

Type-Free Global Intersection Analysis with Linear Displacement Fields

CHENGZHU HE, Xiamen University, China

XUDONG FENG, Zhejiang University, China

ANJUN CHEN, Xiamen University, China

DAN SONG, Tianjin University, China

SHIHUI GUO*, Xiamen University, China

KUI WU, Independent Researcher, United States of America



Fig. 1. Our self-intersection repair method resolves complex entanglements in thin-shell dynamics without topological classification. Top: A square cloth folds its corners inward, creating deep and nested self-intersections. Bottom: A complex garment drops onto the ground, concentrating thousands of intersection contours. By optimizing geometric separation directions directly, our algorithm achieves robust intersection repair, enabling practical, scalable simulation in highly complex intersection scenarios.

*Corresponding author.

Authors' Contact Information: [Chengzhu He](mailto:chengzhuhe@stu.xmu.edu.cn), Xiamen University, Xiamen, China, chengzhuhe@stu.xmu.edu.cn; [Xudong Feng](mailto:xudongfeng18@gmail.com), Zhejiang University, Hangzhou, China, xudongfeng18@gmail.com; [Anjun Chen](mailto:anjunchen@xmu.edu.cn), Xiamen University, Xiamen, China, anjunchen@xmu.edu.cn; [Dan Song](mailto:dan.song@tju.edu.cn), Tianjin University, Tianjin, China, dan.song@tju.edu.cn; [Shihui Guo](mailto:guoshihui@xmu.edu.cn), Xiamen University, Xiamen, China, guoshihui@xmu.edu.cn; [Kui Wu](mailto:kui.wu@gmail.com), Independent Researcher, Los Angeles, CA, United States of America, [walker.kui.wu@gmail.com](mailto:kui.wu@gmail.com).

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

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 1557-7368/2026/9-ART

<https://doi.org/10.1145/3842>

History-independent intersection repair is a critical bottleneck in physical simulation and geometric processing. Existing global intersection analysis (GIA) methods identify invalid regions of relatively small area, but they rely on type-specific heuristics for different contour categories, while local intersection contour minimization (ICM) methods are simpler to deploy but often stall in local minima under deep or nested penetrations. In this paper, we present a novel untangling framework that formulates intersection resolution as a minimum-displacement separation-direction problem. For each intersection contour, we test a set of candidate linear displacement directions using VF/EE/FV ray casting and filter the raw hits using cluster culling based on primitive adjacency to find the valid hits belonging to the penetration region. We rank candidate directions by displacement cost, and refine the most promising direction with projected Newton updates on the sphere. Finally, we convert the valid hit support into standard proximity-style penalty pairs and integrate them into the standard physics simulator. We evaluate our method on diverse contour configurations and self-intersection

benchmarks. Compared with the GIA and ICM baselines, our method resolves intersections more consistently, particularly in challenging cases involving deep and nested layers. The results show that our method is a practical, solver-compatible untangling pipeline for complex mesh intersections.

CCS Concepts: • **Computing methodologies** → **Physical simulation**.

Additional Key Words and Phrases: Untangling, collision response, computational geometry, intersection detection, numerical optimization

ACM Reference Format:

Chengzhu He, Xudong Feng, Anjun Chen, Dan Song, Shihui Guo, and Kui Wu. 2026. Type-Free Global Intersection Analysis with Linear Displacement Fields. *ACM Trans. Graph.* 1, 1 (September 2026), 17 pages. <https://doi.org/10.1145/3842>

1 Introduction

Interactive applications increasingly require *history-independent* methods for repairing initial intersections caused by user editing, scenario generation, retargeting, or previous solver failures. This problem is particularly challenging for cloth and thin-shell geometries, where reliable inside–outside tests are often unavailable, and the intersecting regions may form deep, nested, or self-interlocking configurations.

Classical untangling methods can be broadly divided into global intersection analysis (GIA) and local intersection contour minimization (ICM). GIA methods [Baraff et al. 2003; Cha and Ko 2020, 2022; Wicke et al. 2006] reason about intersection topology by identifying invalid regions through topological splitting and flood-filling. They provide stronger global support construction, but usually require contour-type-specific rules to determine how the flood fill should be performed. As a result, their behavior can be sensitive to missing or duplicated discrete intersections, and the required case analysis becomes difficult to extend to complex, degenerate, or partially non-manifold configurations. In contrast, ICM methods [Volino and Magnenat-Thalmann 2006; Ye and Zhao 2012] optimize local geometric objectives around the intersection contour. They are simple and highly parallelizable, but their reliance on local contour geometry makes them prone to poor local minima, especially under deep or nested penetrations where shortening the contour does not necessarily imply separation.

To bridge this divide, we propose an untangling framework that constructs separation fields directly from contour-seeded ray casting, eliminating the need for discrete contour-type classification and type-specific illegal-region rules. We formulate intersection resolution as a **direction-selection problem** that minimizes the untangling displacement cost, yielding a unified response for both multi-body penetrations and complex self-intersections. Candidate directions are generated from geometric configurations, evaluated based on displacement cost, and further optimized through spherical Newton refinement. To obtain the coherent response regions, we enforce hit adjacency on both mesh sides and apply an adaptive k -ring restriction around the detected contour. We incorporate the resulting penalty pairs into an implicit Newton/PCG solver and validate them on both canonical examples and large-scale self-intersection benchmarks. The source code of our implementation can be found at <https://github.com/ChengzhuUwU/Untangling26>.

2 Related Work

2.1 Physics Simulation

Soft- and rigid-body dynamics. Large-step implicit integration [Baraff and Witkin 1998] and robust contact handling [Bridson et al. 2002] established the practical baseline for production-scale dynamics. Subsequent work on time integration [Gast et al. 2015; Li et al. 2019; Mattos Da Silva et al. 2025] further improved the stability and accuracy of numerical stepping. On the constitutive side, stable material laws and finite-element formulations [Chen et al. 2024a; Hughes 2012; Kim 2020; Kim and Eberle 2020; Macklin and Muller 2021; Smith et al. 2018, 2019; Volino et al. 2009], discrete shell and bending models [Feng et al. 2022; Grinspun et al. 2003; Wang et al. 2023b; Wu and Kim 2023], and specialized cloth, yarn-level, and rod descriptions [Choi and Ko 2005; Feng et al. 2024; Hsu et al. 2025; Liu et al. 2013; Sperl et al. 2020; Yuan et al. 2026] have broadened the range of phenomena that can be captured faithfully. At the solver level, efficient preconditioners and multigrid methods [Chen et al. 2021; Tamstorf et al. 2015; Wang et al. 2018b; Wu et al. 2022], projective and position-based dynamics [Bouaziz et al. 2014; Fratarcangeli et al. 2016; Macklin et al. 2016, 2019; Müller 2008; Müller et al. 2007, 2020], and high-performance nonlinear updates [Chen et al. 2024b; Lan et al. 2023, 2022b; Liu et al. 2017; Tournier et al. 2015; Wang 2015] make large scenes tractable.

Collision detection and contact response. On the detection side, BVH traversal strategies [Karras 2012; Wang et al. 2018a] cull the broad phase, while continuous collision detection (CCD) conservatively identifies the time of impact for moving primitives [Li et al. 2021]. Recent work has tightened CCD guarantees through exact root parity [Wang et al. 2022], bounded floating-point arithmetic [Chen et al. 2026; Wang 2014], inclusion methods [Chen et al. 2024c; Wang et al. 2021], sum-of-squares relaxations [Zhang et al. 2023], and convex-body formulations [Yuan et al. 2025].

On the response side, implicit contact handling [Harmon et al. 2009; Otaduy et al. 2009], impact zone method [Tang et al. 2018; Wang et al. 2023a], and GPU-based repulsion [Wu et al. 2020] bring contact-aware simulation to interactive rates. Optimization-based formulations replace hand-tuned impulses with barrier or penalty objectives: incremental potential contact (IPC) embeds distance barriers in the timestep solve [Li et al. 2020, 2021], and subsequent variants extend the principle to rigid or affine bodies [Ferguson et al. 2021; He et al. 2026; Lan et al. 2022a], alternative barrier shapes [Ando 2024; Huang et al. 2024], non-distance barriers and barrier-free formulations [Lan et al. 2024; Zheng et al. 2025], and faster solvers [Guo et al. 2024, 2026a, 2025a, 2026b; Huang et al. 2025].

2.2 Untangling

Interactive authoring, animation retargeting, and simulation initialization can all introduce intersections without prior penetration-free state history. Removing such intersections is a foundational problem in physics simulation and geometric processing.

Volumetric and rigid-body proxies. Several systems resolve initial state intersections by constructing a proxy representation. Rigid-body packages typically rely on convex separation queries [Gilbert

et al. 1988; van den Bergen 2001] or volumetric decompositions [Mamou et al. 2016; Wei et al. 2022]. For deformable objects, proxy designs include tetrahedral air meshes [Müller et al. 2015], penetration projections [Je et al. 2012; Tang and Kim 2014], voxelized distance fields [Santesteban et al. 2022; Wong et al. 2018], and shortest-path routes to open boundaries [Chen et al. 2023]. While these approaches excel on closed solids or simple boundaries, they fundamentally mismatch generalized thin-shell tangles. A nested cloth cannot be reliably approximated by convex unions or well-defined interior volumes, and its flexible nature requires distributed, topologically varying displacements rather than a single rigid-translation rule.

Global Intersection Analysis (GIA). Early work [Volino et al. 1995] tracked the state of nearby collision pairs in previous frames. However, retaining and selecting the appropriate historical collision information is challenging [Baraff et al. 2003]. In addition, this scheme suffers in reliability when dealing with persistent penetrations. Baraff et al. [2003] introduced history-free cloth untangling by tracing closed intersection contours and performing parallel flood-fills to identify the interior cloth regions. Wicke et al. [2006] extended this scheme to open boundaries by categorizing contours into seven distinct parameter-domain types (e.g., CLOSED, EIGHT, BLI) and applying type-specific extrapolation rules. Subsequent methods [Cha and Ko 2020, 2022] further refined the response logic, but the core design of branching on discrete contour types remains. However, strict contour-type routing suffers from fundamental limitations. Discrete classification is fragile under numerical detection noise, real-world tangles frequently produce multi-loop or non-manifold contours that fall outside predefined category tables, and flood-filling open or non-orientable surfaces often fails to identify the true physical penetration basin. These challenges motivate a unified, type-free response construction operating directly on geometric ray evidence.

Intersection Contour Minimization (ICM). Volino and Magnenat-Thalmann [2006] bypass global partitioning and instead evaluate every intersecting edge-face (EF) pair locally, minimizing the overall contour length to separate the tangled geometry. Later extensions incorporate orientation information and couple the solver with discrete collision response pipelines [Ye et al. 2017, 2015; Ye and Zhao 2012], maintaining high parallelism and broader compatibility. Yet, purely local contour-length minimization inevitably falls into local minima. Symmetric interlocking configurations and multi-layer cloth stacks trap the descent, as shortening the intersection contour does not strictly imply geometric separation. Even when paired with contour-level response, the final optimization steps in these pipelines remain restricted by their local solutions.

Fast and learning-based methods. Implicit and learning-based methods form a third branch beside classical contour analysis. Buffet et al. [2019] model layered garments with implicit untangling operators that provide smooth responses for clothing design workflows, while Guo et al. [2025c] and Guo et al. [2025b] target fast outfit assembly, knot modeling, and pose transfer through progressive repair stages. Recent data-driven works use implicit distance fields [Santesteban et al. 2022], graph neural networks [Lee et al. 2023],

and boundary-volume energy formulations [Grigorev et al. 2024; Liao et al. 2024] to resolve layered cloth intersections under learned or task-specific assumptions. Focusing on explicit geometric repair, Jang et al. [2025] formulate a static self-intersection resolver using a local signed tangent-point energy. These methods are effective in their target regimes, but they do not replace the need for a contour-level, topology-aware construction when the input contains deep or nested intersections.

Overall, previous untangling methods struggle with two recurring challenges that motivate our design: (i) unified handling of diverse contour configurations, avoiding the discontinuous response behavior that arises in the GIA family when mesh degeneracies, detection noise, non-manifold input, or contour-type misclassification cause discrete branching to select the wrong repair rule; and (ii) global convergence under complex penetrations, where contour-minimization in ICM and tangent-point energy formulations stall in local minima.

3 Method

We build on implicit time integration for cloth, thin shells, and deformable bodies, as well as affine and rigid-body dynamics [Ferguson et al. 2021; Lan et al. 2022a], which formulate each time step as an energy minimization. To simplify the exposition, we describe our formulation in terms of nodal positions $\mathbf{x} \in \mathbb{R}^{3N}$; the pipeline applies identically to generalized coordinates $\mathbf{q}$ for rigid and affine bodies (see Supplemental Material). A standard incremental potential (IP) step resolves the next physical state by minimizing:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \left(\frac{1}{2h^2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_{\mathbf{M}}^2 + \Psi(\mathbf{x}) \right) \quad (1)$$

Here h is the time step, $\mathbf{M}$ is the mass matrix, $\tilde{\mathbf{x}} = \mathbf{x}^n + h\mathbf{v}^n + h^2\mathbf{M}^{-1}\mathbf{f}_{\text{ext}}$ is the inertial prediction, $\mathbf{x}^n$ and $\mathbf{v}^n$ are the position and velocity at the beginning of the timestep, $\mathbf{f}_{\text{ext}}$ is the external force, and $\Psi(\mathbf{x})$ aggregates all potential energies (elastic, bending, contact, boundary conditions, etc.).

When the input state contains intersections, continuous collision barriers break down because an interior-point contact barrier requires a penetration-free initial configuration. Untangling repairs the state by injecting history-free separation constraints $C_k(\mathbf{x}) > 0$, pulling the mesh toward a safe configuration Ω_{safe} . However, assigning arbitrary separation routes can violently inflate the inertial or elastic penalties in Eq. 1, inducing unnatural mesh distortion and solver instability. Thus, the geometric untangling route must minimize total induced displacement. For each contour, given a separation direction $\mathbf{r} \in \mathbb{S}^2$ and its (filtered) hit response support $\mathcal{S}(\mathbf{r})$, we evaluate the separation cost using a mass-weighted L_2 displacement metric analogous to the inertial term $\sum_i m_i \|\Delta \mathbf{x}_i\|^2$ in Eq. 1:

$$\begin{aligned} D_C(\mathbf{r}) &= \max_{h \in \mathcal{S}(\mathbf{r})} \delta_h(\mathbf{r}) \\ J(\mathbf{r}) &= \sum_{h \in \mathcal{S}(\mathbf{r})} m_h D_C(\mathbf{r})^2 = \rho D_C(\mathbf{r})^2 \sum_{h \in \mathcal{S}(\mathbf{r})} A_h \end{aligned} \quad (2)$$

where $m_h = \rho A_h$ is the mass associated with hit h , computed from the area A_h of its source and destination primitives and the surface density ρ , and $D_C(\mathbf{r})$ is the maximum ray travel distance of this contour. All hits in $\mathcal{S}(\mathbf{r})$ share the same contour depth D_C ; we treat




Fig. 2. Linear-path intersection evidence. (a) VF casts source vertices along $\mathbf{r}$ toward destination faces; (b) EE sweeps a source edge along $\mathbf{r}$, forming a ruled quadrilateral (parallelogram strip), and detects intersections with destination edges; (c) FV sweeps a source face along $\mathbf{r}$, forming a triangular prism, and tests containment of destination vertices inside the prism volume.

the area density ρ as a uniform constant and drop it. By minimizing $J(\mathbf{r})$, we select separation paths that reconcile geometric untangling with the physical solver's lowest-energy displacement modes.

3.1 Overview of the Direction-Search Pipeline

Given discrete edge-face (EF) intersection pairs, we treat each intersection contour C as a topological seed for constructing a separation response. For each contour, our pipeline alternates between freezing the physical state $\mathbf{x}$ to find the optimal separation direction $\mathbf{r}^* = \arg \min J(\mathbf{r})$, and freezing $\mathbf{r}^*$ to advance $\mathbf{x}$ via the global IP solver:

- (1) **Initialize** directed candidates $\mathbf{r} \in \mathcal{D}_0$ from contour-local geometry;
- (2) **Cast** to find the raw hit set $\mathcal{H}_{\text{raw}}(\mathbf{r})$ via vertex-face (VF), edge-edge (EE), and face-vertex (FV) intersection tests;
- (3) **Filter** raw hits via dual-sided product-complex cluster culling, retaining only the physical component $\mathcal{S}(\mathbf{r})$ attached to the detected contour;
- (4) **Refine** selected candidates using projected Newton or Cauchy steps on $\mathbb{S}^2$, re-evaluating the filtered response objective $J(\mathbf{r})$ at each trial;
- (5) **Assemble** the final optimized hit set into solver constraints $\{C_k\}$ to advance the physical state $\mathbf{x}$ in Eq. 1.

Crucially, the same evaluation path is used for open or closed contours, boundary contacts, self-intersections, and cross-object intersections; our formulation does not rely on classifying intersections into predefined topological categories (Fig. 3).

3.2 Direction Evaluation

An intersection contour C locally separates the participating surfaces into two *logical sides*, which we designate as the *source* and *destination* sides. For cross-body intersections between distinct bodies, these two sides correspond directly to the two separate meshes. For self-intersections within a single mesh, they represent two topologically proximal but locally opposing surface regions (whose exact spatial boundaries are initially unknown and will be discovered by

ray casting). Ray casting is consistently launched from the source side toward the destination side along candidate direction $\mathbf{r} \in \mathbb{S}^2$.

To ensure geometric completeness across the entire displacement field, we collect linear-path intersections using three complementary primitive tests: vertex-face (VF), edge-edge (EE), and face-vertex (FV), as depicted in Fig. 2. Together, these tests exhaustively capture all contact events induced by sweeping the source side along $\mathbf{r}$ into the destination side. The raw evidence set is defined as:

$$\mathcal{H}_{\text{raw}}(\mathbf{r}) = \mathcal{H}_{\text{VF}}(\mathbf{r}) \cup \mathcal{H}_{\text{EE}}(\mathbf{r}) \cup \mathcal{H}_{\text{FV}}(\mathbf{r}) \quad (3)$$

VF Test. For a VF test, let a source vertex p travel along separation direction $\mathbf{r}$ by distance δ to reach an intersection point $\tilde{T} = p + \delta\mathbf{r}$. This point must lie on a destination triangle $T = (T_0, T_1, T_2)$, parameterized by barycentric coordinates (u, v) as $\tilde{T} = (1 - u - v)T_0 + uT_1 + vT_2 = T_0 + u(T_1 - T_0) + v(T_2 - T_0)$. Equating the two representations yields $p + \delta\mathbf{r} = T_0 + u(T_1 - T_0) + v(T_2 - T_0)$, which rearranges to the linear system:

$$\begin{bmatrix} \mathbf{r} & -(T_1 - T_0) & -(T_2 - T_0) \end{bmatrix} \begin{bmatrix} \delta \\ u \\ v \end{bmatrix} = T_0 - p \quad (4)$$

with the geometric containment criteria $u \geq 0, v \geq 0, u + v \leq 1$, and travel distance $\delta \geq 0$.

EE Test. For an EE test between a source edge $E_1 = (p_0, p_1)$ parameterized by $s \in [0, 1]$ as $p(s) = p_0 + s(p_1 - p_0)$ and a destination edge $E_2 = (q_0, q_1)$ parameterized by $t \in [0, 1]$ as $q(t) = q_0 + t(q_1 - q_0)$, the swept segment along $\mathbf{r}$ meets E_2 when $p_0 + s(p_1 - p_0) + \delta\mathbf{r} = q_0 + t(q_1 - q_0)$. Rearranging gives:

$$\begin{bmatrix} \mathbf{r} & (p_1 - p_0) & -(q_1 - q_0) \end{bmatrix} \begin{bmatrix} \delta \\ s \\ t \end{bmatrix} = q_0 - p_0 \quad (5)$$

with parameters $s, t \in [0, 1]$ and $\delta \geq 0$.

FV Test. An FV event tests whether a source face $F = (S_0, S_1, S_2)$ parameterized by barycentric coordinates (u, v) sweeps along $\mathbf{r}$ over a destination vertex q , i.e., $S_0 + u(S_1 - S_0) + v(S_2 - S_0) + \delta\mathbf{r} = q$. Equivalently, casting destination vertex q along the reverse direction $-\mathbf{r}$ toward face F yields $q - \delta\mathbf{r} = S_0 + u(S_1 - S_0) + v(S_2 - S_0)$, which rearranges to:

$$\begin{bmatrix} -\mathbf{r} & -(S_1 - S_0) & -(S_2 - S_0) \end{bmatrix} \begin{bmatrix} \delta \\ u \\ v \end{bmatrix} = S_0 - q \quad (6)$$

with barycentric bounds $u \geq 0, v \geq 0, u + v \leq 1$, and $\delta \geq 0$. Formulating FV as a reverse vertex-ray query reuses the highly optimized VF ray traversal and delivers an order-of-magnitude speedup over explicit face-prism swept volume tests.

Unified Cramer System and Directional Decoupling. The VF, EE, and FV systems in Eqs. (4), (5), and (6) share a unified 3×3 form $[\mathbf{r} \mid \mathbf{a} \mid \mathbf{b}][\delta, \alpha, \beta]^T = \mathbf{d}$, where $\mathbf{a}, \mathbf{b}, \mathbf{d}$ correspond directly to the geometry-dependent columns and right-hand sides of each respective test. Applying Cramer's rule isolates the ray travel depth δ_h :

$$\delta_h(\mathbf{r}) = \frac{\det[\mathbf{d} \mid \mathbf{a} \mid \mathbf{b}]}{\det[\mathbf{r} \mid \mathbf{a} \mid \mathbf{b}]} = \frac{\mathbf{n}_h^T \mathbf{d}}{\mathbf{n}_h^T \mathbf{r}} = \frac{v_h}{\mathbf{n}_h^T \mathbf{r}}, \quad \mathbf{n}_h = \mathbf{a} \times \mathbf{b} \quad (7)$$




Fig. 3. Ray-casting and hit-selection pipeline. (a) The overlap of the source and destination meshes constructs the intersection contours. (b) A candidate direction generates raw VF, EE, and FV hits using the tests in Fig. 2. (c) Hits are grouped by topological adjacency. (d) Only the hits from the contour-anchored cluster are preserved. (e) Response result after applying untangling constraints.

The vector $\mathbf{n}_h = \mathbf{a} \times \mathbf{b}$ is the *characteristic normal* of hit h : it reduces to the oriented face normal for VF and FV pairs and to the edge cross-product normal for EE pairs. For fixed nodal positions $\mathbf{x}$, the numerator $v_h = \mathbf{n}_h^T \mathbf{d}$ is invariant to $\mathbf{r}$. Consequently, $\delta_h(\mathbf{r}) = v_h / (\mathbf{n}_h^T \mathbf{r})$ depends solely on $\mathbf{r}$, which decouples geometric primitive evaluation from directional search and ensures that the separation energy $J(\mathbf{r})$ depends only on $\mathbf{r}$ and is locally differentiable inside each fixed-hit region.

Each surviving hit h stores its distance $\delta_h(\mathbf{r})$, characteristic normal $\mathbf{n}_h$, and response area weight A_h . Following the topological filtering described in the next subsection, the final validated hit set is denoted $\mathcal{S}(\mathbf{r})$. The response objective $J(\mathbf{r})$ and contour depth $D_C(\mathbf{r})$ are then evaluated using Eq. 2. Note that all hits associated with the given contour are subject to the shared maximum depth D_C , and the direction $\mathbf{r}$ is ranked and selected based on the total cost J .

3.3 Topological Cluster Culling

Raw ray casting along a candidate direction $\mathbf{r}$ collects all linear-path intersections across the candidate domain, which often includes unrelated hits on remote mesh regions. We filter these based on a dual-sided hit connectivity graph $G_r = (\mathcal{H}_{\text{raw}}(\mathbf{r}), \mathcal{E}_r)$.

Recall from Sec. 3.2 that the contour partitions the local geometry into two logical sides. Each ray hit h therefore pairs a source primitive $\text{src}(h)$ (swept along $\mathbf{r}$) with a destination primitive $\text{dst}(h)$ (the struck element). To build the edge set $\mathcal{E}_r$, two hits $h_1, h_2 \in \mathcal{H}_{\text{raw}}(\mathbf{r})$ are connected by an adjacency edge $(h_1, h_2) \in \mathcal{E}_r$ if and only if their primitives are mutually adjacent on both the source and destination sides simultaneously. Table 1 formalizes the complete symmetric dual-side adjacency predicates.

We extract the connected components $\{\Gamma_j\}$ of G_r . To isolate the physical separation support from detached geometric debris, the valid component must anchor to the active intersection contour. Let $\mathcal{A}_0$ and $\mathcal{A}_1$ denote the sets of contour seed vertices on the source and destination sides, respectively, and let $\mathcal{N}_1(\cdot)$ denote their one-ring vertex neighborhoods. A hit h is considered *contour-anchored* if its primitives directly touch the seed neighborhood on both logical sides:

$$\mathcal{B}_C(\mathbf{r}) = \left\{ h \in \mathcal{H}_{\text{raw}}(\mathbf{r}) \mid \begin{array}{l} \text{src}(h) \cap \mathcal{N}_1(\mathcal{A}_0) \neq \emptyset \\ \text{dst}(h) \cap \mathcal{N}_1(\mathcal{A}_1) \neq \emptyset \end{array} \right\} \quad (8)$$

Table 1. Dual-side hit adjacency predicates.

Hit Pair	Source-side Condition	Destination-side Condition
$\text{VF}(v, f) \sim \text{VF}(v', f')$	$v \sim_V v'$ (share an edge)	$f = f'$ (same face)
$\text{FV}(f, v) \sim \text{FV}(f', v')$	$f = f'$ (same face)	$v \sim_V v'$ (share an edge)
$\text{VF}(v, f) \sim \text{FV}(f', v')$	$v \in f'$ (vertex incident to face)	$v' \in f$ (vertex incident to face)
$\text{FV}(f', v') \sim \text{VF}(v, f)$	$v \in f'$ (vertex incident to face)	$v' \in f$ (vertex incident to face)
$\text{VF}(v, f) \sim \text{EE}(e_s, e_d)$	$v \in e_s$ (vertex incident to edge)	$e_d \subset f$ (edge incident to face)
$\text{EE}(e_s, e_d) \sim \text{VF}(v, f)$	$v \in e_s$ (vertex incident to edge)	$e_d \subset f$ (edge incident to face)
$\text{FV}(f, v) \sim \text{EE}(e_s, e_d)$	$e_s \subset f$ (edge incident to face)	$v \in e_d$ (vertex incident to edge)
$\text{EE}(e_s, e_d) \sim \text{FV}(f, v)$	$e_s \subset f$ (edge incident to face)	$v \in e_d$ (vertex incident to edge)
$\text{EE}(e_1, e_2) \sim \text{EE}(e'_1, e'_2)$	$e_1 = e'_1$ (same edge)	$e_2 \sim_E e'_2$ (share a vertex)
$\text{EE}(e_1, e_2) \sim \text{EE}(e'_1, e'_2)$	$e_1 \sim_E e'_1$ (share a vertex)	$e_2 = e'_2$ (same edge)

Any component failing to contain at least one contour-anchored hit ($|\Gamma_j \cap \mathcal{B}_C(\mathbf{r})| = 0$) is immediately pruned as detached geometric debris. Among surviving candidate components, we select the unique physical response band $\mathcal{S}(\mathbf{r}) = \Gamma^*$ that maximizes the edge-face anchor coverage:

$$C_C(\Gamma) = |\{e \in \mathcal{C} \mid \text{Endpoints}(e) \subset \text{Support}(\Gamma)\}| \quad (9)$$

The adjacency predicates in Table 1 are codimension-one face adjacencies in the Cartesian product complex $\mathcal{K}_s \times \mathcal{K}_d$ (see the supplemental material for a full topological analysis).

3.4 Candidate Direction Construction

Our goal is to reduce the displacement measure J of each contour in Eq. 2. The variable for one contour is its separation direction $\mathbf{r}$, which is constrained to the unit sphere $\mathbb{S}^2$. A direction first determines the raw VF/EE/FV hits through ray casting, and Cluster Culling then determines the valid set used to evaluate J . Thus, the direction search is the composite problem

$$\begin{aligned} \mathbf{r}^* &= \arg \min_{\mathbf{r} \in \mathbb{S}^2, \mathcal{S}(\mathbf{r}) \neq \emptyset} J(\mathbf{r}) \\ \mathcal{H}_{\text{raw}}(\mathbf{r}) &= \text{Raycast}(\mathcal{C}, \mathbf{r}) \\ \mathcal{S}(\mathbf{r}) &= \text{Cluster}(\mathcal{H}_{\text{raw}}(\mathbf{r})) \end{aligned} \quad (10)$$

Here Raycast returns the tests in Sec. 3.2, and Cluster applies the two-sided connectivity rule in Sec. 3.3. Although $\mathbf{r}$ varies continuously on $\mathbb{S}^2$, discrete hit and cluster transitions render $J(\mathbf{r})$ generally non-convex and only piecewise smooth. Within each fixed-hit basin, however, Eq. (7) makes the energy locally two-dimensional, smooth,

and analytically differentiable on $\mathbb{S}^2$. This motivates a hybrid optimization strategy: we first sample geometric candidate directions to identify promising energy basins, and then perform local spherical optimization within each selected basin.

While our current initial candidate set $\mathcal{D}_0$ is heuristic, an intriguing direction for future work is analyzing surface curvature, geometry, and contour topology to explicitly determine the number and location of candidate energy basins on $\mathbb{S}^2$.

3.4.1 Initial Candidates. Because $J(\mathbf{r})$ is non-convex and non-smooth across hit-component boundaries, we first sample a small set of promising initial directions $\mathcal{D}_0$ before committing to local optimization. A practical heuristic is that the true separation direction often aligns with the shortest topological extent of the intersecting region. We formalize this geometric prior using the spatial distribution of the contour's intersection geometry.

Let $\mathbf{c}_C$ be the length-weighted centroid of the actual EF intersection points, calculated as:

$$\mathbf{c}_C = \frac{\sum_i \ell_i \mathbf{p}_i}{\sum_i \ell_i} \quad (11)$$

Let $\mathcal{A} = \mathcal{A}_0 \cup \mathcal{A}_1$ denote the contour-supporting seed vertices and $A_v > 0$ their rest-area weights. We compute the area-weighted scatter matrix of the supporting vertices around the centroid:

$$\mathbf{S}_w = \sum_{v \in \mathcal{A}} A_v (\mathbf{x}_v - \mathbf{c}_C)(\mathbf{x}_v - \mathbf{c}_C)^\top \quad (12)$$

The eigendecomposition of $\mathbf{S}_w$ yields three orthonormal principal axes $\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2$ ordered by increasing eigenvalue. The minor axis $\mathbf{q}_0$ identifies the direction of minimum weighted variance—the axis along which the combined supporting vertices have the narrowest spatial spread. This axis serves as a strong primary candidate for the shortest separation path, although it remains a heuristic candidate subject to full ray-cast validation. We include all six signed scatter axes as candidates:

$$\mathcal{D}_{\text{scatter}} = \{\pm \mathbf{q}_0, \pm \mathbf{q}_1, \pm \mathbf{q}_2\} \quad (13)$$

To capture highly asymmetric or bending intersections, we supplement these with four additional directions derived from the individual area-weighted centroids of the two sides $\mathbf{c}_0, \mathbf{c}_1$:

$$\mathbf{r}_{\text{side}} = \pm \frac{\mathbf{c}_1 - \mathbf{c}_0}{\|\mathbf{c}_1 - \mathbf{c}_0\|}, \quad \mathbf{r}_{\text{to_center},m} = \frac{\mathbf{c}_C - \mathbf{c}_m}{\|\mathbf{c}_C - \mathbf{c}_m\|}, \quad m \in \{0, 1\} \quad (14)$$

This yields a total of 10 directed candidates in $\mathcal{D}_0$. Each candidate undergoes identical ray casting, topological filtering, and objective evaluation; the initial geometric statistics merely seed the subsequent physics-aware search.

3.5 Direction Selection and Refinement

Once a candidate produces a valid connected hit set, we build a smooth local model inside its fixed-hit region. We freeze hit identity, barycentric coordinates, and characteristic normals, turning ray distance into a differentiable function of direction. This subsection derives the local Newton model on $\mathbb{S}^2$.

Fixed-Hit Surrogate. For an accepted direction $\mathbf{r}_k$, let S_k be its filtered hit set. Each hit $h \in S_k$ stores its current ray distance $\delta_h^k = \delta_h(\mathbf{r}_k)$ and an oriented unit characteristic normal $\mathbf{n}_h$. During one local step, we freeze the hit identity, barycentric coordinates, and $\mathbf{n}_h$. We also freeze the numerator of the Cramer relation by defining $v_h = \delta_h^k(\mathbf{n}_h^\top \mathbf{r}_k)$. The predicted distance at a nearby direction is therefore

$$\delta_h(\mathbf{r}) = \frac{v_h}{\mathbf{n}_h^\top \mathbf{r}}, \quad v_h = \delta_h(\mathbf{r}_k)(\mathbf{n}_h^\top \mathbf{r}_k) \quad (15)$$

Here v_h is only a constant for the current fixed hit set; it is not an additional physical state or optimization variable. Equation 15 is the local sensitivity model obtained from the exact ray intersection in Eq. 7. It also makes the conditioning explicit: as $|\mathbf{n}_h^\top \mathbf{r}|$ approaches zero, a small directional rotation predicts a very long ray, so such trials are rejected when they are re-cast.

The weighting factor μ_h reflects the nodal effort required to enact the separation. Let $w_{h,j}$ be the signed barycentric coefficient of incident vertex j , and let $m_{h,j}$ be its lumped nodal mass. We define the effective inverse mass κ_h and the corresponding effective mass μ_h :

$$\kappa_h = \sum_{j=1}^{m_h} \frac{w_{h,j}^2}{m_{h,j}}, \quad \mu_h = \kappa_h^{-1} \quad (16)$$

This μ_h is the exact analytical effective mass derived by minimizing the kinetic energy $\sum m_j \|\Delta \mathbf{x}_j\|^2$ subject to a unit normal displacement $\sum w_{h,j} \mathbf{n}^\top \Delta \mathbf{x}_j = 1$. By weighting hits with μ_h , the direction search prefers separation routes that move light, unconstrained vertices rather than heavy or stiffly constrained regions, naturally aligning the geometric proxy with the solver's inertial metric.

The actual separation objective $J(\mathbf{r})$ depends on the L^∞ maximum contour depth, rendering it non-smooth with respect to $\mathbf{r}$. To enable smooth gradient-based refinement inside a fixed-hit region while hit identities remain invariant, we formulate an Iteratively Reweighted Least Squares (IRLS) surrogate. We construct a sequence of dynamic quadratic objectives that smoothly interpolate toward the minimax limit. Incorporating a target clearance margin $\epsilon_{\text{safe}} > 0$ (typically 3 mm), the surrogate at local iteration k is:

$$F_k(\mathbf{r}) = \sum_{h \in S_k} \mu_h \omega_h^{(k)} (\delta_h(\mathbf{r}) + \epsilon_{\text{safe}})^2 \quad (17)$$

To capture the deepest penetrations without sacrificing the quadratic structure of the Newton solve, the exponent p_k is scheduled linearly from 2 to p_{max} across local iterations $k = 0, \dots, K-1$:

$$p_k = 2 + (p_{\text{max}} - 2) \frac{k}{\max(K-1, 1)} \quad (18)$$

The corresponding weights $\omega_h^{(k)}$ are frozen from the currently accepted direction $\mathbf{r}_{\text{acc}}$:

$$\omega_h^{(k)} = (\delta_h(\mathbf{r}_{\text{acc}}) + \epsilon_{\text{safe}} + \eta)^{p_k - 2} \quad (19)$$

where $\eta > 0$ prevents division by zero for shallow contacts. Higher exponents p_k concentrate the gradient on dominant penetrations while maintaining an analytic Hessian for the spherical Newton step.

Using Eq. 15, the differential quantities for Eq. 17 are

$$\frac{\partial \delta_h}{\partial \mathbf{r}} = -\frac{\delta_h \mathbf{n}_h}{\mathbf{n}_h^\top \mathbf{r}} \quad (20)$$

The ambient gradient and Hessian are

$$\begin{aligned} \nabla F_k(\mathbf{r}) &= \sum_{h \in \mathcal{S}_k} 2\mu_h \omega_h^{(k)} (\delta_h + \epsilon_{\text{safe}}) \left(-\frac{\delta_h \mathbf{n}_h}{\mathbf{n}_h^\top \mathbf{r}} \right) \\ \nabla^2 F_k(\mathbf{r}) &= \sum_{h \in \mathcal{S}_k} \frac{2\mu_h \omega_h^{(k)}}{(\mathbf{n}_h^\top \mathbf{r})^2} (3\delta_h^2 + 2\epsilon_{\text{safe}} \delta_h) \mathbf{n}_h \mathbf{n}_h^\top \end{aligned} \quad (21)$$

Each hit contributes one rank-one block aligned with its characteristic normal. This is the local curvature missing from the scatter-matrix initialization.

Spherical Newton Step. The direction must remain unit length, so we solve the Newton system in the tangent plane rather than in $\mathbb{R}^3$. Given the ambient gradient $\mathbf{g}$ and Hessian $\mathbf{H}$ from Eq. 21, we first remove the radial component,

$$\mathbf{g}_T = \mathbf{g} - (\mathbf{r}^\top \mathbf{g}) \mathbf{r} \in T_r \mathbb{S}^2 \quad (22)$$

Let $\mathbf{e}_1, \mathbf{e}_2$ be an orthonormal basis of $T_r \mathbb{S}^2$. The projected 2D system is

$$\mathbf{g}_2 = \begin{pmatrix} \mathbf{e}_1^\top \mathbf{g}_T \\ \mathbf{e}_2^\top \mathbf{g}_T \end{pmatrix}, \quad \mathbf{H}_2 = \begin{pmatrix} \mathbf{e}_1^\top \mathbf{H} \mathbf{e}_1 & \mathbf{e}_1^\top \mathbf{H} \mathbf{e}_2 \\ \mathbf{e}_2^\top \mathbf{H} \mathbf{e}_1 & \mathbf{e}_2^\top \mathbf{H} \mathbf{e}_2 \end{pmatrix} \quad (23)$$

When $\mathbf{H}_2$ is positive definite, we solve $\mathbf{H}_2 \delta = -\mathbf{g}_2$; otherwise, we use a Cauchy step along $-\mathbf{g}_2$. With the trust-region option used in our experiments, we clip either step to the current tangent-plane radius. The trial direction is then retracted by normalization,

$$\mathbf{r}_{\text{trial}} = \frac{\mathbf{r} + \delta_1 \mathbf{e}_1 + \delta_2 \mathbf{e}_2}{\|\mathbf{r} + \delta_1 \mathbf{e}_1 + \delta_2 \mathbf{e}_2\|} \quad (24)$$

This update projects the ambient Hessian into the tangent basis; it does not add the full Riemannian correction term. As illustrated in Fig. 4(a), the IRLS surrogate decreases monotonically along accepted local steps within a fixed-hit region. However, as shown in the spherical energy landscape in Fig. 4(b,c), the objective exhibits multiple local valleys and discrete transitions across component boundaries. We then re-evaluate each trial by ray casting and topological selection. The ratio of actual to predicted reduction expands or contracts the trust radius, and a trial is accepted only if the full-contour objective decreases. A trial whose adaptive k -ring does not converge is rejected. Thus, the smooth model proposes a direction, whereas the reconstructed geometric support determines its validity.

3.6 Response Constraint Assembly

The selected direction $\mathbf{r}^*$ and hit set $\mathcal{S}^* = \mathcal{S}(\mathbf{r}^*)$ are converted into pairwise penalty constraints that share the same algebraic form as standard VF and EE contact pairs, enabling the downstream solver to process untangling and collision responses identically.

In standard collision handling, a proximity pair generates a quadratic penalty energy $E^{\text{col}}(\mathbf{x}) = \frac{1}{2} \kappa A \min(C(\mathbf{x}), 0)^2$ active when primitives encroach upon their safety margin. Our untangling framework adopts the same quadratic penalty formulation, $E_h(\mathbf{x}_h) = \frac{1}{2} \kappa \bar{A}_h C_h(\mathbf{x}_h)^2$, while constructing the constraint function C_h directly from the validated ray-hit support.

To prevent large single-iteration displacements from causing excessive geometric distortion or introducing secondary penetrations when continuous collision detection (CCD) is disabled, the single-iteration correction displacement $D_C(\mathbf{r}^*)$ is clamped to a maximum displacement limit (set to 5 mm when CCD is disabled). The untangling constraint is formulated as an affine measure of displacement along the hit's characteristic normal:

$$C_h(\mathbf{x}_h) = \mathbf{n}_h^\top \left(\sum_{j=1}^{m_h} w_{h,j} \mathbf{x}_{h,j} \right) - D_C(\mathbf{r}^*) \quad (25)$$

Here m_h is the number of vertices incident to pair h . The weighted sum in parentheses computes the relative vector from the source element to the destination element. Its scalar coefficients $w_{h,j}$ are the barycentric weights fixed exactly at the moment of ray casting, derived from the u, v, t coordinates in Eqs. (4), (5), and (6). Specifically, for a VF hit ($m_h = 4$) involving vertices $\mathbf{x}_h = [\mathbf{x}_{h,1}^\top, \dots, \mathbf{x}_{h,4}^\top]^\top = [p^\top, T_0^\top, T_1^\top, T_2^\top]^\top$, the weights are $w_{h,j} \in \{-1, 1 - u - v, u, v\}$. For an EE hit with $\mathbf{x}_h = [p_0^\top, p_1^\top, q_0^\top, q_1^\top]^\top$, the weights are $w_{h,j} \in \{-(1-s), -s, 1-t, t\}$.

Let $\bar{A}_h$ be the standard contact-pair area defined as the mean area of the source and destination primitives. The assembled pairwise response energy is:

$$E_h(\mathbf{x}_h) = \frac{1}{2} \kappa \bar{A}_h C_h(\mathbf{x}_h)^2 \quad (26)$$

Because the barycentric weights are fixed, C_h is strictly affine with respect to $\mathbf{x}_h$. Let the constraint Jacobian be denoted by:

$$\mathbf{J}_h = \frac{\partial C_h}{\partial \mathbf{x}_h} = [\mathbf{w}_{h,1} \mathbf{n}_h^\top \quad \dots \quad \mathbf{w}_{h,m_h} \mathbf{n}_h^\top] \quad (27)$$

The affine nature of the constraint implies that its exact Hessian is zero ($\nabla_{\mathbf{x}_h}^2 C_h = 0$). Thus, the pair contributes the following gradient and constant Gauss-Newton Hessian block to the global solver:

$$\nabla E_h = \kappa \bar{A}_h C_h \mathbf{J}_h^\top \quad (28)$$

$$\nabla^2 E_h = \kappa \bar{A}_h \mathbf{J}_h^\top \mathbf{J}_h \quad (29)$$

This distinction reflects our alternating minimization framework: with the nodal configuration $\mathbf{x}$ fixed, the direction search optimizes $\mathbf{r} \in \mathbb{S}^2$ via $J(\mathbf{r})$ (where F_k serves strictly as a local tangent-plane proposal surrogate); once $\mathbf{r}^*$ is frozen, the response energy $E_h(\mathbf{x}_h)$ optimizes the physical nodal positions $\mathbf{x}$ in the global IP solver using standard contact-area weights $\bar{A}_h$.

Finally, the global implicit step augments the standard unconstrained system (cf. Eq. 1) as:

$$\min_{\mathbf{x}} \frac{1}{2h^2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_{\mathbf{M}}^2 + \Psi(\mathbf{x}) + \sum_C \sum_{h \in \mathcal{S}_C^*} E_h(\mathbf{x}_h) \quad (30)$$

We solve this objective in the same Newton/PCG loop as the elastic and collision energies.

Contour-wide depth target. All pairs along a contour share the same depth target $D_C(\mathbf{r}^*)$ for two reasons. First, a uniform target is the physically consistent limit for stiff materials: as the material stiffness grows, the response region approaches a rigid-body translation whose magnitude equals the contour-wide penetration depth. Second, because primitives inside the response band receive identical displacement along $\mathbf{r}^*$, they undergo no relative motion;




Fig. 4. Direction-search landscape and fixed-hit refinement. (a) The IRLS surrogate decreases along accepted local steps while the hit identities are unchanged. (b) An equirectangular view shows that different seeds descend through different fixed-hit regions. (c) The same measurements on the sphere expose multiple valleys and discontinuous transitions, motivating candidate sampling before local refinement.

this eliminates the risk that the correction itself introduces new self-intersections within the band, and avoids the non-uniform overstretching artifacts that arise when neighboring pairs are driven toward inconsistent depths. Adapting the depth per pair according to local material stiffness is a natural extension that we leave as future work.

4 Implementation Details

4.1 Adaptive k-Ring Candidate Restriction

Evaluating ray intersections across the full mesh for every directional trial incurs unnecessary computational overhead. Because valid response supports are topologically anchored to the intersection contour, we restrict ray queries to progressive geodesic neighborhoods around the contour seeds $\mathcal{A}_s$ ($s \in \{0, 1\}$). The complete expansion procedure is summarized in Algorithm 1 and visualized in Fig. 5.

Distance fields. We precompute two distance fields from the seed set $\mathcal{A}_s$ on the mesh 1-skeleton of each side s , independently for the two logical roles. The *hop distance* $\ell_s(v) \in \mathbb{N}$ counts unweighted graph edges along the shortest path from v to the nearest seed, computed by multi-source BFS in $O(|\mathcal{V}_s| + |\mathcal{E}_s|)$. The *rest geodesic distance* $g_s(v) \in \mathbb{R}_{\geq 0}$ accumulates material-space edge lengths along the shortest path, computed by multi-source Dijkstra, providing a metric-aware scale and mesh diameter normalizers (Supplemental Sec. 2.1). Hop distance defines the candidate domain radius $k_s^{(m)}$ and frontier convergence test, providing a combinatorial neighborhood whose extent is independent of local mesh resolution. Before ray casting, the candidate domain is further filtered by contour-side candidate restriction (Sec. 4.2), which uses the deformed Euclidean chord distance to attribute candidate primitives to their owning contour side.

Using the hop-distance field $\ell_s(v)$, we define the candidate primitive domain at expansion phase m with independent radii $k_s^{(m)}$:

$$\begin{aligned} \mathcal{V}_s^{(m)} &= \{v \in \mathcal{V}_s \mid \ell_s(v) \leq k_s^{(m)}\} \\ \mathcal{E}_s^{(m)} &= \{e \in \mathcal{E}_s \mid \max_{v \in e} \ell_s(v) \leq k_s^{(m)}\} \\ \mathcal{F}_s^{(m)} &= \{f \in \mathcal{F}_s \mid \max_{v \in f} \ell_s(v) \leq k_s^{(m)}\} \end{aligned} \quad (31)$$

Ray casting is executed within these bounded domains. At each phase, we inspect the maximum hop distance reached by the accepted valid hit support $\mathcal{S}^{(m)}$:

$$\ell_{\text{frontier},s}^{(m)} = \max \{\ell_s(v) \mid v \in \text{Vertices}_s(\mathcal{S}^{(m)})\}, \quad s \in \{0, 1\} \quad (32)$$

If $\ell_{\text{frontier},s}^{(m)} < k_s^{(m)}$ for both sides, or if the component achieves full contour anchor coverage ($C_C = 1.0$), the valid support has completely detached from the search frontier and expansion terminates. Otherwise, the active radius expands exponentially ($k_s^{(m+1)} = \min(2k_s^{(m)}, k_{s,\text{max}})$).

On 878 initial contours from the benchmark of Jang et al. [2025], adaptive k -rings reduce raw ray hits from a mean of 60,749 to 1,360—a 152.5× geometric-mean reduction and a 99.44% median reduction (Fig. 6)—while maintaining 100% valid-hit parity with full-mesh ray casting.

4.2 Contour-Side Candidate Restriction

Ray casting evaluates rigid linear separation paths $v + tr$. When a trial direction $\mathbf{r}$ is near-tangential to the surface ($\mathbf{r} \cdot \mathbf{n} \approx 0$), a ray launched from the penetration zone travels along the surface manifold for a substantial distance before clearing at an open mesh boundary or the closed mesh part. While traversing the full extent of the sheet yields the mathematically conservative clearance distance for a rigid translation along $\mathbf{r}$, without geometric bounds, such grazing rays generate sprawling hit footprints that over-extend across many rings of primitives towards distant mesh boundaries. If these over-extended hits form an internally connected subgraph in the product




Fig. 5. Role-aware adaptive k -ring progression on the self-penetrating Klein bottle [Jang et al. 2025]. (a–d) Candidate search domains progressively expand outward from the intersection contour. At each stage, cluster-connectivity and monodromy filters cull raw rays cast from the active domain that strike opposing sheets or distant folds (faint black lines), while legitimate local penetration responses are retained (dark green segments). (e) The assembled responses completely resolve all self-intersections without requiring expensive full-mesh ray casting.

ALGORITHM 1: Direction evaluation and response construction under adaptive k -ring expansion.

Input: Contour seeds $\mathcal{A}_0, \mathcal{A}_1$; candidate directions $\mathcal{D}_0$; initial search radii $k_0^{(0)}, k_1^{(0)}$

Output: Optimal separation direction $\mathbf{r}^*$ and validated hit response $\mathcal{S}^*$

for expansion phase $m = 0, 1, \dots$ **do**

 Construct candidate domains $\mathcal{V}_s^{(m)}, \mathcal{E}_s^{(m)}, \mathcal{F}_s^{(m)}$ for $s \in \{0, 1\}$ via Eq. 31;

for $\mathbf{r} \in \mathcal{D}_0$ **do**

$\mathcal{H}_{\text{raw}}(\mathbf{r}) \leftarrow \text{Raycast}(\mathcal{V}_s^{(m)}, \mathcal{E}_s^{(m)}, \mathcal{F}_s^{(m)}, \mathbf{r})$;

$\mathcal{S}(\mathbf{r}) \leftarrow \text{ClusterCulling}(\mathcal{H}_{\text{raw}}(\mathbf{r}), \mathcal{A}_0, \mathcal{A}_1)$;

 Evaluate displacement objective $J(\mathbf{r})$ via Eq. 2;

$\ell_{\text{frontier},s}^{(m)} \leftarrow \max\{\ell_s(v) \mid v \in \text{Vertices}_s(\mathcal{S}(\mathbf{r}))\}$;

end

if $(\ell_{\text{frontier},s}^{(m)} < k_s^{(m)} \text{ for } s \in \{0, 1\}) \vee (C_C = 1.0)$ **then**

return $(\mathbf{r}^*, \mathcal{S}^*) = \arg \min_{\mathbf{r} \in \mathcal{D}_0} J(\mathbf{r})$;

else

$k_s^{(m+1)} \leftarrow \min(2k_s^{(m)}, k_{s,\text{max}})$ for active frontiers;

end

end

complex (Sec. 3.3), cluster culling alone cannot distinguish them from local penetration responses, causing the solver to inject non-local artificial traction into remote, non-intersecting regions.

We establish a localized validity domain for linear separation paths by restricting the candidate domain of each logical role *before* ray casting. For each candidate vertex v and contour side $s \in \{0, 1\}$, we attribute v to side s if its deformed Euclidean distance to the contour seeds satisfies a relative Voronoi criterion:

$$\phi_s(v) = \frac{c_s(v) - c_{1-s}(v)}{c_s(v) + c_{1-s}(v) + \varepsilon} \leq \tau, \quad \text{where } c_s(v) = \min_{a \in \mathcal{A}_s} \|\mathbf{x}_v - \mathbf{x}_a\|_2. \quad (33)$$

The threshold $\tau \in [0, 1)$ interpolates between two regimes: at $\tau = 0$, the mesh is strictly partitioned by the spatial Voronoi bisector $c_0(v) = c_1(v)$, recovering the classical “illegal region” construction [Baraff et al. 2003; Wicke et al. 2006]; at $\tau = 1$, the restriction vanishes. The default $\tau = 0.25$ admits a controlled overlap band




Fig. 6. Candidate restriction via adaptive k -rings. While preserving the same final valid-hit set as full-mesh ray casting on the evaluated benchmark cases [Jang et al. 2025], adaptive k -rings consistently require far fewer raw hits than the unculled baseline, often by one to two orders of magnitude.

around the interface ($c_s(v) \leq \frac{5}{3}c_{1-s}(v)$), providing sufficient topological support for the product-cell adjacency graph while truncating tangential over-extension. Intersecting each adaptive k -ring domain (Eq. 31) with this geometric mask bounds linear separation paths to the contour’s immediate physical neighborhood.

This side-coordinate partition is a heuristic: it effectively suppresses tangential over-extension and spurious cross-layer hits in practice, but does not provide a formal guarantee of correctness. Reformulating the ownership as a smooth field on the deformed surface and replacing it with a discrete combinatorial selection over connected components are directions for future work.

4.3 Conflict Resolution via Contour Merging

Multiple intersecting regions may generate responses over the same surface patch. If the selected hit set of contour C_A reaches the seed region of C_B , their penalty supports overlap; opposing directions can then cancel or distort the separating motion. Figure 7(a) shows this interaction between two intersection contours.

Simple pairwise greedy elimination can suffer from cyclic deadlocks ($A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$) in dense entangled clusters,




Fig. 7. Intersection contour conflicts and their effect on convergence. Columns (a)–(b) show the same dual-BB/II cloth conflict. Without conflict suppression (a), the shared patch is still entangled after 15 iterations because the two contour supports keep competing over the same region. With suppression (b), the conflict disappears after the first iteration. Columns (c)–(d) show that the GIA-style contour classification in (c) converges to a response with an inconsistent lateral pull on the inner patch, whereas our unified geometric formulation in (d) produces a coherent separation field aligned with the global penetration structure.

where interacting contours mutually extinguish each other. We formulate multi-contour interaction as a conflict graph $\mathcal{G}_{\text{conflict}} = (\mathcal{V}_C, \mathcal{E}_{\text{conflict}})$, where an edge $(C_i, C_j) \in \mathcal{E}_{\text{conflict}}$ indicates that the valid hit support of contour C_i overlaps the seed domain of C_j . Each contour node $u \in \mathcal{V}_C$ is weighted by its total penetration area $w(u) = A_{\text{pen}}(u) = \sum_{h \in S_u} A_h$. We formulate multi-contour competition as finding a Maximum Weighted Independent Set (WIS) on $\mathcal{G}_{\text{conflict}}$:

$$\mathcal{S}_{\text{active}}^* = \arg \max_{\mathcal{S} \subseteq \mathcal{V}_C} \sum_{u \in \mathcal{S}} A_{\text{pen}}(u) \quad \text{s.t.} \quad (u, v) \notin \mathcal{E}_{\text{conflict}}, \quad \forall u, v \in \mathcal{S} \quad (34)$$

We approximate Eq. (34) via a greedy priority heuristic: contours are sorted in descending order of $A_{\text{pen}}(u)$, iteratively accepting the dominant unsuppressed contour and disabling its adjacent conflicting neighbors. This greedy selection guarantees that at least one dominant separation field survives within every interacting conflict component without cyclic mutual annihilation.

5 Experiments

5.1 Experimental Setup

Our simulator uses a constant time step of $\Delta t = 1/60$ s and proximity activation thickness $\hat{d}_{\text{prox}} = 1$ mm. In each iteration, our solver and GIA process at most 256 intersection contours, whereas ICM processes all contours simultaneously. By default, we enable continuous collision detection (CCD) with a maximum response distance $d_{\text{resp}} = 1.0$ m. For the four large-scale benchmark scenes (Synthetic, Round Table, Cloth Drop, and Hair Card), where massive contour entanglements cause IPC-style global CCD line search to severely throttle convergence rates, we employ a discrete collision response scheme by disabling CCD and setting $d_{\text{resp}} = 5$ mm. We use linear-elastic stretching [Kim 2020], dihedral-angle bending [Grinspun et al. 2003], and affine-body coupling [Lan et al. 2022a]. The resulting

Gauss–Newton systems are solved by 200 iterations of PCG (Jacobi-preconditioned conjugate gradients) with a parallel block-sparse matrix–vector product [Huang et al. 2025].

We use LBVH [Karras 2012] for broad-phase culling. Primitive pairs that connect the two logical sides of the same contour are excluded from this bound, as blocking those pairs would forbid the relative motion required to remove the penetration. All other VF/EE pairs remain subject to adaptive CCD [Li et al. 2021] and a quadratic proximity energy. Discrete EF intersections are detected geometrically [Möller 1997], connected by topological adjacency and clustered into intersection contours. Detecting and handling near-overlapping configurations [Yang and Jia 2025] is a promising direction for future work.

All experiments were run on a workstation with an Intel Core Ultra 7 265KF CPU, an NVIDIA GeForce RTX 5070 GPU (12 GB), and 32 GB of system memory. Our main simulation pipeline is implemented on the GPU using LuisaCompute [Zheng et al. 2022], including collision detection, linear-system assembly and the PCG solver.

Baseline configurations. We compare against Intersection Contour Minimization (ICM) [Volino and Magnenat-Thalmann 2006] and boundary-aware Global Intersection Analysis (GIA) [Baraff et al. 2003; Wicke et al. 2006]. ICM follows Volino and Magnenat-Thalmann [2006]. Each discrete EF intersection contributes a Gauss–Newton contour-shortening gradient with scaled-arctangent clipping. We use the *global scheme* described in Volino and Magnenat-Thalmann [2006], which aggregates these contributions along each contour before applying the correction. GIA follows the seven contour classes of Wicke et al. [2006]. Contours admitting a closed partition (Closed, Eight, LL, and BB/II) are resolved by flood-filling both sides and selecting the smaller enclosed region as illegal; open BLI variants use a three-ring geodesic neighborhood of the contour. BL/BI contours additionally receive a boundary-directed response

along the segment joining their two boundary endpoints. Following Baraff et al. [2003], proximity VF pairs assigned to the same illegal region attract rather than repel. We evaluate this logical-side relation at each discrete EF intersection and require consistent labels over both edge vertices and all three face vertices, reducing sensitivity to mesh density.

Large final EF counts for GIA reflect a structural failure mode: once a wrong region label or open-contour fallback is selected, the penalty integration can keep creating new intersections instead of monotonically clearing the original set. We also evaluate on the public 59-case benchmark of Jang et al. [2025], which is the closest recent static self-intersection repair dataset with released geometric inputs. Because Instant Self-Intersection Repair (ISIR) [Jang et al. 2025] is designed for closed volume meshes (Instant Body), its formulation is tailored for closed geometries and does not directly apply to open thin-shell surfaces such as cloth. We evaluate ISIR [Jang et al. 2025] on its own 59-case benchmark suite (Table 6), while comparing thin-shell cloth configurations against thin-shell baselines (GIA and ICM).

5.2 Unit Tests and Degenerate Configurations

5.2.1 Canonical Contour Configurations. We first evaluate our algorithm against the seven canonical contour configurations proposed by Wicke et al. [2006] to ensure backward compatibility with standard topology tests. Because our solver expands adaptively across geodesic neighborhoods, it resolves all seven configurations without contour-class-specific branching or patch construction. ICM traps or stalls in several of these configurations, while our response construction reaches zero residual EF pairs in two to four iterations.

5.2.2 Local Optimal Traps. Fig. 9 isolates two cases where baseline methods stall but our method recovers.

In Cloth-Torus, GIA constructs a spurious illegal region that drives the cloth (green, fixed) inward instead of releasing it apart from the torus (red). Our method avoids committing to a region partition and instead searches for a global separating direction with minimum displacement cost, finding the exterior separation route directly.

In Cloth-Sorento, the cloth (green) penetrates deeply into a complex non-convex sculpture (red). ICM’s contour-minimization gradients from opposite branches of the contour nearly cancel, reaching a local minimum of contour length while the penetration remains. Our method instead identifies the full invalid regions.

5.3 Ablation Study

Direction Proposal Family Analysis. To understand the role of the 10 initial direction candidates defined in Eq. (13) and Eq. (14), we analyze the distribution of the selected best initial candidate across the three geometric families: (1) *Scatter axes* ($\mathcal{D}_{\text{scatter}}$, 6 directions along the principal axes $\pm \mathbf{q}_0, \pm \mathbf{q}_1, \pm \mathbf{q}_2$ of $\mathbf{S}_w$), (2) *Mesh center line* ($\mathbf{r}_{\text{side}}$, 2 directions along $\pm(\mathbf{c}_1 - \mathbf{c}_0)$), and (3) *Mesh-to-contour center* ($\mathbf{r}_{\text{to_center}}$, 2 directions along $(\mathbf{c}_C - \mathbf{c}_m)$). Table 3 reports the selection frequencies across different datasets.

For complex self-intersections on the benchmark of Jang et al. [2025], the scatter family overwhelmingly dominates (87.4%), as the local principal variance axis $\mathbf{q}_0$ directly captures the thinnest

cross-section of the self-intersecting fold, whereas global mesh centers nearly coincide. Conversely, for multi-body interpenetration assemblies derived from Thingi10K [Zhou and Jacobson 2016], the mesh center line family provides the dominant separation vector (63.8%), reflecting the global translational offset between the two colliding bodies. The mesh-to-contour directions act as an indispensable complementary mechanism (8.9% and 13.5%) for non-convex or asymmetric entrapments. This confirms that all three families are essential for robustly covering both single-body self-entanglements and multi-body collisions.

Response stiffness sensitivity. Our untangling response is formulated as a soft penalty (Eq. 26) scaled by a stiffness κ . Sweeping $\kappa \in [10^4, 10^{12}]$ across the 59 cases of Jang et al. [2025] (Table 4), too small a κ cannot overcome internal stretching and bending resistance during de-penetration and the solver stalls, while too large a κ degrades the single-precision penalty conditioning and several cases diverge. We find $\kappa = 10^8$, matching the standard collision penalty stiffness [Li et al. 2020], sufficient to resolve essentially all intersections. Exploring adaptive stiffness scheduling or reformulating the soft penalty as a hard geometric constraint via Augmented Lagrangian multipliers is left as future work.

Shape fidelity and volume preservation. Because our primary objective is to reliably resolve complex intersections, we employ a relatively large response penalty stiffness ($\kappa = 10^8$), which can potentially induce unnatural over-deformation or localized surface pinching during aggressive de-penetration. We mitigate this through three complementary strategies: (1) *material stiffness regularization*, increasing resistance against bending and stretching during untangling (e.g., increasing bending stiffness κ_b from 0.05 to 1.0); (2) *bounded response stepping*, restricting the maximum single-iteration correction displacement (e.g., $d_{\text{resp}} = 5$ mm); and (3) *post-untangling relaxation*, executing a small number of standard IPC simulation steps (e.g., 200 steps with active CCD) once all penetrations are cleared to restore rest-state elastic equilibrium. As demonstrated across different combinations in Table 5 and Fig. 10, progressively applying these strategies systematically suppresses localized pinching and yields distinctly improved volume preservation ratios $\mathcal{V}/\mathcal{V}_0$ across all evaluated benchmark categories.

5.4 Large-Scale Benchmark

To stress-test the method, we evaluate four representative large scenes—Synthetic, Round Table, Cloth Drop, and Hair Card—and the 59-case public distribution of Jang et al. [2025]. We compare against ICM [Volino and Magnenat-Thalmann 2006] and GIA [Wicke et al. 2006]. Table 2 records the performance of each method on different cases. All four representative scenes resolve all intersections under our method, whereas ICM and GIA retain persistent intersections over thousands of iterations (Fig. 11).

5.4.1 Representative Cases. Synthetic consists of a square cloth folded inward from the sheet boundaries, with complex and nested intersections. Round Table and Cloth Drop contain densely entangled, non-manifold garment configurations dropped from the air, initialized with 48,686 and 49,124 EF pairs, respectively. Hair Card interweaves numerous ribbons around a fixed head, with an initial




Fig. 8. The seven canonical contour configurations proposed by [Wicke et al. \[2006\]](#). With the same ray-casting and cluster-culling rule, all seven cases reach zero residual EF pairs in two to four untangling iterations.




Fig. 9. Resolution of local-optimal traps. (Top row) On Cloth-Torus, GIA misclassifies the region partition and stalls at 157 residual EF pairs after 300 iterations. (Bottom row) On Cloth-Sorento, ICM contour-minimization gradients cancel inside the narrow cavity, stalling at 739 residual EF pairs after 300 iterations, whereas our method clears all intersections in 1 iteration.

configuration containing 168,914 EF pairs distributed across 10,268 contours. Fig. 12 and Fig. 1 visualize the untangled states of these scenes, in which our solver completely eliminates all intersecting primitive pairs.

5.5 Breakdown Analysis

To maximize parallel throughput, individual ray-casting tasks across different intersection contours—which query the BVH and perform intersection tests from candidate source vertices, source edges, and destination vertices along trial directions—are batched into a single pass with Q total primitive queries. For a BVH with B leaves, H raw hits, and A adjacency relations between raw hits, evaluation requires $O(Q \log B + H)$ time for BVH traversal and $O(Q + H + A)$ time for clustering. Rather than allocating a dense $H \times H$ matrix, we store incident hits using a Compressed Sparse Row (CSR) list ($O(H)$ storage), from which connected hit clusters (*union-find components*) are extracted via parallel disjoint-set union-find in $O(Q + H + A)$ peak memory.

As shown in Fig. 13 on the most computationally intensive initial frame from [Jang et al. \[2025\]](#) (SMPL_019, 2,082 initial EF pairs), ray casting is the primary bottleneck on CPU (74.8% of total time, dominated by edge-edge queries at 73.8%, alongside 15.2% face-vertex and 11.0% vertex-face tests). On GPU, parallel BVH traversal accelerates ray casting by 7.2 $\times$ (dropping from 648.3 ms to 90.3 ms, with 71.0% EE, 16.5% FV, 12.5% VF), shifting the remaining cost to CSR graph building (37.7%) and host evaluation/assembly (34.1%), lowering overall time from 866.3 ms to 542.3 ms. This host evaluation phase comprises CPU-side candidate post-processing and metric verification, including topological hop and geometric distance computation, validity filtering, and continuous penetration volume and displacement objective evaluation. Implementation details of the parallelized ray casting and cluster culling pipeline are provided in the supplemental material. Achieving higher degrees of GPU parallelism—particularly migrating the remaining host-side evaluation and assembly to device kernels—is left as future work.

Table 2. Convergence and performance comparison across benchmark scenes and validation cases. *Init. C* is the initial contour count, *Iter.* the first zero-EF iteration, and *Budget* the baseline iteration limit. Walltimes span their full iteration budgets rather than time to solution.

Scene / Configuration	Input				Ours				ICM			GIA		
	#V	#F	Init. EF	Init. C	Fin. EF	Iter.	Wall GPU (s)	Wall CPU (s)	Fin. EF	Budget	Wall (s)	Fin. EF	Budget	Wall (s)
Large-scale benchmarks														
Synthetic (Fig. 1)	15,647	30,780	6,331	36	0	88	341.410	494.770	691	4,000	251.132	1,977	4,000	4,163.109
Round Table (Fig. 12a)	105,614	210,157	48,686	2,270	0	320	224.420	315.950	1,840	4,000	570.727	13,975	4,000	6,289.804
Cloth Drop (Fig. 1)	111,208	212,636	49,124	3,646	0	48	68.110	83.260	953	1,200	106.902	72,693	1,200	340.585
Hair Card (Fig. 12b)	32,579	41,668	168,914	10,268	0	889	148.780	450.100	1,428	4,000	422.698	55,651	4,000	2,800.250
Canonical seven configurations														
BB/II (Fig. 8)	153	256	16	1	0	2	0.116	0.019	16	300	1.759	10	300	1.844
BI/BI (Fig. 8)	153	256	4	1	0	2	0.107	0.021	4	300	2.285	0	55	0.268
BLI (Fig. 8)	117	192	7	6	0	1	0.130	0.045	25	300	2.555	7	300	1.374
Closed (Fig. 8)	154	256	47	6	0	2	0.182	0.069	59	300	2.790	87	300	2.706
Cross (Fig. 8)	189	320	53	2	0	2	0.153	0.055	53	300	3.075	53	300	2.356
Eight (Fig. 8)	181	320	14	8	0	2	0.439	0.145	14	300	1.124	202	300	3.193
LL (Fig. 8)	145	256	21	10	0	1	0.220	0.133	51	300	3.048	90	300	2.383
Additional validation and illustrative cases														
Pipeline Unit (Fig. 3)	153	256	16	1	0	3	0.116	0.019	16	300	1.759	10	300	1.844
Klein Bottle (Fig. 5)	2,625	5,224	76	2	0	11	1.423	0.560	82	300	5.688	413	300	8.360
Dual-Cloth Conflict (Fig. 7a,b)	8,082	15,776	495	2	0	2	1.212	1.619	453	300	11.885	2	300	10.529
Cloth Cylinder (Fig. 7c,d)	9,235	17,708	380	2	0	5	1.704	1.245	381	300	12.261	946	300	35.386
Bezier Cloth	2,116	4,050	242	2	0	3	0.131	0.109	242	300	6.069	150	300	6.366
Cloth Torus	5,304	10,446	138	1	0	4	0.224	0.213	134	300	8.301	135	300	12.451
Cloth Sorento	14,317	28,468	214	1	0	2	1.432	0.986	194	300	11.679	460	300	30.545




Fig. 10. Visual comparison of shape fidelity and volume preservation across representative cases from the benchmark dataset of Jang et al. [2025]. Progressing from initial tangled inputs (a) through soft de-penetration (b), stiff regularization (c), bounded stepping (d), and bounded stepping followed by 200 standard IPC relaxation steps (e) eliminates localized pinching and fully restores equilibrium surface fidelity.

5.6 Limitations and Failure Cases

While our method reliably resolves complex multi-layer entanglements across the evaluated benchmarks, the contour-wide linear separation path $\mathbf{r}$ carries three structural limitations.

First, a single direction cannot represent spatially varying untangling flows. As illustrated in Fig. 14, continuous twisting (such as the multi-turn rotating cylinder) requires different parts of the

connected hit support to move in opposing directions. Recall from Eq. (7) that the signed depth of hit h is $\delta_h = \mathbf{v}_h / (\mathbf{n}_h^\top \mathbf{r})$, where $\mathbf{v}_h$ is fixed for given nodal positions. When all characteristic normals $\mathbf{n}_h$ within a connected component agree in orientation relative to $\mathbf{r}$, their depths share the same sign, and a single displacement along $\mathbf{r}$ separates every hit simultaneously. Under continuous twisting, however, $\mathbf{n}_h^\top \mathbf{r}$ changes sign across the component: hits on one side of the

Table 3. Selection distribution of the 10 initial direction candidates across geometric families on the ISIR dataset [Jang et al. 2025] and Thing10K [Zhou and Jacobson 2016].

Dataset	Scatter Axes (6)	Mesh Center Line (2)	Mesh-to-Contour (2)
ISIR	87.4%	3.7%	8.9%
Thing10K	22.7%	63.8%	13.5%

Table 4. Ablation study on response stiffness κ across the 59 benchmark cases of Jang et al. [2025]. $\kappa = 10^8$ resolves all cases; both lower and higher stiffness degrade convergence.

κ	Animal (20)	Human (30)	Misc (9)	All (59)
10^4	0/20 (0%)	0/30 (0%)	0/9 (0%)	0/59 (0%)
10^5	1/20 (5%)	0/30 (0%)	2/9 (22%)	3/59 (5%)
10^6	7/20 (35%)	4/30 (13%)	7/9 (78%)	18/59 (31%)
10^7	15/20 (75%)	26/30 (87%)	8/9 (89%)	49/59 (83%)
10^8	20/20 (100%)	30/30 (100%)	9/9 (100%)	59/59 (100%)
10^9	20/20 (100%)	28/30 (93%)	8/9 (89%)	56/59 (95%)
10^{10}	20/20 (100%)	27/30 (90%)	8/9 (89%)	55/59 (93%)
10^{11}	20/20 (100%)	26/30 (87%)	8/9 (89%)	54/59 (92%)
10^{12}	16/20 (80%)	19/30 (63%)	7/9 (78%)	42/59 (71%)

Table 5. Volume preservation ratio $\mathcal{V}/\mathcal{V}_0$ across benchmark cases of Jang et al. [2025] under different shape fidelity stages compared with ISIR.

Subset	Soft ($\kappa_b=0.05$)	Stiff ($\kappa_b=1.0$)	Bounded ($\kappa_b=1.0$)	Bounded + 200 IPC	ISIR
Animal	0.980	0.984	1.004	1.009	1.034
Human	0.910	0.901	0.977	1.006	1.070
Misc	0.844	0.841	0.992	1.002	1.002
Mean	0.911	0.909	0.991	1.006	1.035

Table 6. Evaluation on the benchmark dataset of Jang et al. [2025]. We report subset repair success rate, mean convergence iterations, final residual EF counts, and total subset walltime under both GPU and CPU execution paths compared with ISIR [Jang et al. 2025].

Subset	#Cases	Avg. Init.	ISIR			Ours				
			Succ.	Iters	Fin. EF	Succ.	Iters	Fin. EF	GPU (s)	CPU (s)
Animal	20	481.6	15%	65.3	653.3	100%	12.30	0	28.5	47.0
Human	30	857.5	0%	237.3	966.0	100%	14.67	0	46.2	80.1
Misc	9	499.8	80%	33.5	59.0	100%	31.67	0	21.8	31.4
All	59	675.5	25.45%	146.0	707.8	100%	16.51	0	96.5	158.5

twist yield positive depths (separation), while hits on the opposite side yield negative depths (deeper penetration under the same r). No choice of r can make all depths positive when the characteristic normals span both hemispheres. Our scale-invariant centrifugal radial basis $\mathbf{r}_{\text{radial}}$ handles rotationally symmetric multi-lobe folds through conformal expansion, but continuous helical entanglements fundamentally demand a spatially varying deformation flow field.

Second, the candidate domain can grow disproportionately to the actual intersection region. For multi-body intersections, even a




Fig. 11. EF-pair convergence on the four large-scale scenes.

small contour between two objects seeds adaptive k -ring expansion (Sec. 4.1) on both bodies; the resulting candidate set may span the entirety of the opposing mesh, producing separation paths far larger than the local penetration warrants. For self-intersections, cluster culling on a single connected mesh can chain hits along the surface into a fully connected component that covers the entire sheet rather than the local entanglement. The intrinsic contour-side candidate restriction (Sec. 4.2) suppresses much of this spatial leakage in practice, but it remains a heuristic: no formal guarantee prevents the candidate support from exceeding the physically relevant region, and the method does not provide global convergence guarantees.

These open challenges motivate future research into spatially continuous untangling flow fields and tighter candidate-domain bounds with provable locality.

6 Conclusions and Future Work

We present a type-free geometric untangling framework for resolving intersecting initial states without classifying contours into predefined topological categories. By treating intersection contours as topological seeds, our pipeline constructs coherent physical separation supports via multi-type ray casting, filters geometric debris through dual-sided product-cell clustering, optimizes separating directions on the unit sphere, and converts the validated hits into standard penalty constraints for implicit Newton-Raphson simulation. Coupled with adaptive k -ring expansion, support validity filtering, and contour merging, our method provides a unified, continuous response construction for both multi-body collisions and self-intersections. Extensive experiments on multi-body collisions and self-intersections—including multi-layer folds, topological singularities, and large-scale thin-shell configurations—show that the method achieves practical convergence on benchmarks where classical GIA and ICM baselines stall or diverge. Adaptive k -ring candidate restriction (Sec. 4.1) reduces ray query volume by one to two orders of magnitude without altering the valid-hit set, while conflict resolution via contour merging (Sec. 4.3) and intrinsic contour-side candidate restriction (Sec. 4.2) extend stability to dense multi-contour



Fig. 12. Untangling results on representative large-scale benchmarks.




Fig. 13. Execution breakdown across CPU (left) and GPU (right) on the initial frame (SMPL_019) from Jang et al. [2025].

scenes where unrestricted candidate domains would otherwise cause spatial leakage or contradictory separation forces. Future work will investigate spatially continuous untangling flow fields, material-stiffness-aware depth blending, fully device-resident GPU pipeline implementations, and codimensional geometric untangling across broader design and animation workflows.

Acknowledgments

This work is supported by National Natural Science Foundation of China (62472364), Fujian Association for Science and Technology (FJKX-2025XRH01), Fujian Provincial Natural Science Foundation of China (2026J002001), and Xiamen Natural Science Foundation of China (3502Z202672001). Shihui Guo is the corresponding author.

References

- Ryoichi Ando. 2024. A cubic barrier with elasticity-inclusive dynamic stiffness. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–13.
- David Baraff and Andrew Witkin. 1998. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. ACM, New York, NY, USA, 43–54.
- David Baraff, Andrew Witkin, and Michael Kass. 2003. Untangling cloth. *ACM Transactions on Graphics (TOG)* 22, 3 (2003), 862–870.




Fig. 14. A representative failure mechanism under complex topological entanglement. The lower part shows the intersection contours. Under the contour-wide linear path, local adjacency can over-connect narrow regions near loop vertices and produce a massive connected hit set. Leakage-aware range selection can limit this growth, but it cannot replace the missing spatially varying separation field.

- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: fusing constraint projections for fast simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Robert Bridson, Ronald Fedkiw, and John Anderson. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. ACM, New York, NY, USA, 594–603.
- Thomas Buffet, Damien Rohmer, Loic Barthe, Laurence Boissieux, and Marie-Paule Cani. 2019. Implicit untangling: A robust solution for modeling layered clothing. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- Ick-Hoon Cha and Hyeon-Seok Ko. 2020. Tanglement resolution in clothing simulation with explicit convergence. *IEEE Transactions on Visualization and Computer Graphics* 28, 7 (2020), 2764–2775.
- Ick-Hoon Cha and Hyeon-Seok Ko. 2022. BLI-resolver: resolving the boundary-loop-interior type intersections for clothing simulation. *The Visual Computer* 38, 4 (2022), 1359–1368.
- Anka He Chen, Ziheng Liu, Yin Yang, and Cem Yuksel. 2024b. Vertex Block Descent. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–16.

- He Chen, Elie Diaz, and Cem Yuksel. 2023. Shortest path to boundary for self-intersecting meshes. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–15.
- Honglin Chen, Hsueh-Ti Derek Liu, David IW Levin, Changxi Zheng, and Alec Jacobson. 2024a. Stabler neo-hookean simulation: Absolute eigenvalue filtering for projected Newton. In *ACM SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery, New York, NY, USA, 1–10.
- Jiong Chen, Florian Schäfer, Jin Huang, and Mathieu Desbrun. 2021. Multiscale Cholesky preconditioning for ill-conditioned problems. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13.
- Xuwen Chen, Junyu Wang, Cheng Yu, Xingyu Ni, Meng Zhang, Bin Wang, Mengyu Chu, and Baoquan Chen. 2026. Floating-Point Robustness in Parametric Surface Continuous Collision Detection: From Algorithm to Benchmarking. *ACM Transactions on Graphics* 45, 4, Article 141 (2026), 14 pages.
- Xuwen Chen, Cheng Yu, Xingyu Ni, Mengyu Chu, Bin Wang, and Baoquan Chen. 2024c. A Time-Dependent Inclusion-Based Method for Continuous Collision Detection between Parametric Surfaces. *ACM Transactions on Graphics* 43, 6, Article 223 (2024), 11 pages.
- Kwang-Jin Choi and Hyeon-Seok Ko. 2005. Stable but responsive cloth. In *ACM SIGGRAPH 2005 Courses*. ACM, New York, NY, USA, 1–es.
- Xudong Feng, Wenchao Huang, Weiwei Xu, and Huamin Wang. 2022. Learning-based bending stiffness parameter estimation by a drape tester. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–16.
- Xudong Feng, Huamin Wang, Yin Yang, and Weiwei Xu. 2024. Neural-assisted homogenization of yarn-level cloth. In *ACM SIGGRAPH 2024 Conference Papers*. ACM, New York, NY, USA, 1–10.
- Zachary Ferguson, Minchen Li, Teseo Schneider, Francisca Gil-Ureta, Timothy Langlois, Chenfanfu Jiang, Denis Zorin, Danny M Kaufman, and Daniele Panozzo. 2021. Intersection-free rigid body dynamics. *ACM Transactions on Graphics* 40, 4 (2021), 1–16.
- Marco Fratarcangeli, Valentina Tibaldo, Fabio Pellacini, et al. 2016. Vivace: a practical gauss-seidel method for stable soft body dynamics. *ACM Trans. Graph.* 35, 6 (2016), 214–1.
- Theodore F Gast, Craig Schroeder, Alexey Stomakhin, Chenfanfu Jiang, and Joseph M Teran. 2015. Optimization integrator for large time steps. *IEEE transactions on visualization and computer graphics* 21, 10 (2015), 1103–1115.
- Elmer G Gilbert, Daniel W Johnson, and S Sathiya Keerthi. 1988. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal on Robotics and Automation* 4, 2 (1988), 193–203.
- Artur Grigorev, Giorgio Becherini, Michael Black, Otmar Hilliges, and Bernhard Thomaszewski. 2024. Contourcraft: Learning to resolve intersections in neural multi-garment simulations. In *ACM SIGGRAPH 2024 conference papers*. ACM, New York, NY, USA, 1–10.
- Eitan Grinspun, Anil N Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. ACM, New York, NY, USA, 62–67.
- Dewen Guo, Minchen Li, Yin Yang, Sheng Li, and Guoping Wang. 2024. Barrier-Augmented Lagrangian for GPU-based Elastodynamic Contact. *ACM Trans. Graph.* 43, 6, Article 225 (Nov. 2024), 17 pages. doi:10.1145/3687988
- Dewen Guo, Zixuan Lu, Zhiyong He, Yuqi Meng, Bohan Wang, Lei Lan, Weiwei Xu, Chenfanfu Jiang, and Yin Yang. 2026a. JGS2-GQ: Training-free 2nd Jacobi with Gaussian Quadrature. *ACM Trans. Graph.* 45, 4 (2026), 1–12.
- Dewen Guo, Ran Tian, Sinuo Liu, Guoping Wang, and Sheng Li. 2025a. Diagonal Hessian Proxy for Efficient Elastic Simulation Using Peridynamics. *IEEE Transactions on Visualization and Computer Graphics* 31, 9 (2025), 6466–6483. doi:10.1109/TVCG.2024.3516480
- Dewen Guo, Zhendong Wang, Minchen Li, Sheng Li, Guoping Wang, Huamin Wang, Chenfanfu Jiang, and Yin Yang. 2026b. Heterogeneous Subspace Corrections for GPU Deformable Multibody Dynamics. *ACM Trans. Graph.* 45, 4 (2026), 1–14.
- Dewen Guo, Zhendong Wang, Zegao Liu, Sheng Li, Guoping Wang, Yin Yang, and Huamin Wang. 2025b. Fast Physics-Based Modeling of Knots and Ties using Templates. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers (SIGGRAPH Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 43, 9 pages. doi:10.1145/3721238.3730622
- Dewen Guo, Zhendong Wang, Zegao Liu, Sheng Li, Guoping Wang, Yin Yang, and Huamin Wang. 2025c. Progressive Outfit Assembly and Instantaneous Pose Transfer. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*. ACM, New York, NY, USA, 1–12.
- David Harmon, Etienne Vouga, Breannan Smith, Rasmus Tamstorf, and Eitan Grinspun. 2009. Asynchronous contact mechanics. In *ACM SIGGRAPH 2009 papers*. Association for Computing Machinery, New York, NY, USA, 1–12.
- Zhiyong He, Dewen Guo, Minghao Guo, Yili Zhao, Wojciech Matusik, Hao Su, Chenfanfu Jiang, Peter Yichen Chen, and Yin Yang. 2026. M-ABD: Scalable, Efficient, and Robust Multi-Affine-Body Dynamics. *ACM Trans. Graph.* 45, 4, Article 99 (July 2026), 19 pages. doi:10.1145/3811276
- Jerry Hsu, Tongtong Wang, Kui Wu, and Cem Yuksel. 2025. Stable Cosserat Rods. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*. Association for Computing Machinery, New York, NY, USA, 1–10.
- Kemeng Huang, Floyd M Chitalu, Huancheng Lin, and Taku Komura. 2024. GIPC: Fast and stable gauss-newton optimization of IPC barrier energy. *ACM Transactions on Graphics* 43, 2 (2024), 1–18.
- Kemeng Huang, Xinyu Lu, Huancheng Lin, Taku Komura, and Minchen Li. 2025. StiffGIPC: Advancing GPU IPC for stiff affine-deformable simulation. *ACM Transactions on Graphics* 44, 3 (2025), 1–20.
- Thomas JR Hughes. 2012. *The finite element method: linear static and dynamic finite element analysis*. Courier Corporation, North Chelmsford, MA, USA.
- Wonjong Jang, Yuchool Jung, Gyeongmin Lee, and Seungyong Lee. 2025. Instant Self-Intersection Repair for 3D Meshes. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–14.
- Changsoo Je, Min Tang, Youngeun Lee, Minkyoun Lee, and Young J Kim. 2012. Poly-Depth: Real-time penetration depth computation using iterative contact-space projection. *ACM Transactions on Graphics (TOG)* 31, 1 (2012), 1–14.
- Tero Karras. 2012. Maximizing parallelism in the construction of BVHs, octrees, and k-d trees. In *Proceedings of the Fourth ACM SIGGRAPH/Eurographics Conference on High-Performance Graphics*. Association for Computing Machinery, New York, NY, USA, 33–37.
- Theodore Kim. 2020. A Finite Element Formulation of Baraff-Witkin Cloth. *Computer Graphics Forum* 39, 8 (2020), 171–179.
- Theodore Kim and David Eberle. 2020. Dynamic deformables: implementation and production practicalities. In *ACM SIGGRAPH 2020 Courses*. Association for Computing Machinery, New York, NY, USA, 1–182.
- Lei Lan, Danny M Kaufman, Minchen Li, Chenfanfu Jiang, and Yin Yang. 2022a. Affine body dynamics: Fast, stable & intersection-free simulation of stiff materials. arXiv preprint arXiv:2201.10022.
- Lei Lan, Minchen Li, Chenfanfu Jiang, Huamin Wang, and Yin Yang. 2023. Second-order Stencil Descent for Interior-point Hyperelasticity. *ACM Transactions on Graphics* 42, 4 (2023), 1–13.
- Lei Lan, Zixuan Lu, Jingyi Long, Chun Yuan, Xuan Li, Xiaowei He, Huamin Wang, Chenfanfu Jiang, and Yin Yang. 2024. Efficient GPU Cloth Simulation with Non-distance Barriers and Subspace Reuse. arXiv preprint arXiv:2407.01234.
- Lei Lan, Guanqun Ma, Yin Yang, Changxi Zheng, Minchen Li, and Chenfanfu Jiang. 2022b. Penetration-free projective dynamics on the GPU. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–16.
- Dohae Lee, Hyun Kang, and In-Kwon Lee. 2023. Clothcombo: Modeling inter-cloth interaction for draping multi-layered clothes. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–13.
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy R 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.* 39, 4 (2020), 49.
- Minchen Li, Ming Gao, Timothy Langlois, Chenfanfu Jiang, and Danny M. Kaufman. 2019. Decomposed Optimization Time Integrator for Large-Step Elastodynamics. *ACM Transactions on Graphics* 38, 4 (2019), 1–19.
- Minchen Li, Danny M Kaufman, and Chenfanfu Jiang. 2021. Codimensional incremental potential contact. *ACM Transactions on Graphics* 40, 4 (2021), 1–24.
- Zhouyingcheng Liao, Sinan Wang, and Taku Komura. 2024. SENC: Handling self-collision in neural cloth simulation. In *European Conference on Computer Vision*. Springer, Springer, Cham, Switzerland, 385–402.
- Tiantian Liu, Adam W Bargteil, James F O'Brien, and Ladislav Kavan. 2013. Fast simulation of mass-spring systems. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–7.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-newton methods for real-time simulation of hyperelastic materials. *ACM Transactions on Graphics (TOG)* 36, 3 (2017), 1–16.
- Miles Macklin and Matthias Muller. 2021. A constraint-based formulation of stable neo-hookean materials. In *Proceedings of the 14th ACM SIGGRAPH conference on motion, interaction and games*. Association for Computing Machinery, New York, NY, USA, 1–7.
- Miles Macklin, Matthias Müller, and Nuttapon Chentanez. 2016. XPBD: position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games*. ACM, New York, NY, USA, 49–54.
- Miles Macklin, Kier Storey, Michelle Lu, Pierre Terdiman, Nuttapon Chentanez, Stefan Jeschke, and Matthias Müller. 2019. Small steps in physics simulation. In *Proceedings of the 18th annual ACM siggraph/eurographics symposium on computer animation*. ACM, New York, NY, USA, 1–7.
- Khaled Mamou, E Lengyel, and A Peters. 2016. Volumetric hierarchical approximate convex decomposition. *Game engine gems* 3 (2016), 141–158.
- Leticia Mattos Da Silva, Silvia Sellán, Natalia Pacheco-Tallaj, and Justin Solomon. 2025. Variational Elastodynamic Simulation. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*. ACM,

- New York, NY, USA, 1–11.
- Tomas Möller. 1997. A fast triangle-triangle intersection test. *Journal of graphics tools* 2, 2 (1997), 25–30.
- Matthias Müller. 2008. Hierarchical position based dynamics. In *Workshop on Virtual Reality Interaction and Physical Simulation (VRIPHYS)*. The Eurographics Association, Goslar, DEU, 1–10.
- Matthias Müller, Nuttapon Chentanez, Tae-Yong Kim, and Miles Macklin. 2015. Air meshes for robust collision handling. *ACM Transactions on Graphics (ToG)* 34, 4 (2015), 1–9.
- Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position based dynamics. *Journal of Visual Communication and Image Representation* 18, 2 (2007), 109–118.
- Matthias Müller, Miles Macklin, Nuttapon Chentanez, Stefan Jeschke, and Tae-Yong Kim. 2020. Detailed rigid body simulation with extended position based dynamics. *Computer Graphics Forum* 39, 8 (2020), 101–112.
- Miguel A Otaduy, Rasmus Tamstorf, Denis Steinemann, and Markus Gross. 2009. Implicit contact handling for deformable objects. *Computer Graphics Forum* 28, 2 (2009), 559–568.
- Igor Santesteban, Miguel Otaduy, Nils Thuerey, and Dan Casas. 2022. ULNeF: Untangled layered neural fields for mix-and-match virtual try-on. *Advances in Neural Information Processing Systems* 35 (2022), 12110–12125.
- Breannan Smith, Fernando De Goes, and Theodore Kim. 2018. Stable neo-hookean flesh simulation. *ACM Transactions on Graphics (TOG)* 37, 2 (2018), 1–15.
- Breannan Smith, Fernando De Goes, and Theodore Kim. 2019. Analytic eigensystems for isotropic distortion energies. *ACM Transactions on Graphics (TOG)* 38, 1 (2019), 1–15.
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized yarn-level cloth. *ACM Trans. Graph.* 39, 4 (2020), 48.
- Rasmus Tamstorf, Toby Jones, and Stephen F McCormick. 2015. Smoothed aggregation multigrid for cloth simulation. *ACM Trans. Graph. (TOG)* 34, 6 (2015), 245.
- Min Tang and Young J Kim. 2014. Interactive generalized penetration depth computation for rigid and articulated models using object norm. *ACM Transactions on Graphics (TOG)* 33, 1 (2014), 1–15.
- Min Tang, Tongtong Wang, Zhongyuan Liu, Ruofeng Tong, and Dinesh Manocha. 2018. I-Cloth: Incremental collision handling for GPU-based interactive cloth simulation. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–10.
- Maxime Tournier, Matthieu Nesme, Benjamin Gilles, and François Faure. 2015. Stable constrained dynamics. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–10.
- Gino van den Bergen. 2001. Proximity queries and penetration depth computation on 3D game objects. Game Developers Conference (GDC). Introduces the Expanding Polytope Algorithm (EPA) for penetration-depth recovery.
- Pascal Volino, Martin Courchesne, and Nadia Magnenat Thalmann. 1995. Versatile and efficient techniques for simulating cloth and other deformable objects. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, New York, NY, USA, 137–144.
- Pascal Volino and Nadia Magnenat-Thalmann. 2006. Resolving surface collisions through intersection contour minimization. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 1154–1159.
- Pascal Volino, Nadia Magnenat-Thalmann, and Francois Faure. 2009. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Transactions on Graphics* 28, 4 (2009), Article–No.
- Bolun Wang, Zachary Ferguson, Xin Jiang, Marco Attene, Daniele Panozzo, and Teseo Schneider. 2022. Fast and exact root parity for continuous collision detection. *Computer Graphics Forum* 41, 2 (2022), 355–363.
- Bolun Wang, Zachary Ferguson, Teseo Schneider, Xin Jiang, Marco Attene, and Daniele Panozzo. 2021. A Large Scale Benchmark and an Inclusion-Based Algorithm for Continuous Collision Detection. *ACM Transactions on Graphics* 40, 5, Article 188 (2021), 16 pages.
- Huamin Wang. 2014. Defending continuous collision detection against errors. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.
- Huamin Wang. 2015. A chebyshev semi-iterative approach for accelerating projective and position-based dynamics. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–9.
- Tianyu Wang, Jiong Chen, Dongping Li, Xiaowei Liu, Huamin Wang, and Kun Zhou. 2023a. Fast GPU-based two-way continuous collision handling. *ACM Transactions on Graphics* 42, 5 (2023), 1–15.
- Xinlei Wang, Min Tang, Dinesh Manocha, and Ruofeng Tong. 2018a. Efficient BVH-based collision detection scheme with ordering and restructuring. *Computer Graphics Forum* 37, 2 (2018), 227–237.
- Zhendong Wang, Longhua Wu, Marco Fratarcangeli, Min Tang, and Huamin Wang. 2018b. Parallel multigrid for nonlinear cloth simulation. *Computer Graphics Forum* 37, 7 (2018), 131–141.
- Zhendong Wang, Yin Yang, and Huamin Wang. 2023b. Stable discrete bending by analytic eigensystem and adaptive orthotropic geometric stiffness. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–16.
- Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. 2022. Approximate convex decomposition for 3d meshes with collision-aware concavity and tree search. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–18.
- Martin Wicke, Hermes Lanker, and Markus Gross. 2006. Untangling cloth with boundaries. In *Proc. of Vision, Modeling, and Visualization*, Vol. 349. VMV, Germany, 356.
- Audrey Wong, David Eberle, and Theodore Kim. 2018. Clean cloth inputs: Removing character self-intersections with volume simulation. In *ACM SIGGRAPH 2018 Talks*. ACM, New York, NY, USA, 1–2.
- Botao Wu, Zhendong Wang, and Huamin Wang. 2022. A GPU-based multilevel additive schwarz preconditioner for cloth and deformable body simulation. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–14.
- Haomiao Wu and Theodore Kim. 2023. An eigenanalysis of angle-based deformation energies. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 6, 3 (2023), 1–19.
- Longhua Wu, Botao Wu, Yin Yang, and Huamin Wang. 2020. A safe and fast repulsion method for GPU-based cloth self collisions. *ACM Transactions on Graphics (TOG)* 40, 1 (2020), 1–18.
- Jieyin Yang and Xiaohong Jia. 2025. Overlap Region Extraction of Two NURBS Surfaces. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–20.
- Juntao Ye, Guanghui Ma, Liguang 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.
- 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*. IEEE, IEEE, Piscataway, NJ, USA, 352–357.
- 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*. Eurographics Association, Aire-la-Ville, Switzerland, 311–316.
- Chun Yuan, Zixuan Lu, Haoyang Shi, Dewen Guo, Huamin Wang, Chenfanfu Jiang, Zherong Pan, Kui Wu, and Yin Yang. 2026. Interactive Yarn-level Knitwear with Nested Douglas-Rachford Splitting. *ACM Trans. Graph.* 45, 4, Article 103 (July 2026), 22 pages. doi:10.1145/3811277
- Xiaodi Yuan, Fanbo Xiang, Yin Yang, and Hao Su. 2025. C5D: A Robust and Efficient Continuous Collision Detection Framework for Convex Shapes. *ACM Transactions on Graphics* 44, 4 (2025), 1–14.
- Paul Zhang, Zoë Marschner, Justin Solomon, and Rasmus Tamstorf. 2023. Sum-of-Squares Collision Detection for Curved Shapes and Paths. In *ACM SIGGRAPH 2023 Conference Proceedings*. ACM, New York, NY, USA, Article 76, 11 pages.
- Juntian Zheng, Zhao Feng Luo, and Minchen Li. 2025. Robust and Efficient Penetration-Free Elastodynamics without Barriers. *ACM Transactions on Graphics* 44, 4 (2025), 1–14.
- Shaokun Zheng, Zhiqian Zhou, Xin Chen, Difei Yan, Chuyan Zhang, Yuefeng Geng, Yan Gu, and Kun Xu. 2022. Luisarender: A high-performance rendering framework with layered and unified interfaces on stream architectures. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–19.
- Qingnan Zhou and Alec Jacobson. 2016. Thing10k: A dataset of 10,000 3d-printing models. arXiv preprint arXiv:1605.04797.