

Robust Containment Queries over Collections of Rational Parametric Curves via Generalized Winding Numbers

JACOB SPAINHOUR, University of Colorado Boulder, USA
 DAVID GUNDERMAN, Indiana University School of Medicine, USA
 KENNETH WEISS, Lawrence Livermore National Laboratory, USA

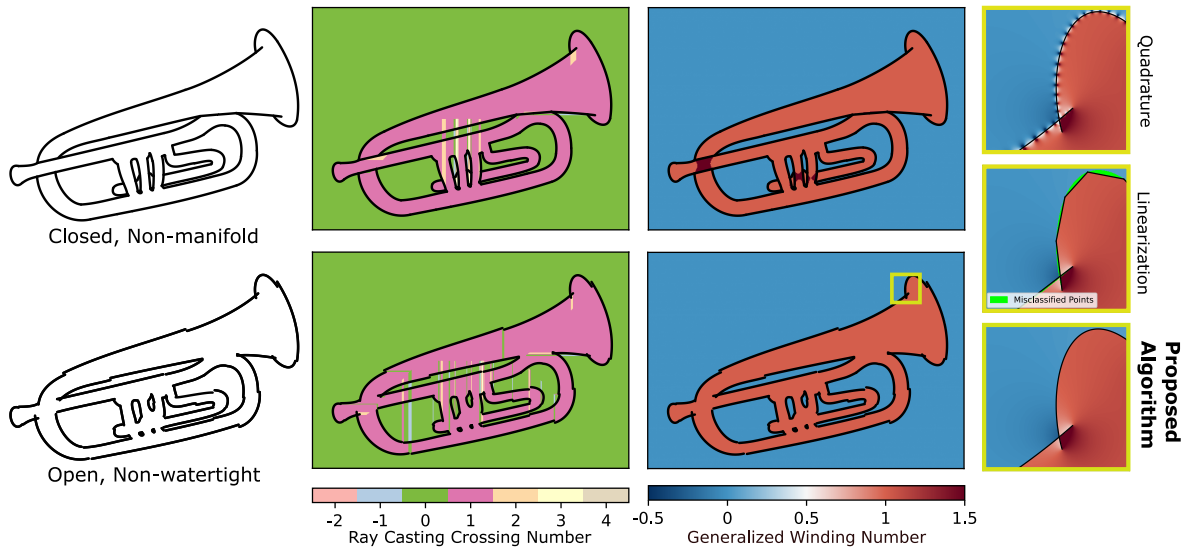


Fig. 1. The point containment query is an important predicate in many applications, yet their evaluation is limited by undesirable, but often unavoidable features of the bounding geometry such as non-manifold or non-watertight edges (left). This causes catastrophic errors for traditional approaches to deciding containment, in particular those that use ray casting (middle left). Because the field of generalized winding numbers degrades smoothly around geometric errors, they are well-suited for use in robust containment queries (middle right). In this paper, we present a framework for evaluating exact generalized winding numbers for arbitrary collections of curved objects, as well as a novel point-in-curved-polygon algorithm that facilitates their efficient calculation (right).

Point containment queries for regions bound by watertight geometric surfaces, i.e., closed and without self-intersections, can be evaluated straightforwardly with a number of well-studied algorithms. When this assumption on domain geometry is not met, such methods are either unusable, or prone to misclassifications that can lead to cascading errors in downstream applications. More robust point classification schemes based on generalized winding numbers have been proposed, as they are indifferent to these imperfections. However, existing algorithms are limited to point clouds and collections of linear elements. We extend this methodology to encompass more general curved shapes with an algorithm that evaluates the winding number scalar field over unstructured collections of rational parametric curves. In particular, we evaluate the winding number for each curve independently, making the derived containment query robust to how the curves are arranged. We ensure geometric fidelity in our queries by treating each curve as equivalent to an adaptively constructed polyline that provably has the same generalized winding number at the point of interest. Our algorithm is numerically stable for points that are arbitrarily close to the model, and explicitly treats points that are coincident with curves. We demonstrate the improvements

in computational performance granted by this method over conventional techniques as well as the robustness induced by its application.

CCS Concepts: • **Computing methodologies** → **Shape analysis; Parametric curve and surface models.**

Additional Key Words and Phrases: Winding number, point containment query, rational Bézier curves, robust geometry processing

1 INTRODUCTION

In the fields of computer graphics, Computer Aided Geometric Design (CAGD) and Computer Aided Engineering (CAE), it is common for geometric objects to be expressed using a boundary representation (B-Rep). Often, these B-Reps are defined via non-uniform rational B-spline (NURBS) curves in 2D and surfaces in 3D. Such objects are invaluable in the design of CAD models, as they allow for precise geometric control of the object’s boundary [Piegl and Tiller 1996]. However, it can still be desirable in many contexts, such as in animation or simulation, to treat the interior volume of these objects explicitly [McAdams et al. 2011].

Authors’ addresses: Jacob Spainhour, University of Colorado Boulder, USA, jacob.spainhour@colorado.edu; David Gunderman, Indiana University School of Medicine, USA, dgunderm@purdue.edu; Kenneth Weiss, Lawrence Livermore National Laboratory, USA, kweiss@llnl.gov.

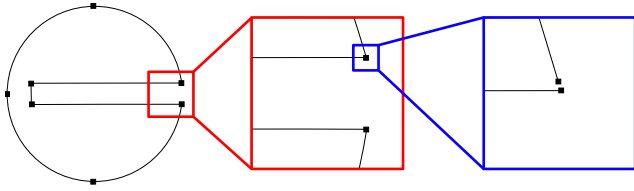


Fig. 2. Geometric errors can be visually imperceptible, but can still cause a shape to have no topological interior.

Such a task motivates development of the *containment query*, a geometric predicate that returns whether an arbitrary spatial location is contained within an arbitrary shape. The rapid evaluation of containment queries is quite well understood for many classes of shapes, such as polygons and regions bounded by Bézier curves [Carvalho and Cavalcanti 1995; Nishita et al. 1990; Patrikalakis and Maekawa 2010]. However, such methods assume the boundary is *watertight*, having no gaps between connected components. In reality, the vast majority of B-Reps are created by hand within some form of CAD software, leading to so-called “messy” CAD geometry.

A straightforward source of these problems is human error during design or construction, but more subtle issues can be introduced through varying tolerances within and between different software tools. This results in boundary models with human-imperceptible, but numerically significant gaps and overlaps between individual components of the boundary [Marussig and Hughes 2017] (see Figure 2). Even in situations prioritizing visibility over strict geometric fidelity, such as in computer graphics and animation, care must be taken to ensure that these errors are properly accounted for. For example, the *dynamic planar map* heuristic used by Adobe Illustrator’s Live Paint Bucket tool allows users to set fills for closed regions bounded by intersecting collections of Bézier curves and exposes a global threshold parameter to connect gaps between curves in the collection [Asente et al. 2007].

As an additional class of motivating examples, we consider applications in multiphysics simulation, for which containment classification errors can lead to cascading problems and incorrect results. The mathematical underpinnings of containment within non-watertight shapes is dubious, which leads to a tenuous theoretical foundation for established methods that assume geometric continuity. This is problematic, as errors are frequently inconsistent within local regions, with even arbitrarily close points receiving different classifications. More practically, these errors can result in unexpected and fatal errors in simulation software, wasting computational resources and developer time [Sederberg et al. 2008].

At the same time, physical applications often prevent approximating curved geometry with piecewise linear segments. Although this would allow for treatment of the messy—but now linear—geometry with established methods, such an approximation can produce sizable errors in the relevant numerical method [Sevilla et al. 2008]. As a concrete example, multimaterial Arbitrary Lagrangian-Eulerian (ALE) methods require additional preprocessing to initialize the volume fraction of each material in a computational cell [Barlow et al. 2016; Hirt et al. 1974]. This, in turn, necessitates classifying each point in a quadrature rule along the background grid as inside

or outside the volume bounded by the corresponding shape [Weiss et al. 2016]. It is critical that this identification takes into account the full curvature of the bounding geometry, as piecewise linear approximations can cause quadrature points near material boundaries to be misclassified, unpredictably compounding the error in the calculation.

Such geometric errors can be resolved with techniques in surface repair [Bischoff et al. 2005; Mezentsev and Woehler 1999; Park et al. 2021], which alter the B-Rep itself to an approximated surface that lacks these imperfections. However, doing so by hand can be impractical for the typical number of imperfections in such a model, and automated procedures during geometry pre-processing can inadvertently remove important model features. Instead, we require a containment query that is *robust* to such artifacts, in the sense that an appropriate inside/outside classification is always returned, independent of size and frequency of these boundary errors.

To address these concerns, we present in this work three principal contributions:

- We extend the theoretical framework of generalized winding numbers [Jacobson et al. 2013] to the context of an unstructured collection of 2D rational parametric curves, from which one can derive a robust containment query that indicates whether an arbitrary point should be considered interior to non-watertight geometry.
- We address the issue of defining and evaluating generalized winding numbers for points that are coincident with the boundary. This discussion lies outside the scope of much of the existing literature, as such points require unique consideration in the context of curved geometry.
- We provide a novel algorithm for evaluating integer winding numbers with respect to curved geometry, which is the principal component of our technique for evaluating generalized winding numbers on such shapes. We demonstrate that this calculation is robust to non-watertight geometry and considerably more performant than state-of-the-art techniques.

Altogether, these contributions allow for robust containment queries of objects implied by messy 2D CAD geometry, composed of collections of parametric curves (see Figure 1). Our algorithm has been implemented in Axom, a BSD-licensed open-source library of Computer Science infrastructure for HPC applications [Capps et al. 2024].

2 BACKGROUND AND RELATED WORK

Even in the ideal and watertight setting, methods of evaluating containment vary wildly depending on the numerical representation of the bounding geometry. As simple examples, the operation can be performed trivially for axis-aligned quadrilaterals through the evaluation of inequalities, for spheres by computing a distance to the center, and for triangles using Barycentric coordinates. More complicated procedures are necessary for point containment in arbitrary polygons [Haines 1994]. These algorithms can generally be broken down into two categories (see Figure 3).

Ray Casting algorithms determine containment by extending a ray from the specified query point out to infinity, and counting the

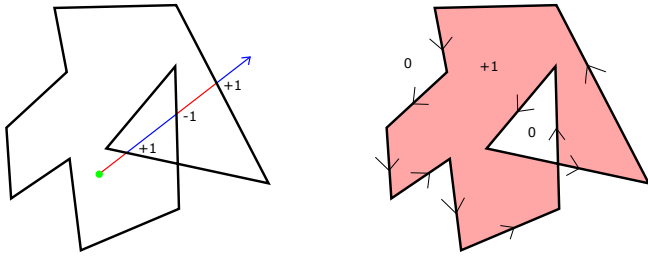


Fig. 3. Given an arbitrary query point, containment can be determined with a ray casting algorithm that counts (signed) *intersections* between the polygon and a ray extending from the query (left), or a winding number algorithm that counts *revolutions* of the curve around the query (right).

number of times the ray intersects the domain boundary to produce a *crossing number* [Shimrat 1962]. As a point moves along this ray, it will alternate between interior and exterior, and therefore an odd (even) crossing number indicates the point is interior (exterior). Computing intersections between the given ray and the linear edges of a polygon can be done very efficiently, with many graphics processing units (GPUs) having dedicated ray-tracing cores that can perform the operation massively in parallel [Zellmann et al. 2022]. However, these methods are sensitive to the specific ray that is extended, needing to account for special cases when the ray intersects a polygon vertex or entire edge, which can be unnecessarily more costly [Edelsbrunner and Mücke 1990; Sommariva and Vianello 2022]. *Winding number* algorithms are a class of algorithms that are more indifferent to these issues, and, as such, are the type from which we derive our proposed method. In a winding number algorithm, we instead count the number of revolutions around the query point made by a particle traveling on the domain boundary [Hormann and Agathos 2001]. This value can be considered a generalization of the crossing number, and can be used to determine containment with the same even-odd rule. By virtue of not being dependent on any particular extended ray, these algorithms implicitly handle edge-cases that are problematic for ray casting. This means even in the context of simple, watertight geometry, such algorithms can be considered a more robust approach to containment queries.

Containment of arbitrary points in more general curved shapes is a much more difficult problem, and has been studied extensively in the context of 2D vector graphics. As an example, we refer to the process of *linearization*, in which a curve is approximated by connected linear segments, an approach that is fairly common in both modern graphics engines and physics codes. In the former, information about the targeted display can be leveraged to ensure that the approximating polyline is constructed with sufficiently high resolution such that it is indistinguishable to a viewer [Kilgard 2020]. Combined with the fact that the polyline must be rendered with some predetermined thickness, it is therefore theoretically possible (although rarely enforced) for containment queries on linearized watertight geometries on a rendered display to have no *observable* misclassifications.

On the other hand, this type of polyline approximation of curved geometry is known to cause severe and often unexpected errors in the modeling of physical systems through PDEs and finite element

analysis [Sevilla et al. 2008]. This is a particular problem in the case where points of interest are allowed to be placed arbitrarily close to the curve, as then no guarantee can be made as to how many refinements in the approximation are necessary to achieve perfect geometric fidelity. In place of such a linearization, the use of NURBS to represent boundaries in a finite element analysis has been quite successful in reducing or even eliminating geometric error [Sevilla et al. 2008, 2011]. As a result of this, we consider operating directly on curves without any approximation to be a necessary component of our proposed algorithm so as to preserve the accuracy of an underlying numerical method.

We similarly desire an algorithm that can reach this level of geometric fidelity for a very general class of curves. For example, cubic Bézier curves are the highest order required by postscript and SVG formats, as the level of detail they afford is typically sufficient for design applications. However, applications in computational multiphysics can necessitate an algorithm that works for *rational* Bézier curves of *arbitrary* order. This is because the NURBS shapes that are used by more sophisticated modeling software can always be decomposed into a collection of (in principle arbitrary order) rational Bézier primitives through the process of Bézier extraction [Piegl and Tiller 1996; Thomas et al. 2015]. However as we will see, the proposed algorithm is ultimately indifferent to the order and/or rationality of the curve in question.

Containment within regions defined by Bézier curves has also been studied in the context of 2D parametric trimming curves for 3D NURBS surface patches. In this sub-problem, it must be determined if a point in the 2D parameter space of the 3D surface is contained within the (often, but not necessarily watertight) trimming curves. Because this and other contemporary 2D containment problems are frequently solved with ray casting algorithms [Marussig and Hughes 2017], we note that they can largely be divided into two categories. In the algebraic approach, intersections between rays and arbitrary curves can be computed using typical root-finding techniques in parameter space [Patrikalakis and Maekawa 2010], although standard bisection and gradient descent methods can only identify at most a single point of intersection. In contrast, geometric approaches take place in physical space, decomposing the curve into subcurves until intersections with the ray can be identified [Farin 2001]. The current state-of-the-art method, *Bézier clipping* [Nishita et al. 1990; Sederberg and Nishita 1990], lies at the intersection of these two approaches, where recursive algebraic subdivision is accelerated using the convex hull property of a Bézier curve. However, such methods are known to possibly report incorrect intersections and suffer from inefficiency in the presence of multiple spatially equivalent intersection points [Efremov et al. 2005].

These issues are worsened by geometric errors present in the shape. As an example, if the arbitrary ray is by chance extended through a very small gap in the model, then the point must be classified as “exterior” despite this being unintuitive and contrary to the intentions of the mesh designer. Non-manifold edges in a shape are similarly problematic for ray casting methods. These issues can be somewhat mitigated by extending several rays and taking a consensus, but such approaches impose an additional computational burden while still producing noisy and potentially inconsistent classifications for nearby points [Nooruddin and Turk 2003]. We

now show that, in addition to being indifferent to the specific edge cases introduced by ray casting, winding number algorithms can be extended such that they are also robust to more general issues introduced by non-watertight, non-manifold, curved geometry.

3 GENERALIZED WINDING NUMBERS

The generalized winding number is an extension of the standard integer-valued winding number to (potentially) non-watertight regions. While integer winding numbers partition the enclosed regions of a domain, the generalized winding number generates a harmonic scalar field that smoothly degrades in the presence of discontinuities and self-intersections. This allows for the calculation of robust containment queries in the presence of messy geometry. As an example, a simple rule for handling the fractional values that result from imperfections in the bounding geometry is to round them to their nearest integer.

Such generalizations of winding numbers have recently shown great utility in geometry processing applications. In the context of both (linear) triangle meshes [Jacobson et al. 2013] and point clouds [Barill et al. 2018], the resulting inside/outside classification can be used to generate volumetric triangle/tetrahedral meshes out of completely unstructured geometric data. Additionally, the smooth degradation of the winding number field naturally introduces desirable locality properties, which in turn lead to more robust and performant algorithms in, for example, exact mesh booleans [Tretner et al. 2022]. More recently, these particular winding number algorithms have been used in the context of immersogeometric analysis to simulate fluid flow around geometric objects defined by unstructured point clouds [Balu et al. 2023]. It is speculated in Gunderman [2021] that one could extend the calculation of generalized winding numbers by utilizing their own novel quadrature schemes for such objects [Gunderman et al. 2020], but preliminary tests of this approach are limited by the accuracy of numerical integration. Despite this recent surge in use, methods that compute *exact* generalized winding numbers for more general curved geometric objects are, to our knowledge, currently absent from the literature.

Nevertheless, the problem of computing generalized winding numbers for curved shapes has close analogues in the fields of both computer graphics and boundary integral equations. In the former, a connection can be drawn to diffusion curves, a primitive in vector graphics that also produces a scalar field defined by the solution of a differential equation with respect to boundary curves (albeit with different data prescribed on the curve) [Jacobson et al. 2013]. These applications have the advantage of being solved on a grid generated by a rasterized domain, meaning that a global solution can be obtained efficiently using a geometric multi-grid method [Orzan et al. 2008]. Sawhney and Crane [2020] recently introduced a grid-free approach in which random walks and projections onto curved geometry are used to solve certain classes of PDEs, including that which implicitly defines the generalized winding number, focusing on improving speed and locality at the cost of some degree of accuracy and consistency. Within the context of boundary integral equations, calculating winding numbers is a special case of the more general problem of integrating double layer potentials

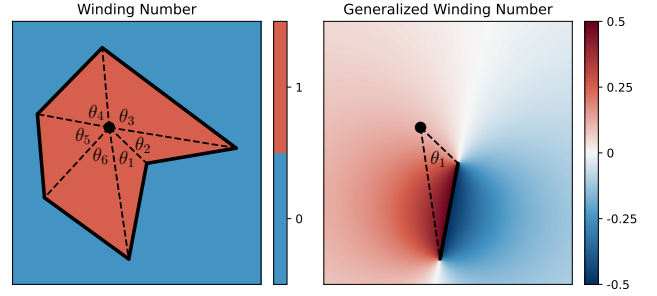


Fig. 4. (left) The winding number for a piecewise linear shape can be computed by summing angles subtended by each segment. (right) Winding numbers for linear segments can be computed at arbitrary points independently of the remaining shape.

with a constant density function. However, such algorithms are typically applied to closed domains with simpler (although still curved) types of geometry, precluding general use on arbitrary, messy CAD models [Klinteberg and Barnett 2020]. Furthermore, our proposed algorithm relies exclusively on geometric principles, eliminating the need to use costly and potentially inaccurate quadrature schemes.

We begin by reviewing the formal description of a generalized winding number in \mathbb{R}^2 . Given an oriented curve $\Gamma \subset \mathbb{R}^2$ and query point $q \in \mathbb{R}^2 \setminus \Gamma$, the winding number w describes the (potentially incomplete) number of times the curve travels counterclockwise around the point. In the case where Γ is a closed curve, the scalar field $w_\Gamma(q)$ is integer-valued, and induces a partition of \mathbb{R}^2 that can be interpreted as containment in the region bounded by Γ . In the following discussion, it is notationally convenient to assume that the query point is located at the origin, and that the curve Γ has been appropriately translated by q . The scalar field is invariant to such global translations, and so the winding number of Γ at q , $w_\Gamma(q)$, is identical to the winding number at the origin after translating the curve, $w_{\Gamma-q}(0)$.

When the curve is piecewise linear with $\Gamma = \cup_i L_i$, the winding number at q can be computed as the sum of *signed* angles θ_i subtended by each linear component L_i , given by

$$w_{\Gamma-q}(0) := \frac{1}{2\pi} \sum_i \theta_i. \quad (1)$$

Furthermore, the winding number of each individual component of the piecewise linear shape (and of straight lines in general) is given by

$$w_{L_i-q}(0) = \frac{1}{2\pi} \theta_i, \quad (2)$$

again where θ_i is the angle subtended by L_i at q (see Figure 4).

We can consider the case for a more general collection of arbitrary curves analogously. Let $\Gamma = \cup_i \Gamma_i$ be a collection of rational Bézier curves that constitute the (likely not watertight) boundary of a 2D CAD model. In this case, the winding number of a point q with

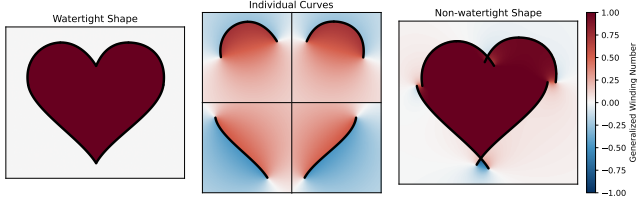


Fig. 5. The winding number with respect to closed shapes (left) is always an integer. However, by independently summing contributions from each curved component (center), the generalized winding number can be computed over collections of curves containing gaps and overlaps (right).

respect to a single curve Γ_i is given by

$$w_{\Gamma-q}(0) := \frac{1}{2\pi} \int_{\Gamma-q} d\theta = \frac{1}{2\pi} \sum_i \int_{\Gamma_i-q} d\theta \quad (3)$$

by the properties of the integration. This is notable, as it describes how the winding number for each individual component Γ_i can be computed completely independently of every other curve, with

$$w_{\Gamma_i-q}(0) := \frac{1}{2\pi} \int_{\Gamma_i-q} d\theta. \quad (4)$$

It is from this property of the generalized winding number that we derive both the robustness of our winding number calculation, and the robustness of the derived containment query. While the *collection* of Bézier curves may be messy in the sense that pairs of endpoints may not perfectly align, generalized winding numbers on each *individual* curve are well-defined and can be computed stably. At the same time, the sum of these values at a given point is exactly the generalized winding number of the entire shape, which can be used to produce intuitive and robust containment queries for non-watertight shapes (see Figure 5). Most commonly, this is done by rounding this value to the nearest integer, and applying to it a conventional even-odd or non-zero rule to make the classification. Importantly, each of these conventions treat points with a zero winding number as exterior.

Framed in this way, the remaining task is to evaluate generalized winding numbers for a single, arbitrary rational Bézier curve. As an initial approach, one could consider direct evaluation of the integral through numerical quadrature. When the curve is given parametrically as $(x(t), y(t)) \in \mathbb{R}^2$, $t \in [0, 1]$, as is the case for (rational) Bézier curves, we can evaluate this formula more practically in terms of Cartesian coordinates, as

$$w_{\Gamma}(q) := \frac{1}{2\pi} \int_0^1 \frac{x(t)y'(t) - x'(t)y(t)}{x^2(t) + y^2(t)} dt. \quad (5)$$

Direct evaluation of this integral through standard techniques is notoriously difficult. In particular, when Γ is close to the query point (or rather the origin under this translated notation) then the near-singular behavior of the integrand causes Gaussian quadrature to become unstable, as seen in Figure 6. Although some error is to be expected for this type of numerical integration, this error rapidly becomes unacceptable as the query point approaches the curve. Perhaps more concerningly, there is no immediate way to verify

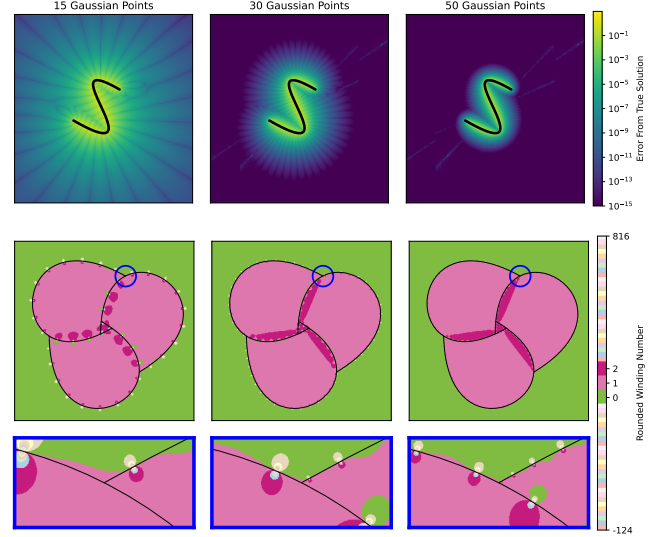


Fig. 6. (top) The absolute error (log scale) in Gaussian quadrature used to compute the generalized winding number on a cubic Bézier curve with Equation 5, evaluated with 15-, 30- and 50-nodes. (middle) Using Gaussian quadrature to compute generalized winding numbers over a shape leads to unacceptable errors. (bottom) Close-up of highlighted region.

that the value returned by the method is accurate. We can see how this would affect a containment query on a closed shape in Figure 6. In this example, the only correct values for the rounded winding number are 0, 1, and 2. However, even the most accurate quadrature scheme shown produces values in the vicinity of the quadrature nodes that are completely meaningless.

4 GENERALIZED WINDING NUMBERS FOR CURVED GEOMETRY

While these integration formulae are useful from a theoretical perspective, their inherent instability necessitates an approach that avoids evaluating them directly. In their place, we develop a framework that computes winding numbers based only on geometric properties of the curve. To this end, we make heavy use of a particular kind of closing curve for each shape. Given an (in principle arbitrary) parametric curve Γ , we define the linear segment $C : [0, 1] \rightarrow \mathbb{R}^2$ by $C(t) = \Gamma(1)(1-t) + \Gamma(0)t$ as the *linear closure* of our curve.

As the name suggests, the union of this closure C and Γ will always be a properly oriented, closed curve, and thus partitions \mathbb{R}^2 into discrete enclosing regions. Thus, winding numbers with respect to the total curve $\Gamma \cup C$ are integers, such that

$$w_{\Gamma}(a) + w_C(a) = w_{\Gamma \cup C}(a) \in \mathbb{Z}. \quad (6)$$

Most importantly, the winding number of such a closure can always be computed *exactly* without the need to appeal to quadrature by Equation 4, as the angle subtended by the introduced linear segment can be evaluated through a (relatively) inexpensive arccosine. In this way, we can always solve the problem of the *generalized* winding number using a solution to the *integer* winding number problem on the closed curve, as subtracting away the contribution of this linear

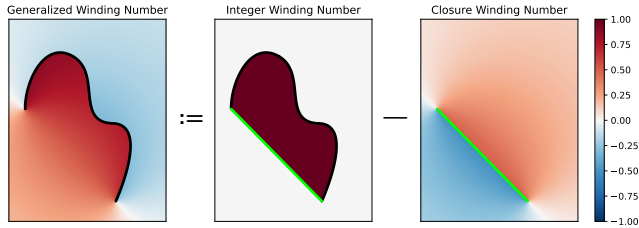


Fig. 7. The unknown winding number with respect to a curved shape (left) can be computed by finding the *integer* winding number of the closed figure (center) and subtracting away the contribution of its closure (right).

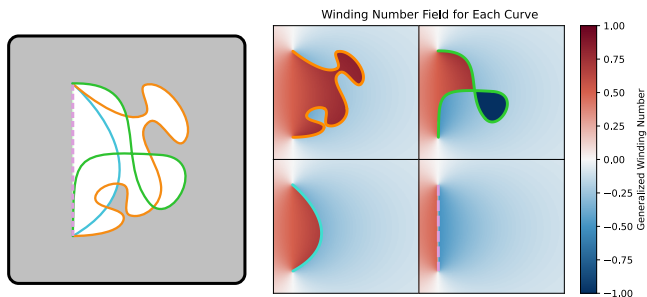


Fig. 8. (left) All three curves are closed by the same dashed line, and the shaded region is exterior to all three closed shapes. (right) This means that the winding number field generated by each curve (and the closure, up to orientation) is *identical* in the shaded region.

closure can be done independently of the original curve geometry. While we present our own algorithm to compute this integer winding number in Section 5 that outperforms known alternatives, we note that this strategy is fully compatible with more conventional techniques for computing containment queries in closed regions bounded by curves, e.g., Bézier clipping [Nishita et al. 1990].

This principle is particularly useful in the case when $w_{\Gamma \cup C}(q) = 0$, i.e., the query point is located outside the closed shape $\Gamma \cup C$, or outside the convex hull of Γ itself, where we have $w_{\Gamma}(q) = -w_C(q)$. This usage is introduced in Jacobson et al. [2013], where it is used in a hierarchical evaluation of the winding number for collections of linear triangular facets. An important consequence of this for curved geometry is that for query points that are far enough away, we can reverse the orientation of C and treat the shape as *equivalent* to its straight line closure, and compute the winding number of the entire curved segment *immediately* and *exactly*.

We further extend this principle to evaluate generalized winding numbers for points that are arbitrarily close to the curve with the same assurances of exact accuracy. In such cases, Γ can be replaced with a piecewise linear approximation $\tilde{\Gamma}$ so long as the *integer* winding number at q remains unchanged between $\Gamma \cup C$ and $\tilde{\Gamma} \cup C$ (see Figure 8). Doing so necessarily leaves the *generalized* winding number at q unchanged as well, as we have

$$w_{\Gamma}(q) = w_{\Gamma \cup C}(q) - w_C(q) = w_{\tilde{\Gamma} \cup C}(q) - w_C(q) = w_{\tilde{\Gamma}}(q).$$

It is from these observations that we derive an algorithm for computing *exact* generalized winding numbers over a collection of curves. For both far and near query points, we construct appropriate linearizations of our curves that are guaranteed to have the same generalized winding number at the point of interest as their curved counterparts. For a given point, the vast majority of curves in the model will be considered far, and their contribution to the generalized winding number can be computed with a single arccosine evaluation. For the few curves that are close enough such that the linear closure itself is an insufficient approximation, we construct an approximating polyline that provably generates the same winding number. In such cases, we evaluate the integer winding number for the closed polygon, which can be done without reference to any trigonometric functions using the point-in-polygon algorithm in Hormann and Agathos [2001] and subtract away the contribution of the linear closure.

In either case, we are able to compute the *exact* generalized winding number for an arbitrary point using only a single evaluation of arccosine for each curve in the shape. All that remains is to ensure that this polyline is adaptively constructed such that evaluating the integer winding number of the polyline is efficient.

5 WINDING NUMBER ALGORITHMS

We now describe our complete algorithm for solving the generalized winding number problem for a rational Bézier curve at a given query point (Algorithm 1). The core of this algorithm is described in Algorithm 2, which computes the *integer* winding number for a closed rational Bézier curve without reference to ray casting.

In brief, given such a curve Γ , we adaptively construct a polyline approximate $\tilde{\Gamma}$ which has the same winding number at the point of interest. We then apply a standard point-in-polygon algorithm to this closed polyline to compute their shared integer winding number and subtract the contribution of their shared closure.

For this procedure to work, we require that our polyline $\tilde{\Gamma}$ satisfies $w_{\Gamma \cup C}(q) = w_{\tilde{\Gamma} \cup C}(q)$. The simplest way to ensure this is for q to be located outside both closed shapes $\Gamma \cup C$ and $\tilde{\Gamma} \cup C$, as then this shared integer winding number is equal to 0. Because this cannot be guaranteed in general, we recursively bisect the curve into components $\{\Gamma_i\}_i$ until q is located outside the closed shape for each.

There are a number of ways to ensure that for each component, $w_{\Gamma_i \cup C_i}(q) = 0$. For example the convex hull property of a Bézier curve states that the curve Γ_i is completely contained within the convex hull of its control nodes. Given a description of this convex hull, classic point-in-polygon algorithms such as Hormann and Agathos [2001] can be applied to test for containment. This leads to the following outline for our algorithm:

- If the query point is located outside this convex hull, then it is guaranteed to be outside the closed shape $\Gamma_i \cup C_i$ and we replace the component Γ_i with the reversal of its linear closure C_i , as they now provably have the same generalized winding number.
- Otherwise, we bisect the Bézier curve and repeat the algorithm on each half. As a base case, we check whether the

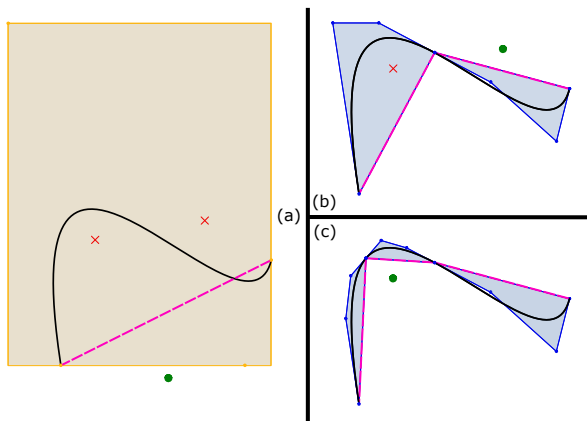


Fig. 9. Three iterations of the approximating polyline algorithm. In (a), one point is outside the bounding box, and its winding number can be computed from the dashed closure. In (b), after a bisection we can compute the winding number for an additional point. We repeat the process in (c) to compute the winding number for the remaining point.

curve is approximately linear (Algorithm 3) as such segments can be added directly to the polyline.

Once this approximating polyline is fully constructed, we close it and use the same point-in-polygon algorithm, `PolygonWindingNumber`, to compute the integer winding number for the closed polygon. The contribution of this closure is then subtracted, and we are left with the exact generalized winding number for the curve. Figure 9 illustrates this algorithm for three successively closer query points to a Bézier curve.

To improve computational efficiency in the case where subdivision is expensive and exact geometric accuracy is not needed, we optionally provide in Algorithm 3 an early stopping criterion. We choose this particular heuristic because it is easier to evaluate than the distance between the query point and the curve, and has a more intuitive spatial interpretation than a simple bound on the maximum number of curve subdivisions. In practice, Bézier curves are smooth nearly everywhere on their interior, and so the recursive bisection step needs to be done relatively few times before the query point is found to be outside the convex hulls of each component, even for points very close to the curve. Importantly, this causes the accuracy of Algorithm 1 and runtime of the procedure to both be reasonably insensitive to this numerical tolerance, as the algorithm is likely to terminate for these smooth curves well before there has been sufficient refinement for subdivisions to be considered linear. Indeed, in the numerical examples we provide, we set this value to machine epsilon and find no significant loss in performance (See Section 7.2).

We note that it can be problematic if the query point is located directly on the curve, as in such cases the winding number is not formally defined. Since we need our algorithm to be compatible to such input points, we discuss this case extensively in Section 6.

Despite the simplicity of the above procedure, directly computing containment within a Bézier curve’s convex hull can be prohibitively expensive. Instead, we first determine if the query is contained in

ALGORITHM 1: `WindingNumberCurve`: Evaluate the generalized winding number for an arbitrary rational parametric Bézier curve

Input: Γ : Rational parametric Bézier curve Γ

q : Query point

Output: w_{Γ} : The winding number evaluated at q

/ Store the linear closure of Γ */*

$C(t) \leftarrow (1-t)\Gamma(1) + t\Gamma(0)$

$w_C \leftarrow (1/2\pi) \times (\text{Signed angle subtended by } \overrightarrow{qC(0)} \text{ and } \overrightarrow{qC(1)})$

if $q \notin \text{BoundingBox}(\Gamma)$ **then**

 | **return** $-w_C$

else

 | **return** `IntegerWindingNumberCurve`($\Gamma \cup C, q$) $- w_C(q)$

end

an axis-aligned bounding box that encompasses the convex hull. Although this can have a much larger area, the bounding box containment query is inexpensive and most points will be exterior to it. If it is not, we use a test from Ma and Hewitt [2003] to check if the control polygon (defined by a Bézier curve’s control nodes) is already simple and convex (see Algorithm 4). If so, we perform the point-in-polygon test to determine if the query lies outside the convex hull of the control polygon. For efficiency, our implementation notes that when a curve is simple and convex, its subcomponent curves will be as well.

Altogether, the proposed algorithm is given by Algorithm 1. We note that this procedure is not meaningfully restricted to rational Bézier curves. Their focus throughout this work reflects their simplicity during calculation through manipulation of their control nodes, which itself justifies their ubiquity in application. More generally, the algorithm is applicable to any curve for which it is possible to construct a bounding box and a linear closure. In particular, extension to collections of NURBS curves is straightforward.

6 GENERALIZED WINDING NUMBERS FOR COINCIDENT POINTS

The generalized winding number is, in a strict mathematical sense, undefined for points located directly on the curve. This poses a practical problem during the implementation of a containment query, as they are often executed massively in parallel for large clusters of points without any *a priori* knowledge of their position relative to the boundary. Furthermore, our desire for exact geometric fidelity precludes us from applying the standard trick of perturbing such points randomly to place them definitively on one side of the singular boundary. In spite of this, it is necessary for compatibility with respect to downstream applications that the algorithm definitively return a value that adheres to a convention that is mathematically justified and intuitive to the caller of `WindingNumberCurve`.

An important aspect of the winding number scalar field is that it is harmonic, being the unique solution of a Laplacian operator with boundary data prescribed by each side of our curves, enforcing a jump discontinuity across them. Therefore, the simplest convention for the winding number of coincident points is to take it as the average value across this jump discontinuity. For linear segments, this boundary data enforces that the winding number approaches $+1/2$

ALGORITHM 2: IntegerWindingNumberCurve Evaluate the integer winding number for a rational parametric Bézier curve closed by a linear segment. We use an algorithm from [Hormann and Agathos 2001] for PolygonWindingNumber.

Input: $\Gamma \cup C$: Closed rational parametric Bézier curve
 q : Query point
Output: $w_{\Gamma \cup C}$ The integer winding number evaluated at q

```

 $\tilde{\Gamma} \leftarrow \{\}$  // Initialize the polyline approximation
 $w_{\Gamma \cup C} \leftarrow 0$  // Initialize the winding number

Push  $\Gamma$  onto an empty stack.
while the stack is not empty do
   $\Gamma_0 \leftarrow \text{StackPop}$ 
   $C_0(t) \leftarrow (1-t)\Gamma_0(0) + t\Gamma_0(1)$ 

  /* Check for coincidence at the endpoints */
  if isSimpleConvex( $\Gamma_0$ ) and ( $q = \Gamma_0(0)$  or  $q = \Gamma_0(1)$ ) then
    /* Track the contribution of coincident points */
     $w_{\Gamma \cup C} \text{ += ConvexEndpointWindingNumber}(\Gamma_0, q)$ 
     $\tilde{\Gamma} \leftarrow \tilde{\Gamma} \cup \{C_0\}$  // Add to the polyline
  else
    if {isSimpleConvex( $\Gamma_0$ ) and  $q \notin \text{ControlPolygon}(\Gamma_0)$ }
    or isApproximatelyLinear( $C_0$ )
    then
       $\tilde{\Gamma} \leftarrow \tilde{\Gamma} \cup \{C_0\}$  // Add to the polyline
    else
       $\Gamma_1, \Gamma_2 \leftarrow \text{Bisection}(\Gamma_0)$ 
      StackPush( $\Gamma_1, \Gamma_2$ )
    end
  end
end

 $P \leftarrow \tilde{\Gamma} \cup \{C\}$  // Close the polyline, forming a polygon
return PolygonWindingNumber( $P, q$ ) +  $w_{\Gamma \cup C}$ 

```

ALGORITHM 3: isApproximatelyLinear: Return true if each control point of a rational parametric Bézier curve is within a given tolerance of the closure.

Input: Γ : Rational parametric Bézier curve
 ϵ : User tolerance

```

 $C_0(t) \leftarrow (1-t)\Gamma_0(0) + t\Gamma_1(1)$ 
for each control node  $P_i$  of  $\Gamma$  do
  if SquaredDistance( $C_0, P_i$ )  $\geq \epsilon$  then
    return false
  end
end
return true

```

from one side and $-1/2$ on the other. This means that the winding number should be exactly 0 at every coincident point along linear segments. Similarly, this convention would ensure that the winding number along a more general closed shape is a fixed half-integer value along its length, changing only across self-intersections. However, this presents a unique problem for open, curved shapes, as

ALGORITHM 4: isSimpleConvex: Return true if the polygon of control points of a rational parametric Bézier curve is simple and convex (Adapted from [Ma and Hewitt 2003])

Input: Γ : Rational parametric Bézier curve with control nodes P_0, \dots, P_n

```

for  $i = 1, \dots, (n-1)$  do
  /* Store the linear segment connecting the nodes */
   $S(t) \leftarrow (1-t)P_{i-1} + tP_{i+1}$ 
  if  $i \leq n/2$  then
    if  $P_i$  and  $P_n$  are on the same side of  $S$  then
      return false
    end
  else
    if  $P_i$  and  $P_0$  are on the same side of  $S$  then
      return false
    end
  end
end
return true

```

the winding number is no longer constant along the length of the curve. This underscores the importance of a single convention for coincident points, as the winding number for individual curves must be computed without any knowledge of the other components that make up the shape, which can be unintuitive if that shape is itself closed. Nevertheless, we can evaluate the winding number at a coincident point for a single curve just as before, by computing the half-integer winding number of the closed shape and subtracting the contribution of the closing line.

This convention for a coincident winding number is well-justified in a mathematical context as well. As explained by Jacobson et al. [2013], the generalized winding number can be understood in the context of ray casting as the *average* number of intersections from rays cast in all directions from q . When q is on a straight line, only two of the uncountably infinite possible directions will intersect the line, resulting in an average of zero intersections. Furthermore, the discontinuity in Equation 5 is, in fact, removable, and evaluating the remaining integrand agrees with this interpretation of coincident winding numbers.

Despite this conceptual clarity, evaluating coincident winding numbers is complicated in practice, as it is computationally expensive to perform the projection necessary to identify when a point is located exactly on a Bézier curve [Ma and Hewitt 2003]. The exception to this is the endpoints, which are interpolated exactly by the first and last control points of the curve. At such points, the coincident winding number is equal to the signed angle spanned by the non-coincident endpoint, and a *tangent* vector at the coincident endpoint (see Figure 10). Furthermore, computing such a tangent is trivial at the endpoint of a Bézier curve, as it is defined by the endpoint and the adjacent control node. Interior points become endpoints in linear time by the repeated bisections of our algorithm. As this spanned angle must account for full revolutions of this span around the tangent vector, our algorithm only applies this edge-case to curves with simple and convex control polygons

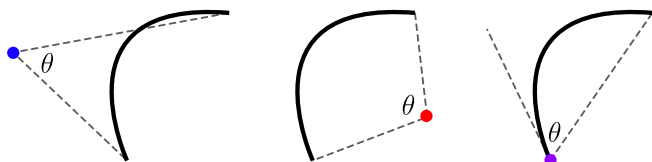


Fig. 10. For points located outside the closed curve (left, center), the winding number is proportional to the angle subtended by the endpoints. For points located directly on an endpoint (right), we define the winding number as proportional to the angle subtended between the non-coincident endpoint and a tangent vector.

ALGORITHM 5: `ConvexEndpointWindingNumber`: Return the convention for “coincident winding numbers” as outlined in Section 6

Input: Γ : Rational parametric Bézier curve

q : Query point

Output: w_p : The winding number evaluated at the endpoint of the curve

```

if  $q = \Gamma(0)$  then
  return  $\frac{1}{2\pi} \times (\text{Signed angle subtended by } \overrightarrow{q\Gamma(1)} \text{ and } \overrightarrow{\Gamma'(0)})$ .
else
  return  $\frac{1}{2\pi} \times (\text{Signed angle subtended by } \overrightarrow{q\Gamma(0)} \text{ and } -\overrightarrow{\Gamma'(1)})$ .
end

```

(see Algorithm 5). This particular criterion is typically met after very few bisections, and ensures that no full revolutions can occur on the local subcurve.

7 NUMERICAL EXPERIMENTS AND RESULTS

7.1 Robustness of Containment Queries on Curved Geometry

We first show the utility of a generalized winding number approach to containment queries in the context of messy CAD geometry. We compare our approach to conventional ray casting for containment, which remains the de-facto standard for exact containment in curved watertight geometry. Consider the geometry of Figure 11, which is composed of 87 linear segments and 477 cubic Bézier curves. We remove the marked curve in Figure 11 and determine containment at each pixel in the image using a simple ray casting algorithm that extends the ray to the right of the pixel and counts the number of intersections with the shape. As expected, the use of a ray casting algorithm causes catastrophic issues, in the sense that the containment errors are located at a great distance from the actual geometric error. This makes the model unusable in the applications of interest without specific oversight and (often user-driven) correction of these errors.

We then apply our generalized winding number algorithm to the same geometry in Figure 12, evaluating at each pixel the winding number field generated by the bounding geometry. In this example, points which we expect to be interior to this shape have winding numbers close to 1, while the value for points we expect to be exterior is close to 0. We see that deleting the curve does degrade the field, but that the most severe effects are *localized* to the site of

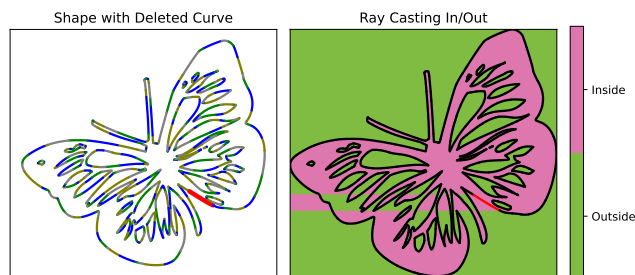


Fig. 11. (left) A watertight geometric shape, with the exception of the curve in red that is removed during calculation. (right) Ray casting classifies points as interior (pink) and exterior (green). As expected, the geometric error leads to numerous errors in containment queries.

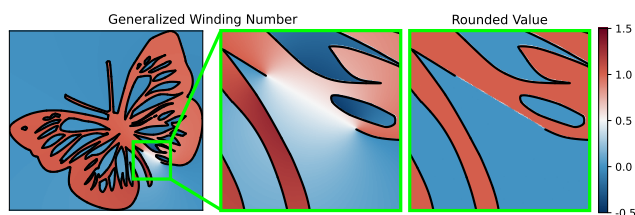


Fig. 12. (left) The generalized winding number computed using Algorithm 1. (center) The winding number in the region around the deleted curve degrades smoothly. (right) Rounding the winding number produces an approximation that better conforms to the designer’s intuition.

the geometric error. In some sense, the locality of this degradation is unintuitive, as containment is necessarily a *global* property of points relative to bounding geometry, i.e., the winding number at every point is dependent on every curve. Importantly, however, the influence of distant curves becomes rapidly negligible. While the exact value of the winding number now a non-integer value (nearly) everywhere in the domain, away from the deleted curve the difference from the nearest integer is nowhere large enough to have an adverse effect on the actual determination of containment.

A natural mapping of winding numbers to containment classifications is to round each value to the nearest integer and apply a non-zero rule. For example, in Figure 12 we see that the *rounded* winding number produces a clear boundary between regions along the $1/2$ isocurve. While this implied boundary lacks the curvature of the deleted curve, it is a reasonable approximation and nevertheless allows for the surface to be used immediately in applications. This behavior is especially desirable when geometric errors arise from small gaps between components of the mesh, as the resulting containment query reflects the simplest closure of the shape permitted by the provided geometric data.

Through Figure 13, we reiterate the two important kinds of robustness in this work. The shape in this example is intentionally deformed such that no pairs of adjacent edges are connected. From a distance, i.e., the scale at which the model geometry is visually “good enough,” there are functionally no irregularities in the shape’s

appearance. This mirrors the realistic setting, in which such imperceptible tolerances are unexpected and remain catastrophic for conventional containment algorithms. As before we apply our algorithm to each pixel in the bounding region, and plot the winding number field alongside its difference from the nearest integer.

The first type of robustness is derived from our specific computation of the generalized winding number through Algorithm 1, which is completely robust to non-watertight geometry. Because we only ever consider individual curves without any reference to their overall arrangement, both the computational performance and accuracy of our calculation is completely unaffected by the non-watertightness of the shape. Furthermore, our adaptive algorithm ensures that the winding number is computed stably at arbitrary distances from individual curves. On the other hand, *any* containment query derived from generalized winding numbers is robust to geometric artifacts in the boundary. This is because although even small errors have far-reaching influence on the scalar field of winding numbers, this influence degrades rapidly and gracefully away from the sources of these errors. Even at this frequency of topological errors in this example, the resulting values of the field still very clearly partition the shape into the expected interior and exterior.

It is noted in Jacobson et al. [2013] that a useful perspective on the fractional value of the winding number is as a measure of confidence for the derived containment query. For example, points for which the winding number is close to an integer can be considered “more likely” to be classified accurately according to the unknown, but presumably watertight shape that lacks any of the present boundary artifacts. On the other hand, points with winding numbers near half-integer values are often close to the boundaries that are only implied by the rounding heuristic, and correlate with containment classifications that are less informed by existing surface elements. In Figure 13, we specifically identify each point in the image with a winding number in the range $[0.25, 0.75]$, which are the only points that we can therefore consider to have an “uncertain” classification according to our rounding heuristic. From this, we can see that the *vast* majority of points in this image are classified with high confidence.

Of course, the rounding heuristic is not the only way one can determine containment from a fractional winding number, as access to the full winding number field through our algorithm permits treatment of containment through standard segmentation strategies. For example, Jacobson et al. [2013] meshes the interior of the shape, utilizing an energy minimization technique to enforce additional smoothness in the resulting segmentation. Additionally, the above usage of the winding numbers as a measure of confidence can be made more rigorous through further statistical analysis. In Sellán and Jacobson [2022], the generalized winding number field is used as a prior for a number of statistical distributions, including likelihood of containment. This rigor is desirable for certain applications, as the fractional value of the winding number alone cannot be interpreted as a probability for the accuracy of the derived containment query. Because we are ultimately concerned with efficient evaluation of the winding number in this work, we use only the rounding heuristic to determine containment in our examples.

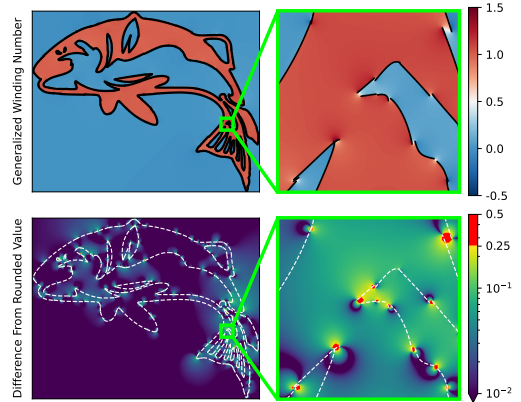


Fig. 13. (top) Generalized winding number for a shape that is intentionally deformed to slightly separate adjacent curves. (bottom) Absolute difference (log scale) between computed winding number and rounded winding number. Points for which the difference is greater than 0.25 are highlighted, as this indicates some degree of uncertainty in the classification. As we see, such points are very sparsely distributed throughout the domain.

7.2 Algorithm Performance

We now consider the computational performance of Algorithms 1 and 2. In each of the following examples, we evaluate our algorithm in the context where exact geometric fidelity is prioritized, and set all numerical tolerances to zero. For comparison, we consider state-of-the-art techniques that have been adapted to collections of parametric curves. Perhaps the most common winding number-based approach for watertight geometry would be to discretize the shape into approximating polygons with linear edges and to apply a standard point-in-polygon algorithm to the polygonal shapes to determine containment. If watertightness cannot be guaranteed, the generalized winding number approaches of Jacobson et al. [2013]; Trettner et al. [2022] admit a straightforward adaptation to a linearized shape.

In demonstration of this fact, consider Figure 14, where we compare the performance of our serial implementation of the proposed adaptive strategy to such a linearized method. Specifically, we consider refinements of each curves in the shape into a fixed number of linear segments. We note that the shape considered is non-watertight and non-manifold, making it a particularly challenging case for conventional containment queries. Nevertheless, the fractional winding number field clearly partitions the space according to an intuitive understanding of its interior.

By comparing wall-clock runtime, we see that the use of a fixed level of refinement of the curve leads to a more efficient calculation of the generalized winding number as expected, as curve subdivisions can be computed during a preprocessing step and reused. While the analogous overhead in the adaptive algorithm can be similarly mitigated through hash maps that cache the results of curve subdivisions, we consider further implementation optimizations for collections of curves to be future work, and instead focus here on efficient and accurate calculation for individual curves. Indeed, when geometric fidelity is not a concern, using a fixed linearization

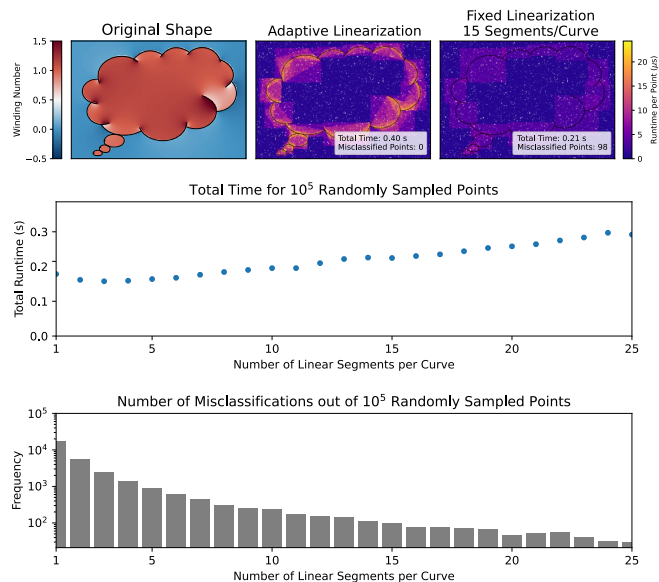


Fig. 14. Given a non-watertight, non-manifold shape, we observe the practical effect of curve linearization on geometric accuracy and runtime. We compare our adaptive strategy to those that uses a uniform, piecewise linearization of the curve at fixed levels of refinement. We record the total time spent to evaluate the generalized winding number for 10^5 points randomly sampled from a uniform distribution, and the number of these points which are ultimately misclassified. We see that a fixed linearization can offer better computational performance, as subdivisions of the curve can be efficiently precomputed. However, we note that even at high levels of refinement there are still a considerable number of misclassifications, which is unacceptable in many downstream applications of interest.

to compute generalized winding numbers with Algorithm 1 is a reasonable approach, although it may involve a considerable memory footprint for larger models.

Importantly, such a linearized method offers no guarantees for geometric fidelity. By comparing the level of refinement against an approximate likelihood that points are ultimately misclassified, we see that although the rate of misclassification decreases with increased linear refinement, it does so at a rate that we consider to be far too slow to be practical for many applications of interest, especially for downstream applications that are sensitive to *any* misclassifications. This is also an issue for methods that sample the bounding geometry non-deterministically, such as those of Barill et al. [2018]; Nooruddin and Turk [2003]; Sawhney and Crane [2020]. In contrast, the proposed method is able to achieve perfect geometric fidelity for points that are arbitrarily close to the boundary. In place of these fixed linearization methods, we turn our attention to those that operate directly on curved geometry without approximation.

There remains the issue that the methods known to the authors that operate directly on curved geometry à la Bézier clipping [Nishita et al. 1990] can do so only in the watertight case. This regrettably means that no direct comparison for the performance Algorithm 1 can be made at this time. In its place, we note that our procedure for computing the exact generalized winding number

for any open, curved object requires applying an *arbitrary* integer winding number algorithm on the closed shape and subsequently subtracting out the contribution of the closure of the curve. It is in this context that we demonstrate the performance of Algorithm 2, which is one such algorithm that computes the integer winding number of a closed, curved object. We show that Algorithm 2 is considerably more performant than alternatives utilizing crossing numbers for containment queries on closed, curved geometry, as such methods necessarily introduce inefficiencies through a reliance on ray casting.

We first consider a common procedure for ray casting that applies a geometric *binary search* to compute ray-Bézier curve intersections [Farin 2001]. In this algorithm, we extend from the query a ray in the direction of the closest edge of a bounding box, and recursively bisect the curve until guarantees for potential intersections can be made. If a subcurve is sufficiently linear (see Algorithm 3) such that an intersection is guaranteed, then the signed intersection can be recorded. If the ray does not intersect the subcurve’s bounding box, then it can be discarded. This procedure converges to points of intersection linearly, and is capable of identifying multiple intersections with the same curve.

Bézier clipping [Nishita et al. 1990; Sederberg and Nishita 1990] is a similar procedure, discarding sections of the curve that are guaranteed to have no intersections with the ray. However, the convergence of this approach to points of intersection is quadratic, as the curve is instead split into three subcurves at each iteration.

To make the appropriate comparison to Algorithm 2 in the context of exact winding numbers for non-watertight geometry, we utilize each *within* the framework described in Algorithm 1. That is to say, we iterate over each open curve in the model and close it, compute the integer winding number with Algorithm 2 or one of the above alternate techniques, and subtract out the contribution of the closure. We note that, while these alternate techniques themselves are well-known in the context of ray-tracing and containment of closed shapes, computing generalized winding numbers in this way is, to our knowledge, itself a novel application of their use.

In comparing these three algorithms, there are a number of specific implementation details that govern the wall-time of their evaluation, in particular the use of a spatial index to efficiently handle far away points. However, we are concerned with their efficiency in the near-curve regime, where the number of *curve evaluations* (most commonly occurring as a result of a curve bisection) is an effective proxy for computational efficiency between the methods. Furthermore, the computation of the exact fractional value for the generalized winding number only involves the evaluation of a single arccosine for each curve in the model, and this cost is equivalent across all three methods.

Our first example in Figure 15 uses a relatively simple shape featuring at most polynomial, cubic Bézier curves. To compare the three algorithms (binary search, Bézier clipping, and the proposed approach presented in Algorithm 2), we randomly sample 250,000 points from a bounding box around the shape using a uniform distribution, and count the number of curve evaluations on each that needs to be made to determine containment with perfect geometric fidelity. Again, we emphasize that all three methods are being used here to compute the generalized winding number, which they do

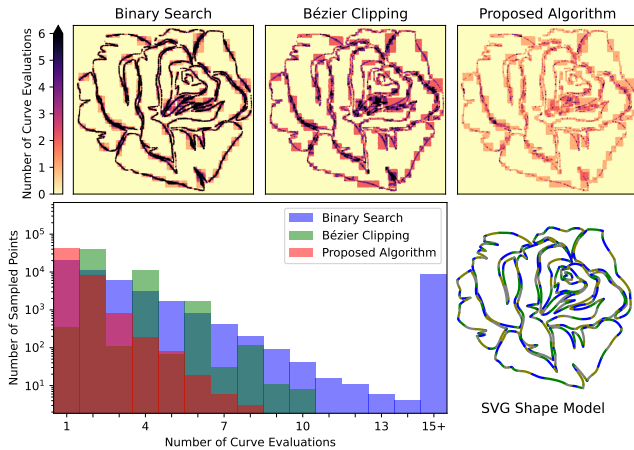


Fig. 15. We compare the cost of a conventional geometric bisection method, Bézier clipping, and our proposed method on a watertight model. We uniformly sample 250,000 points from a bounding box and measure the number CAD model evaluations, i.e., curve subdivisions, that are necessary to evaluate the winding number at each point to full geometric precision. In the histogram, we omit the number of sampled points that require zero function evaluations, as this occurs whenever the point is outside each curve’s bounding box, occurring at the same frequency for each method.

by computing the integer winding number of the closed shape for each curve in the model.

As expected, the binary search ray casting algorithm performs the poorest, with a considerable number of points requiring more than 15 curve subdivisions to make an evaluation. This is because the terminating condition of the naive algorithm is linearity of the curve component near the intersection, for which the number of bisections increases with curvature. This behavior is improved considerably with Bézier clipping, which performs more sophisticated subdivision (at greater computational expense) by considering the convex hull of the curve. However, the proposed algorithm improves on this further, with no points requiring more than 8 function evaluations, and most requiring only one or two.

Furthermore, we recognize that for the vast majority of points, *zero* function evaluations are required for each algorithm. This is because the first containment check performed in each is one to an axis-aligned bounding box for the entire curve, and most points lie outside the bounding box for most curves. For this reason, integer winding numbers for most combinations of points and curves are known to be zero without any curve evaluations.

One explanation for the relatively poor performance of the two ray casting algorithms is that such procedures are unnecessarily informative. Not only do ray casting algorithms *identify* potential intersections, but they also inherently *locate* the intersections, information which is rarely needed when the ray itself is chosen arbitrarily. In contrast, by not needing to compute the location of these intersections, the proposed algorithm is able to converge with considerably fewer curve evaluations, despite only doing so linearly. This is particularly relevant in the case where the curves of interest

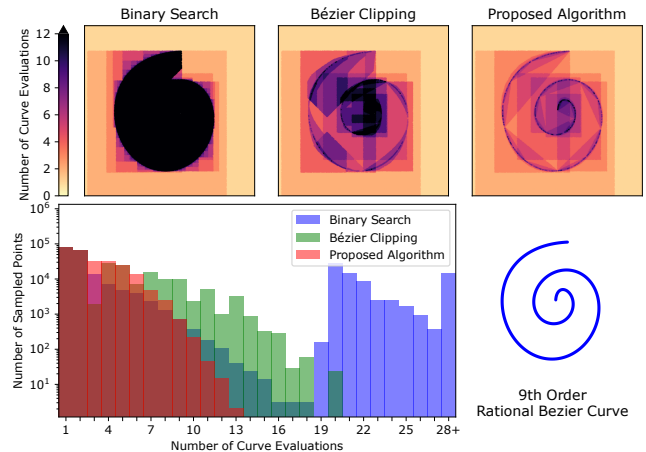


Fig. 16. For higher-order curves, additional curve subdivisions are necessary to reach perfect geometric fidelity. The proposed algorithm still has superior performance in such cases, as the location of the various intersections need not be computed directly.

contain multiple overlapping components, self-intersections, or regions of significant curvature, the latter of these cases being tested directly on a rational curve in Figure 16. In these examples, the ray casting algorithms perform particularly poorly, as any ray extended will intersect the curve multiple times, incurring additional cost with each.

Finally, we consider the sensitivity of our procedure for computing generalized winding numbers to the numerical tolerance introduced in Algorithm 3. We do this by considering performance and accuracy as a function of this tolerance, respectively measured in terms of the number of curve evaluations required to determine containment and the number of misclassifications caused by the final layer of implied linearization. These metrics are evaluated over 10^5 points sampled from a bounding box for the closed shape in Figure 17, which for demonstration purposes depicts a 9th order, rational curve that far exceeds the complexity of those typically used in practice. We see that invoking Algorithm 3 to prematurely terminate the procedure does marginally reduce the overall computational cost, but that this performance improvement is closely correlated with losses in classification accuracy. This effect is amplified on shapes with such complex curvature, as it becomes more likely that the implied linearization imposed by the tolerance will produce misclassifications. At the same time, we observe that even for very intricate shapes, the performance of Algorithm 2 rapidly stabilizes along both metrics. As a result, our practical preference is to set this tolerance to zero even in cases where geometric accuracy is not a priority.

As is always the case when generalized winding numbers are used to determine containment, there is still an assumption that the collection of curves are reasonably *oriented*. Because the scalar field generated by a single curve is computed completely independently of every other curve by design, a single reversed curve interferes with the contribution of surrounding curves, functionally reversing nearby containment queries, as can be observed in Figure 18. While

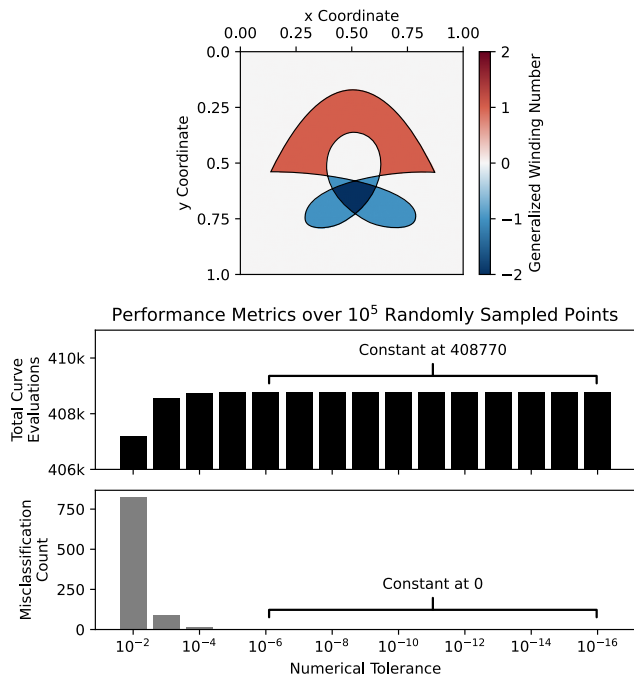


Fig. 17. Beyond a certain threshold, the adaptive nature of the proposed algorithm causes it to be insensitive to the numerical tolerance in Algorithm 3. In practice, we set this tolerance to zero, as this ensures exact geometric fidelity while not incurring a large additional computational burden.

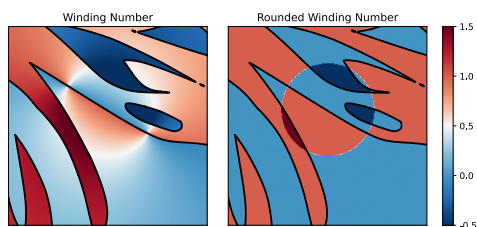


Fig. 18. Generalized winding numbers generated by the shape in Figure 12, but with the orientation of the highlighted curve reversed rather than removed entirely. The reversal of this curve causes surrounding containment classifications to be swapped within a local region.

the induced classification errors are still local to the geometric error, and would no doubt cause catastrophic errors *without* the use of generalized winding numbers, they nevertheless represent a type of geometric error unaccounted for by the current method.

8 CONCLUDING REMARKS

In this paper, we proposed a generalized winding number algorithm for collections of parametric curves, addressing the inherent challenges for curved geometry that are not shared by their linear counterparts. For the vast majority of points that are considered “sufficiently far” from a given curve, we circumvent the need for expensive linearization of the shape by treating the curve directly

as a single segment connecting its endpoints. For all other points, we demonstrated that our recursive algorithm outperforms existing ray casting algorithms adapted for this context. Finally, we formalize a procedure for handling winding numbers of points that are coincident with the curve, further increasing the practicality of our algorithm. In each case, this algorithm makes minimal reference to the often near-singular integrals from which the winding number is theoretically derived, instead operating strictly according to geometric and trigonometric principles. Nevertheless, we are able to *exactly* compute generalized numbers for curved shapes in two dimensions. Whereas the focus of this work has been on efficiently and robustly computing winding numbers for individual curves within a collection, follow-up work will consider accelerating the overall query workflow using a spatial index, as in Barill et al. [2018]; Jacobson et al. [2013]; Weiss et al. [2016] and exploiting the inherent parallelism within Algorithm 1 via threaded and/or GPU implementation.

In future work, we are interested in developing a generalized winding number for curved surfaces in \mathbb{R}^3 , but anticipate considerable challenges in developing an efficient numerical implementation. While much of the mathematical underpinnings remains the same, we cannot merely replace polar angles with their solid angle analogues. For example, while a line can close any curve in 2D, there are no simple geometric primitives that can close an arbitrary surface in 3D, or at the very least none that have direct formulas for their winding numbers. Nevertheless, the potential impacts of such a technique is highly motivating in spite of the challenges, and its pursuit will likely have theoretical foundations in the work presented here.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their thoughtful consideration and insightful feedback. This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

REFERENCES

- P. Asente, M. Schuster, and T. Pettit. 2007. Dynamic planar map illustration. *ACM Transactions on Graphics* 26, 3 (July 2007), 30:1–10. <https://doi.org/10.1145/1276377.1276415>
- A. Balu, M. R. Rajanna, J. Khristy, F. Xu, A. Krishnamurthy, and M.-C. Hsu. 2023. Direct immersogeometric fluid flow and heat transfer analysis of objects represented by point clouds. *Computer Methods in Applied Mechanics and Engineering* 404 (Feb. 2023), 115742. <https://doi.org/10.1016/j.cma.2022.115742>
- G. Barill, N. Dickson, R. Schmidt, D. I. W. Levin, and A. Jacobson. 2018. Fast Winding Numbers for Soups and Clouds. *ACM Transactions on Graphics* 37, 4, Article 43 (July 2018), 12 pages. <https://doi.org/10.1145/3197517.3201337>
- A. J. Barlow, P.-H. Maire, W. J. Rider, R. N. Rieben, and M. J. Shashkov. 2016. Arbitrary Lagrangian–Eulerian methods for modeling high-speed compressible multimaterial flows. *J. Comput. Phys.* 322 (2016), 603–665. <https://doi.org/10.1016/j.jcp.2016.07.001>
- S. Bischoff, D. Pavic, and L. Kobbelt. 2005. Automatic restoration of polygon models. *ACM Transactions on Graphics* 24, 4 (Oct. 2005), 1332–1352. <https://doi.org/10.1145/1095878.1095883>
- A. Capps, R. Carson, B. Corbett, N. Elliott, J. Essman, B. Gunney, B. Han, C. Harrison, R. Hornung, M. Larsen, A. Moody, E. Pauli, R. Settigast, L. Taylor, K. Weiss, C. White, B. Whitlock, M. Yang, and G. Zagaris. 2017–2024. *Axom: CS infrastructure components for HPC applications*. <https://doi.org/10.11578/dc.20210915.2> <https://github.com/llnl/axom>.
- P. C. P. Carvalho and P. R. Cavalcanti. 1995. IL2 - Point in Polyhedron Testing Using Spherical Polygons. In *Graphics Gems V*, A. W. Paeth (Ed.). Academic Press, Boston, 42–49. <https://doi.org/10.1016/B978-0-12-543457-7.50015-2>
- H. Edelsbrunner and E. P. Mücke. 1990. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.* 9, 1 (Jan. 1990), 66–104. <https://doi.org/10.1145/77635.77639>

- A. Efremov, V. Havran, and H.-P. Seidel. 2005. Robust and Numerically Stable Bézier Clipping Method for Ray Tracing NURBS Surfaces. In *Proceedings of the 21st Spring Conference on Computer Graphics* (Budmerice, Slovakia) (SCCG '05). Association for Computing Machinery, New York, NY, USA, 127–135. <https://doi.org/10.1145/1090122.1090144>
- G. E. Farin. 2001. *Curves and surfaces for CAD: A practical guide*. Morgan Kaufmann.
- D. Gunderman. 2021. *High-Order Spatial Discretization and Numerical Integration Schemes for Curved Geometries*. Ph. D. Dissertation. Department of Applied Mathematics, University of Colorado Boulder.
- D. Gunderman, K. Weiss, and J. A. Evans. 2020. Spectral mesh-free quadrature for planar regions bounded by rational parametric curves. *Computer-Aided Design* 130 (2020). <https://doi.org/10.1016/j.cad.2020.102944>
- E. Haines. 1994. I.4. - Point in Polygon Strategies. In *Graphics Gems*, P. S. Heckbert (Ed.). Academic Press, 24–46. <https://doi.org/10.1016/B978-0-12-336156-1.50013-6>
- C.W. Hirt, A. A. Amsden, and J. L. Cook. 1974. An Arbitrary Lagrangian-Eulerian computing method for all flow speeds. *J. Comput. Phys.* 14, 3 (1974), 227–253. [https://doi.org/10.1016/0021-9991\(74\)90051-5](https://doi.org/10.1016/0021-9991(74)90051-5)
- K. Hormann and A. Agathos. 2001. The point in polygon problem for arbitrary polygons. *Computational Geometry* 20, 3 (2001), 131–144. [https://doi.org/10.1016/S0925-7721\(01\)00012-8](https://doi.org/10.1016/S0925-7721(01)00012-8)
- A. Jacobson, L. Kavan, and O. Sorkine. 2013. Robust Inside-Outside Segmentation using Generalized Winding Numbers. *ACM Trans. Graph.* 32, 4 (2013), 1–12. <https://doi.org/10.1145/2461912.2461916>
- M. J. Kilgard. 2020. Polar stroking: New theory and methods for stroking paths. *ACM Transactions on Graphics* 39, 4 (Aug. 2020), 145:1–15. <https://doi.org/10.1145/3386569.3392458>
- L. Klinteberg and A. H. Barnett. 2020. Accurate quadrature of nearly singular line integrals in two and three dimensions by singularity swapping. *BIT Numerical Mathematics* 61 (07 2020), 1–36. <https://doi.org/10.1007/s10543-020-00820-5>
- Y. L. Ma and W. T. Hewitt. 2003. Point inversion and projection for NURBS curve and surface: Control polygon approach. *Computer Aided Geometric Design* 20, 2 (2003), 79–99. [https://doi.org/10.1016/S0167-8396\(03\)00021-9](https://doi.org/10.1016/S0167-8396(03)00021-9)
- B. Marussig and T. Hughes. 2017. A Review of Trimming in Isogeometric Analysis: Challenges, Data Exchange and Simulation Aspects. *Archives of Computational Methods in Engineering* 25 (06 2017), 1–69. <https://doi.org/10.1007/s11831-017-9220-9>
- A. McAdams, Y. Zhu, A. Selle, M. Empey, R. Tamstorf, J. Teran, and E. Sifakis. 2011. Efficient Elasticity for Character Skinning with Contact and Collisions. *ACM Trans. Graph.* 30, 4, Article 37 (July 2011), 12 pages. <https://doi.org/10.1145/2010324.1964932>
- A. A. Mezentsev and T. Woehler. 1999. Methods and Algorithms of Automated CAD Repair for Incremental Surface Meshing. In *IMR*. Citeseer, 299–309.
- T. Nishita, T. W. Sederberg, and M. Kakimoto. 1990. Ray Tracing Trimmed Rational Surface Patches. *SIGGRAPH Comput. Graph.* 24, 4 (sep 1990), 337–345. <https://doi.org/10.1145/97880.97916>
- F. S. Nooruddin and G. Turk. 2003. Simplification and repair of polygonal models using volumetric techniques. *IEEE Transactions on Visualization and Computer Graphics* 9, 2 (2003), 191–205. <https://doi.org/10.1109/TVCG.2003.1196006>
- A. Orzan, A. Bousseau, H. Winnemöller, P. Barla, J. Thollot, and D. Salesin. 2008. Diffusion Curves: A Vector Representation for Smooth-Shaded Images. *ACM Transactions on Graphics* 27, 3 (Aug. 2008), 92:1–8. <https://doi.org/10.1145/1399504.1360691>
- M. A. Park, R. Haines, N. J. Wyman, P. A. Baker, and A. Loseille. 2021. Boundary Representation Tolerance Impacts on Mesh Generation and Adaptation. In *AIAA AVIATION 2021 FORUM*. <https://doi.org/10.2514/6.2021-2992>
- N. Patrikalakis and T. Maekawa. 2010. *Shape Interrogation for Computer Aided Design and Manufacturing*. <https://doi.org/10.1007/978-3-642-04074-0>
- L. Piegl and W. Tiller. 1996. *The NURBS book*. Springer Science & Business Media.
- R. Sawhney and K. Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Transactions on Graphics* 39, 4 (2020).
- T. W. Sederberg, G. T. Finnigan, X. Li, H. Lin, and H. Ipson. 2008. Watertight trimmed NURBS. *ACM Transactions on Graphics* 27, 3 (Aug. 2008), 1–8. <https://doi.org/10.1145/1360612.1360678>
- T. W. Sederberg and T. Nishita. 1990. Curve intersection using Bézier clipping. *Computer-Aided Design* 22, 9 (1990), 538–549. [https://doi.org/10.1016/0010-4485\(90\)90039-F](https://doi.org/10.1016/0010-4485(90)90039-F)
- S. Sellán and A. Jacobson. 2022. Stochastic Poisson Surface Reconstruction. *ACM Transactions on Graphics* 41, 6, Article 227 (Nov. 2022), 12 pages. <https://doi.org/10.1145/3550454.3555441>
- R. Sevilla, S. Fernández-Méndez, and A. Huerta. 2008. NURBS-enhanced finite element method for Euler equations. *International Journal for Numerical Methods in Fluids* 57, 9 (2008), 1051–1069. <https://doi.org/10.1002/ld.1711>
- R. Sevilla, S. Fernández-Méndez, and A. Huerta. 2011. 3D NURBS-enhanced finite element method (NEFEM). *Internat. J. Numer. Methods Engrg.* 88, 2 (2011), 103–125. <https://doi.org/10.1002/nme.3164>
- M. Shimrat. 1962. Algorithm 112: Position of Point Relative to Polygon. *Commun. ACM* 5, 8 (Aug. 1962), 434. <https://doi.org/10.1145/368637.368653>
- A. Sommariva and M. Vianello. 2022. inRS: Implementing the indicator function of NURBS-shaped planar domains. *Applied Mathematics Letters* 130 (2022), 108026. <https://doi.org/10.1016/j.aml.2022.108026>
- D. C. Thomas, M. A. Scott, J. A. Evans, K. Tew, and E. J. Evans. 2015. Bézier projection: A unified approach for local projection and quadrature-free refinement and coarsening of NURBS and T-splines with particular application to isogeometric design and analysis. *Computer Methods in Applied Mechanics and Engineering* 284 (2015), 55–105. <https://doi.org/10.1016/j.cma.2014.07.014> Isogeometric Analysis Special Issue.
- P. Trettner, J. Nehring-Wirxel, and L. Kobbelt. 2022. EMBER: Exact Mesh Booleans via Efficient & Robust Local Arrangements. *ACM Trans. Graph.* 41, 4, Article 39 (July 2022), 15 pages. <https://doi.org/10.1145/3528223.3530181>
- K. Weiss, G. Zagaris, R. Rieben, and A. Cook. 2016. Spatially accelerated shape embedding in multimaterial simulations. In *Proceedings 25th International Meshing Roundtable (IMR '16)*, S. Canann (Ed.). Washington, D.C.
- S. Zellmann, D. Seifried, N. Morrical, I. Wald, W. Usher, J. P. Law-Smith, S. Walch-Gassner, and A. Hinkenjann. 2022. Point Containment Queries on Ray-Tracing Cores for AMR Flow Visualization. *Computing in Science & Engineering* 24, 02 (March 2022), 40–51. <https://doi.org/10.1109/MCSE.2022.3153677>