Extremum Sum Proximity Potential

FEDERICO SICHETTI, NYU Courant Institute of Mathematical Sciences, USA
ZIZHOU HUANG, Roblox, USA
DANIELE PANOZZO, NYU Courant Institute of Mathematical Sciences, USA
ENRICO PUPPO, University of Genoa, Italy
DENIS ZORIN, New York University, USA

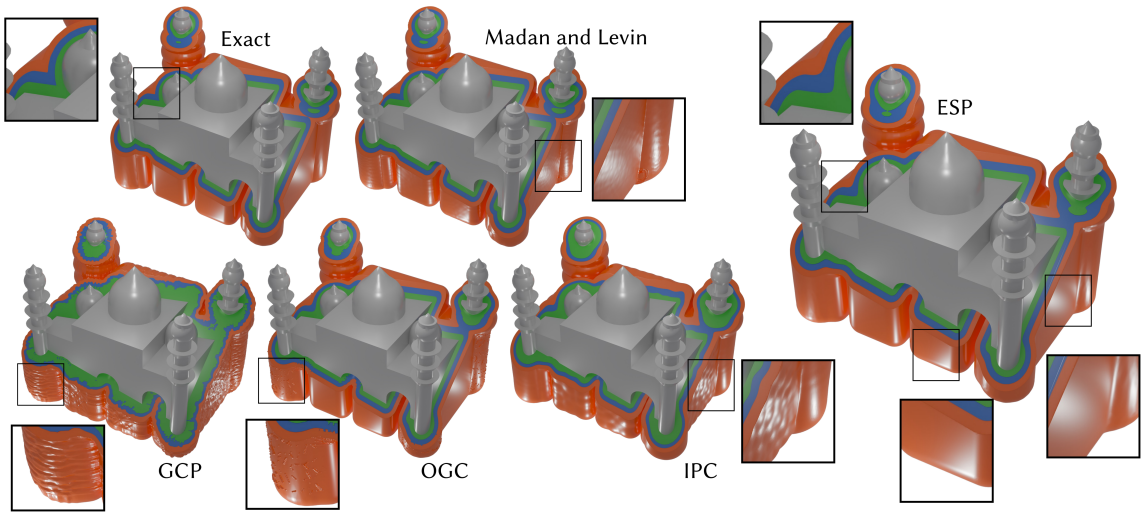


Fig. 1. Level sets of the exact distance-based potential (Exact), GCP [Huang et al. 2025], OGC [Chen et al. 2025], IPC [Li et al. 2020], Madan-Levin smoothed distance potential [Madan and Levin 2022], and the potential proposed in this paper (ESP). Note the normal discontinuities near the medial axis for Exact, Smooth but bumpy surface for GCP and IPC, smaller unevenness for Madan-Levin, and discontinuities for OGC. The insets for the same locations for ESP show the absence of these artifacts.

Distance-based barrier potentials and smooth distance functions are fundamental in geometry processing and physics-based simulation. In this paper, we propose a general, simple definition of a differentiable proximity potential, *extremum sum potential* (ESP), obtained as a signed combination of critical values of the distance function. Surprisingly, this potential is differentiable, while remaining qualitatively similar to the closest-point distance function. The resulting potential shares all key features with known potentials, including barrier behavior, absence of spurious forces, mesh resolution independence etc. Viewed as a distance approximation, the inverse potential defines a smooth distance-like field that avoids sharp medial-axis artifacts.

CCS Concepts: • **Computing methodologies** → **Mesh models**; **Volumetric models**; **Physical simulation**.

Authors' Contact Information: Federico Sichetti, NYU Courant Institute of Mathematical Sciences, New York, USA, f.sichetti@nyu.edu; Zizhou Huang, Roblox, San Mateo, USA, zizhou@nyu.edu; Daniele Panozzo, NYU Courant Institute of Mathematical Sciences, New York, USA, panozzo@nyu.edu; Enrico Puppo, University of Genoa, Genoa, Italy, enrico.puppo@unige.it; Denis Zorin, Courant Institute School of Mathematics, Computing and Data Science, New York University, New York, NY, USA, dzorin@cs.nyu.edu.

This work is licensed under a Creative Commons Attribution 4.0 International License.
SA Conference Papers '26, Kuala Lumpur, Malaysia
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2842-6/2026/12
https://doi.org/10.1145/3829340.3842312

Additional Key Words and Phrases: Repulsive potentials, Distance functions, Contact

ACM Reference Format:
Federico Sichetti, Zizhou Huang, Daniele Panozzo, Enrico Puppo, and Denis Zorin. 2026. Extremum Sum Proximity Potential. In *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26), December 01–04, 2026, Kuala Lumpur, Malaysia.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3829340.3842312

1 Introduction

Distance functions and distance-based force potentials play a central role in geometry processing, physics-based simulation, and optimization. In geometry processing, distance fields are widely used for deformations and sculpting, shape analysis, interference detection, meshing, and other applications. In contact and collision handling, barrier potentials derived from distance functions provide a principled way to enforce non-interpenetration. Geometric barrier potentials and smooth distance functions are closely related: the inverse of a (positive) barrier potential defines a smooth approximation of distance, and conversely, a negative power or logarithm of a point-to-surface distance can be used as a barrier.

Despite their ubiquity, designing a distance function or potential remains a significant challenge due to a set of competing requirements: retaining qualitative similarity to the (non-smooth) Euclidean distance function, differentiability, and efficient evaluation. Prior work has proposed a range of differentiable distance functions and barrier potentials, each balancing these features and the needs of specific applications.

In this paper, we present a new formulation of a *differentiable* proximity potential and an associated distance function, based on a sum of barrier-function values at all *extrema* of the distance to the surface, not just minima, with signs in the summation determined by the Morse indices of the extrema. Remarkably, this simple formulation results in a potential that reduces to the surface distance close to a smooth surface, and at the same time is globally differentiable, both for smooth surfaces and meshes.

We show that for meshes, the proposed formulation coincides with the heuristic for approximating local maxima of a localized contact potential [Li et al. 2023]. Its integrated version satisfies the key contact energy properties proposed by Huang et al. [2025], while being substantially simpler, and offers additional properties useful for smooth offset construction and for reducing artifacts in contact handling (Section 5).

In summary, this paper makes the following contributions:

- a general formulation of a differentiable proximity potential as an alternating sum over critical points of the distance function;
- theoretical analysis of its properties, with emphasis on smooth distance field construction and contact handling;
- a discretized version of the potential, including high-order quadrature support, retaining all key properties;
- applications of the potential behavior to offset surfaces, and simulations involving rigid and elastic objects.

Our results suggest that a single, elementary potential construction, with parameter-independent level set structure, captures the core geometric properties needed for both smooth distance fields and robust contact.

2 Related work

We focus on the most closely related recent constructions, referring to Huang et al. [2025] (supplement) for a detailed survey of the contact literature, barrier formulations in particular.

A general concept of *Euler integration*, which uses the Euler characteristic of subsets as a measure, was introduced for integer-valued functions in [Schapira 1995; Viro 1988]. Baryshnikov and Ghrist [2010] generalize it to continuous functions, Morse functions in particular, and demonstrate that it reduces to the signed sums of values at critical points that we consider in this work. These ideas were applied to sensor networks and shape analysis [Baryshnikov and Ghrist 2009; Crawford et al. 2020; Ghrist et al. 2018]. However, the differentiability of the resulting functions and other properties we examine here were not considered.

Barrier energies often obtained by integrating a pointwise potential, provide a common foundation for penetration-free simulation, popularized in graphics by Incremental Potential Contact (IPC) and a body of follow-up work [Ferguson et al. 2021; Huang et al. 2024; Li et al. 2020, 2023, 2021; Shen et al. 2024].

The mesh potential derived in this paper surprisingly leads to the same point-wise expression, which was proposed by Li et al. [2023] as a heuristic "smoothed max/min" interpretation, under a strong localization assumption, with the mesh potential extent bounded by the adjacent edge lengths. Our main contribution is to show that the underlying continuum functional is an alternating-sign sum over distance critical points (minima, saddles, maxima) and requires no localization assumption. This reinterpretation directly guides discretization and quadrature, preserving key properties (positivity, barrier behavior, and near-point behavior) without tying the interaction radius to mesh edge length.

Huang et al. [2025] introduce the *geometric contact potential* (GCP) and provide a systematic set of desiderata for barrier potentials in the contact setting. They derive a relatively complex potential achieving these mathematical properties, based on the idea of "smoothing out" local minima of the distance function. The proposed potential has similar properties, eliminating the same type of artifacts in the contact context, but does not require the smoothing construction to achieve differentiability.

Offset geometric contact (OGC) [Chen et al. 2025] provides an elegant geometric formulation for a barrier potential for contact that has some of the properties of GCP, and works more naturally for piecewise-linear surfaces; its main deficiency is the lack of differentiability or even continuity in concave regions and close to sharp convex corners. It can be interpreted as considering only local minima in the formulation we propose. Remarkably, the extremum-sum expression, while remaining close to these formulations, is differentiable.

Other potential methods. Du et al. [2024] propose a fully C^1 surface representation in combination with piecewise-linear meshes: they base their contact potential upon distances between the discretization nodes and the smooth surface, reducing artifacts for sliding contact. Since the barrier is evaluated only on node–surface pairs, closest pairs that are not sampled (particularly on edges) can vanish at finite energy. The method does not guarantee intersection-free iterations, which is critical for IPC's robustness [Li et al. 2020] and retained by our formulation; see Huang et al. [2025] for a detailed comparison.

More broadly, continuum barrier models based on surface/volume integrals have also been explored in computational mechanics and adapted for self-contact [Alaydin et al. 2021; Kamensky et al. 2018; Sauer and De Lorenzis 2013]. These formulations are attractive for mesh independence in principle, but in practice typically require explicit self-exclusion neighborhoods and careful localization; moreover, point-pair quadrature discretizations may fail to preserve exact non-penetration on coarse meshes.

Offset surface construction. The level sets of our potential are smooth offset surfaces; see Zint et al. [2025] for a recent overview of explicit offset construction, which produces meshes with combinatorial guarantees (watertightness, no self-intersections, correct topology), whereas we obtain a closed-form differentiable field (Figure 10).

Smoothed distance functions. A related line of work constructs smooth distance approximations as implicit fields, primarily by smoothing the min/max operation itself (e.g., soft-min / LogSumExp / Kreisselmeier–Steinhauser distances), often augmented with weighting and acceleration to address bulge artifacts and scaling to large co-dimensional inputs. Most relevantly, Madan and Levin [2022] combine a LogSumExp smooth minimum with per-primitive blending weights and a Barnes–Hut acceleration; while satisfying many of the same criteria as the construction we consider, it weighs co-dimensional features through parameters tuned per input, which makes the field discretization-dependent, and low values of the smoothing parameter make the weights overcompensate, producing creases in the level sets. Smoothness moreover holds for the exact sum, but not for what is evaluated: as the authors observe, the near/far switch of the Barnes–Hut expansion introduces (small) jumps in the field. Vaillant et al. [2013] use a smooth distance function based on RBFs to resolve self-contact in skin deformation. The zero isosurface of this function is not guaranteed to lie exactly on the input surface, thus is not useful for our purposes.

PDE-based distance approximations. A complementary family of methods obtains smooth distance-like fields as solutions of a partial differential equation posed over the ambient domain. Classical *fast marching* [Sethian 1996; Tsitsiklis 1995] solves the eikonal equation $\|\nabla u\| = 1$ on a grid; the *heat method* [Crane et al. 2013, 2017] recovers geodesic distance from a short-time heat diffusion; Belyaev and Fayolle [2015] unifies a number of variational and PDE-based constructions, including screened-Poisson and p-Poisson approximations that converge to the eikonal solution [Belyaev and Fayolle 2020; Fayolle and Belyaev 2018]; the *biharmonic distance* of Lipman et al. [2010] produces a globally smooth, intrinsic distance-like function from the Green's function of the bi-Laplacian, and similar variational formulations continue to be developed [Edelstein et al. 2023]. Point-wise evaluation is possible in some instances using Monte Carlo sampling instead of solving the PDE on an explicit space discretization [Sawhney and Crane 2020; Sawhney et al. 2023].

Since these methods compute solutions numerically, up to a prescribed tolerance, they do not satisfy important properties of a distance potential, such as positivity, monotonicity, and barrier property (see Section 3).

Repulsive surfaces and self-avoidance energies. Repulsive surface energies – such as tangent-point energies [Sassen et al. 2024; Strzelecki and von der Mosel 2013; Yu et al. 2021a,b] – are related to the geometric contact potential for a suitable choice of parameters. These energies are natural for defining self-avoiding smooth surfaces as minimizers, but are typically global, expensive without localization, and less directly suited to mesh contact barriers in the presence of sharp features (see Huang et al. [2025, Section 2.3] for a detailed discussion). The formulation we describe achieves a similar objective and admits discretizations that preserve the essential properties.

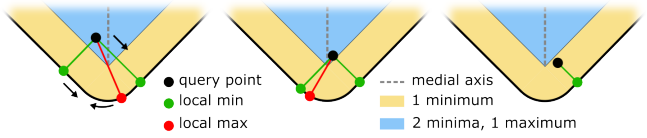


Fig. 2. Extrema of the distance from a point (black dot) to an open curve. In the blue region, there are two local minima and one local maximum; as the point approaches the boundary of the blue region, one minimum and one maximum collapse; in the yellow region, there is just one local minimum. In the expression for the extremum-sum potential, the collapsing minimum and maximum have different signs, and the difference vanishes differentiably.

3 Pointwise proximity potential

3.1 Motivating example

Our goal is to define a pointwise *proximity potential* $P(q)$, depending on a curve or surface M, that informally pushes points away from M, and is infinite for points on M. In many contexts (geometric design and shape optimization, and physical simulation), differentiability is highly desirable, and rules out the obvious solutions building directly on the closest-point distance.

As a motivating example, we consider a smooth curve $f(t)$ in the plane depicted in Figure 2. For this curve, the standard Euclidean point-to-curve distance, and, as a consequence, any potential depending on it, is not differentiable, with a gradient discontinuity at the medial axis. One possible remedy conceptually similar to the one used by Chen et al. [2025], is to combine potentials with respect to local minima of the distance function, resulting in a smooth transition near the medial axis. However, while addressing the medial axis problem, this leads to a discontinuity at points where the number of local minima changes. Another, more complex one, along the lines of Madan and Levin [2022], Huang et al. [2025] and Yu et al. [2021a] is to replace the minimal distance with a smooth kernel, and integrate over the surface, requiring surface quadrature, and making it difficult to ensure that the numerical approximation maintains the properties of the integral formulation.

Our key observation is that one can eliminate the gradient discontinuity if we consider all *extrema* of the distance (i.e., maxima and minima in the case of curves), and combine the potentials corresponding to these, with the opposite signs forced by the Morse indices of extrema that vanish and appear in pairs. In this way, when a pair of extrema appears, their initial contribution to the potential is zero, with the contribution increasing as they move apart. While continuity is intuitive in this case, differentiability and the barrier property (i.e., growth to infinity) on the surface are less obvious. It turns out that these properties hold.

We start with the smooth surface formulation, as the formulas are most transparent in this case. Our practical implementation is based on the specialization of the general formulation to meshes described below.

3.2 Smooth curves and surfaces

Suppose the curve is given in the parametric form as $f = f(s)$. We consider all extrema of the distance function from a point q to

the curve. Let s_i be the solutions of the condition for the squared-distance function $\|q - f(s)\|^2 = r(q,s)^2$ to have an extremum, i.e., $(q-f(s))\cdot f'(s) = 0$. Assuming a curve in a general position, each s_i is an isolated simple minimum or a maximum (we omit the inflection points: an inflection is a coincident minimum–maximum pair, whose two oppositely signed terms cancel exactly, so its net contribution is zero), i.e. the distance is a Morse function. Denote these sets as $S_{\min}$, $S_{\max}$. Finally, let $b(r)$ be a *barrier function*: monotonically decreasing, diverging to $+\infty$ as $r \to 0$, and twice differentiable (e.g., $b(r) = r^{-p}$ or $b(r) = -\ln r$); differentiability alone suffices for the results below, while the second derivative is used by the Newton solver. Then the formula for the proposed proximity potential is

$$P(q) = \sum_{s_i \in S_{\min}} b(r(q,s_i)) - \sum_{s_i \in S_{\max}} b(r(q,s_i)) \tag{1}$$

The intuition for why this potential is at least continuous (differentiability is proved in Theorem 3.1) is that for varying q, the extrema of $r(q,s)$ viewed as a function of s may appear and disappear only in pairs (maximum, minimum), and these appear at the same point (Figure 2), and move away continuously from each other, so initially these cancel each other in the sum, and their contribution increases continuously. While we exclude isolated points for which the distance function may have non-isolated extrema (this requires a circular arc being a part of the curve or a surface), the potential can be defined at such points by continuity. An appealing property of this formula is that it yields an explicit expression for piecewise-linear curves, yielding a "natural" discretization. Also note that the numbers of local minima and maxima on a closed curve always match (but not on an open curve as in Figure 2), and for each distance minimum, we can assign a matching maximum that is further away, which shows that the potential is positive if $b(r)$ is positive and decreasing (this is not the case for the surfaces, as discussed below).

Smooth surfaces. The formula also generalizes to smooth surfaces. We use s now to denote 2D parametric coordinates on the surface, and let s_i be the roots of the system of equations $(q-f(s))\cdot t_j(s) = 0$, where $j = 0, 1$, and t_j are a pair of independent tangent vectors at a point. Assuming isolated *simple* extrema only, i.e., ones with a non-degenerate Hessian (a *complex* extremum, such as a monkey saddle, is decomposed into simple ones, which are all added to the resulting set), each s_i can be classified as a minimum, maximum or a saddle. Denote these sets as $S_{\min}$, $S_{\max}$, S_{saddle}. Then the formula is:

$$P(q) = \sum_{s_i \in S_{\min}} b(r(q,s_i)) - \sum_{s_i \in S_{\text{saddle}}} b(r(q,s_i)) + \sum_{s_i \in S_{\max}} b(r(q,s_i)) \tag{2}$$

Equivalently, each critical point enters with the sign $(-1)^{\text{ind}(s_i)}$ of its Morse index; signed sums of this form are exactly those arising in Euler integration [Baryshnikov and Ghrist 2010] (Section 2). For a localized $b(r)$ vanishing for $r > \hat{d}$, where $\hat{d}$ is the *barrier localization parameter*, the potential defined above coincides with $b(d(q,f))$, where $d(q,f) = \min_s r(q,s)$ is the Euclidean distance from q to the surface, if the potential extent is less than the minimal curvature radius. For a convex object, there is a single minimum and maximum, so if $b(r)$ support is smaller than the object size, the resulting potential again coincides with $b(d(q,f))$ (in the case of convex objects, $d(q,f)$ is differentiable).

The following theorem holds for the potentials defined above, assuming the distance function critical points are isolated, and with the additional technical assumption that the points q for which the distance function critical points are not simple (i.e., the Hessian of the distance function is singular) are isolated and the Hessian of the distance function has rank deficiency 1. This assumption is not really necessary, but it makes the proof more transparent.

Theorem 3.1. *Let $f(s)$ be a smooth curve or surface, satisfying the assumptions above, and $P_f(q)$ be a potential defined by Equation* (1) *or* (2) *respectively. Assume that $b(r)$ is differentiable, monotonically decreasing, and $b(r) \to +\infty$ as $r \to 0$. $P_f(q)$ has the following properties:*

- *It is a C^1 function of q away from f;*
- *$P_f(q) \to \infty$ on f;*
- *The limit of its gradient at points of f is along the normal direction to the curve/surface.*
- *If $b(r)$ is positive, then $P_f(q) > 0$ for all q for curves (but not surfaces).*

The proof is provided in Section S1 of the supplemental document. An additional potential term required for positivity on surfaces is also discussed in Section S3 of the supplemental document.

3.3 Piecewise-linear curves and surfaces

Next, we show that the definitions above can be applied to piecewise-linear curves and surfaces, i.e., polylines and meshes.

Polylines. Let $p_0, \ldots p_N$ be points connected into a polyline M. For each segment (p_i, p_{i+1}) let $r_i^e(q)$ be the distance to segment i from q, and $r_i^v(q)$ the distance to p_i. For a given point q, consider the set of all local minima and maxima of $r(q,x)$ for x varying on M.

We assume that there is a strict order between distances to any two adjacent vertices. If the distances are exactly equal, the ties are broken using, e.g., vertex ordering based on indices (if values are equal, the value at the vertex with a larger index is considered larger, using the simulation of simplicity approach [Bremer et al. 2004]). Note that the choice of tie-breaking order does not affect the value of the potential and is just required to make the definition well-defined (by equivalence to Equation (5)).

A local minimum can be: (1) a closest point on an edge in its interior; (2) a vertex v for which the closest point to q on incident edges is v itself. A local maximum is always a vertex v for which two adjacent edges both have closest points to q different from v. The following formula is the analog of the smooth definition above:

$$P(q) = \sum_{i \in E_{\min}} b(r_i^e(q)) + \sum_{i \in V_{\min}} b(r_i^v(q)) - \sum_{i \in V_{\max}} b(r_i^v(q)) \tag{3}$$

where $E_{\min}$ is the set of edges for which the closest point to q is in the interior, $V_{\max}$ and $V_{\min}$ are the sets of local-maximum and local-minimum vertices, respectively (see Figure 3 left). Note that the sets of $E_{\min}$, $V_{\max}$ and $V_{\min}$ are small, so the summation is over just a few points. If $b(r)$ vanishes for $r > \hat{d}$, then in most cases, this coincides with evaluating it at the nearest point.

Meshes. Throughout, we consider linear triangle surface meshes embedded in $\mathbb{R}^3$. To define a similar potential for meshes, we need to classify different types of local extrema (Figure 4). For simplicity, we consider closed meshes only, but the definition extends easily to meshes with boundaries (Section S4 of the supplemental document).

- For *faces*, interior points can only be minima. $F_{\min}$ is the set of faces containing a local minimum.
- For *edges*, if an interior point x is closest to q on the edge, there are three possibilities for incident faces: x may be closest to q on both faces, or on neither, or on only one of them. The first two cases yield minima and saddles, respectively; in the third case, the edge is not extremal.
- For *vertices*, the situation is more complicated: a minimum, maximum, or a complex saddle is possible.

For a vertex x, let $e_0, \dots e_{k-1}$ be incident edges, enumerated counterclockwise, with two sequential edges bounding a face f_i. If all edges have x as the closest point, then it is a minimum; if none have it as the closest point, it is a maximum. Otherwise, it may be a saddle. We determine whether it is a saddle in the following way. Split the edges into sequential subsequences $e_{j_m}, e_{j_m+1} \dots e_{j_{m+1}}$, such that for all edges in the subsequence, the vertex is the closest point in the edge, or it is not. We refer to the former subsequences as +-subsequences. The faces between edges of the +-subsequence also have their closest point at x, but if at least one edge of a face is not in a +-subsequence, then it has the closest point elsewhere. Suppose there are k^+ distinct +-subsequences. If $k^+ = 1$, then the vertex is not extremal (the neighborhood is split into two parts, one above the vertex and one below); otherwise, it is a saddle, and $k^+ - 1$ is the saddle multiplicity (see Figure 3 right).

With this classification in hand, we define the potential as follows:

$$\begin{aligned} P(q) = &\left(\sum_{i \in F_{\min}} b(r_i^f(q)) + \sum_{i \in E_{\min}} b(r_i^e(q)) + \sum_{i \in V_{\min}} b(r_i^v(q)) \right) \\ &- \left(\sum_{i \in E_{\text{saddle}}} b(r_i^e(q)) + \sum_{i \in V_{\text{saddle}}} (k^+(v_i) - 1) b(r_i^v(q)) \right) \\ &+ \sum_{i \in V_{\max}} b(r_i^v(q)) \end{aligned} \tag{4}$$

We summarize the properties of these potentials in the following theorem:

Theorem 3.2. *Assume $b(r)$ is differentiable, monotonically decreasing, and $b(r) \to +\infty$ as $r \to 0$. The potential defined by Equation (3) or Equation (4) has the following properties:*

- *It is a C^1 function of q away from the polyline/mesh; for q fixed, $P(q)$ is a C^1 function of vertex positions p_i. It is moreover C^2 away from a set of measure zero, namely the loci where the closest point on an element transitions between its interior and its boundary.*
- *$P_f \to +\infty$ as q approaches the polyline/mesh.*
- *If q lies on a segment (p, p'), where p is a point on the polyline/mesh, consisting of points for which p is the closest point on the polyline/mesh, then as q approaches p, the direction of the gradient of the potential approaches the direction of $p' - p$.*

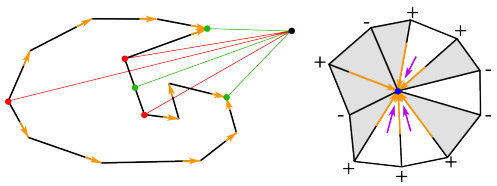


Fig. 3. Left: Cancellation in Equation 5, relating it to the signed sum over extrema. Minima are shown in green, maxima in red. The arrows on edges indicate the vertex at which the distance to the edge is minimal; the potential terms for these vertex-edge pairs cancel. Right: A monkey saddle at the blue vertex of a mesh results from $k^+ = 3$ alternating +-subsequences.

- *The potential can be expressed as the following sum over* all *edges and vertices for polylines:*

$$P(q) = \sum_{i=0}^{N-1} b(r_i^e(q)) - \sum_{i=1}^{N-1} b(r_i^v(q)) \tag{5}$$

where r_i^e and r_i^v are the distances to the edges and vertices, respectively. For meshes, the formula is:

$$P(q) = \sum_f b(r_f(q)) - \sum_e b(r_e(q)) + \sum_v b(r_v(q)) \tag{6}$$

The proof of the theorem is included in Section S2 of the supplemental document.

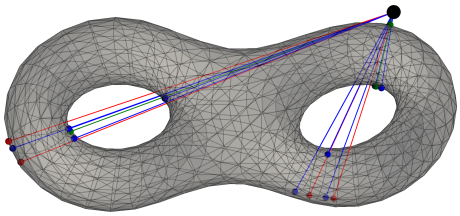


Fig. 4. Extrema of a distance function from a point (black) to a double torus mesh. Minima are shown in green, maxima in red, saddles in blue.

The expression (5) is the same expression introduced heuristically by Li et al. [2023], as an approximation to the maximum potential. Here, we demonstrate that the resulting potential is equal to the difference of the sums of maxima and minima above, so, in the special case of a single maximum (e.g., convex objects, or having a single maximum locally), it indeed coincides with the maximum. The alternating-sign formula provides a convenient mathematical tool, considerably simplifying the proof of the potential properties. It is also the form we evaluate in practice: (5) and (6) sum over all elements and require no classification of extrema, so the separate $E_{\min}$ and $V_{\min}$ terms of (3) and (4), which are convenient for stating the geometric meaning of the potential, never need to be formed explicitly.

We further note that the piecewise-linear potentials (3) and (4) converge to the smooth potentials (1) and (2), if increasingly fine polylines and meshes are sampled from a smooth curve or surface, as long as the normals of the polyline or mesh converge to the normals of the curve or surface. Figure 5 shows the convergence of level sets

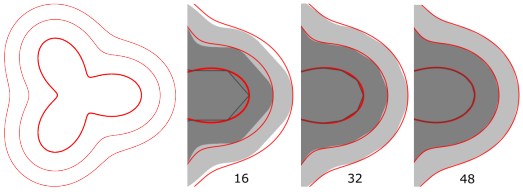


Fig. 5. Convergence of a level set of a potential with respect to a polyline approximation to a curve to the limit level set. A trefoil curve (red), shown with two level sets, is approximated with progressively finer polylines consisting of 16, 32, and 48 segments (dark gray).

in 2D; we leave the formal proof for future work, but include further convergence evidence in Section S7 of the supplemental material.

Like IPC [Li et al. 2020] and OGC [Chen et al. 2025], and unlike the pointwise version of GCP [Huang et al. 2025], the proposed potential admits an exact closed-form pointwise expression.

Algorithmic details. To evaluate the potential for polylines and meshes, we need to find extrema of the distance function. The criteria for face, edge, and vertex extrema above directly lead to an algorithm. A face contains an extremum of the distance function, only if q projects to its interior, which reduces to 3 halfspace tests. For an edge, we need to check if q projects to its interior (two halfspace tests), and whether for both incident faces the closest point is on the edge or not (two halfplane tests in triangle planes). Finally, for a vertex, we need to check for which incident edges the vertex is the closest point to q (one halfspace test for each edge). So algorithmically, extremum determination reduces to halfspace and halfplane tests, which can be implemented using robust predicates. If the barrier $b(r)$ has a finite extent, we only need to check elements within its radius $\hat{d}$ from the evaluation point q. Section S10 of the supplemental document gives the algorithms in pseudocode.

4 Repulsive energy

The pointwise potential is already useful for, e.g., smooth offset construction and computing point-surface interaction. To apply ESP to contact or intersection avoidance, we need to compute its integral over the polyline/mesh itself. However, since the evaluation point p now lies *on* f, the distance $r(p,\cdot)$ has a trivial minimum of value zero at p itself, which makes the potential infinite and carries no proximity information, as we only care about proximity of the p to *distinct* parts of the surface. We thus define $\mathring{P}_f(p)$ as the potential of Equations (3) and (4) with the trivial minimum at p excluded from the alternating sum. Dropping this term adds a caveat to the barrier property that we discuss in Section S5 of the supplemental material, showing how it does not prevent using our potential for simulation. We define the continuum version of the energy for a curve or a surface f as

$$\mathcal{I} = \int_f \mathring{P}_f(p)dp \tag{7}$$

where the integral is taken over the curve/surface. In the following we will implicitly overload P_f to indicate $\mathring{P}_f$ whenever the evaluation point lies on f.

We assume here that $b(r)$ diverges slowly enough as $r \to 0$ for $\mathcal{I}$ to be defined; Section S6 of the supplemental document quantifies this condition and the resulting trade-off with the barrier property.

Because the potential does not have a closed-form integral for most barrier functions $b(r)$, numerical quadrature is needed. At the same time, as discussed by Huang et al. [2025], we would like the discretization to retain the key properties of the continuum potential.

For a mesh M, $\mathring{P}_M(p)$ is realized equivalently by excluding all primitives containing p from the sum defining $P_f(p)$. The energy has the following properties directly obtained from the properties of $P_M(p)$: (1) it is differentiable with respect to q and with respect to mesh vertices; (2) it is repulsive close enough to contact, i.e., forces acting on the vertices point away from the surface.

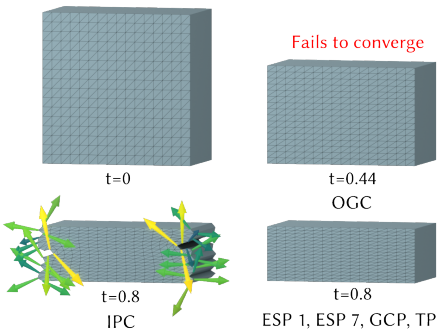


Fig. 6. Spurious tangential forces under compression: the IPC potential repels the compressed surface, causing crumpling. Like GCP, ESP (orders 1 and 7, indistinguishable here) and TP (the tangent-point energy, Section S8.1) remove this artifact.

In particular, as we show in Figure 6, similar to GCP, OGC and the tangent-point energy (TP) [Yu et al. 2021a], and in contrast to IPC, the energy minimizes artificial forces if points on a mesh are compressed together.

Quadrature design. The discretization should retain differentiability and the barrier property; additionally, we want to minimize discretization-related artifacts.

To design a quadrature rule for the integral (7) that preserves these properties, we start with a standard quadrature rule on triangular faces (resp. polyline edges in 2D), and write the energy as an elementwise sum

$$\mathcal{I} = \sum_i 2A_i^0 \int_{s\in T^{\text{ref}}} P_M(g_i(s))ds$$

where the summation is over all triangles of the mesh, T^{ref} is a standard reference triangle, g_i is the affine map from the reference triangle to the i-th triangle of M, and A_i^0 is the area of that triangle in the *rest* configuration (its length, in 2D). Integrating against the rest measure rather than the deformed one keeps the quadrature weights independent of the unknowns, so that no area derivatives enter the gradient and Hessian; the potential P_M is still evaluated on the deformed mesh. The energy we minimize is therefore the integral of P_M against the rest measure, which agrees with (7) at

the rest configuration. We always use rules that include the triangle vertices (vertex nodes are what make the barrier property hold for vertex-face contact), and have positive weights.

Using fixed quadrature points with respect to T^{ref} is efficient and allows for an arbitrary quadrature order, but it does not preserve the barrier property in 3D, as contact may happen at any interior edge point, which therefore needs to be included as a quadrature point. The barrier property, differentiability, and accurate integration are difficult to combine, as these additional quadrature points are not fixed with respect to the reference triangles, and their reference-triangle position depends non-differentiably on the mesh vertex positions (cf. [Huang et al. 2025]). To address this problem, we treat two types of quadrature points differently.

The first set consists of the quadrature points of a quadrature rule of arbitrary order, fixed with respect to the reference triangle, used to accurately approximate the integral in the flat regime; the second set includes the closest points between edges, and guarantees that the barrier property holds for all contact types.

Near-far potential split. Our goal is to avoid placing extra quadrature points on edges until they become necessary, that is, when the surface approaches edge-edge contact.

We assume a localized barrier function $b(r)$ vanishing smoothly for $r > \hat{d}$. For the purpose of using two separate quadratures, we pick a near-field cutoff $\bar{d} = \alpha\hat{d}$ with $0 < \alpha < 1$ and a C^2 smooth step function $G_{\bar{d}} : [0, \infty) \to [0, 1]$ satisfying $G_{\bar{d}}(d) = 1$ for $d \le \frac{1}{2}\bar{d}$ and $G_{\bar{d}}(d) = 0$ for $d \ge \bar{d}$.

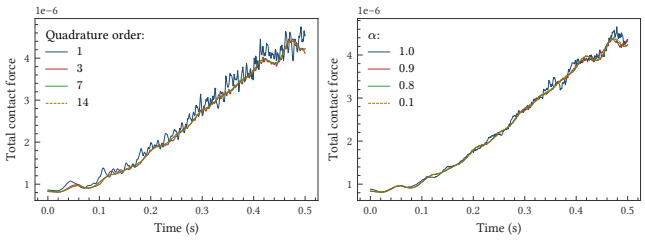


Fig. 7. Left: total force for the half-pipe example at different quadrature orders. Right: sensitivity to the near/far partition parameter $\alpha = \bar{d}/\hat{d}$ in the same setting; below $\alpha \approx 0.8$ the near quadrature has no influence.

The role of α is to keep the edge–edge quadrature out of the ordinary surface-contact regime. Suppose that in sustained surface contact the separation settles at some fraction of $\hat{d}$; choosing $\bar{d}$ below that separation leaves P_{near} inactive there, so sliding is integrated by the fixed high-order rule alone, and the contact-specific points activate only in genuine edge–edge approach. Any α meeting this condition works: smaller values localize the near term further, though a value that is too small can make the problem harder to solve. We use $\alpha = 0.05$ in our experiments as a relatively conservative choice, though the usable range depends on the scene and barrier stiffness: Figure 7 (right) shows that for the half-pipe example, one may use values as high as $\alpha \approx 0.8$.

Then we decompose the kernel and the potential:

$$b_{\text{near}}(d) = G_{\bar{d}}(d)\, b(d), \quad b_{\text{far}}(d) = \big(1 - G_{\bar{d}}(d)\big)\, b(d) \tag{8}$$

and $P_{\text{near}}(q)$ and $P_{\text{far}}(q)$ are defined by replacing b with b_{near} and b_{far} respectively in the pointwise potential formulas (5) and (6). As these are linear in b, $P_{\text{near}} + P_{\text{far}} = P$.

Near-term quadrature. For an edge e with endpoints v_0, v_1, let $N(e) = \{e' \in E : d(e, e') < \bar{d},\ e' \nparallel e\}$, where $d(\cdot,\cdot)$ is the edge-edge distance, and let $(q_{e,e'}, q'_{e,e'})$ be the closest pair of points on e and e'. For every $e' \in N(e)$ whose closest pair lies in the interior of both edges we place a quadrature point at $q_{e,e'}$, and write $Q_E(e)$ for the resulting set. The position of the closest pair is not C^0 across parallel configurations [Li et al. 2020], where the vertex quadrature points already ensure divergence, and not C^1 where it reaches an endpoint; we therefore scale the weight of these points by a factor vanishing smoothly at both. This is the mollifier of Huang et al. [2025]; more details are given in Section S9 of the supplemental document.

The complete set $Q(f)$ of quadrature points for a face f is the union of the set of fixed quadrature points $Q_F(f)$ and that of edge points $Q_E(e_i)$ for the three triangle edges e_i, $i = 1, 2, 3$.

We define unnormalized weights as follows. For a fixed quadrature point $q \in Q_F(f)$, $w(q) = \omega_q$ is the weight of the underlying rule, positive by assumption. For an edge quadrature point $q \in Q_E(e_i)$ arising from the pair (e, e'),

$$w(q) = \big(S(d(e, e'))\, \mu(e, e')\big)^k, \tag{9}$$

where S is a C^2 decreasing step with $S(0) = 1$ and $S(d) = 0$ for $d \ge \bar{d}$ (a scaled cubic B-spline, given in Section S9 of the supplemental), and $\mu(e, e') \in [0, 1]$ is the mollifier above. The exponent $k \ge 1$ makes wP_{near} vanish where an endpoint transition coincides with contact; it depends on the barrier and is given there too.

We compute the integral of the near term as

$$\mathcal{I}_{\text{near}} = \sum_f \frac{A_f^0}{\sum_{q \in Q(f)} w(q)} \sum_{q \in Q(f)} w(q) P_{\text{near}}(q),$$

that is, the rest-area-scaled weighted average of P_{near} over $Q(f)$. This sum depends differentiably on the position of the vertices of the mesh: the prefactor A_f^0 is a constant, the fixed points depend affinely on the vertices, and each edge point enters multiplied by a weight that vanishes to sufficient order wherever its position ceases to be smooth. It maintains the barrier property because whichever point realizes the closest distance to the mesh carries a positive weight. Finally, the complete discretized energy is the sum of $\mathcal{I}_{\text{near}}$ and the integral $\mathcal{I}_{\text{far}}$ of P_{far}, obtained from the same expression with $Q_F(f)$ in place of $Q(f)$, i.e., using only the points fixed with respect to the reference triangle.

While the pointwise mesh potential coincides with that of Li et al. [2023] (Section 3.3), we treat the discrete edge-edge contact differently to reduce sliding artifacts: they approximate edge-surface distances by edge-edge distances for efficiency, whereas we evaluate the full pointwise potential at the quadrature point; moreover, we split the barrier into near and far parts.

4.1 Repulsive energy properties

Huang et al. [2025] proposed a list of properties desirable for a contact/repulsive energy that both the discrete and the underlying continuum energies need to satisfy. Our potential satisfies the key

properties of finiteness, barrier and differentiability, and removes the resolution-dependent rest-state forces, which we discuss in turn below; Table S1 of the supplemental document compares them across methods. The supplementary also discusses an important "no free lunch" limitation for the limit continuum potential (Section S6).

Finiteness. This property requires the energy to be finite for all non-self-intersecting meshes. For the discrete energy, this is achieved automatically.

Barrier property. This property requires that the energy grows to infinity as the surface/mesh configuration approaches contact; in the discrete case, this means that either a vertex-face or an edge-edge distance for two non-intersecting edges decreases to zero. From the energy expression, as long as the quadrature points include all vertex-face and edge-edge closest points, this also follows immediately.

Differentiability. The energy is C^1 with respect to the surface control points. This follows from the differentiability of the pointwise potential and the expressions for the variable quadrature weights for the near quadrature.

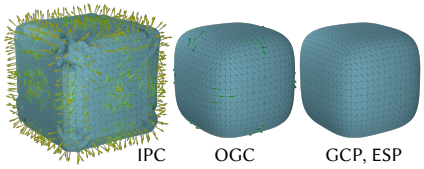


Fig. 8. Rest-state spurious forces: IPC repels without deformation once the barrier extent $\hat{d}$ exceeds the edge lengths; ESP (order 1) is free of this even at large $\hat{d}$. Removing rest forces where the rest state is genuinely near contact instead requires an adaptive $\hat{d}$, as in GCP. OGC has small forces from its edge-edge formulation.

No rest-state forces. Unlike IPC, the potential produces no rest-state repulsion when $\hat{d}$ exceeds the local edge length (Figure 8), so $\hat{d}$ need not be tied to the mesh resolution. Forces at rest remain only where distinct parts of the surface lie within $\hat{d}$ of each other in the undeformed configuration, being genuinely close to contact; removing those without globally reducing $\hat{d}$ requires making it variable and adjusting it to the rest state, in the same way as it is done by Huang et al. [2025].

5 Evaluation

Our tests are implemented on the IPC C++ infrastructure [Ferguson et al. 2020] and use Schneider et al. [2019] for physical modeling. All potentials are evaluated in the same Newton solver, with identical line-search, CCD [Li et al. 2020], and stopping criteria per experiment, isolating the choice of potential; per-experiment parameters are in Section S11 of the supplemental document.

Comparisons. We evaluate ESP relative to other potentials in two ways: (1) as a smooth distance-like function for offset construction, preserving the qualitative behavior of the Euclidean distance while creating smooth offset surfaces, and (2) as a contact potential that retains desirable properties of prior constructions while eliminating some artifacts.

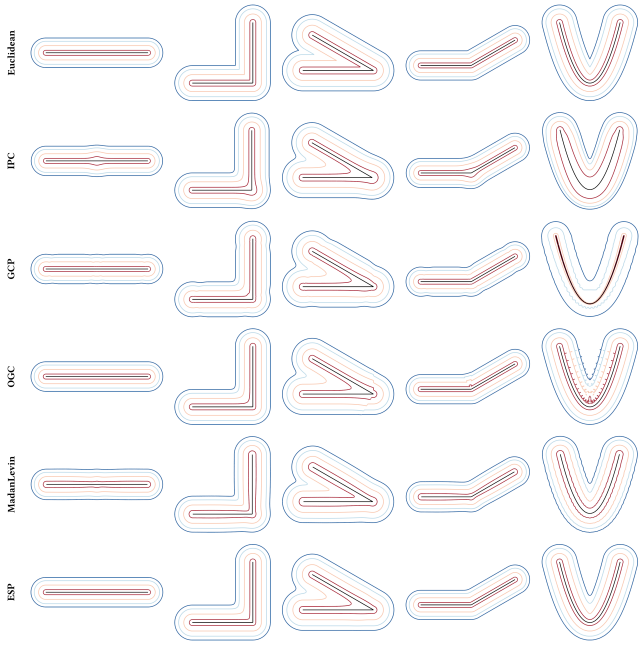


Fig. 9. 2D offsets from basic polyline configurations for different potentials, similar to Figure 4.2 of [Li et al. 2023]; all methods use the same barrier function. ESP level sets avoid bumps in flat areas (cf. GCP, Madan–Levin), bulges at corners (cf. IPC, GCP), and value and gradient discontinuities (cf. Euclidean distance, OGC).

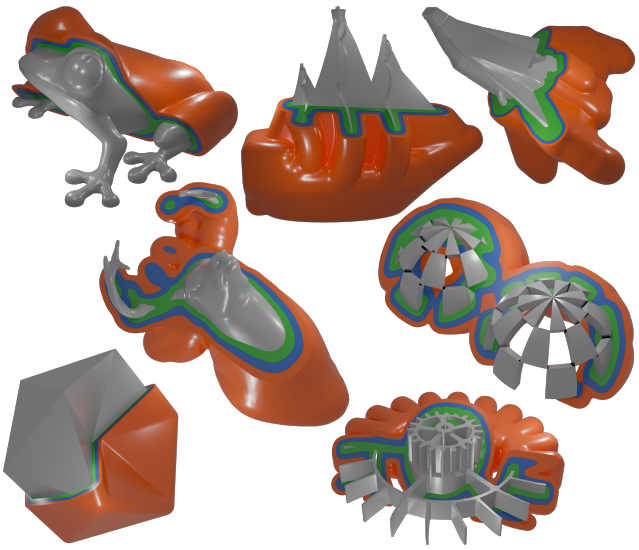


Fig. 10. Smooth offset surface examples obtained with the ESP potential.

Comparisons of offset behavior can be seen, for all variants, in Figures 1 and 9. Next, we discuss all distance function/proximity potentials one by one.

Compared to a barrier function applied to the *exact distance*, the resulting potential is differentiable everywhere, including at the medial axis, while otherwise behaving similarly.

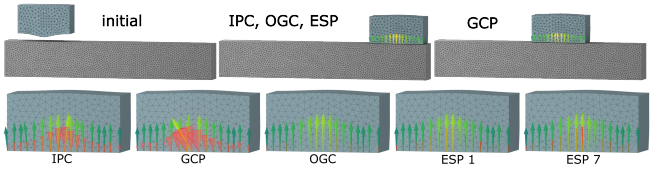


Fig. 11. Deformable sliding under gravity. Top: IPC, OGC, and ESP (orders 1 and 7) agree; spurious forces slow GCP down. Bottom: zoom at $t = 2.46$s; each red sector spans a vertex's force directions over all time steps, with radius scaled by the mean force magnitude.

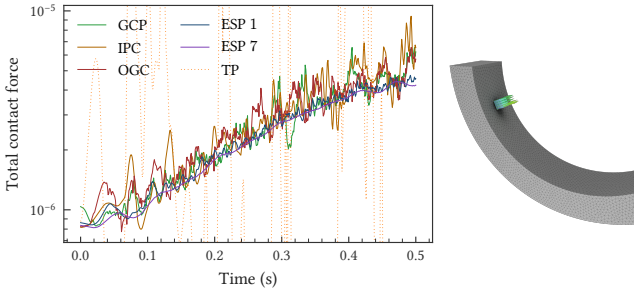


Fig. 12. Sliding on a half-pipe: all methods oscillate more in total force than ESP (shown at orders 1 and 7); high-order ESP is the most realistic. TP oscillates far more than the rest and fails to sustain sliding, tumbling down the pipe instead (Section S8.1).

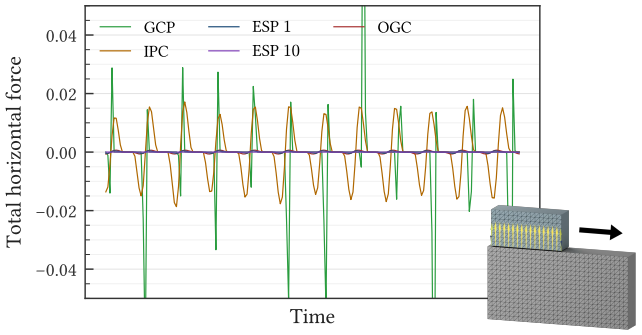


Fig. 13. A rigid object sliding on top of another: spurious horizontal forces appear for IPC and GCP. OGC and ESP at quadrature order 1 have small forces, negligible for ESP at order 10.

The *IPC potential* suffers from bulges where contributions from multiple elements overlap, notably in straight/planar regions (Figures 1 and 9); in contact, this translates into tangential force artifacts in sliding motion (Figures 11, 12, and 13). As discussed by Huang et al. [2022] and Chen et al. [2025], it also suffers from repulsion artifacts under tangential compression (Figure 6) and rest-state repulsion if the potential extent exceeds mesh edge lengths (Figure 8).

The *GCP potential*, while explicitly differentiable and free of tangential compression and rest-state artifacts, has an even more uneven potential surface from the interaction of multiple mollifiers, producing similar tangential force artifacts (Figures 11 to 13); see supplemental video.

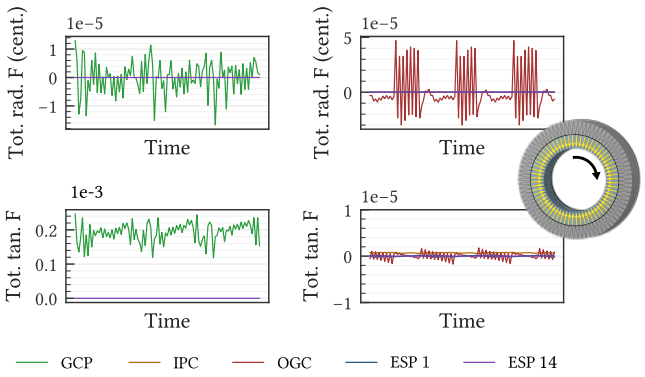


Fig. 14. Rotating ring: total radial and tangential forces. GCP is plotted separately, as its oscillations are much larger. ESP tangential forces are negligible at both orders 1 and 14. Radial forces are mean-centered to show oscillations; in the top-right plot IPC and ESP nearly overlap.

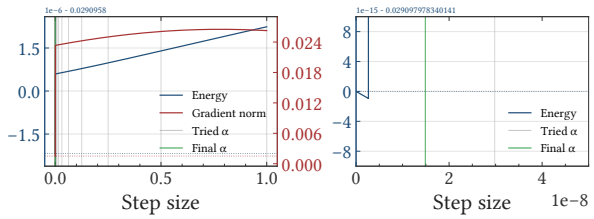


Fig. 15. Newton/line-search stalling for OGC, in the simulation of Figure 16 (ESP at order 1, run with the OGC two-stage barrier, does not exhibit this failure). Left: energy and gradient norm along the backtracking steps α, seeking an energy below the initial value (dotted line). Right: zooming near zero, such values do exist, but the energy discontinuity forces termination at the minimal step (10^{-8} here).

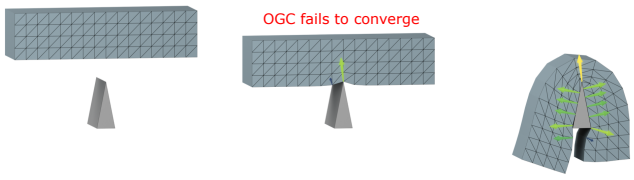


Fig. 16. Soft object falling onto a spike: OGC stalls at the center frame; all other methods, including ESP at order 1, complete the simulation with similar results.

The *OGC potential*, while identical to ours in convex areas, is discontinuous in concave areas, producing significant level set artifacts even for moderate concavity: Figure 14 shows force jumps across these discontinuities. The resulting instability can break numerical methods relying on energy/gradient continuity (Figures 15, 16, and 11).

Finally, Madan and Levin [2022] yields well-behaved surfaces, but these may still have artifacts in planar areas, depending on parameters (Figures 1 and 9).

Potential properties. Similar to GCP and tangent-point (TP) kernels [Yu et al. 2021a], ESP reduces repulsive artifacts under tangential compression and in the rest shape (Figures 6 and 8), while being substantially simpler. The high-genus correction makes it approximate the exact distance at large distances if desired (Figure S1 of the supplemental). For the total energy, the near-far partition is essential for capturing sliding without artifacts and naturally supports high-order quadrature for large contact areas (Figure 7).

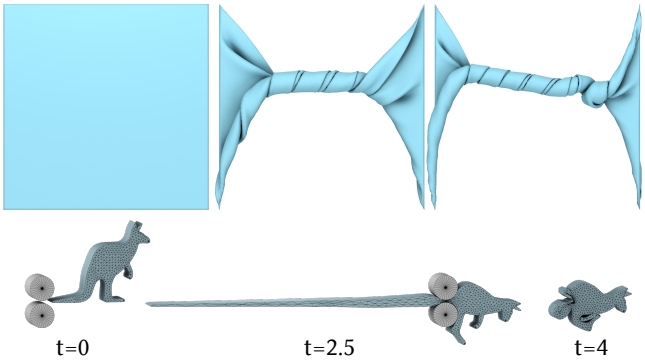


Fig. 17. Top: "mat twist" robustness test from [Li et al. 2020] with ESP at order 1. Bottom: the tail of the 2D kangaroo dragged through the small hole, also with ESP; both the order-1 and order-14 configurations of Table S3 complete the test, and are shown in the supplemental video.

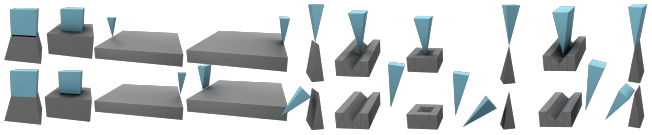


Fig. 18. Contact based on ESP passes all Erleben contact tests [Erleben 2018].

Finally, the near quadrature gives it the same level of robustness as IPC and related potentials (Figures 17 and 18).

6 Conclusion

We have demonstrated that a remarkably simple formulation, summing distance extrema with signs rather than considering only the closest point, yields a differentiable proximity potential with desirable properties: high-quality smooth offset construction and contact simulation, while retaining the robustness and guarantees of prior contact formulations. The pointwise potential for meshes has analytic expressions requiring few barrier function evaluations.

Limitations and future work. While local extrema can be identified efficiently and robustly, acceleration methods could be designed for barrier functions with large or infinite support. Our formulation reverts to low-order integration when contact spans few elements, and the same definition extends naturally to higher-order surfaces.

Availability. Our implementation and data are available at github.com/fsichetti/esp.

Acknowledgments

This work was supported by NSF grants OAC-2411349 and IIS-2313156, as well as the NYU IT High Performance Computing resources, services, and staff expertise.

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

M.D. Alaydin, D.J. Benson, and Y. Bazilevs. 2021. An updated Lagrangian framework for Isogeometric Kirchhoff–Love thin-shell analysis. *Computer Methods in Applied Mechanics and Engineering* 384 (2021), 113977.

Yuliy Baryshnikov and Robert Ghrist. 2009. Target Enumeration via Euler Characteristic Integrals. *SIAM J. Appl. Math.* 70, 3 (2009), 825–844. doi:10.1137/070687293

Yuliy Baryshnikov and Robert Ghrist. 2010. Euler Integration over Definable Functions. *Proceedings of the National Academy of Sciences of the United States of America* 107, 21 (2010), 9525–9530. doi:10.1073/pnas.0910927107

Alexander G. Belyaev and Pierre-Alain Fayolle. 2015. On Variational and PDE-Based Distance Function Approximations. *Computer Graphics Forum* 34, 8 (2015), 104–118. doi:10.1111/cgf.12611

Alexander G. Belyaev and Pierre-Alain Fayolle. 2020. An ADMM-based Scheme for Distance Function Approximation. *Numerical Algorithms* 84, 3 (2020), 983–996. doi:10.1007/s11075-019-00789-5

Peer-Timo Bremer, Herbert Edelsbrunner, Bernd Hamann, and Valerio Pascucci. 2004. A Topological Hierarchy for Functions on Triangulated Surfaces. *IEEE Transactions on Visualization and Computer Graphics* 10, 4 (2004), 385–396. doi:10.1109/TVCG.2004.3

Anka He Chen, Jerry Hsu, Ziheng Liu, Miles Macklin, Yin Yang, and Cem Yuksel. 2025. Offset Geometric Contact. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–21.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2013. Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow. *ACM Transactions on Graphics* 32, 5 (2013), 152:1–152:11. doi:10.1145/2516971.2516977

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2017. The Heat Method for Distance Computation. *Commun. ACM* 60, 11 (2017), 90–99. doi:10.1145/3131280

Lorin Crawford, Anthea Monod, Andrew X. Chen, Sayan Mukherjee, and Raúl Rabadán. 2020. Predicting Clinical Outcomes in Glioblastoma: An Application of Topological and Functional Data Analysis. *J. Amer. Statist. Assoc.* 115, 531 (2020), 1139–1150. doi:10.1080/01621459.2019.1671198

Yinwei Du, Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2024. Robust and Artefact-Free Deformable Contact with Smooth Surface Representations. *Computer Graphics Forum* (2024). doi:10.1111/cgf.15187

Michal Edelstein, Nestor Guillen, Justin Solomon, and Mirela Ben-Chen. 2023. A Convex Optimization Framework for Regularized Geodesic Distances. In *ACM SIGGRAPH 2023 Conference Proceedings*. ACM, 1–11. doi:10.1145/3588432.3591523

Kenny Erleben. 2018. Methodology for Assessing Mesh-Based Contact Point Methods. *ACM Transactions on Graphics* 37, 3, Article 39 (July 2018), 30 pages.

Pierre-Alain Fayolle and Alexander G. Belyaev. 2018. p-Laplace Diffusion for Distance Function Estimation, Optimal Transport Approximation, and Image Enhancement. *Computer Aided Geometric Design* 67 (2018), 1–20. doi:10.1016/j.cagd.2018.09.003

Zachary Ferguson et al. 2020. IPC Toolkit. https://ipc-sim.github.io/ipc-toolkit/. https://ipc-sim.github.io/ipc-toolkit/

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 (Proceedings of SIGGRAPH)* 40, 4, Article 183 (July 2021), 16 pages.

Robert Ghrist, Rachel Levanger, and Huy Mai. 2018. Persistent Homology and Euler Integral Transforms. *Journal of Applied and Computational Topology* 2, 1–2 (2018), 55–60. doi:10.1007/s41468-018-0017-1

Kemeng Huang, Floyd M. Chitalu, Huancheng Lin, and Taku Komura. 2024. GIPC: Fast and stable Gauss-Newton optimization of IPC barrier energy. *ACM Trans. Graph.* (jan 2024). doi:10.1145/3643028 Just Accepted.

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

Zizhou Huang, Davi Colli Tozoni, Arvi Gjoka, Zachary Ferguson, Teseo Schneider, Daniele Panozzo, and Denis Zorin. 2022. Differentiable solver for time-dependent deformation problems with contact. arXiv:2205.13643 [cs.GR]

David Kamensky, Fei Xu, Chung-Hao Lee, Jinhui Yan, Yuri Bazilevs, and Ming-Chen Hsu. 2018. A contact formulation based on a volumetric potential: Application to isogeometric simulations of atrioventricular valves. *Computer Methods in Applied Mechanics and Engineering* 330 (2018), 522–546.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Transactions on Graphics* 39, 4, Article 49 (July 2020), 20 pages.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2023. Convergent Incremental Potential Contact. arXiv:2307.15908 [math.NA]

Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021. Codimensional incremental potential contact. *ACM Trans. Graph.* 40, 4, Article 170 (jul 2021), 24 pages. doi:10.1145/3450626.3459767

Yaron Lipman, Raif M. Rustamov, and Thomas A. Funkhouser. 2010. Biharmonic Distance. *ACM Transactions on Graphics* 29, 3 (2010), 27:1–27:11. doi:10.1145/1805964.1805971

Abhishek Madan and David IW Levin. 2022. Fast evaluation of smooth distance constraints on co-dimensional geometry. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–17.

Josua Sassen, Henrik Schumacher, Martin Rumpf, and Keenan Crane. 2024. Repulsive Shells. *ACM Trans. Graph.* 43, 4 (2024).

Roger A Sauer and Laura De Lorenzis. 2013. A computational contact formulation based on surface potentials. *Computer Methods in Applied Mechanics and Engineering* 253 (2013), 369–395.

Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Transactions on Graphics* 39, 4 (2020), 123:1–123:18. doi:10.1145/3386569.3392374

Rohan Sawhney, Bailey Miller, Ioannis Gkioulekas, and Keenan Crane. 2023. Walk on Stars: A Grid-Free Monte Carlo Method for PDEs with Neumann Boundary Conditions. *ACM Transactions on Graphics* 42, 4 (2023), 80:1–80:20. doi:10.1145/3592398

Pierre Schapira. 1995. Tomography of Constructible Functions. In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, 11th International Symposium, AAECC-11 (Lecture Notes in Computer Science, Vol. 948)*. Springer, 427–435. doi:10.1007/3-540-60114-7_35

Teseo Schneider, Jérémie Dumas, Xifeng Gao, Denis Zorin, and Daniele Panozzo. 2019. PolyFEM. https://polyfem.github.io/.

James A. Sethian. 1996. A Fast Marching Level Set Method for Monotonically Advancing Fronts. *Proceedings of the National Academy of Sciences* 93, 4 (1996), 1591–1595. doi:10.1073/pnas.93.4.1591

Xing Shen, Runyuan Cai, Mengxiao Bi, and Tangjie Lv. 2024. Preconditioned Nonlinear Conjugate Gradient Method for Real-time Interior-point Hyperelasticity. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) *(SIGGRAPH '24)*. Association for Computing Machinery, New York, NY, USA, Article 96, 11 pages. doi:10.1145/3641519.3657490

Paweł Strzelecki and Heiko von der Mosel. 2013. Tangent-Point Repulsive Potentials for a Class of Non-smooth m-dimensional Sets in $\mathbb{R}^n$. Part I: Smoothing and Self-avoidance Effects. *Journal of Geometric Analysis* 23, 3 (2013), 1085–1139.

John N. Tsitsiklis. 1995. Efficient Algorithms for Globally Optimal Trajectories. *IEEE Trans. Automat. Control* 40, 9 (1995), 1528–1538. doi:10.1109/9.412624

Rodolphe Vaillant, Loïc Barthe, Gaël Guennebaud, Marie-Paule Cani, Damien Rohmer, Brian Wyvill, Olivier Gourmel, and Mathias Paulin. 2013. Implicit skinning: real-time skin deformation with contact modeling. 32, 4, Article 125 (July 2013), 12 pages.

Oleg Ya. Viro. 1988. Some Integral Calculus Based on Euler Characteristic. In *Topology and Geometry—Rohlin Seminar*. Lecture Notes in Mathematics, Vol. 1346. Springer, Berlin, 127–138. doi:10.1007/BFb0082775

Chris Yu, Caleb Brakensiek, Henrik Schumacher, and Keenan Crane. 2021a. Repulsive Surfaces. *ACM Trans. Graph.* 40, 6 (2021).

Chris Yu, Henrik Schumacher, and Keenan Crane. 2021b. Repulsive Curves. *ACM Trans. Graph.* 40, 2 (2021).

Daniel Zint, Zhouyuan Chen, Yifei Zhu, Denis Zorin, Teseo Schneider, and Daniele Panozzo. 2025. Topological Offsets. *ACM Transactions on Graphics* (July 2025). doi:10.1145/3731157