogram is colored by the height of the bars.

Additionally, we compare the distortion against AT on a challenging pair of inputs with significant non-isometric deformation. Figure 22 shows that while AT produces a valid bijective map, our method produces a more aligned correspondence with lower symmetric Dirichlet energy. While not explicitly enforced, our method better respects the intrinsic symmetries of the shape (*e.g.*, preserving the left-right symmetry of the human model).

## 5.5 Implementation and Parameters

The main parameter of our method is the Ginzburg-Landau parameter $\lambda$, weighting the potential energy. As discussed in Section 4.2 and Appendix C, setting $\lambda \approx 100\lambda_{A\times B}$ generally produces high-quality results. However, we observed that examples requiring significant untangling or large deformations from the initialization benefit from a simple annealing scheme, using a smaller $\lambda$ first to align global features before refining with the default value. When using landmarks, we set equal widths for the Gaussian kernels, $\sigma_A = \sigma_B$, setting them depending on the landmarks distribution. We acknowledge that this parameter is sensitive to the specific configuration: the optimal $\sigma$ depends both on the geometric scale and on the distribution of the landmarks. While the precise value of this parameter does not dramatically change the correspondences, we select this parameter manually to ensure the potential wells are sufficiently wide to attract the zero set. We expect that future work on adaptive pinning potentials will allow this parameter to be determined automatically. We refer the reader to the supplementary material for a complete list of parameters used in each figure.

For all examples, we used L-BFGS to minimize the Ginzburg-Landau energy, terminating when the norm of the projected gradient fell below $10^{-5}$ or after a maximum of 1000 iterations. For reproducibility, comprehensive pseudocode can be found in the supplementary material, and C++, MATLAB, and Python implementations of our method will be released upon publication.

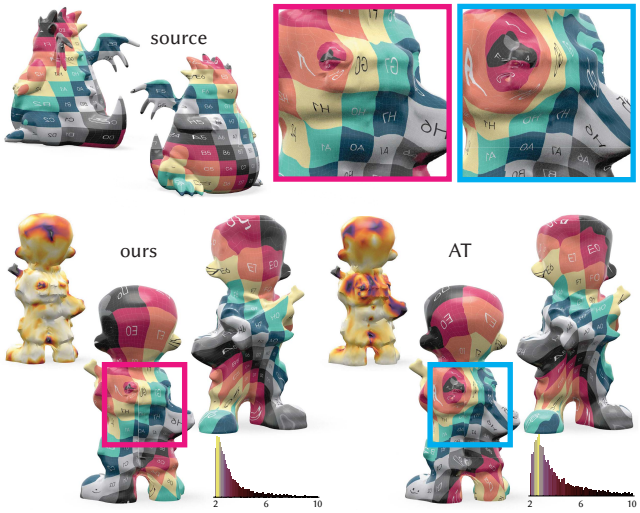


Fig. 22. Distortion is inevitable between non-isometric shapes. Compared with the adaptive triangulations (AT) method [Schmidt et al. 2023], our method produces dramatically less distorted correspondences (as measured by the symmetric Dirichlet energy) that better respect the intrinsic symmetries of the model.

## 6 Limitations and Future Work

While our algorithm offers new perspectives and significant advantages in terms of correspondence quality relative to prior work, it also presents certain challenges and suggests avenues for future research.

*Performance.* Without a multiresolution hierarchy, the runtime to compute an implicit minimal surface scales slightly superlinearly with the product of the vertex counts (Figure 23). On a typical pair of models with $|V_A| \approx |V_B| \approx 500$, our algorithm takes approximately 20 seconds, while for a pair of models with with $|V_A| \approx |V_B| \approx 5000$ our algorithm takes approximately 1 hour. Our code is implemented in C++; timings are reported on an Intel i7-14700K CPU with 64GB of RAM. While the direct application of our method is significantly slower than remeshing based approaches such as [Schmidt et al. 2023], it is competitive with ISM approach of [Schmidt et al. 2020], which reports a runtime of approximately 3 hours for models with $|V_A| \approx |V_B| \approx 4000$.

The runtime of our algorithm is dominated by the eigenvalue initialization and the minimization of the Ginzburg-Landau functional; the remaining operations, including the construction of surface connections and overlay mesh extraction, are negligible. For the former, we employ a naïve LOBPCG solver, which could be significantly accelerated by incorporating algebraic or geometric multigrid preconditioners [Xu and Zikatanov 2017] adapted to the tensor-product structure of the operator. Similarly, for the energy minimization, we rely on a generic off-the-shelf algorithm (L-BFGS). Using a custom

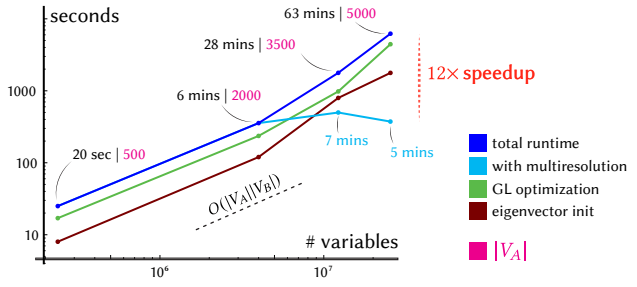


Fig. 23. Computing implicit minimal surfaces scales roughly linearly in the number of variables, with the eigenvalue initialization taking a shade less than half the total runtime. By adopting a multiresolution acceleration scheme, we decrease the effective runtime significantly.

solver that uses the structure of the Ginzburg-Landau energy to exploit parallelization would likely provide significant speedups. These costs are partially mitigated by our multiresolution hierarchy (Section 4.5), which shifts most of the computation to coarser meshes. For the example in Figure 12 ($|V_A| \approx |V_B| \approx 5000$) we obtained a 12× speedup and computed a distortion minimizing bijective correspondence in only 5 minutes.

Memory limitations place a hard constraint on the input mesh size: although the mass and stiffness matrices are never assembled in the product space, the section $z$ requires $|V_A| \times |V_B|$ storage, exceeding capacity for pairs with greater than ≈70k vertices each. By using a decimated proxy mesh to compute the correspondence there are further opportunities to decouple the runtime and memory requirements from the input resolution.

*Bijectivity.* The non-convexity of the Ginzburg-Landau functional means that the optimization may converge to local minima. While theoretical work suggests that critical points—beyond just global minima—can describe minimal surfaces, it remains unclear whether *all* critical points correspond to such surfaces. Consequently, we cannot guarantee that our algorithm will always produce a bijective and distortion-minimizing mapping. Nevertheless, unlike other methods, our approach ensures an orientation-preserving mapping, which proves highly efficient for surfaces exhibiting isometric symmetries.

In a few situations where the input surfaces have spurious regions of high curvature, long thin features, and/or inconsistent landmark constraints we found that energy minimization would attract the singularities from multiple slices to the same point. These singularity "sinks" result in invalid correspondences; though in these cases, we were able to produce distortion minimizing correspondences by suitably modifying the Ginzburg-Landau parameter. Further theoretical work would need to be developed to characterize (and constrain) when these implicit minimal surfaces are graphs over both factors to exclude these failure modes.

*Generalization.* Extending the algorithm to general surfaces, not limited to genus zero surfaces with boundaries, is a compelling direction for future work. The topology of the implicit surfaces, encoded in the curvature of the connection, needs to be modified so that it represents surfaces that are the graphs of bijective correspondences. In addition to total curvature $2\pi$ on each slice, additional curvature must be concentrated, according to the map homotopy type, on the closed two-dimensional surfaces corresponding to the product of homology generators. Beyond changing the topology of the implicit surfaces, additional complications arise when minimizing the Ginzburg-Landau functional that may obstruct extracting a correspondence. For instance, the energy density of critical points of the Ginzburg-Landau functional may not concentrate around a singularity as $\varepsilon \to 0$, instead spreading out over the entire space according to a harmonic one-form.

The implicit representation we studied can also be generalized to describing codimension-3 manifolds, replacing complex line bundles with rank-3 vector bundles. Looking for implicit codimension-3 minimal surfaces in the product of two volumetric domains would be the natural generalization of our method to volumetric correspondences.


Finally, we mention that beyond bijective correspondences, our method can compute *covering maps* instead of bijections. For instance, by picking a connection on $A$ which has total curvature $4\pi$ instead of $2\pi$, we can compute a double covering map where every point of $B$ is covered by two points of $A$.

## References


S. Mazdak Abulnaga, Oded Stein, Polina Golland, and Justin Solomon. 2023. Symmetric volume maps: Order-invariant volumetric mesh correspondence with free boundary. *ACM Transactions on Graphics (TOG)* 42, 3 (2023), 1–20. https://doi.org/10.1145/3572897

Amandine Aftalion, Etienne Sandier, and Sylvia Serfaty. 2001. Pinning phenomena in the Ginzburg-Landau model of superconductivity. *Journal de mathématiques pures et appliquées* 80, 3 (2001), 339–372. https://doi.org/10.1016/S0021-7824(00)01180-6

Noam Aigerman, Shahar Z Kovalsky, and Yaron Lipman. 2017. Spherical Orbifold Tutte Embeddings. *ACM Trans. Graph.* 36, 4 (2017), 90. https://doi.org/10.1145/3072959.3073615

Noam Aigerman and Yaron Lipman. 2015. Orbifold Tutte Embeddings. *ACM Trans. Graph.* 34, 6 (2015), 190–1. https://doi.org/10.1145/2816795.2818099

Noam Aigerman and Yaron Lipman. 2016. Hyperbolic Orbifold Tutte Embeddings. *ACM Trans. Graph.* 35, 6 (2016), 217–1. https://doi.org/10.1145/2980179.2982412

Noam Aigerman, Roi Poranne, and Yaron Lipman. 2014. Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–12. https://doi.org/10.1145/2601097.2601158

Noam Aigerman, Roi Poranne, and Yaron Lipman. 2015. Seamless surface mappings. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–13. https://doi.org/10.1145/2766921

Giovanni Alberti, Sisto Baldo, and Giandomenico Orlandi. 2005. Variational convergence for functionals of Ginzburg-Landau type. *Indiana University Mathematics Journal* (2005), 1411–1472. https://www.jstor.org/stable/24902326

N. El Amrani, P. Roetzer, and F. Bernard. 2025. High-Resolution 3D Shape Matching with Global Optimality and Geometric Consistency. *Computer Graphics Forum* 44, 5 (2025), e70208. https://doi.org/10.1111/cgf.70208

Alex Baden, Keenan Crane, and Misha Kazhdan. 2018. Möbius Registration. *Computer Graphics Forum (SGP)* 37, 5 (2018). https://doi.org/10.1111/cgf.13503

M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* 61, 2 (2005), 139–157. https://doi.org/10.1023/B:VISI.0000043755.93987.aa

A. Bobenko and B. Springborn. 2007. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete & Computational Geometry* 38, 4 (2007). https://doi.org/10.1007/s00454-007-9006-1

Lucas Brifault, David Cohen-Steiner, and Mathieu Desbrun. 2025. Efficient and Scalable Spatial Regularization of Optimal Transport. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*. ACM, Article 120, 10 pages. https://doi.org/10.1145/3757377.3763976

Giacomo Canevari, Federico Luigi Dipasquale, and Giandomenico Orlandi. 2023. The Yang–Mills–Higgs functional on complex line bundles: Γ-convergence and the London equation. *Archive for Rational Mechanics and Analysis* 247, 6 (2023), 104.

https://doi.org/10.1007/s00205-023-01933-1
Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. 2018. Shape from metric. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 63. https://doi.org/10.1145/3197517.3201276
Keenan Crane, Mathieu Desbrun, and Peter Schröder. 2010. Trivial Connections on Discrete Surfaces. *Computer Graphics Forum (SGP)* 29, 5 (2010), 1525–1533. https://doi.org/10.1111/j.1467-8659.2010.01761.x
Camillo De Lellis. 2014. The regularity of minimal surfaces in higher codimension. *Current Developments in Mathematics* (2014), 153–229.
Guido De Philippis and Alessandro Pigati. 2022. Non-degenerate minimal submanifolds as energy concentration sets: A variational approach. *Communications on Pure and Applied Mathematics* (2022). https://doi.org/10.1002/cpa.22193
Mathieu Desbrun, Anil N Hirani, Melvin Leok, and Jerrold E Marsden. 2005. Discrete exterior calculus. arXiv:math/0508341 [math.DG]
Nicolas Donati, Etienne Corman, Simone Melzi, and Maks Ovsjanikov. 2022. Complex Functional Maps: A Conformal Link Between Tangent Bundles. *Computer Graphics Forum* 41, 1 (2022), 317–334. https://doi.org/10.1111/cgf.14437
M. Eisenberger, Z. Lähner, and D. Cremers. 2019. Divergence-Free Shape Correspondence by Deformation. *Computer Graphics Forum (SGP)* 38, 5 (2019), 1–12. https://doi.org/10.1111/cgf.13785
Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. 2020. Smooth Shells: Multi-Scale Shape Registration With Functional Maps. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 12262–12271. https://doi.org/10.1109/CVPR42600.2020.01228
D. Ezuz, B. Heeren, O. Azencot, M. Rumpf, and M. Ben-Chen. 2019a. Elastic Correspondence between Triangle Meshes. *Computer Graphics Forum* 38, 2 (2019), 121–134. https://doi.org/10.1111/cgf.13624
Danielle Ezuz, Justin Solomon, and Mirela Ben-Chen. 2019b. Reversible harmonic maps between discrete surfaces. *ACM Transactions on Graphics (TOG)* 38, 2 (2019), 1–12. https://doi.org/10.1145/3202660
Herbert Federer and Wendell H. Fleming. 1960. Normal and Integral Currents. *Annals of Mathematics* 72, 3 (1960), 458–520. https://doi.org/10.2307/1970227
Mark Gillespie, Nicholas Sharp, and Keenan Crane. 2021. Integer coordinates for intrinsic geometry processing. *ACM Transactions on Graphics (TOG)* 40, 6 (2021). https://doi.org/10.1145/3478513.3480522
Philipp Gunz and Philipp Mitteroecker. 2013. Semilandmarks: A Method for Quantifying Curves and Surfaces. *Hystrix the Italian Journal of Mammalogy* 24, 1 (May 2013), 103–109. https://doi.org/10.4404/hystrix-24.1-6292
Qi-Xing Huang, Bart Adams, Martin Wicke, and Leonidas J. Guibas. 2008. Non-rigid registration under isometric deformations. *Computer Graphics Forum (SGP)* (2008), 1449–1457. https://doi.org/10.1111/j.1467-8659.2008.01285.x
Robert L Jerrard and Halil Mete Soner. 2002. The Jacobian and the Ginzburg-Landau energy. *Calculus of Variations and Partial Differential Equations* 14, 2 (2002), 151–191. https://doi.org/10.1007/s005260100093
T. Kanai, H. Suzuki, and F. Kimura. 1997. 3D geometric metamorphosis based on harmonic map. In *Proceedings The Fifth Pacific Conference on Computer Graphics and Applications*. 97–104. https://doi.org/10.1109/PCCGA.1997.626179
Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. 2011. Blended intrinsic maps. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–12. https://doi.org/10.1145/2010324.1964974
Felix Knöppel. 2020. Riemann Surfaces: Complex Analysis from a Differential Geometric Viewpoint. (2020). Lecture Notes.
Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally optimal direction fields. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10. https://doi.org/10.1145/2461912.2462005
Felix Knöppel and Ulrich Pinkall. 2016. *Complex line bundles over simplicial complexes and their applications*. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-50447-5_6
A. V. Knyazev, M. E. Argentati, I. Lashuk, and E. E. Ovtchinnikov. 2007. Block Locally Optimal Preconditioned Eigenvalue Xolvers (BLOPEX) in Hypre and PETSc. *SIAM Journal on Scientific Computing* 29, 5 (2007), 2224–2239. https://doi.org/10.1137/060661624
Maximilian Kohlbrenner, Singchun Lee, Marc Alexa, and Misha Kazhdan. 2023. Poisson Manifold Reconstruction — Beyond Co-dimension One. *Computer Graphics Forum (SGP)* 42, 5 (2023), e14907. https://doi.org/10.1111/cgf.14907
Vladislav Kraevoy and Alla Sheffer. 2004. Cross-parameterization and compatible remeshing of 3D models. *ACM Transactions on Graphics (TOG)* 23, 3 (2004), 861–869. https://doi.org/10.1145/1015706.1015811
Hao Li, Robert W. Sumner, and Mark Pauly. 2008b. Global Correspondence Optimization for Non-Rigid Registration of Depth Scans. *Computer Graphics Forum (SGP)* 27, 5 (2008), 1421–1430. https://doi.org/10.1111/j.1467-8659.2008.01282.x
Xin Li, Yunfan Bao, Xiaohu Guo, Miao Jin, Xianfeng Gu, and Hong Qin. 2008a. Globally optimal surface mapping for surfaces with arbitrary topology. *IEEE Transactions on Visualization and Computer Graphics* 14, 4 (2008), 805–819. https://doi.org/10.1109/TVCG.2008.32
Fanghua Lin and Tristan Riviere. 1999. Complex Ginzburg-Landau equations in high dimensions and codimension two area minimizing currents. *Journal of the European Mathematical Society* 1, 3 (1999), 237–311. https://doi.org/10.1007/s100970050008
O. Litany, E. Rodolà, A. M. Bronstein, M. M. Bronstein, and D. Cremers. 2016. Non-rigid puzzles. *Computer Graphics Forum (SGP)* 35, 5 (Aug 2016), 135–143. https://doi.org/10.1111/cgf.12970
Nathan Litke, Marc Droske, Martin Rumpf, and Peter Schröder. 2005. An image processing approach to surface matching. In *Proceedings of the Third Eurographics Symposium on Geometry Processing (SGP '05)*. Eurographics Association, 207–es. https://doi.org/10.2312/SGP/SGP05/207-216
Beibei Liu, Yiying Tong, Fernando De Goes, and Mathieu Desbrun. 2016. Discrete Connection and Covariant Derivative for Vector Field Analysis and Design. *ACM Transactions on Graphics (TOG)* 35, 3, Article 23 (March 2016), 17 pages. https://doi.org/10.1145/2870629
Manish Mandad, David Cohen-Steiner, Leif Kobbelt, Pierre Alliez, and Mathieu Desbrun. 2017. Variance-minimizing transport plans for inter-surface mapping. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–14. https://doi.org/10.1145/3072959.3073671
Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. 2019. ZoomOut: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–14. https://doi.org/10.1145/3355089.3356524
Takashi Michikawa, Takashi Kanai, Masahiro Fujita, and Hiroaki Chiyokura. 2001. Multiresolution Interpolation Meshes. In *Proceedings of the 9th Pacific Conference on Computer Graphics and Applications (PG '01)*. IEEE Computer Society, 60. https://doi.org/10.1109/PCCGA.2001.962858
John Willard Milnor and James D Stasheff. 1974. *Characteristic Classes*. Number 76 in Annals of Mathematics Studies. Princeton University Press.
Luciano Modica and S Mortola. 1977. Un esempio di Γ-convergenza. *Bollettino della Unione Matematica Italiana B* 14 (1977), 285–299.
Luca Morreale, Noam Aigerman, Vladimir Kim, and Niloy J. Mitra. 2021. Neural Surface Maps. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 4637–4646. https://doi.org/10.1109/CVPR46437.2021.00461
Liviu Nicolaescu and Nikhil Savale. 2017. The Gauss-Bonnet-Chern theorem: a probabilistic perspective. *Trans. Amer. Math. Soc.* 369, 4 (2017), 2951–2986. https://doi.org/10.1090/tran/6895
Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–11. https://doi.org/10.1145/2185520.2185526
Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas, Frederic Chazal, and Alex Bronstein. 2017. Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses*. ACM, Article 5, 62 pages. https://doi.org/10.1145/3084873.3084877
David Palmer, Albert Chern, and Justin Solomon. 2024. Lifting Directional Fields to Minimal Sections. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–20. https://doi.org/10.1145/3658198
David Palmer, Dmitriy Smirnov, Stephanie Wang, Albert Chern, and Justin Solomon. 2022. DeepCurrents: Learning Implicit Representations of Shapes with Boundaries. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 18644–18654. https://doi.org/10.1109/CVPR52688.2022.01811
Davide Parise, Alessandro Pigati, and Daniel Stern. 2024. Convergence of the self-dual $U(1)$-Yang–Mills–Higgs energies to the $(n-2)$-area functional. *Communications on Pure and Applied Mathematics* 77, 1 (2024), 670–730. https://doi.org/10.1002/cpa.22150
Chao Peng and Sabin Timalsena. 2016. Fast mapping and morphing for genus-zero meshes with cross spherical parameterization. *Computers & Graphics* 59 (2016), 107–118. https://doi.org/10.1016/j.cag.2016.06.001
Alessandro Pigati and Daniel Stern. 2021. Minimal submanifolds from the abelian Higgs model. *Inventiones mathematicae* 223 (2021), 1027–1095.
Roman Poya, Rogelio Ortigosa, and Theodore Kim. 2023. Geometric optimisation via spectral shifting. *ACM Transactions on Graphics (TOG)* 42, 3 (2023), 1–15. https://doi.org/10.1145/3585003
Jing Ren, Simone Melzi, Maks Ovsjanikov, and Peter Wonka. 2020. Maptree: Recovering multiple solutions in the space of maps. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–17. https://doi.org/10.1145/3414685.3417800
Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. 2021. Discrete optimization for shape matching. *Computer Graphics Forum (SGP)* 40, 5 (2021), 81–96. https://doi.org/10.1111/cgf.14359
Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. 2018. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–16. https://doi.org/10.1145/3272127.3275040
Paul Roetzer and Florian Bernard. 2025. Fast Globally Optimal and Geometrically Consistent 3D Shape Matching. https://doi.org/10.48550/arXiv.2504.06385 arXiv:2504.06385 [cs]

Paul Roetzer, Paul Swoboda, Daniel Cremers, and Florian Bernard. 2022. A Scalable Combinatorial Solver for Elastic Geometrically Consistent 3D Shape Matching. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 428–438. https://doi.org/10.1109/CVPR52688.2022.00052

Yusuf Sahillioğlu. 2020. Recent advances in shape correspondence. *The Visual Computer* 36, 8 (2020), 1705–1721. https://doi.org/10.1007/s00371-019-01760-0

Patrick Schmidt, Janis Born, Marcel Campen, and Leif Kobbelt. 2019. Distortion-minimizing injective maps between surfaces. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–15. https://doi.org/10.1145/3355089.3356519

Patrick Schmidt, Marcel Campen, Janis Born, and Leif Kobbelt. 2020. Inter-surface maps via constant-curvature metrics. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 119–1. https://doi.org/10.1145/3386569.3392399

P. Schmidt, D. Pieper, and L. Kobbelt. 2023. Surface Maps via Adaptive Triangulations. *Computer Graphics Forum* 42, 2 (2023), 103–117. https://doi.org/10.1111/cgf.14747

John Schreiner, Arul Asirvatham, Emil Praun, and Hugues Hoppe. 2004. Inter-surface mapping. *ACM Transactions on Graphics (TOG)* 23, 3 (2004), 870–877. https://doi.org/10.1145/1015706.1015812

Andrei Sharf, Marina Blumenkrants, Ariel Shamir, and Daniel Cohen-Or. 2006. Snappaste: an interactive technique for easy mesh composition. *The Visual Computer* 22, 9 (2006), 835–844. https://doi.org/10.1007/s00371-006-0068-5

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. Navigating Intrinsic Triangulations. *ACM Transactions on Graphics (TOG)* 38, 4 (2019). https://doi.org/10.1145/3306346.3322979

Rui Shi, Wei Zeng, Zhengyu Su, Jian Jiang, Hanna Damasio, Zhonglin Lu, Yalin Wang, Shing-Tung Yau, and Xianfeng Gu. 2016. Hyperbolic harmonic mapping for surface registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 5 (2016), 965–980. https://doi.org/10.1109/TPAMI.2016.2567398

Justin Solomon, Andy Nguyen, Adrian Butscher, Mirela Ben-Chen, and Leonidas Guibas. 2012. Soft Maps Between Surfaces. *Computer Graphics Forum (SGP)* 31, 5 (2012), 1617–1626. https://doi.org/10.1111/j.1467-8659.2012.03167.x

Kenshi Takayama. 2022. Compatible intrinsic triangulations. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–12. https://doi.org/10.1145/3528223.3530175

Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. 2012. Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *IEEE Transactions on Visualization and Computer Graphics* 19, 7 (2012), 1199–1217. https://doi.org/10.1109/TVCG.2012.310

Alex Tsui, Devin Fenton, Phong Vuong, Joel Hass, Patrice Koehl, Nina Amenta, David Coeurjolly, Charles DeCarli, and Owen Carmichael. 2013. Globally Optimal Cortical Surface Matching with Exact Landmark Correspondence. In *Information Processing in Medical Imaging*. Springer Berlin Heidelberg, 487–498. https://doi.org/10.1007/978-3-642-38868-2_41

Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. 2011. A Survey on Shape Correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. https://doi.org/10.1111/j.1467-8659.2011.01884.x

Matthias Vestner, Roee Litman, Emanuele Rodola, Alex Bronstein, and Daniel Cremers. 2017. Product Manifold Filter: Non-rigid Shape Correspondence via Kernel Density Estimation in the Product Space. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 6681–6690. https://doi.org/10.1109/CVPR.2017.707

Stephanie Wang and Albert Chern. 2021. Computing minimal surfaces with differential forms. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14. https://doi.org/10.1145/3450626.3459781

Steffen Weißmann, Ulrich Pinkall, and Peter Schröder. 2014. Smoke rings from smoke. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–8. https://doi.org/10.1145/2601097.2601171

H Whitney. 1957. Geometric Integration Theory Princeton Univ. *Press, Princeton, NJ* (1957), 62.

Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers. 2011a. Geometrically Consistent Elastic Matching of 3D Shapes: A Linear Programming Solution. In *2011 International Conference on Computer Vision*. IEEE, Barcelona, Spain, 2134–2141. https://doi.org/10.1109/ICCV.2011.6126489

Thomas Windheuser, Ulrich Schlickwei, Frank R. Schimdt, and Daniel Cremers. 2011b. Large-Scale Integer Linear Programming for Orientation Preserving 3 *D* Shape Matching. *Computer Graphics Forum* 30, 5 (Aug. 2011), 1471–1480. https://doi.org/10.1111/j.1467-8659.2011.02021.x

Jinchao Xu and Ludmil Zikatanov. 2017. Algebraic multigrid methods. *Acta Numerica* 26 (2017), 591–721. https://doi.org/10.1017/S0962492917000083

Yang Yang, Xiao-Ming Fu, Shuangming Chai, Shi-Wei Xiao, and Ligang Liu. 2018. Volume-enhanced compatible remeshing of 3D models. *IEEE Transactions on Visualization and Computer Graphics* 25, 10 (2018), 2999–3010. https://doi.org/10.1109/TVCG.2018.2861396

Yang Yang, Wen-Xiang Zhang, Yuan Liu, Ligang Liu, and Xiao-Ming Fu. 2020. Error-bounded compatible remeshing. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 113–1. https://doi.org/10.1145/3386569.3392434

Laurence C Young. 1963. Some extremal questions for simplicial complexes V. The relative area of a Klein bottle. *Rendiconti del Circolo Matematico di Palermo* 12, 3 (1963), 257–274. https://doi.org/10.1007/BF02851262

## A Riemannian Geometry in the Product of Surfaces

Consider a pair of oriented smooth surfaces $A$ and $B$ with Riemannian metrics $g_A$ and $g_B$, respectively. The product space $A \times B$ is a Riemannian 4-manifold endowed with the metric $g_{A\times B}(X_A \oplus X_B, Y_A \oplus Y_B) \coloneqq g_A(X_A, Y_A) + g_B(X_B, Y_B)$.

Proposition A.1. *Given $\varphi : A \to B$ the area of its graph $\Sigma_\varphi$ can be expressed as the integral of the singular values $\sigma_1, \sigma_2$ of the deformation $d\varphi$:*

$$\mathrm{Area}(\Sigma_\varphi) = \int_A \sqrt{(1+\sigma_1^2)(1+\sigma_2^2)}\ \mathrm{vol}_A\,.$$

Proof. The graph can be parameterized by $F_\varphi : A \to \Sigma_\varphi, F_\varphi(p) \coloneqq (p, \varphi(p))$, and so $\mathrm{Area}(\Sigma_\varphi)$ can be computed by integrating the area form induced by the parameterization:

$$\begin{aligned}
\mathrm{Area}(\Sigma_\varphi) &= \int_A \det\left(F_\varphi^\star g_{A\times B}\right)^{\frac{1}{2}} \mathrm{vol}_A \\
&= \int_A \det\left(I + g_A^{-1}\varphi^\star g_B\right)^{\frac{1}{2}} \mathrm{vol}_A \\
&= \int_A \left[1 + \det\left(g_A^{-1}\varphi^\star g_B\right) + \mathrm{tr}\left(g_A^{-1}\varphi^\star g_B\right)\right]^{\frac{1}{2}} \mathrm{vol}_A \\
&= \int_A \left[1 + \sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2\right]^{\frac{1}{2}} \mathrm{vol}_A\,.
\end{aligned}$$

□

As expected the change for variables to $B$ leads to the same function applied to the inverse singular values:

$$\begin{aligned}
\mathrm{Area}(\Sigma_\varphi) &= \int_B \left[1 + \sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2\right]^{\frac{1}{2}} \frac{\det\left((\varphi^{-1})^\star g_B\right)^{\frac{1}{2}}}{\det\left(g_B\right)^{\frac{1}{2}}} \mathrm{vol}_B \\
&= \int_B \left[1 + \sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2\right]^{\frac{1}{2}} \frac{1}{\sigma_1\sigma_2} \mathrm{vol}_B \\
&= \int_B \left[\frac{1}{\sigma_1^2\sigma_2^2} + 1 + \frac{1}{\sigma_2^2} + \frac{1}{\sigma_1^2}\right]^{\frac{1}{2}} \mathrm{vol}_B\,.
\end{aligned}$$

## B The Product-Space Dirichlet Energy

Here we discretize the smooth expression of Dirichlet energy in Equation 6 to obtain a discrete Dirichlet energy on the product space. We consider a connection $\nabla$ that is the tensor product of connections $\nabla^A$ and $\nabla^B$ on $A$ and $B$, respectively. For convenience, we reproduce Equation 6 below:

$$\mathcal{E}^D(z) = \int_A \left(\int_B |\nabla z|_B^2\ \mathrm{vol}_B\right)\mathrm{vol}_A + \int_B \left(\int_A |\nabla z|_A^2\ \mathrm{vol}_A\right)\mathrm{vol}_B$$

Discretizing the second term yields:

$$\int_{A\times B} |\nabla^A z(p,q)|^2 \operatorname{vol}_{A\times B}(p,q) = \sum_{\substack{ijk\in F_A \\ abc\in F_B}} \underbrace{\int_{ijk}\int_{abc} \Big| \sum_{\substack{u\in ijk \\ v\in abc}} Z_{u,v}\nabla^A\phi_u^A(p)\phi_v^B(q)\Big|^2 \operatorname{vol}_A(p)\operatorname{vol}_B(q)}_{X_{ijk,abc}}.$$

The integrand evaluates to

$$X_{ijk,abc} = \sum_{\substack{m\in ijk \\ n\in abc}} \bar{Z}_{m,n}\Bigg[\sum_{\substack{u\in ijk \\ v\in abc}} \left(\int_{ijk} (\nabla^A\phi_m^A)^\star(p)\nabla^A\phi_u^A(p)\operatorname{vol}_A(p)\right)\cdot \left(Z_{u,v}\int_{abc}\bar{\phi}_n^B(q)\phi_v^B(q)\operatorname{vol}_B(q)\right)\Bigg],$$

where we recognize the integrals as the components of connection Laplacian of $A$ and the mass matrix of $B$. Hence, the connection Laplacian on $A\times B$ has a tensor-product structure:

$$L^\nabla_{A\times B} = L^\nabla_A \otimes M^\nabla_B + M^\nabla_A \otimes L^\nabla_B. \tag{35}$$

Using a matrix representation $Z\in\mathbb{C}^{V_A\times V_B}$ of the complex field $z : A\times B\to\mathbb{C}$, the discretization of Equation 6 only depends on left and right multiplications of the connection Laplacian and mass matrices defined on each manifold:

$$\mathcal{E}^D(Z) = \frac{1}{2}\langle Z, L^\nabla_A Z (M^\nabla_B)^\top\rangle_{\mathbb{C}} + \frac{1}{2}\langle Z, M^\nabla_A Z (L^\nabla_B)^\top\rangle_{\mathbb{C}}.$$

## C The Ginzburg-Landau Parameter

For the following result, we consider a Hermitian positive-definite matrix $A\in\mathbb{C}^{d\times d}$ along with a diagonal matrix $M\in\mathbb{R}^{d\times d}$—for our application, they will be the connection Laplacian and mass matrix on a four-dimensional cell complex, respectively.

Lemma C.1. *Let $\lambda_0 > 0$ be the smallest eigenvalue of the generalized eigenvalue problem $Az = \lambda Mz$ . Consider the discrete energy*

$$e(z) = \tfrac{1}{2}\langle Az, z\rangle + \tfrac{\lambda}{4}\langle Mu, u\rangle,$$

*where $u_i = 1-|z_i|^2$ for every vertex $i = 1,\dots,d$. If $0<\lambda\le\lambda_0$ then the only critical point of $e$ is given by $z\equiv 0$.*

Proof. Taking the scalar product between a stationary point $\bar{z}$ and the gradient $\nabla e(\bar{z})$, we obtain:

$$\langle \bar{z}, \nabla e(\bar{z})\rangle = \langle \bar{z}, A\bar{z}\rangle - \lambda\langle \bar{z}, M\bar{z}\rangle + \lambda\langle |\bar{z}|^2, M|\bar{z}|^2\rangle = 0. \tag{36}$$

Using the fact that the smallest eigenvalue of $A$ is greater than $\lambda$, we conclude that $\bar{z}$ must be equal to zeros:

$$0 \le \langle \bar{z}, A\bar{z} - \lambda M\bar{z}\rangle = -\lambda\langle |\bar{z}|^2, M|\bar{z}|^2\rangle = 0 \le 0 \tag{37}$$

□

In the continuous setting, the connection Laplacian $\Delta^\nabla = \Delta^\nabla_A \otimes I_B + I_A\otimes\Delta^\nabla_B$ on $A\times B$ decomposes as a sum of connection Laplacians on each shape. Its minimal eigenvalue is therefore the sum of the minimal eigenvalues on each shape: $\lambda_0 = \lambda^\nabla_A + \lambda^\nabla_B$, where $\lambda^\nabla_A$ and $\lambda^\nabla_B$ are the minimal eigenvalues of $\Delta^\nabla_A$ and $\Delta^\nabla_B$ respectively (defined with respect to the corresponding volume form on each shape).

In practice, however, due to differing discretization choices, we use $L^\nabla_{A\times B} = L^\nabla_A\otimes M^\nabla_B + M^\nabla_A\otimes L^\nabla_B$ for the Dirichlet energy, but $M = M_A\otimes M_B$ for the potential energy, using lumped mass matrices $M_A$ and $M_B$. The generalized eigenvalue problem $Az = \lambda Mz$ then doesn't strictly separate. Motivated by the continuous setting, we approximate the stability threshold using $\lambda = \lambda^\nabla_A + \lambda^\nabla_B$ where $\lambda^\nabla_A$ and $\lambda^\nabla_B$ are the smallest eigenvalues of $L^\nabla_A$ with respect to $M^\nabla_A$, and $L^\nabla_B$ with respect to $M^\nabla_B$, respectively. As both mass matrices approximate the same continuous volume form, we expect this value to provide a close approximation of the true stability threshold.

## D Alternative Surface Connections

We describe two discrete connections with curvature $\Omega = \frac{1}{2}\Omega^{LC}$: one based on concentrating and redistributing the curvature of the Levi-Civita connection and another based on spin structures. Below, $M$ is a triangle mesh describing one of the two input surfaces, with $r^{LC}$ and $\Omega^{LC}$ as in Section 3.4.1.

While we obtained similar correspondences irrespective of the choice of connection, the available multiresolution schemes differ (Section 4.5). For instance, to prolongate the sections of the bundles constructed below requires keeping track of parallel transport maps across the hierarchy of meshes. Further refinements and generalizations of our approach may be possible by changing the surface connections and their curvature.

*A Vector Field Connection.* We can modify the discrete Levi-Civita connection by computing an "offset connection" $\tilde{r}$ so that $r = \tilde{r}\, r^{LC}$ is compatible with $\Omega$. Similar to our discretization of the skyscraper bundle, we first select an arbitrary face $f_0$, and define

$$\tilde{\Omega} = \Omega^{LC} - 2\pi(\chi(M)-1)\delta_{f_0},$$

where $\delta_{f_0}$ is the Kronecker delta on faces. By construction, $\tilde{\Omega}$ sums to $2\pi$ on $M$. Due to the $2\pi$ ambiguity in the curvature of a discrete complex line bundle, it is also compatible with $r^{LC}$. The offset connection can now be computed by solving a Poisson equation: we set $\tilde{r}_{ij} = e^{\imath\tilde{\rho}_{ij}}$ where $\tilde{\rho}$ solves

$$\arg\min_{\tilde{\rho}} \|\tilde{\rho}\|^2 \quad \text{s.t.} \quad \mathrm{d}_1\tilde{\rho} = \Omega - \tilde{\Omega}. \tag{38}$$

*Spin Connections.* A canonical choice on genus zero surfaces is given by a discrete spin connection [Chern et al. 2018], which is given by computing a square-root of the Levi-Civita connection

$$r^{\text{spin}}_{ij} = \pm(r^{LC}_{ij})^{1/2}.$$

Unlike the other constructions we discussed, the construction does not depend on the choice of an arbitrary face. The signs of the square-root, however, must be chosen appropriately so that $r^{\text{spin}}$ is compatible with $\Omega$. A simple spanning tree based algorithm to choosing the signs is given in [Chern et al. 2018, Algorithm 3].

## E Edge-Edge Intersections

We identify the intersection between an edge of $A$ and an edge of $B$ by locating zeros in "edge-edge" faces of $A\times B$. On such a face complex section values are interpolated using the tensor product of the basis functions along each edge,

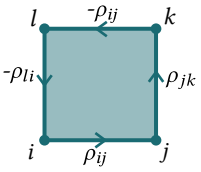

so a point $(s,t)$ in the face is a zero if and only if

$$0 = (1-s)(1-t)e^{\imath(s\rho_{ij}+t\rho_{jk})}z_i + s(1-t)e^{\imath(-(1-s)\rho_{ij}+t\rho_{jk})}z_j + ste^{\imath(-(1-s)\rho_{ij}-(1-t)\rho_{jk})}z_k + (1-s)te^{\imath(s\rho_{ij}-(1-t)\rho_{jk})}z_l.$$

Dividing through by $z_i e^{\imath(s\rho_{ij}+t\rho_{jk})}$, we obtain

$$0 = (1-s)(1-t) + s(1-t)\frac{z_j}{r_{ij}z_i} + st\frac{z_k}{r_{ij}r_{jk}z_i} + (1-s)t\frac{z_l}{r_{jk}z_i}.$$

Now $t$ is the solution to the following quadratic equation and $s$ is the associated quotient:

$$t^2\,\mathrm{Im}(a\bar{c}) + t\,\mathrm{Im}(a\bar{d} - c\bar{b}) + \mathrm{Im}(b\bar{d}) = 0 \quad s = -\frac{\mathrm{Re}(ct+d)}{\mathrm{Re}(at+b)},$$

where

$$b = -1 + r_{jk}^{-1}\frac{z_j}{z_i},\ c = -1 + r_{jk}^{-1}\frac{z_l}{z_i},\ d = 1$$
$$a = -1 - b - c + r_{jk}^{-1}r_{jk}^{-1}\frac{z_k}{z_i}.$$