Untangling Surfaces via Shape and Mesh Repulsion

JIŘÍ MINARČÍK[*], Carnegie Mellon University, United States of America and Resistant AI, Czech Republic
MICHAEL LIU[*], Carnegie Mellon University, United States of America
KEENAN CRANE, Carnegie Mellon University, United States of America and Roblox Research, United States of America
MINCHEN LI, Carnegie Mellon University, United States of America and Genesis AI, United States of America

Fig. 1. Inspired by Gian Lorenzo Bernini's *Bust of Medusa*, we construct a 3D model based on an existing Poisson reconstruction, modifying it so that many snakes are represented as separate mesh components rather than being merged into a single surface. We evaluate our method on this challenging model, which contains over 215k vertices and 431k faces and exhibits highly complex topology. Self-intersecting surface with intersections highlighted *(left)* and embedded surface obtained without changing mesh connectivity *(right)*.

Self-intersections are widespread in surface meshes and invalidate downstream simulation, fabrication, and learning pipelines. Existing approaches typically treat self-intersections as local collision events, but embeddedness (i.e., lack of self-intersections) is a global geometric property that cannot be enforced through local reasoning alone. We introduce an energy-based framework that enforces surface embeddedness simultaneously at the shape and mesh levels, based on the insight that successful untangling requires accounting for both global shape-level interactions and local mesh-level interactions. A shape-level energy captures global entanglement independent of discretization, while a mesh-level penalty regularizes local discrete interactions. Together, these energies enable reliable removal of self-intersections without changing mesh connectivity and apply to a broad class of geometries, including surfaces with boundary, non-manifold configurations, immersion failures, and multi-object scenes. Compared to prior state-of-the-art methods, our approach resolves self-intersections across challenging datasets, enabling reliable downstream processing of surface meshes.

CCS Concepts: • **Computing methodologies → Shape modeling**; • **Mathematics of computing → Continuous optimization**.

Additional Key Words and Phrases: Surface embeddedness, Self-intersection removal, Energy-based methods, Geometric optimization, Mesh processing

ACM Reference Format:
Jiří Minarčík, Michael Liu, Keenan Crane, and Minchen Li. 2026. Untangling Surfaces via Shape and Mesh Repulsion. *ACM Trans. Graph.* 45, 4, Article 163 (July 2026), 15 pages. https://doi.org/10.1145/3811382

---

[*] Both authors contributed equally to this research.

Authors' addresses: Jiří Minarčík, jiri@minarcik.com, Carnegie Mellon University, Pittsburgh, United States of America and Resistant AI, Prague, Czech Republic; Michael Liu, appledorem.g@gmail.com, Carnegie Mellon University, Pittsburgh, United States of America; Keenan Crane, kmcrane@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, United States of America and Roblox Research, San Mateo, United States of America; Minchen Li, minchernl@gmail.com, Carnegie Mellon University, Pittsburgh, United States of America and Genesis AI, Pittsburgh, United States of America.

ACM 0730-0301/2026/7-ART163
https://doi.org/10.1145/3811382

1 INTRODUCTION

Surface meshes are commonly used to represent the boundaries of physical shapes. In simulation, fabrication, and robotics, downstream algorithms implicitly assume that a surface separates interior from exterior, while geometry processing and learning-based pipelines often rely on representations that cannot encode self-intersections at all, such as signed distance fields or volumetric occupancy grids. Self-intersections violate these assumptions and invalidate many standard methods. As a result, self-intersecting meshes routinely break collision handling, produce non-manufacturable geometry, and corrupt data-driven pipelines. At the same time, large-scale 3D datasets assembled from scanning, reconstruction, and modeling pipelines contain self-intersections at scale [Deitke et al. 2023a,b] and are increasingly used for training and evaluation in geometry processing and generative modeling, including recent neural shape synthesis frameworks [Xiang et al. 2024]. As a result, enforcing surface validity is a practical challenge.

Despite decades of work, robustly removing self-intersections remains difficult. Many existing approaches treat self-intersections as local collision events and resolve them through discrete detection, projection, or constraint-based handling. While effective in simple cases, such methods often struggle when self-intersection involves multiple distant surface regions, complex entanglement, or highly degenerate geometry. Energy-based alternatives relying on repulsion or singular interaction kernels can be effective at preventing self-intersections, but are typically ill-suited for removing them: when initialized from intersecting configurations, these energies often become infinite, undefined, or numerically unstable. Practical implementations therefore rely on ad hoc techniques to weaken or regularize repulsion, frequently leading to unwanted geometric artifacts, residual intersections, or failure to converge. Even recent state-of-the-art methods illustrate that reliably eliminating all self-intersections remains challenging [Jang et al. 2025]. Robust self-intersection removal requires energies that are aggressively repulsive yet finite, and that can simultaneously address large-scale self-intersection structure and local geometric degeneracies.

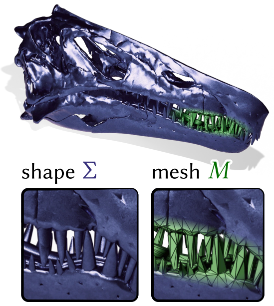

Fig. 2. Self-intersections may be introduced by simple mesh operations such as decimation (inspired by Figure 2 from [Sacht et al. 2013]).

To address these challenges, we introduce an energy-based framework that enforces surface embeddedness by coupling complementary *shape-level* and *mesh-level* components. At the shape level, interactions capture large-scale proximity across the surface, while the mesh-level formulation operates locally on discrete triangle interactions. We propose a new self-contact energy at the shape level designed specifically to detect and penalize large-scale self-intersection, capturing long-range entanglement independently of discretization. A key feature of this energy is that it can automatically adapt to the spatial scales most relevant for untangling a given model, providing the strongest effective "push" at each stage of the untangling process (Subsection 3.1.4). At the mesh level, we adapt Minkowski-based penalties [Minarčík et al. 2024] to polygonal face interactions in three dimensions, converting discrete triangle–triangle conflicts into smooth, differentiable terms suitable for continuous optimization. Our 3D Minkowski penalty formulation enables efficient evaluation of penetration depths and gradients. While the shape-level energy drives large-scale untangling, the mesh-level penalty resolves small-scale geometric degeneracies inherent to discrete meshes. Once self-intersections are eliminated, embeddedness can be maintained using standard intersection-aware optimization to further reduce deformation and restore surface appearance while preserving the untangled configuration.

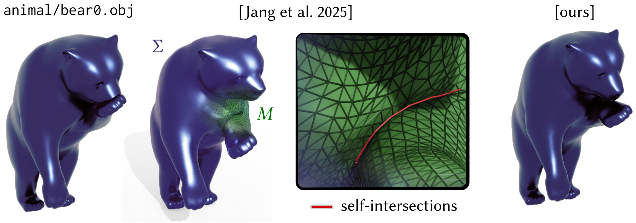

Fig. 3. The `animal/bear0.obj` mesh appears visually repaired after applying the method of Jang et al. [Jang et al. 2025] *(middle)*, yet non-adjacent triangle–triangle intersections persist in tightly folded crease regions at the mesh level, as revealed in the zoomed view. Our method removes these residual intersections and produces a fully embedded surface *(right)*.

Our formulation removes self-intersections without changing mesh connectivity and applies to a broad class of inputs, including non-manifold geometry, difficult corner cases such as cross-caps, and multi-object scenes with fixed components. Our method demonstrates consistent repair across challenging datasets and eliminates all self-intersections in the benchmark. By explicitly targeting embeddedness rather than collision resolution, our method provides a practical and robust framework for restoring valid surface geometry. We provide an open-source implementation of our method and all datasets used in this paper at https://github.com/Simulation-Intelligence/Mesh-Untangling.

2 STATE OF THE ART AND LIMITATIONS

Physics-based collision handling. Self-intersection handling is well studied in physics-based simulation using continuous or discrete collision detection coupled with penetration depth estimation and collision response [Baraff et al. 2003; Volino and Magnenat-Thalmann 2006; Ye et al. 2015]. Closely related interference-aware modeling methods maintain non-intersection during interactive deformation [Harmon et al. 2011]. These methods rely on temporal coherence or volumetric representations and are therefore not directly applicable to the repair of static surface meshes, where neither collision history nor a well-defined interior volume is available.

Cloth untangling. A substantial literature addresses untangling in cloth and garment simulation via global intersection analysis, contour minimization, and topology-based case classification [Baraff et al. 2003; Buffet et al. 2019; Cha and Ko 2022a,b; Wicke et al. 2006;

Ye et al. 2017; Ye and Zhao 2012]. These approaches are embedded in garment simulation pipelines and rely on assumptions such as layering, boundary conditions, orientation, or collision history, and target inter-object or boundary-driven collisions within a simulation context rather than general self-intersections of a single static surface mesh.

Volumetric and decomposition-based approaches. Several methods resolve self-intersections by converting surfaces into volumetric or cell-complex representations and applying volumetric processing or simulation [Li and Barbič 2018; Sacht et al. 2013]. While effective for producing simulation-ready representations, these pipelines are computationally expensive and not designed for fast, minimal-deformation surface repair. Moreover, they rely on a consistent notion of interior and exterior and locally well-defined surface normals, and therefore may break down for non-orientable surfaces (e.g., Möbius strips) or immersion failures (e.g., Whitney umbrellas). Related work enforces injectivity of maps, including 3D deformation processing [Fang et al. 2021; Li et al. 2020] and 2D parameterization [Du et al. 2021], but does not directly address self-intersection removal in static surface meshes.

Flow-based surface regularization. Flow-based surface regularization methods apply local geometric energies to smooth surface geometry, as in conformalized mean curvature flow (cMCF) [Sacht et al. 2013], sometimes followed by an explicit reverse flow intended to restore geometric detail. Such approaches, however, smooth geometry rather than explicitly enforcing embeddedness. In particular, because the driving energies are local, the forward flow is insensitive to large-scale self-intersection. As a result, configurations with multiple connected components may freely evolve into mutually intersecting shapes under the forward flow (see the inset example in Subsection 5.3). Moreover, the flow can drift far from the input geometry and pass into a different isotopy class, so that although the reverse flow may produce an embedded surface, it cannot recover one that remains close to the original shape or preserves the original configuration (see Subsection 5.5).

Energy-based repulsion. Repulsive formulations penalize spatially close but intrinsically distant surface elements and are well suited for intersection avoidance on embedded surfaces. Such double-integral repulsion energies originate in the knot energy literature [O'Hara 1991] and later inspired tangent-point energies and related smooth contact potentials in geometry processing [Huang et al. 2025; Sassen et al. 2024; Yu et al. 2021]. However, such energies are typically singular or ill-conditioned on non-embedded configurations, making them poorly suited for repair when initialized from self-intersecting inputs. Several methods therefore attempt to regularize or localize repulsion, for example by weakening interaction strength or restricting forces to detected intersections [Jang et al. 2025]. While these strategies yield fast and practical algorithms, residual self-intersections often persist in tightly folded regions (Fig. 3).

Summary and positioning. Our method enforces embeddedness by coupling a global *shape-level* self-contact energy with a complementary *mesh-level* intersection penalty. The formulation operates directly on static surface geometry, avoids volumetric discretization, and does not assume a simulation context or temporal coherence. It applies without modification to multi-object scenes and higher-dimensional ambient spaces (e.g., 2D surfaces in $\mathbb{R}^4$). We provide no theoretical guarantee that the method will always produce an embedded surface. In practice, however, it consistently resolves all self-intersections across a wide range of challenging inputs and substantially outperforms recent state-of-the-art methods (Section 5).

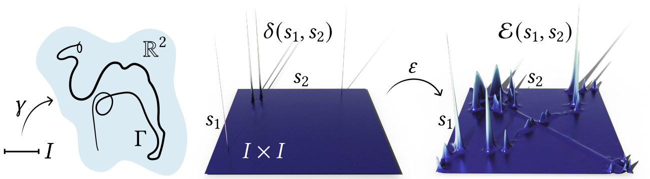

Fig. 4. Integrand $\mathcal{E}(s_1, s_2)$ of the Gaussian self-contact energy (1) for a planar curve Γ. Peaks in $I \times I$ correspond to self-intersections and converge to scaled Dirac distributions $\delta(s_1, s_2)$ as $\varepsilon \to 0$. Increasing ε smooths and spatially distributes these peaks, spreading gradient information along the curve and enabling globally aware untangling dynamics.

3 METHOD

We first describe the shape-level and mesh-level energies before combining them into a single untangling objective. Once self-intersections are eliminated, we apply standard intersection-aware optimization to reduce deformation while preserving the untangled configuration.

Notation. We distinguish between energies defined on continuous shapes and their discrete counterparts. We denote by $\Sigma \subset \mathbb{R}^3$ an underlying surface and by $M = (V, F)$ a triangular mesh approximating Σ. Energies defined on continuous shapes are written using parentheses, e.g., $\mathcal{E}(\Sigma)$, while energies evaluated on meshes are written using brackets, e.g., $\mathcal{E}[M]$. We assume Σ is a compact surface endowed with an area measure, not necessarily embedded, orientable, or even manifold (in particular, no tangent space or smooth structure is required). In the discrete mesh setting, embeddedness corresponds to the absence of self-intersections. Throughout this section, we use the term *self-contact* in a geometric sense to denote zero-separation interactions between geodesically distant regions of a shape.

3.1 Gaussian Self-Contact Energy (Shape Level)

We introduce a Gaussian self-contact energy designed specifically for removing self-intersections. In contrast to classical repulsive energies that are singular on intersecting configurations and primarily suited for prevention, this formulation remains finite while strongly penalizing close proximity between intrinsically distant regions of the shape. The Gaussian term provides a global signal for untangling and is complemented by a mesh-level Minkowski penalty (Section 3.2) that enforces separation at the discrete level.

3.1.1 Curves. We begin with planar curves to build intuition for the surface energy. Let $\Gamma \subset \mathbb{R}^2$ be a rectifiable planar curve parametrized by $\gamma : I \to \mathbb{R}^2$, allowing corners and self-intersections. In this setting, any such curve can be continuously untangled into an embedding.

We define the Gaussian self-contact energy on curves as

$$\mathcal{E}_{\mathrm{G}}(\Gamma;\varepsilon) = \iint_{I\times I} \frac{w(s_1,s_2)}{\varepsilon} \exp\left(-\frac{\|\gamma(s_1)-\gamma(s_2)\|^2}{\varepsilon^2}\right) ds_1\, ds_2. \quad (1)$$

Here s_1 and s_2 denote arc-length parameters along Γ. The normalization by ε ensures consistent scaling under uniform rescaling of the curve, matching the one-dimensional Hausdorff measure of Γ. The intrinsic weight function $w : I \times I \longrightarrow \mathbb{R}_+$ is symmetric, nonnegative, and vanishes on the diagonal ($s_1 = s_2$); throughout we use $w(s_1, s_2) = d_\Gamma(s_1, s_2)$, the intrinsic arc-length distance along the curve, which suppresses contributions from neighboring points. As illustrated in Fig. 4, the energy penalizes pairs of points that are close in Euclidean space but far along the curve, similar in spirit to classic knot energies such as the Möbius energy [Freedman et al. 1994] and tangent-point energies [Buck and Orloff 1995]. Like these formulations, self-intersections contribute strongly, while the energy remains finite on intersecting configurations and becomes increasingly steep as distinct branches approach one another.

3.1.2 Surfaces. For a surface $\Sigma \subset \mathbb{R}^3$ with area measure $d\mathcal{H}^2$, we define the Gaussian self-contact energy

$$\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon) = \iint_{\Sigma\times\Sigma} \frac{w(x,y)}{\varepsilon^p} \exp\left(-\frac{\|x-y\|^2}{\varepsilon^2}\right) d\mathcal{H}^2(x)\, d\mathcal{H}^2(y). \quad (2)$$

The parameter $\varepsilon > 0$ is the Gaussian kernel width; as $\varepsilon \longrightarrow 0$ the energy concentrates on self-intersecting regions of Σ, while larger ε distributes the contribution over a broader neighborhood (cf. Fig. 4). This interaction scale is induced by the kernel and its bandwidth, rather than being an intrinsic global property of the energy itself. Here $p > 0$ denotes the normalization exponent; for surfaces we set $p = 2$, matching the two-dimensional area measure so that the energy has consistent scaling under uniform rescaling of Σ. The intrinsic weight function $w : \Sigma\times\Sigma \longrightarrow \mathbb{R}_+$ is symmetric, nonnegative, and vanishes on the diagonal ($x = y$), suppressing contributions from pairs of nearby points along the surface; in the discrete setting (Section 3.5.2), we use a simple adjacency-based definition.

3.1.3 Global Adaptive Bandwidth. To select a suitable value of ε, we treat the bandwidth as a variational parameter and analyze the dependence of $\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon)$ on the scale parameter. For any fixed surface Σ, the energy vanishes as ε approaches 0 and $+\infty$: at very small scales, interactions are suppressed by the Gaussian kernel, while at very large scales they are washed out by normalization. Consequently, $\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon)$ necessarily attains a positive global maximum at an intermediate scale. This maximum provides a useful signal for identifying a characteristic interaction scale of the surface. To find this maximum, we differentiate with respect to ε,

$$\frac{d}{d\varepsilon}\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon) = \varepsilon^{-(p+3)}\left(2\,\mathcal{M}_2(\varepsilon) - p\,\varepsilon^2\mathcal{M}_0(\varepsilon)\right), \quad (3)$$

where $\mathcal{M}_k(\varepsilon)$ denote weighted Gaussian kernel moments,

$$\mathcal{M}_k(\varepsilon) = \iint_{\Sigma\times\Sigma} \|x-y\|^k\, w(x,y)\, \exp\left(-\frac{\|x-y\|^2}{\varepsilon^2}\right) d\mathcal{H}^2(x)\, d\mathcal{H}^2(y),$$

used here for $k = 0$ and $k = 2$. Critical points $\varepsilon^\star$ satisfy $\frac{d}{d\varepsilon}\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon^\star) = 0$, leading to the implicit scale-selection condition

$$(\varepsilon^\star)^2 = \frac{2}{p}\,\frac{\mathcal{M}_2(\varepsilon^\star)}{\mathcal{M}_0(\varepsilon^\star)} =: g(\varepsilon^\star)^2.$$

Local maxima of $\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon)$ correspond to stable fixed points of the relation $\varepsilon = g(\varepsilon)$, while local minima give rise to unstable fixed points. Depending on the geometry, $\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon)$ may admit multiple local maxima, each associated with a distinct interaction scale. We therefore compute candidate scales via fixed-point iteration $\varepsilon^{k+1} = g(\varepsilon^k)$. Details of the derivation and its discrete counterpart are provided in Appendix B. The update map g is nonnegative and bounded by $\operatorname{diam}(\Sigma)$, so the iteration remains within a geometrically meaningful range and converges locally near stable fixed points.

3.1.4 Local Adaptive Bandwidth. A single global bandwidth ε enforces a uniform interaction scale across the surface. In practice, surfaces often contain features at very different scales: thin folds or narrow handles may coexist with large smooth regions. Using one global scale forces these features to interact in the same way. We therefore consider spatially varying bandwidths $\varepsilon(x)$ and restrict attention to constructions that preserve symmetry of pairwise interactions and admit a well-defined energy. Among several possible constructions, we adopt a symmetric formulation in which a per-point bandwidth field $\varepsilon(x)$ is combined into a pairwise bandwidth $\varepsilon(x,y) = f(\varepsilon(x),\varepsilon(y))$. The combiner f is chosen to preserve symmetry, positivity, and scale invariance. This yields the energy

$$\mathcal{E}_{\mathrm{G}}(\Sigma;\varepsilon) = \iint_{\Sigma\times\Sigma} \frac{w(x,y)}{\varepsilon(x,y)^p} \exp\left(-\frac{\|x-y\|^2}{\varepsilon(x,y)^2}\right) d\mathcal{H}^2(x)\, d\mathcal{H}^2(y).$$

This formulation reduces to the global model when $\varepsilon(x) \equiv \varepsilon$ and preserves the variational structure of the energy. Alternative local constructions and their limitations are analyzed in Appendix B. We use the symmetric formulation throughout.

3.1.5 Discrete Formulation. We discretize the symmetric local bandwidth formulation on a triangular mesh $M = (V, F)$ with vertices $v_i \in \mathbb{R}^3$ and associated vertex areas A_i. The continuous bandwidth field $\varepsilon(x)$ is represented by per-vertex values $\{\varepsilon_i\}_{i=1}^N$, and pairwise distances are $r_{ij} = \|v_i - v_j\|$. Continuous weights $w(x,y)$ are evaluated discretely as w_{ij}.

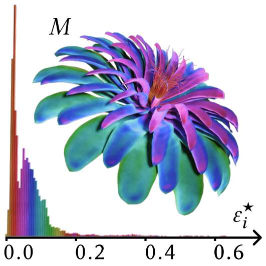

Fig. 5. Per-vertex optimal bandwidths $\varepsilon_i^\star$ mapped onto the mesh, with a histogram of their distribution. Fine central features yield smaller $\varepsilon_i^\star$.

Discrete Energy. We define pairwise bandwidths using the quadratic mean combiner $\varepsilon_{ij}^2 = \frac{1}{2}(\varepsilon_i^2 + \varepsilon_j^2)$. The discrete energy reads

$$\mathcal{E}_{\mathrm{G}}[M;\boldsymbol{\varepsilon}] = \sum_{i,j} A_i A_j w_{ij}\, \varepsilon_{ij}^{-p} \exp\left(-r_{ij}^2 \varepsilon_{ij}^{-2}\right). \quad (4)$$

To emphasize locality and reduce computational cost, sums over j may be restricted to a neighborhood $\mathcal{N}_i$.

Discrete bandwidth selection. For fixed geometry, optimal bandwidths are obtained by maximizing $\mathcal{E}_{\mathrm{G}}[M;\boldsymbol{\varepsilon}]$ with respect to the per-vertex values ε_i. In the symmetric local formulation, stationarity leads to a coupled nonlocal variational condition that is cumbersome to express explicitly in the continuous setting. After discretization, however, this condition admits a simple fixed-point form, closely mirroring the global adaptive bandwidth selection,

$$(\varepsilon_i^\star)^2 = \frac{4}{p}\,\frac{\sum_{j\in\mathcal{N}_i} b_{ij}\, r_{ij}^2}{\sum_{j\in\mathcal{N}_i} b_{ij}} - \frac{\sum_{j\in\mathcal{N}_i} b_{ij}\,(\varepsilon_j^\star)^2}{\sum_{j\in\mathcal{N}_i} b_{ij}} =:\; g_i(\boldsymbol{\varepsilon}^\star),$$

where $b_{ij} = A_i A_j w_{ij}\, \varepsilon_{ij}^{-(p+4)} \exp(-r_{ij}^2 \varepsilon_{ij}^{-2})$ and $\varepsilon_{ij} = f(\varepsilon_i, \varepsilon_j)$. As in the global case, optimal scales are obtained as fixed points of an update map $\varepsilon^2 = g(\varepsilon)$, but here the updates are coupled across vertices through neighborhood interactions. For comparison, the independent local construction yields N decoupled fixed-point relations of the same form. The derivation of this coupled system from the continuous variational formulation is given in Appendix B. In practice, we solve this system by damped Jacobi iteration, $(\varepsilon_i^{k+1})^2 = (1-\theta)(\varepsilon_i^k)^2 + \theta\, g_i(\boldsymbol{\varepsilon}^k)$, using $\theta = 0.25$. The effect of bandwidth adaptation on convergence is examined in Section 5.6.

3.2 Minkowski Penalty (Mesh Level)

The Gaussian self-contact energy operates at the level of the underlying shape and is effective at resolving large-scale entanglement between distant surface regions. However, on discrete meshes, self-intersections also arise from *local geometric degeneracies* such as thin spikes, near-coplanar folds, or poorly shaped triangles. To stabilize the optimization in these regimes, we introduce a complementary *mesh-level* penalty that acts directly on pairs of triangles. This term adapts the Minkowski-based penalty framework introduced in [Minarčík et al. 2024] to polygonal face interactions in three dimensions, yielding a smooth, differentiable measure of triangle–triangle separation suitable for continuous optimization. In particular, we provide a 3D triangle–triangle formulation enabling efficient evaluation of a smooth separation measure and its gradients.

Minkowski difference and separation. Let $A, B \subset \mathbb{R}^3$ be two triangles. Their *Minkowski difference* is the convex polytope

$$A - B = \{a - b \mid a \in A,\ b \in B\}. \tag{5}$$

A fundamental property of the Minkowski difference is

$$A \cap B \neq \emptyset \quad\Longleftrightarrow\quad 0 \in A - B. \tag{6}$$

Thus, intersection testing reduces to determining whether the origin lies inside a convex polytope. More importantly, the signed distance of the origin to $A - B$ provides a natural *penetration-depth-like* measure of overlap that depends only on relative configuration, without requiring volumetric inside-outside information (Figure 6).

Half-space formulation. A tempting approach to evaluating the signed distance to the Minkowski difference $A - B$ is to reason directly about its vertices. Since $A - B$ is the convex hull of all differences $a_i - b_j$, one could in principle construct this polytope explicitly and then compute its signed distance function. In practice, however, this route is unnecessarily indirect: explicit polytope construction introduces combinatorial complexity, and recovering the signed distance from such a representation requires additional geometric queries. Since our goal is only the signed distance to $A - B$, we bypass this intermediate step and instead represent the Minkowski difference as an intersection of supporting half-spaces,

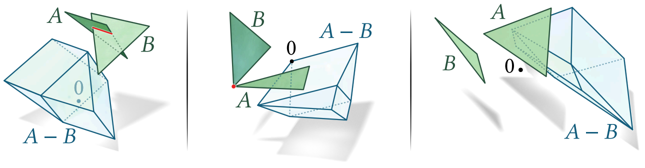

Fig. 6. Intersection, tangency, and separation of triangles A and B correspond to the origin lying inside, on the boundary of, or outside the Minkowski difference $A - B$ *(left to right)*. The resulting signed distance provides a smooth separation measure for mesh-level contact.

$$A - B \;=\; \bigcap_{n \in \mathcal{N}} H_n, \tag{7}$$

where each half-space $H_n = \{\, x \in \mathbb{R}^3 \mid \langle n, x\rangle \le h_n \,\}$ is defined by a unit normal direction n and support value $h_n = \max_{i,j} \langle n, a_i - b_j\rangle$.

The set of directions $\mathcal{N}$ encodes all possible face orientations of the Minkowski difference (Figure 7). Any face normal of $A - B$ must arise either from a triangle face of A or $-B$, or from an edge–edge interaction between A and B. Accordingly, $\mathcal{N}$ consists of the two triangle face normals together with all cross products $e_A \times e_B$ between edges e_A of A and edges e_B of B, normalized to unit length.

Signed distance at the origin. For a supporting half-space $H_n = \{x \mid \langle n, x\rangle \le h_n\}$ with unit normal n, we define the signed distance

$$\phi(x, H_n) = \langle n, x\rangle - h_n.$$

Using the finite set of candidate normals $\mathcal{N}$ (Figure 7), we evaluate the signed distance to the Minkowski difference at the origin by taking the maximum over these half-spaces,

$$\phi(0, A - B) := \max_{n \in \mathcal{N}} \phi(0, H_n). \tag{8}$$

For intersecting triangle pairs, $\phi(0, A - B) < 0$, while non-negative values indicate tangency or separation. Although (8) is defined using a finite normal set, it is exact in the regime relevant to the penalty: when $0 \in A - B$, the closest point lies on a supporting face whose normal is contained in $\mathcal{N}$, and the maximizing half-space recovers the correct signed distance. Gradients follow directly from the active half-space and are piecewise smooth.

Penalty and gradients. Let $\mathcal{P}$ denote the set of unordered pairs of non-adjacent triangles in the mesh,

$$\mathcal{P} := \bigl\{(A, B) \bigm| A, B \text{ are non-adjacent triangles of } M\bigr\}. \tag{9}$$

The mesh-level Minkowski penalty is defined as

$$\mathcal{E}_{\mathrm{M}}[M] = \sum_{(A,B)\in\mathcal{P}} \max\{0, -\phi(0, A - B)\}. \tag{10}$$

Only intersecting triangle pairs contribute to the sum, for which $0 \in A - B$ and $\phi(0, A - B) < 0$. In this regime, the signed distance at the origin is realized by a supporting face whose normal lies in the candidate set $\mathcal{N}$, so the half-space formulation recovers the correct penetration depth. Gradients follow from the active supporting half-space and are piecewise smooth, enabling efficient optimization.

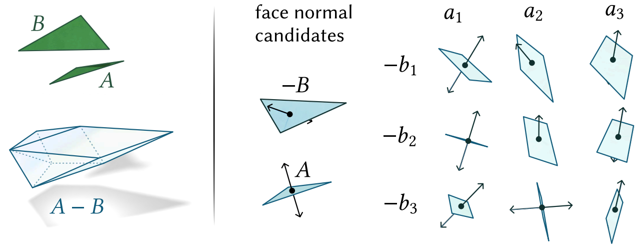

Fig. 7. Faces of the polytope $A-B$ admit only a small set of normal directions $\mathcal{N}$. These arise either from triangle faces of A or $-B$, or from edge–edge parallelogram directions induced by pairs of edges from A and B.

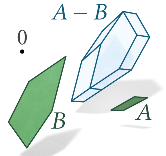

General polygonal faces. Although described for triangles, the mesh-level Minkowski penalty applies equally to meshes with general convex polygonal faces, such as quadrilaterals or mixed polygonal discretizations. The construction depends only on the relative configuration of the faces: the Minkowski difference $A - B$ remains a convex polytope, and the half-space surrogate is obtained by considering candidate normals induced by edge directions of the two faces. Compared to the triangular case, this enlarges the normal set $\mathcal{N}$ with no change to the formulation, evaluation, or optimization.

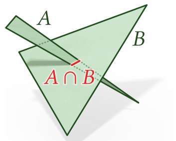

Why not intersection length? A natural alternative is to penalize the length of the intersection curve between intersecting triangles. While geometrically well defined and independent of volumetric information, *intersection length measures only the size of the contact set, not the distance to a non-intersecting configuration.* For instance, a thin triangle orthogonally piercing a larger one produces an arbitrarily short intersection segment despite requiring a large displacement to resolve the intersection. As a result, minimizing intersection length promotes tangential sliding or thinning rather than fast separation. Moreover, length-based gradients depend on the combinatorial structure of the intersection and are unstable under topological changes, requiring costly case-dependent differentiation.

3.3 Untangling Objective

We formulate surface untangling as the minimization of a single objective combining complementary shape-level, mesh-level, and regularization terms:

$$\mathcal{E}_{\text{untangle}}[M] = \lambda_{\text{G}}\,\mathcal{E}_{\text{G}}[M;\boldsymbol{\varepsilon}] + \lambda_{\text{M}}\,\mathcal{E}_{\text{M}}[M] + \lambda_{\text{reg}}\,\mathcal{E}_{\text{reg}}[M]. \tag{11}$$

Here, $\mathcal{E}\text{G}[M;\boldsymbol{\varepsilon}]$ denotes the discrete Gaussian self-contact energy with an adaptive bandwidth field $\boldsymbol{\varepsilon}$, $\mathcal{E}\text{M}[M]$ denotes the mesh-level Minkowski penalty, and $\mathcal{E}_{\text{reg}}[M]$ is a regularization term encouraging local shape preservation. The Gaussian self-contact energy drives large-scale untangling of the surface, enabling removal of global entanglement. The Minkowski penalty stabilizes local triangle–triangle interactions near folds, thin features, and other discrete degeneracies. As a regularizer, we use an as-rigid-as-possible (ARAP) energy [Alexa et al. 2000], which discourages unnecessary distortion and guides the surface toward a nearby configuration once intersections are resolved.

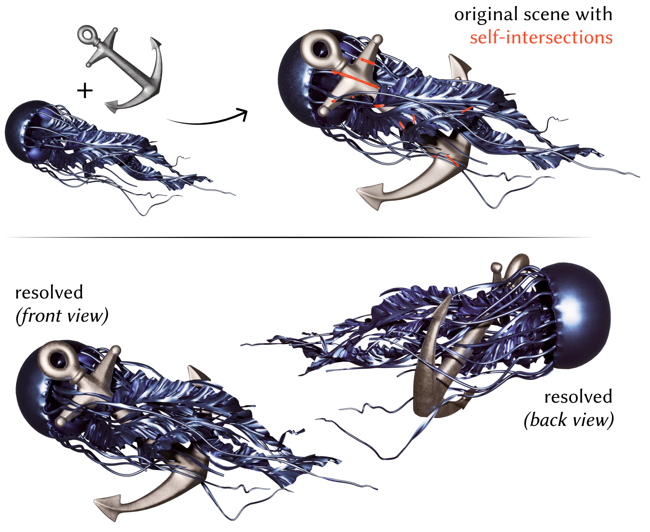

Fig. 8. A self-intersecting jellyfish mesh is optimized with a rigid anchor whose vertices are fixed by zeroing their gradients. The optimization removes intersections within and between objects while respecting non-moving geometry, enabling constrained scene arrangement.

3.4 Recovery Objective

The untangling process can introduce unwanted deformation to the mesh. To obtain an intersection-free mesh that is close to the initial shape and inspired by Guaranteed Globally Injective 3D Deformation Processing (IDP) [Fang et al. 2021], we perform shape recovery, or the reverse flow, with continuous collision detection (CCD) applied in every optimization step. Similar to the untangling objective, we formulate the recovery process as the minimization of the following objective:

$$\mathcal{E}_{\text{recovery}}[M] = \tilde{\lambda}_{\text{reg}}\,\tilde{\mathcal{E}}_{\text{reg}}[M] + \lambda_{\text{IPC}}\,\mathcal{E}_{\text{IPC}}[M], \tag{12}$$

where $\tilde{\mathcal{E}}_{\text{reg}}[M]$ denotes the regularization energy for the reverse flow. Similar to IDP, it consists of an ARAP energy computed per triangle face, a bending energy, and a quadratic penalty, while $\mathcal{E}_{\text{IPC}}[M]$ denotes the IPC barrier energy [Li et al. 2020] for avoiding interpenetration.

3.5 Optimization and Implementation Details

We summarize the key optimization and implementation choices in this section; further details are provided in the supplementary material.

3.5.1 Solver. For the untangling stage, we use the L-BFGS solver [Liu and Nocedal 1989] with $\beta_0 = 10^{-4}$ for the Armijo condition parameter. We employ a backtracking line search that enforces the Armijo sufficient-decrease condition, initialized with step size $\alpha = 1$ and iteratively halved until the condition is satisfied or $\alpha < 10^{-10}$, at which point the line search terminates. The solver terminates when no non-adjacent triangle–triangle intersections are detected.

For the reverse-flow stage, we use the idea of artificial time-stepping in IDP, and we run for fixed 500 Newton steps. To provide strong non-intersection guarantees, we employ an intersection-aware backtracking line search in the style of IPC. Each Newton step is solved using a preconditioned conjugate gradient method with a relative tolerance of 10^{-6} and a diagonal Jacobi preconditioner.

3.5.2 Discrete Intrinsic Weight Function. In the discrete setting, we use a deliberately simple intrinsic weight function w, with $w_{ij} = 0$ for adjacent vertices and $w_{ij} = 1$ otherwise. This yields compact closed-form gradients and efficient optimization.

3.5.3 Explicit Gradients. To avoid automatic differentiation and improve performance, we derive explicit closed-form expressions for the gradient of the Gaussian self-contact energy. Quadrature weights $a_{ij} = A_i A_j w_{ij}$ are treated as constant during differentiation and are recomputed periodically to reflect geometric changes. Let $\{v_i\}_{i=1}^N \subset \mathbb{R}^3$ denote the mesh vertices, and let $r_{ij} = \|v_i - v_j\|$ denote their pairwise distances. The gradient with respect to a vertex v_i is

$$\nabla_{v_i} \mathcal{E}_{\mathrm{G}}[M; \varepsilon] = -\frac{2}{\varepsilon^{p+2}} \sum_{j \neq i} a_{ij} \exp\left(-r_{ij}^2 \varepsilon^{-2}\right) (v_i - v_j),$$

which corresponds to Gaussian-weighted pairwise forces that increase separation between spatially close regions.

3.5.4 Neighborhood Truncation. To accelerate evaluation of the Gaussian self-contact energy, we restrict pairwise interactions to a finite neighborhood. Specifically, for each vertex v_i we consider only vertices v_j satisfying $r_{ij} \leq \mathcal{N}$ and discard all pairs with $r_{ij} > \mathcal{N}$. Since the Gaussian kernel decays exponentially in r_{ij}/ε, contributions from distant pairs are numerically negligible, and this truncation does not affect optimization behavior while significantly reducing computational cost. In practice, $\mathcal{N}$ is chosen as a fixed multiple of the bandwidth ε.

3.5.5 Scaling and Units. Under uniform scaling $M \mapsto XM$, pairwise distances and vertex areas scale as $r_{ij} \mapsto X r_{ij}$ and $A_i \mapsto X^2 A_i$. For surfaces ($p = 2$) with bandwidth scaled as $\varepsilon \mapsto X\varepsilon$, the Gaussian self-contact energy scales as X^3 when using an intrinsic weight function w, and as X^2 for scale-free weights. The Minkowski penalty $\mathcal{E}_{\mathrm{M}}[M]$ scales linearly. In our experiments, we address differing scaling behavior by rescaling all meshes to a common unit scale.

4 APPLICATIONS

Once surface embeddedness can be reliably enforced, a wide class of otherwise difficult problems reduces to tractable geometric optimizations. We present representative applications spanning practical geometry processing as well as mathematically motivated constructions arising in geometric modeling and topology.

4.1 Dataset Cleaning

Self-intersecting meshes are pervasive in large-scale datasets of 3D shapes used in geometry processing and generative modeling, and violate the assumption that a surface represents a valid physical shape, corrupting downstream simulation, fabrication, and learning-based pipelines [Xiang et al. 2024]. Our method enables dataset cleaning by removing all self-intersections across diverse benchmarks spanning generative models, animation assets, and analytic constructions (subsection 5.2), while remaining effective on non-manifold meshes, immersion failures, and severely degenerate geometries (see subsection 5.4).

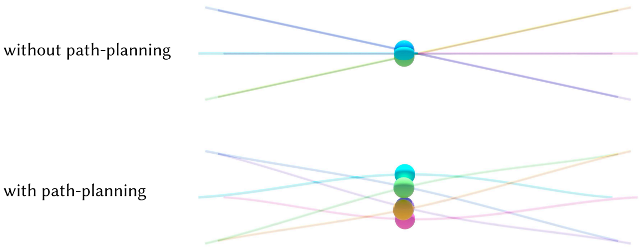

Fig. 9. Initial planned routes for six agents (*top*) and the corresponding untangled routes (*bottom*).

4.2 Constrained Scene Arrangement

Our formulation naturally extends to scenes composed of multiple independent meshes, including configurations with prescribed non-moving geometry. By zeroing gradients of selected vertices, individual components can be treated as rigid anchors while the remaining geometry is optimized jointly. This enables the removal of intersections within individual meshes as well as between distinct objects while respecting fixed constraints (Fig. 8). Note that this capability is not supported by [Sacht et al. 2013], whose conformal mean curvature flow treats meshes independently and does not account for other objects, allowing convergence to configurations with multiple overlapping spheres and preventing fixed geometry. This behavior follows from enforcing embeddedness at the level of the combined scene rather than locally or per object.

4.3 3D Multi-Agent Path Planning

With anchored constraints, our framework naturally extends to multi-agent 3D path planning: given an initial route for each agent with prescribed start and goal positions, we jointly optimize the routes to eliminate inter-route intersections while penalizing deviations from the initial plans. To demonstrate this idea, we model six agents as spheres and represent each agent's swept path as a cylindrical volume. We anchor the two circular end caps to fix the start and goal positions. We then apply our untangling and recovery pipeline to these cylinder volumes, yielding non-intersecting trajectories while keeping each route close to its initial plan. The initial route and the resulting route are shown in (Fig. 9), and we observe that the agents' routes have zero intersections in a spiral shape while staying close to its originally planned route. We note that our method does not guarantee intersection-free agent motion, since the bounding volumes may deform during optimization. Nevertheless, it may be possible in future work to impose additional constraints so that each volume remains a uniform bound on the agent's swept region.

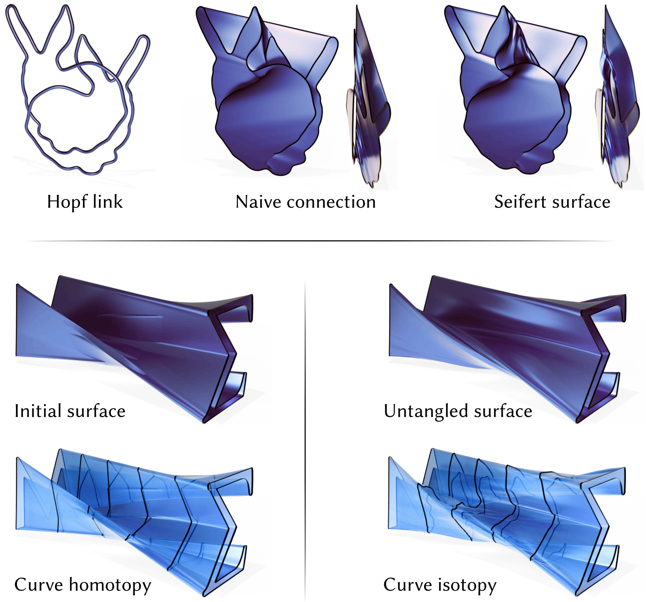

Fig. 10. We demonstrate how surface untangling with fixed boundary can be used in two related geometric constructions. In the top row, a Hopf-linked pair of planar curves is connected by a naive spanning surface, which is untangled into an embedded Seifert surface realizing the same prescribed boundary. In the bottom row, we construct an explicit isotopy between two embedded planar curves by first connecting them with a naive annular surface and then untangling it into an embedded annulus, whose planar cross-sections trace the isotopy.

4.4 Construction of Embedded Seifert Surfaces

A *Seifert surface* for a knot or link $\Gamma \subset \mathbb{R}^3$ is a compact, oriented, embedded surface Σ satisfying $\partial\Sigma = \Gamma$. Such surfaces are fundamental objects in low-dimensional topology, where they define classical invariants such as knot genus, and they arise naturally in geometric modeling and scientific visualization as explicit spanning surfaces for knotted curves and fields. A classical result due to Seifert gives an explicit combinatorial construction from knot diagrams by resolving crossings into Seifert circles and attaching twisted bands [Seifert 1934]. While this algorithm guarantees the existence of a spanning surface and fixes its topology, the resulting geometry is largely determined by the band structure. Subsequent work has shown how such surfaces can be smoothed and improved geometrically while preserving embeddedness [van Wijk and Cohen 2005], but there is limited control over the surface shape when the boundary curve is prescribed as a smooth embedding in $\mathbb{R}^3$. We demonstrate how our energy-based framework can be used as a geometric optimization stage in this setting. Given a prescribed boundary Γ and an initial spanning surface Σ_0 satisfying $\partial\Sigma_0 = \Gamma$ that may contain self-intersections, we optimize an energy combining untangling terms with curvature-based regularization while enforcing the boundary constraint $\partial\Sigma_t = \Gamma$ throughout the evolution. The optimization removes self-intersections and produces an embedded surface Σ realizing the same boundary. Constructing an initial spanning surface with the desired topology remains a separate problem; our method addresses the geometric stage of transforming a given surface into an embedded Seifert surface with fixed boundary. In Fig. 10, we demonstrate this procedure on a Hopf link.

4.5 Explicit Isotopies Between Embedded Curves

Beyond static surface repair, our method can also be viewed as operating on surfaces that encode a continuous deformation in one higher dimension. We illustrate this using isotopies between planar curves, which can be represented as embedded surfaces in $\mathbb{R}^3$. Let $\Gamma_1, \Gamma_2 \subset \mathbb{R}^2$ be closed embedded curves. An explicit isotopy between them can be described by a continuous map

$$\gamma : S^1 \times [0,1] \to \mathbb{R}^2,$$

such that $\gamma(\cdot, 0)$ and $\gamma(\cdot, 1)$ parametrize Γ_1 and Γ_2, respectively, and $\gamma(\cdot, t)$ is embedded for all $t \in [0,1]$. By the Jordan–Schönflies theorem, any two embedded closed planar curves admit an ambient deformation in the plane. We construct such deformations by embedding Γ_1 and Γ_2 as boundary curves in parallel planes in $\mathbb{R}^3$ and initializing an annular surface Σ_0 with boundary $\partial\Sigma_0 = \Gamma_1 \sqcup \Gamma_2$, which may be self-intersecting. Minimizing an energy that enforces surface embeddedness while keeping the boundary fixed yields an embedded annulus Σ. Intersections are resolved at the level of the full space–time surface, and planar cross-sections of Σ directly produce an explicit isotopy between Γ_1 and Γ_2, as shown in Fig. 10. In the planar setting, curve shortening flow preserves embeddedness and, by Grayson's theorem, evolves any embedded closed curve to a convex one [Grayson 1987], suggesting an alternative route to isotopy construction. However, such flows impose strong geometric smoothing and provide limited control over intermediate configurations, feature preservation, and pointwise correspondence along the curve. In contrast, the surface-based formulation enforces embeddedness globally in space–time and supports additional geometric constraints through the energy, making it suitable for controlled deformation trajectories.

Higher-Dimensional Isotopies. The same viewpoint extends naturally beyond planar curves. Given two embeddings $K_1, K_2 \subset \mathbb{R}^3$ of the same knot, a deformation between them can be represented as a surface embedded in one higher dimension. In this setting, isotopies correspond to embedded cobordisms in $\mathbb{R}^4$, where the additional dimension plays the role of a deformation parameter. While existence results guarantee that such deformations exist, constructing explicit and geometrically controlled isotopies between knot embeddings remains a challenging and active area of research [Döhrer et al. 2025]. The surface-based formulation presented here reframes knot deformations as an embeddedness problem in one higher dimension.

5 RESULTS

Our results show that enforcing embeddedness at both the shape and mesh levels yields complete removal of self-intersections in cases where prior state-of-the-art methods leave residual defects. Beyond such comparisons, we demonstrate robustness on a wide range of challenging inputs arising in practice, including non-manifold, non-orientable, non-immersed, and multi-component geometries. We

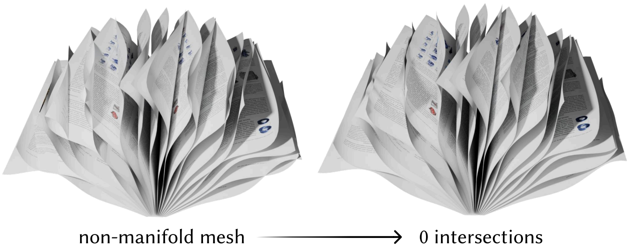

Fig. 11. A book model composed of multiple pages welded along a shared spine, forming non-manifold junctions and extensive self-intersections *(left)* is untangled while preserving the non-manifold connectivity *(right)*.

further present ablations that clarify the roles and tradeoffs of the individual components of our formulation.

Implementation. We implement our system in C++ and CUDA (Toolkit 13.1) and evaluate it on Ubuntu 24.04. All benchmark experiments for comparisons are run on a workstation equipped with an Intel Core i9-13900KF CPU (24 cores, up to 5.8 GHz) and an NVIDIA GeForce RTX 4090 GPU. Our implementation additionally relies on NVIDIA Thrust and cuSPARSE for parallel primitives and sparse linear algebra.

5.1 Comparison to Instant Repair

The closest related work and current state-of-the-art method is *Instant Repair* [Jang et al. 2025], which addresses the same problem of removing self-intersections in static surface meshes. Their method adapts Tangent Point Energy, originally designed for self-intersection prevention, to the repair setting. While their approach resolves a large fraction of intersections, most meshes in their dataset are not fully repaired. In practice, partial solutions are insufficient, as mesh repair is a preprocessing step for pipelines that require fully embedded surfaces. Empirically, unresolved cases concentrate in regions of fine-scale geometric complexity, including small mesh degeneracies, creases, and high-curvature areas, where smooth global energies struggle to enforce local disentanglement due to cancellation of diffused gradients. In contrast, our method explicitly targets such configurations and resolves all self-intersections in the dataset. Furthermore, since we resolve all self-intersections, we can obtain fully embedded meshes with small deformation by applying the reverse flow. Finally, unlike Tangent Point Energy formulations, our objective remains well-defined for non-manifold intermediate states and immersion failures such as cross-caps, as it does not rely on the existence of a well-defined tangent space.

5.2 Benchmark Overview

Beyond the ISIR benchmark, we evaluate our method on a set of application-driven datasets spanning generated meshes, animation assets, and mathematically constructed surfaces. Together, these datasets cover a wide range of geometric complexity, articulation, and self-intersection patterns, allowing us to assess robustness across both practical and extreme scenarios. Below, we briefly describe each dataset.

SMAL. A parametric articulated animal model with consistent mesh topology across poses and shapes. We evaluate five categories (lions, cows, dogs, hippos, and horses), totaling 49 meshes. Large articulation-induced deformations frequently introduce self-intersections, making this dataset representative of non-rigid animation scenarios.

SCAPE. [Dragomir et al. 2005] A data-driven human body model that produces meshes of human bodies in a range of articulated poses. The dataset contains 71 meshes; to focus on untangling behavior, we retain the 68 meshes that exhibit self-intersections. We normalize this dataset so that the maximum size of its bounding box is 1.

Microsoft Rocketbox Avatar Library. [Gonzalez-Franco et al. 2020] A collection of fully rigged, production-style character meshes with substantial variation in shape and pose. We retain the largest connected component of each mesh, yielding 140 single-component models. Intersection-free meshes are included to reflect realistic, unfiltered pipeline settings.

Analytic Surfaces. A curated set of 26 analytically constructed surfaces exhibiting extreme self-intersection behavior, including non-orientable and highly degenerate configurations (see Section A and Figs. 20, 21, 19).

Throughout the benchmark, we compare against the open-source implementation of Instant Repair (ISIR). To ensure a meaningful comparison, we increase its iteration cap to 15,000 per case, select parameters to maximize its success rate on each dataset, and report its best result over the run. This brings ISIR's runtime into the same order of magnitude as ours and reflects its strongest attainable performance under generous computational budgets.

For our method, we use the same L-BFGS configuration for the untangling stage and the same Newton solver and energy-weight settings for the reverse-flow stage across all datasets. Within each dataset, we use a shared choice of weights for λ_M, λ_{reg} and neighborhood size N in the majority of cases, and adjust them only for a small number of extreme examples. Across all benchmarks, the method consistently produces final surfaces with zero non-adjacent triangle-triangle intersections (Table 1). In practice, shared parameter settings resolve all cases in several datasets (e.g., SCAPE and Analytic Surfaces) and suffice for the majority of meshes in the remaining benchmarks.

5.3 Comparison to Conformalized Mean Curvature Flow

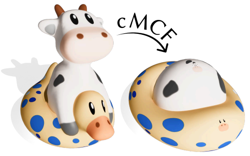

Conformalized mean curvature flow (cMCF) [Sacht et al. 2013] is a natural first idea for addressing self-intersections, but it is fundamentally not designed for this task. Any removal of self-intersections is an indirect byproduct of geometric smoothing rather than an explicit notion of embeddedness, and the method is known to introduce severe distortions, including global shape collapse and orientation flips (Fig. 14). As acknowledged in the original work, cMCF is formulated for closed genus-zero surfaces and does not naturally extend to surfaces with boundary (Fig. 12), non-manifold geometry (Fig. 11), or constrained regions. Moreover, cMCF treats meshes independently:

Table 1. Benchmark Results: Our Method vs. *Instant Repair*.

Dataset Name	Method	Dataset Size	Fully Resolved Meshes	Unique Intersecting Triangles	Time (s)
ISIR Benchmark	Ours	60	**60**	**0**	2217.20
ISIR Benchmark	*Instant Repair*	60	11	8331	2815.47
SMAL	Ours	49	**49**	**0**	1409.23
SMAL	*Instant Repair*	49	0	27892	2406.82
Analytic Surfaces	Ours	26	**26**	**0**	1062.61
Analytic Surfaces	*Instant Repair*	26	6	22975	627.79
Microsoft Rocketbox	Ours	140	**140**	**0**	4678.35
Microsoft Rocketbox	*Instant Repair*	*140	22	11335	4983.71
SCAPE	Ours	68	**68**	**0**	1706.11
SCAPE	*Instant Repair*	68	0	5193	4605.40

*Two models containing quad-faces in Microsoft Rocketbox are ignored for *Instant Repair*.

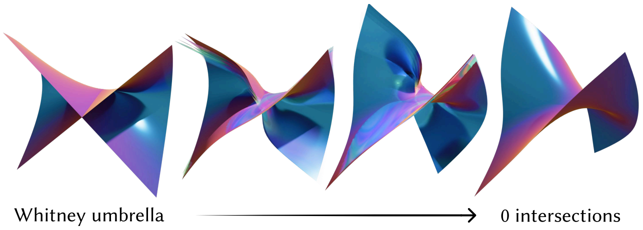

Fig. 12. Minimizing Gaussian self-contact energy untangles surfaces with singularities. Starting from a Whitney umbrella with a crosscap *(left)*, the flow removes self-intersections and produces an embedded surface *(right)*.

in multi-object scenes, surfaces do not interact and may converge to mutually intersecting configurations, making it unsuitable for constrained or multi-object untangling problems (Fig. 8).

5.4 Robustness to Challenging Geometry

Our method operates without assuming manifoldness or immersion of the input surface and therefore handles non-manifold meshes and singular geometry. This includes meshes with multi-face junctions and sheet-like intersections (Figure 11), surfaces with immersion failures such as the Whitney umbrella with a cross-cap singularity (Figure 12), geometrically degenerate but topologically simple surfaces such as the cylindrical constructions in Figure 21, as well as other examples throughout the paper that contain non-manifold configurations, including the multi-object jellyfish scene (Figure 8). This robustness follows from the formulation of the untangling energy: both the Gaussian self-contact energy and the Minkowski penalty depend only on embedded distances between surface elements, while the regularization term relies solely on local adjacency information, without requiring consistent normals, curvature, orientability, or a manifold differential structure.

5.5 Why the Untangling Phase Is Necessary

One might ask whether reverse flow could be applied starting from a simple embedded surface consistent with the mesh topology (e.g., a sphere or torus), bypassing the energy optimization used to remove self-intersections. Reverse flow, however, is constrained by isotopy: it can recover the target surface only when initialized from an embedding in the same isotopy class. Initializing from a simple embedded surface in a different isotopy class is topologically obstructed and cannot converge to the desired geometry (Figure 13). Moreover, for self-intersecting inputs the isotopy class is not defined a priori, since the surface is not embedded. The energy optimization therefore plays a necessary role by resolving self-intersections while remaining close to the input geometry, producing a valid embedded initialization from which reverse flow can meaningfully operate.

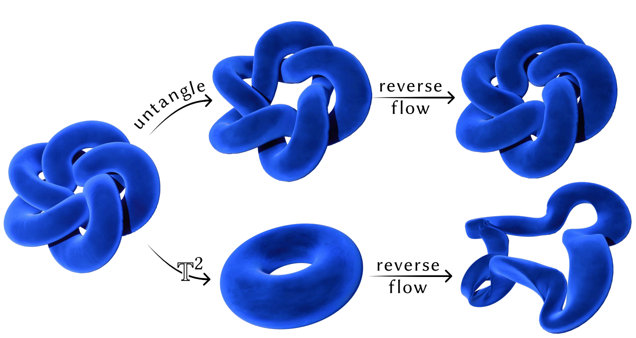

Fig. 13. The untangle phase yields a nearby embedding from which reverse flow can recover the original shape *(top)*. In contrast, when initialized from a simple shape in a different isotopy class, such as a torus, reverse flow cannot reach the target surface, regardless of optimization *(bottom)*.

5.6 Adaptive Bandwidth Ablation

Bandwidth Scale. The bandwidth parameter ε plays a decisive role in determining not only the rate of untangling but also the qualitative outcome of the deformation. Figure 14 illustrates this effect on the self-intersecting *Flying Saucer* surface introduced by Sacht et al. [2013]. For large bandwidths ($\varepsilon \gg \varepsilon^\star$), self-intersections are resolved through a global inside-out flip of the surface, consistent with the behavior observed for conformalized mean curvature flow. In contrast, using the optimal bandwidth $\varepsilon^\star$ removes all self-intersections without inducing a global orientation flip, producing a closed surface with less global deformation.

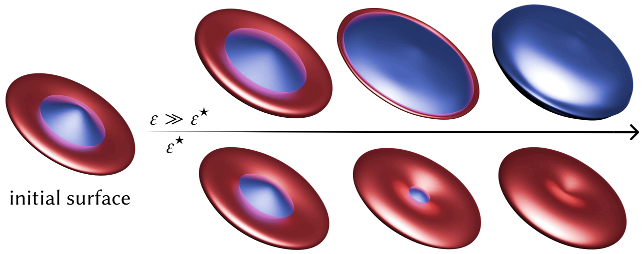

Fig. 14. Starting from the self-intersecting *Flying Saucer* [Sacht et al. 2013] *(left)*, a large bandwidth ($\varepsilon \approx 10\,\varepsilon^\star$) resolves intersections by flipping the surface orientation, as reported for cMCF. Using the optimal bandwidth $\varepsilon^\star$ instead removes all self-intersections while preserving orientation.

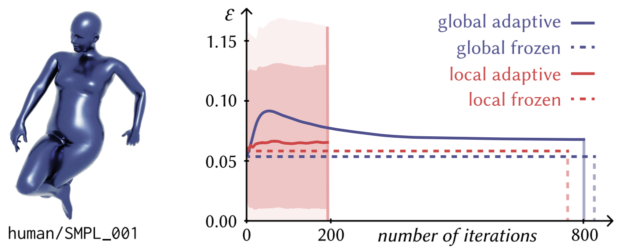

Fig. 15. Evolution of the bandwidth parameter ε for an example mesh `human/SMPL_001`. Local adaptive bandwidth converges in 195 iterations, compared to 773 iterations for frozen local ε, while global bandwidths require 808 iterations when adaptive and 835 when frozen.

Bandwidth Adaptation. We compare bandwidth strategies in which ε is either held fixed throughout optimization (*frozen*) or updated at each iteration based on the current geometry (*adaptive*), considering both local (per-vertex) and global (single scalar) formulations. We evaluate these variants on meshes from [Jang et al. 2025], measuring the number of solver iterations required to reach an embedded configuration. Median iteration counts are 225 for local-adaptive, 245 for global-adaptive, 1036 for local-frozen, and 1504 for global-frozen bandwidths. Adaptive bandwidths therefore converge in substantially fewer iterations, with local adaptation performing best. Bandwidth updates incur a per-iteration overhead (below 20%), which is offset by the reduction in total iterations (Fig. 15).

5.7 Do We Need Both Interaction Energies?

Robust surface untangling requires reasoning about self-intersections at both the level of the continuous shape and the discrete mesh. Figures 16 and 17 illustrate complementary failure modes when only one interaction energy is used. Mesh-level penalties efficiently resolve local collisions but cannot induce the coordinated global motion required in deeply interpenetrating configurations, while shape-level energies capture large-scale entanglement yet may overlook fine-scale mesh degeneracies. Combining both terms is therefore essential for robust untangling across geometric scales.

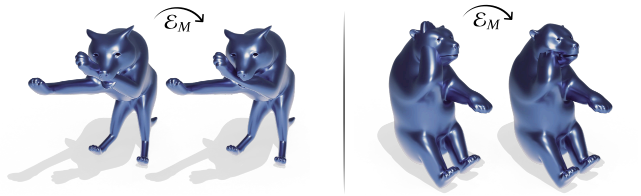

Fig. 16. The `wolf3` mesh from [Jang et al. 2025] can be untangled using the Minkowski penalty $\mathcal{E}_M$ alone, as local interactions are sufficient to resolve its self-intersections *(left)*. In contrast, the `bear3` mesh exhibits a deep penetration of the paw into the head that cannot be resolved by $\mathcal{E}_M$ alone, since purely local forces fail to induce the required global motion *(right)*; a global shape energy term is required. Overall, 33 out of 60 meshes in this dataset are successfully untangled using $\mathcal{E}_M$ alone.

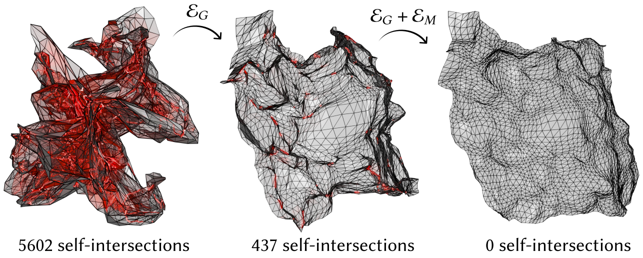

Fig. 17. To stress test the limitations of shape-level interaction energies, we generate an extremely crumpled surface by applying high-frequency noise to a planar square mesh. Minimizing the Gaussian self-intersection energy $\mathcal{E}_G$ alone substantially reduces large-scale penetrations but fails to eliminate all self-intersections due to locally degenerate mesh configurations *(middle)*. Adding the mesh-level Minkowski penalty $\mathcal{E}_M$ resolves the remaining intersections, yielding a fully embedded surface *(right)*.

5.8 Effect of Regularization Strength

The regularization weight λ_{reg} controls the trade-off between shape fidelity and the efficiency of untangling. As shown in Fig. 18, larger values preserve the original geometry more closely but require more iterations to eliminate all self-intersections, while smaller values allow faster untangling at the cost of larger intermediate deformation. In practice, we intentionally use weak regularization during the forward untangling phase, permitting substantial deformation to resolve complex self-intersections. Because the method includes a reverse flow that maintains embeddedness and does not alter mesh connectivity, the original geometry can subsequently be recovered within the same optimization framework.

6 LIMITATIONS AND FUTURE WORK

Topological Obstructions. Our approach, like any method based on continuous deformation, is limited by topology. In particular, surfaces that admit no embedding in $\mathbb{R}^3$, such as the Klein bottle or the real projective plane $\mathbb{RP}^2$, cannot be transformed into an embedded configuration by any geometric process. In these cases, self-intersections reflect a fundamental obstruction rather than removable geometric entanglement.

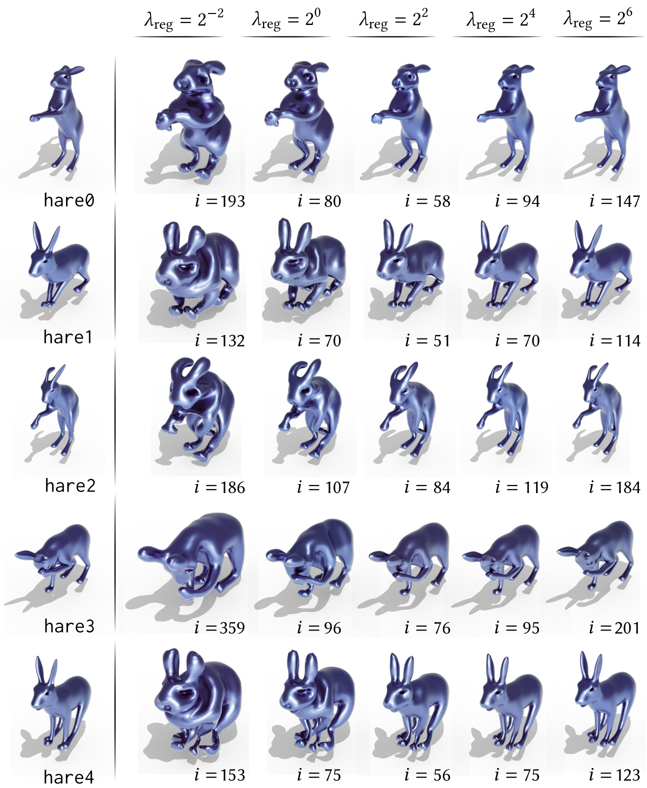

Fig. 18. Effect of the regularization weight λ_{reg} on untangling behavior. Input meshes with self-intersections are shown *(left)*, and untangled results obtained with increasing values of λ_{reg} are shown *(right)*. The number of iterations i required to eliminate all self-intersections is reported for each run. Very small λ_{reg} can also slow convergence due to excessive deformation.

Extreme Polyhedral Geometries. To stress test the untangling framework, we applied our method to highly constrained high-genus polyhedral surfaces, motivated by recent constructions of symmetric polyhedral embeddings of regular maps with genus $2 \leq g \leq 14$ [Bokowski and H. 2025]. Starting from known embedded configurations reported in [Bokowski and H. 2025], the method recovers a nearby embedded surface after small perturbations *(inset top)*, for example for the `R3.2_D3` surface under small random noise. In contrast, when initialized far from any known embedding, such as from random vertex positions for the `R3.1_T` topology, the optimization does not converge to an embedded realization *(inset bottom)*, although the number of self-intersections is substantially reduced (from 243 to 108). The constructions in [Bokowski and H. 2025] rely on specialized numerical optimization and symmetry exploitation; incorporating an untangling energy with suitable regularization could complement such approaches and potentially accelerate or systematize similar searches.

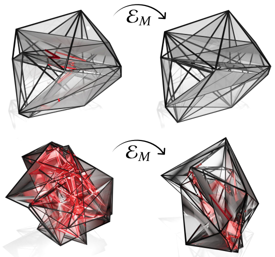

Higher-Dimensional and Space-Time Extensions. Several directions opened by this work lie beyond the three-dimensional setting. As discussed in Sec. 4.5, interpreting one spatial dimension as a deformation parameter allows isotopies to be formulated as embedded surfaces in a higher dimension. For example, in knot theory, deformations between two embeddings of the same knot in $\mathbb{R}^3$ correspond to embedded cobordisms in $\mathbb{R}^4$. Although existence results guarantee that such deformations exist, constructing explicit and geometrically controlled knot deformations remains difficult and is an active area of research [Döhrer et al. 2025]. More broadly, treating surfaces as space–time objects connects naturally to swept volumes and deformation trajectories, suggesting potential applications in shape evolution and motion planning.

7 CONCLUSION

We presented an energy-based framework for enforcing embeddedness by coupling complementary *shape-level* and *mesh-level* energies within a single variational formulation. The shape-level term detects and penalizes large-scale self-intersections independently of mesh resolution, while the mesh-level penalty regularizes local discrete intersections. Together, these components enable consistent removal of self-intersections without changing mesh connectivity across a broad range of inputs, including highly entangled surfaces, non-manifold geometry, and multi-object scenes with fixed components. More broadly, reliable enforcement of embeddedness transforms many otherwise ill-posed surface manipulation and deformation problems into tractable geometric optimizations.

ACKNOWLEDGMENTS

Minchen Li acknowledges partial support from a Junior Faculty Startup Fund from Carnegie Mellon University and gift funding from Genesis AI. Jiří Minarčík acknowledges support from the Fulbright–Masaryk Scholarship. Keenan Crane acknowledges support from NSF awards 2212290 and 2504890.

REFERENCES

Marc Alexa, Daniel Cohen-Or, and David Levin. 2000. As-rigid-as-possible shape interpolation. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., USA, 157–164. https://doi.org/10.1145/344779.344859

Thomas F. Banchoff. 1974. Triple points and surgery of immersed surfaces. *Proc. Amer. Math. Soc.* 46, 3 (1974), 407–413. https://doi.org/10.1090/S0002-9939-1974-0377897-1

David Baraff, Andrew Witkin, and Michael Kass. 2003. Untangling cloth. In *ACM SIGGRAPH 2003 Papers* (San Diego, California) *(SIGGRAPH '03)*. Association for Computing Machinery, New York, NY, USA, 862–870. https://doi.org/10.1145/1201775.882357

Jürgen Bokowski and Kevin H. 2025. Polyhedral Embeddings of Triangular Regular Maps of Genus g, 2 <= g <= 14, and Neighborly Spatial Polyhedra. *Symmetry* 17, 4 (2025). https://doi.org/10.3390/sym17040622

Gregory Buck and Jeremey Orloff. 1995. A simple energy function for knots. *Topology Appl.* 61, 3 (1995), 205–214. https://doi.org/10.1016/0166-8641(94)00024-W

Thomas Buffet, Damien Rohmer, Loïc Barthe, Laurence Boissieux, and Marie-Paule Cani. 2019. Implicit untangling: a robust solution for modeling layered clothing. *ACM Trans. Graph.* 38, 4, Article 120 (jul 2019), 12 pages. https://doi.org/10.1145/3306346.3323010

Ick-Hoon Cha and Hyeong-Seok Ko. 2022a. BLI-resolver: resolving the boundary-loop-interior type intersections for clothing simulation. *The Visual Computer* 38 (2022), 1359–1368. https://doi.org/10.1007/s00371-021-02279-z

Ick-Hoon Cha and Hyeong-Seok Ko. 2022b. Tanglement Resolution in Clothing Simulation With Explicit Convergence. *IEEE Transactions on Visualization and Computer Graphics* 28, 7 (July 2022), 2764–2775. https://doi.org/10.1109/TVCG.2020.3039566

Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani,

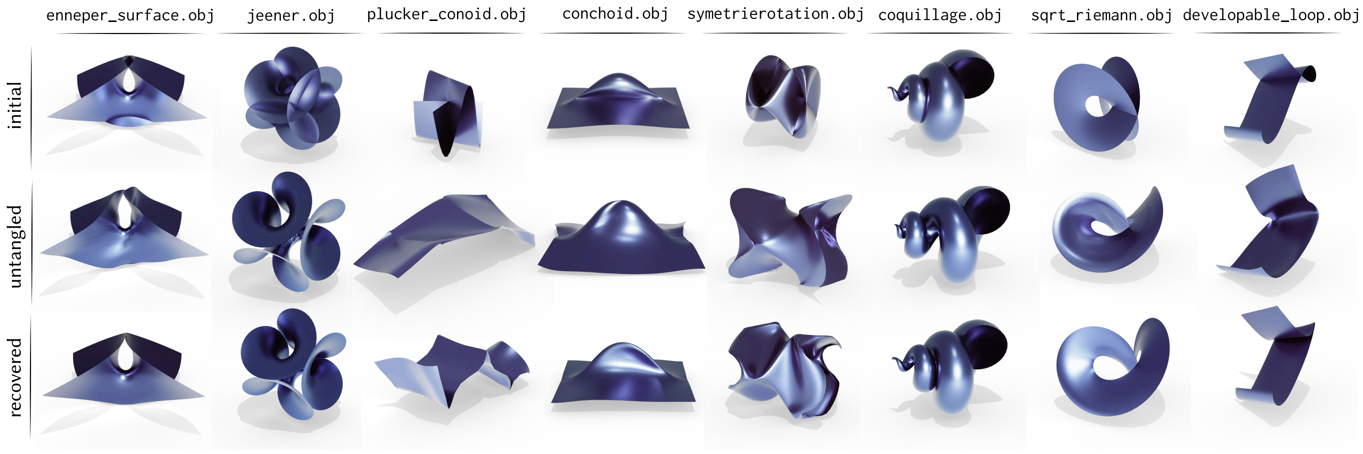

Fig. 19. Representative subset of our mathematical surface dataset. Each column shows one surface in three states: the original self-intersecting input *(top)*, the fully untangled embedded configuration *(middle)*, and the final result after reversing the flow while preserving embeddedness *(bottom)*. The examples span diverse surface types and qualitatively different intersection configurations, illustrating the variety of untangling scenarios covered by the benchmark.

Ludwig Schmidt, and Ali Farhadi. 2023a. Objaverse-XL: A Universe of 10M+ 3D Objects. *arXiv preprint arXiv:2307.05663* (2023).

Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. 2023b. Objaverse: A Universe of Annotated 3D Objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* 13142–13153.

Elias Döhrer, Philipp Reiter, and Henrik Schumacher. 2025. On a Complete Riemannian Metric on the Space of Embedded Curves. (2025). arXiv:2501.16647 [math.DG]

Anguelov Dragomir, Srinivasan Praveen, Koller Daphne, Thrun Sebastian, Rodgers Jim, and Davis James. 2005. SCAPE. *ACM Transactions on Graphics* 24, 3 (07 2005), 408–416. https://doi.org/10.1145/1073204.1073207

Xingyi Du, Danny M. Kaufman, Qingnan Zhou, Shahar Z. Kovalsky, Yajie Yan, Noam Aigerman, and Tao Ju. 2021. Optimizing global injectivity for constrained parameterization. *ACM Trans. Graph.* 40, 6, Article 260 (Dec. 2021), 18 pages. https://doi.org/10.1145/3478513.3480556

Yu Fang, Minchen Li, Chenfanfu Jiang, and Danny M. Kaufman. 2021. Guaranteed Globally Injective 3D Deformation Processing. *ACM Trans. Graph. (SIGGRAPH)* 40, 4, Article 75 (2021).

Michael H. Freedman, Zheng-Xu He, and Zhenghan Wang. 1994. Möbius energy of knots and unknots. *Annals of Mathematics (2)* 139, 1 (1994), 1–50. https://doi.org/10.2307/2946626

Martin Golubitsky and Victor Guillemin. 1973. *Stable Mappings and Their Singularities.* Graduate Texts in Mathematics, Vol. 14. Springer-Verlag, New York. https://doi.org/10.1007/978-1-4615-7904-5

Mar Gonzalez-Franco, Eyal Ofek, Ye Pan, Angus Antley, Anthony Steed, Bernhard Spanlang, Antonella Maselli, Domna Banakou, Nuria Pelechano, Sergio Orts-Escolano, Veronica Orvalho, Laura Trutoiu, Markus Wojcik, Maria V. Sanchez-Vives, Jeremy Bailenson, Mel Slater, and Jaron Lanier. 2020. The Rocketbox Library and the Utility of Freely Available Rigged Avatars. *Frontiers in Virtual Reality* Volume 1 - 2020 (2020). https://doi.org/10.3389/frvir.2020.561558

Matthew A. Grayson. 1987. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry* 26, 2 (1987), 285–314. https://doi.org/10.4310/jdg/1214441371

David Harmon, Daniele Panozzo, Olga Sorkine, and Denis Zorin. 2011. Interference-aware geometric modeling. In *Proceedings of the 2011 SIGGRAPH Asia Conference* (Hong Kong, China) *(SA '11)*. Association for Computing Machinery, New York, NY, USA, Article 137, 10 pages. https://doi.org/10.1145/2024156.2024171

Zizhou Huang, Maxwell Paik, Zachary Ferguson, Daniele Panozzo, and Denis Zorin. 2025. Geometric Contact Potential. *ACM Trans. Graph.* 44, 4, Article 158 (2025), 24 pages. https://doi.org/10.1145/3731142

Wonjong Jang, Yucheol Jung, Gyeongmin Lee, and Seungyong Lee. 2025. Instant Self-Intersection Repair for 3D Meshes. *ACM Trans. Graph.* 44, 4, Article 143 (July 2025), 14 pages. https://doi.org/10.1145/3731427

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Trans. Graph. (SIGGRAPH)* 39, 4, Article 49 (2020).

Yijing Li and Jernej Barbič. 2018. Immersion of Self-Intersecting Solids and Surfaces. *ACM Transactions on Graphics* 37, 4 (2018), 45:1–45:14. https://doi.org/10.1145/3197517.3201327

Dong C. Liu and Jorge Nocedal. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical Programming* 45, 1 (1989), 503–528. https://doi.org/10.1007/BF01589116

Jiří Minarčík, Sam Estep, Wode Ni, and Keenan Crane. 2024. Minkowski Penalties: Robust Differentiable Constraint Enforcement for Vector Graphics. *ACM SIGGRAPH 2024 Conference Proceedings* 43, 4 (2024). https://doi.org/10.1145/3641519.3657495

Jun O'Hara. 1991. Energy of a knot. *Topology* 30, 2 (1991), 241–247. https://doi.org/10.1016/0040-9383(91)90022-5

Ulrich Pinkall. 1985. Regular Homotopy Classes of Immersed Surfaces. *Topology* 24, 4 (1985), 421–434.

Leonardo Sacht, Alec Jacobson, Daniele Panozzo, Christian Schüller, and Olga Sorkine-Hornung. 2013. Consistent Volumetric Discretizations Inside Self-Intersecting Surfaces. *Computer Graphics Forum (proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing)* 32, 5 (2013), 147–156.

Josua Sassen, Henrik Schumacher, Martin Rumpf, and Keenan Crane. 2024. Repulsive Shells. *ACM Transactions on Graphics* 43, 4 (2024), 140:1–140:22. https://doi.org/10.1145/3658174 Best Paper Award, SIGGRAPH 2024.

H. Seifert. 1934. Über das Geschlecht von Knoten. *Math. Ann.* 110 (1934), 571–592.

Jarke J. van Wijk and Arjeh M. Cohen. 2005. Visualization of the genus of knots. In *Proceedings IEEE Visualization 2005.* IEEE Computer Society Press, 567–574.

Pascal Volino and Nadia Magnenat-Thalmann. 2006. Resolving surface collisions through intersection contour minimization. *ACM Trans. Graph.* 25, 3 (jul 2006), 1154–1159. https://doi.org/10.1145/1141911.1142007

Hassler Whitney. 1944. The self-intersections of a smooth n-manifold in $2n$-space. *Annals of Mathematics, Second Series* 45, 2 (1944), 220–246. https://doi.org/10.2307/1969265

Martin Wicke, Hermes Lanker, and Markus Gross. 2006. Untangling cloth with boundaries. In *Proc. of Vision, Modeling, and Visualization*, Vol. 349. 356.

Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. 2024. Structured 3D Latents for Scalable and Versatile 3D Generation. *arXiv preprint arXiv:2412.01506* (2024).

Juntao Ye, Guanghui Ma, Liguo Jiang, Lan Chen, Jituo Li, Gang Xiong, Xiaopeng Zhang, and Min Tang. 2017. A Unified Cloth Untangling Framework Through Discrete Collision Detection. *Computer Graphics Forum* 36, 7 (2017), 217–228. https://doi.org/10.1111/cgf.13287 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13287

Juntao Ye, Timo R. Nyberg, and Gang Xiong. 2015. Fast Discrete Intersection Detection for Cloth Penetration Resolution. In *2015 IEEE International Conference on Multimedia Big Data.* 352–357. https://doi.org/10.1109/BigMM.2015.22

Juntao Ye and Jing Zhao. 2012. The intersection contour minimization method for untangling oriented deformable surfaces. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation.* 311–316.

Chris Yu, Caleb Brakensiek, Henrik Schumacher, and Keenan Crane. 2021. Repulsive Surfaces. *ACM Transactions on Graphics* 40, 6 (2021), 268:1–268:19. https://doi.org/10.1145/3478513.3480521

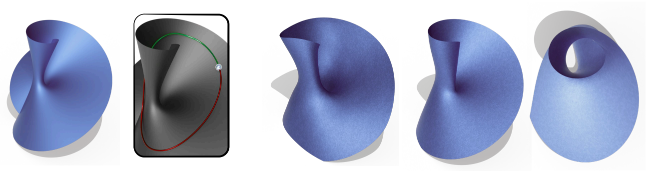

initial *non-orientable* surface embedded reverse flow *(front and top)*

Fig. 20. Extending the parameter domain of a Möbius surface yields a severely self-intersecting, intrinsically non-orientable mesh for which inside–outside reasoning is undefined. Our method recovers an embedded surface in $\mathbb{R}^3$ without altering topology.

A MATHEMATICAL SURFACE DATASET

We propose a mathematical surface dataset of 26 meshes designed to systematically expose challenging self-intersection phenomena. The dataset spans orientable and non-orientable surfaces, surfaces with and without boundary, varying genus, and multiple boundary components. Beyond topological diversity, the meshes exhibit qualitatively different intersection structures, including isolated double curves, branching intersection networks, triple points, and intersection endpoints at singularities (Fig. 19).

Surfaces beyond the immersion regime. Classical results in immersion theory, such as classification up to regular homotopy [Pinkall 1985] and parity laws for triple points [Banchoff 1974], provide strong global invariants but apply only to immersed surfaces. In contrast, surfaces arising in geometric modeling and mathematical visualization routinely violate immersion assumptions, exhibiting stable singularities such as cross-caps and creases. Singularity theory following Whitney [Whitney 1944] and Mather, as presented for example in Golubitsky–Guillemin [Golubitsky and Guillemin 1973], characterizes the local structure of these singularities but does not constrain their global arrangement, allowing complex self-intersection configurations to arise naturally on a single surface. These phenomena are deliberately represented in the dataset.

B VARIATIONAL ANALYSIS OF ADAPTIVE BANDWIDTH

We treat the bandwidth parameters as variational degrees of freedom and derive optimal values from stationarity of the discrete Gaussian self-contact energy. The resulting conditions show that the bandwidth corresponds to a weighted root-mean-square of pairwise distances, yielding an intrinsic interaction scale.

B.1 Global Bandwidth (Discrete Derivation)

The discrete Gaussian interaction energy is

$$\mathcal{E}_\varepsilon[M] = \sum_{i,j} A_i A_j w_{ij} \, \frac{1}{\varepsilon^p} \exp\left(-\varepsilon^{-2} r_{ij}^2\right),$$

where $r_{ij} = \|v_i - v_j\|$. For fixed geometry, the bandwidth ε controls which pairwise distances contribute significantly to the energy. Small values emphasize only very close pairs, while large values distribute the interaction over a broader range of distances. Differentiating with respect to ε gives

$$\frac{d}{d\varepsilon}\mathcal{E}_\varepsilon = \frac{1}{\varepsilon^{p+3}} \sum_{i,j} A_i A_j w_{ij} \, \exp\left(-\varepsilon^{-2} r_{ij}^2\right) \left(2r_{ij}^2 - p\,\varepsilon^2\right).$$

Setting $\frac{d}{d\varepsilon}\mathcal{E}_\varepsilon = 0$ yields the stationarity condition

$$(\varepsilon^\star)^2 = \frac{2}{p} \frac{\sum_{i,j} r_{ij}^2 A_i A_j w_{ij} \exp[-(\varepsilon^\star)^{-2} r_{ij}^2]}{\sum_{i,j} A_i A_j w_{ij} \exp[-(\varepsilon^\star)^{-2} r_{ij}^2]}.$$

The optimal scale $\varepsilon^\star$ equals $\sqrt{2/p}$ times the Gaussian-weighted RMS of pairwise distances, consistent with $\varepsilon^2 = 2\mathcal{M}_2/(p\mathcal{M}_0)$. A damped fixed-point iteration is

$$\varepsilon_{k+1}^2 \leftarrow (1-\theta)\varepsilon_k^2 + \theta \frac{2}{p} \frac{\sum_{i,j} r_{ij}^2 A_i A_j w_{ij} \exp(-\varepsilon_k^{-2} r_{ij}^2)}{\sum_{i,j} A_i A_j w_{ij} \exp(-\varepsilon_k^{-2} r_{ij}^2)}.$$

B.2 Fixed-Point Interpretation

The update defines a nonlinear map $\varepsilon = g(\varepsilon)$. Since pairwise distances are bounded by $\mathrm{diam}(M)$, the map is nonnegative and bounded, and the iteration remains within a geometrically meaningful range. In practice, a damped update with $\theta < 1$ is used for stability.

B.3 Local Bandwidth Formulations

To capture spatial variation, we introduce per-vertex bandwidths ε_i and consider ways of localizing the global variational problem.

(I) Independent local bandwidth. A direct extension assigns to each vertex i a local energy depending only on interactions centered at i,

$$\mathcal{E}_i(\varepsilon_i) = \frac{A_i}{\varepsilon_i^p} \sum_j A_j w_{ij} \exp\left(-\varepsilon_i^{-2} r_{ij}^2\right),$$

and optimizes $\mathcal{E} = \sum_i \mathcal{E}_i$ with respect to each ε_i independently. Differentiating $\mathcal{E}_i$ with respect to ε_i gives

$$\partial_{\varepsilon_i}\mathcal{E}_i = \varepsilon_i^{-(p+3)} \sum_j A_i A_j w_{ij} \exp\left(-\varepsilon_i^{-2} r_{ij}^2\right) \left(2r_{ij}^2 - p\varepsilon_i^2\right),$$

and setting $\partial_{\varepsilon_i}\mathcal{E}_i = 0$ yields the corresponding stationarity condition

$$(\varepsilon_i^\star)^2 = \frac{2}{p} \frac{\sum_j r_{ij}^2 A_j w_{ij} \exp[-(\varepsilon_i^\star)^{-2} r_{ij}^2]}{\sum_j A_j w_{ij} \exp[-(\varepsilon_i^\star)^{-2} r_{ij}^2]}.$$

This produces decoupled updates, but the interaction is asymmetric in (i, j) and does not arise from a symmetric pairwise energy.

(II) Symmetric local bandwidth (used). To preserve symmetry, we instead define pairwise bandwidths $\varepsilon_{ij} = f(\varepsilon_i, \varepsilon_j)$ and optimize the symmetric energy

$$\mathcal{E}[M; \boldsymbol{\varepsilon}] = \sum_{i,j} A_i A_j w_{ij} \, \varepsilon_{ij}^{-p} \exp\left(-\varepsilon_{ij}^{-2} r_{ij}^2\right).$$

In this formulation, each ε_i influences all incident pairs (i, j), leading to a coupled variational problem. The derivative with respect to ε_i involves contributions from all such pairs and yields a nonlinear system whose explicit form is derived in the following subsection.

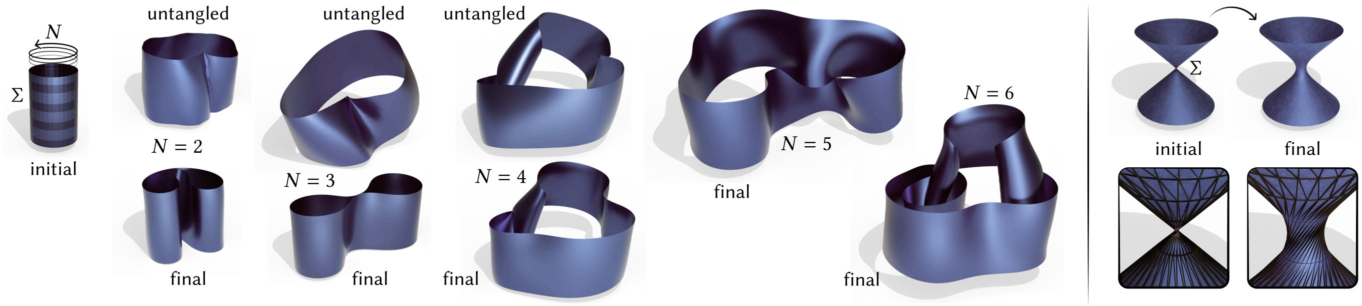

Fig. 21. Our method resolves two extreme cylindrical degeneracies: a cylinder wrapped N times around its axis, producing dense self-intersections *(left)*, and a cylinder with a smoothly inverted vertical radius that collapses an entire cross section to a single point *(right)*. After resolution, the singular region locally opens into a shape reminiscent of a ruled hyperboloid *(bottom right)*, while the rest of the surface undergoes only minimal deformation.

Table 2. Bandwidth models for the Gaussian self-contact energy: a global scalar fixed-point iteration, independent per-vertex updates (I), a coupled per-vertex system (II), and a closed-form pairwise solution that yields a diverging energy (III). Model (II) is used in all experiments.

	MODEL	BANDWIDTH	ENERGY
	Global	$\varepsilon^\star = g(\varepsilon^\star)$	$\sum_{i,j} \mathcal{E}_{ij}(\varepsilon^\star)$
(I)	Independent local	$\varepsilon_i^\star = g_i(\varepsilon_i^\star)$	$\sum_i \mathcal{E}_i(\varepsilon_i^\star)$
(II)	**Symmetric local**	$\varepsilon_i^\star = g_i(\boldsymbol{\varepsilon}^\star)$	$\sum_{i,j} \mathcal{E}_{ij}(f(\varepsilon_i^\star, \varepsilon_j^\star))$
(III)	Pairwise optimal	$\varepsilon_{ij}^\star \propto r_{ij}$	*diverges*

(III) Pairwise optimal bandwidth. A limiting construction treats each ε_{ij} as an independent variable and maximizes the energy with respect to each pair. Differentiating with respect to ε_{ij} gives

$$\partial_{\varepsilon_{ij}}\left(\varepsilon_{ij}^{-p}\exp\left(-\varepsilon_{ij}^{-2}r_{ij}^2\right)\right)=\varepsilon_{ij}^{-(p+3)}\exp\left(-\varepsilon_{ij}^{-2}r_{ij}^2\right)\left(2r_{ij}^2-p\varepsilon_{ij}^2\right),$$

and setting this to zero yields $(\varepsilon_{ij}^\star)^2 = \frac{2}{p}r_{ij}^2$. Substituting this back into the energy eliminates the Gaussian kernel and produces a singular Riesz interaction proportional to r_{ij}^{-p}, which diverges as $r_{ij} \longrightarrow 0$ and is unsuitable for regularized self-contact.

B.4 Symmetric Formulation: Discrete Derivation

Differentiating $\mathcal{E}$ with respect to ε_i gives

$$\partial_{\varepsilon_i}\mathcal{E}=\sum_j A_iA_jw_{ij}\varepsilon_{ij}^{-(p+4)}\exp\left(-\varepsilon_{ij}^{-2}r_{ij}^2\right)\left(2r_{ij}^2-p\varepsilon_{ij}^2\right)\cdot\partial_{\varepsilon_i}\varepsilon_{ij}^2.$$

For the quadratic mean, $\varepsilon_{ij}^2 = \frac{1}{2}(\varepsilon_i^2 + \varepsilon_j^2)$ and $\partial_{\varepsilon_i}\varepsilon_{ij}^2 = \varepsilon_i$, so substitution yields the stationarity condition

$$\varepsilon_i\sum_j A_iA_jw_{ij}\varepsilon_{ij}^{-(p+4)}\exp\left(-\varepsilon_{ij}^{-2}r_{ij}^2\right)\left(2r_{ij}^2-p\varepsilon_{ij}^2\right)=0.$$

Introducing the weights $b_{ij} = A_iA_jw_{ij}\,\varepsilon_{ij}^{-(p+4)}\exp(-\varepsilon_{ij}^{-2}r_{ij}^2)$ allows the expression to be written in the compact form

$$(\varepsilon_i^\star)^2=\frac{4}{p}\frac{\sum_j b_{ij}r_{ij}^2}{\sum_j b_{ij}}-\frac{\sum_j b_{ij}(\varepsilon_j^\star)^2}{\sum_j b_{ij}}.$$

The corresponding damped Jacobi iteration is

$$(\varepsilon_i^{k+1})^2\leftarrow(1-\theta)(\varepsilon_i^k)^2+\theta\left(\frac{4}{p}\frac{\sum_j b_{ij}^k r_{ij}^2}{\sum_j b_{ij}^k}-\frac{\sum_j b_{ij}^k(\varepsilon_j^k)^2}{\sum_j b_{ij}^k}\right).$$

B.5 Comparison of Bandwidth Models

The global formulation yields a single RMS interaction scale and cannot resolve spatial variation. The independent local formulation decouples updates but breaks symmetry and lacks a consistent pairwise interpretation. The symmetric local formulation preserves symmetry and variational structure while allowing spatial adaptation, producing a coupled nonlinear system. The pairwise optimal formulation reduces to a singular Riesz interaction and diverges at collisions. The symmetric local formulation is the only one that preserves symmetry, admits a variational interpretation, and yields finite interaction. In all cases, the effective interaction scale is set by the kernel bandwidth, not by the algebraic form of the energy.