Optimal Curve Straightening is $\exists\mathbb{R}$-Complete

Jeff Erickson

University of Illinois, Urbana-Champaign

Abstract

We prove that the following problem has the same computational complexity as the existential theory of the reals: Given a generic self-intersecting closed curve γ in the plane and an integer m, is there a polygon with m vertices that is isotopic to γ? Our reduction implies implies two stronger results, as corollaries of similar results for pseudoline arrangements. First, there are isotopy classes in which every m-gon with integer coordinates requires $2^{\Omega(m)}$ bits of precision. Second, for any semi-algebraic set V, there is an integer m and a closed curve γ such that the space of all m-gons isotopic to γ is homotopy equivalent to V.

1   Introduction

The image of any self-intersecting closed curve γ in the plane has a natural structure as a (not necessarily simple) 4-regular plane graph, whose vertices are the points where γ intersects itself. After subdividing if necessary to remove any loops or parallel edges, the classical Wagner-Fáry-Stein-Stojaković theorem implies that there is a combinatorially equivalent plane map whose edges are straight line segments. Thus, any curve with n self-intersections can be continuously deformed, without ever changing its intersection pattern, into a polygon with at most $O(n)$ vertices. This type of continuous deformation is called an *isotopy*.

While this $O(n)$ bound is tight in the worst case, it is far from optimal for all curves. For example, straightening the curve in Figure 1 by finding a piecewise-linear embedding of its image graph yields the decagon on the left, but in fact the curve is isotopic to the self-intersecting hexagon on the right. As a more extreme example, consider any polygon in which *every* pair of edges intersects, for example, the regular star polygon with Schäfli symbol $\{m, \lfloor m/2 \rfloor\}$ for some odd integer m. Interpreting such a polygon as a curve and then converting that curve back to a polygon by straightening its image graph roughly *squares* the number of vertices.
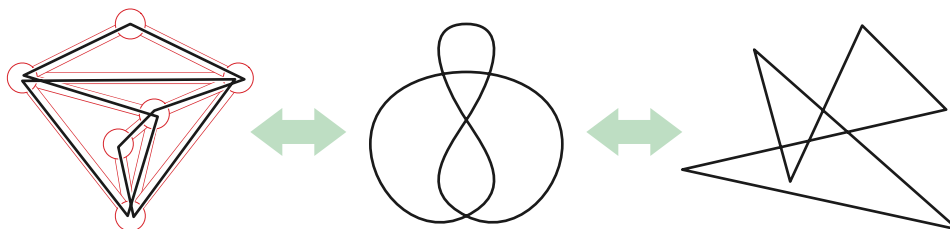

Figure 1. Two polygons isotopic to the same closed curve.

This quadratic gap between best and worst cases motivates the following question: How quickly can we find a polygon with the minimum number of vertices that is isotopic to a given closed curve? This is a special case of *minimum-segment drawings* problem proposed by Dujmović *et al.* [26], which asks for

the minimum number of line segments whose union is isomorphic, as a straight-line planar graph, to a given plane graph.

The main result of this paper is that deciding whether a given closed curve is isotopic to an *arbitrary* polygon with a given number of vertices has the same computational complexity as *the existential theory of the reals*—the problem of deciding whether a given system of multivariate polynomial equalities and inequalities has a real solution. In standard computational complexity terms, our result implies that optimal curve-straightening is both NP-hard and in PSPACE.

In Section 3, we prove that curve-straightening is hard using a simple reduction from the stretchability problem for pseudoline arrangements, which was proved equivalent to the existential theory of the reals by Mnëv [47, 48]. Figure 2 shows a typical example of our reduction. Our argument is similar to, but both simplifies and strengthens, a proof by Durocher *et al.* [27] that computing minimum-segment drawings of plane graphs is ∃ℝ-hard, which relies on a reduction by Bose *et al.* from pseudoline stretchability to recognizing arrangement graphs [12].

(a)   (b)
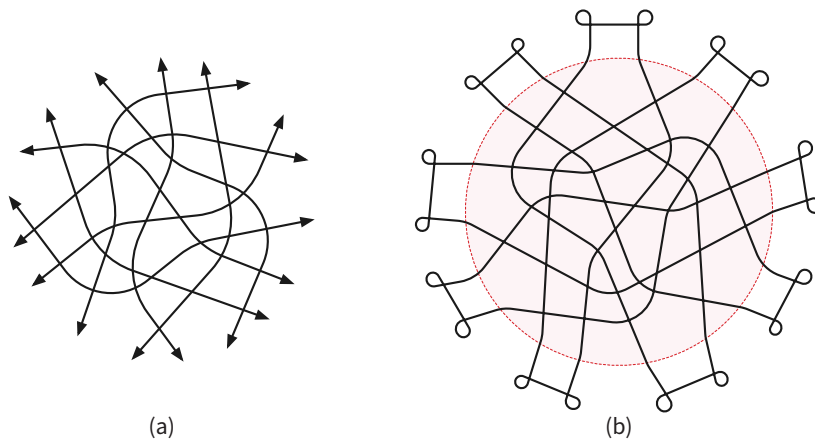
Figure 2. (a) Ringel's [56] minimal non-stretchable arrangement of 9 pseudolines. (b) A closed curve derived from Ringel's arrangement that is not isotopic to any 36-gon.

Together with Mnëv's universality theorem for pseudoline arrangements [47, 48], our reduction implies a similar universality result for polygons: the set of all polygons isotopic to a given curve and with a given number of vertices can have essentially arbitrary topology. Similarly, results of Goodman *et al.* [34, 35] imply that an integer polygon isotopic to a given curve could require doubly-exponentially large coordinates in the worst case.

Our reduction *to* the existential theory of the reals in Section 4 is surprisingly more delicate than our hardness proof, for two reasons. First, we must be much more explicit about how the isotopy class of the input curve is represented; we use a variant of codes proposed by Gauss [30, 31] and Dowker and Thistlethwaite [25], with some additional assumptions to establish an embedding in the plane instead of the sphere. More significantly, there is no direct correspondence between the vertex coordinates of the polygon and the conditions that guarantee isotopy. (By contrast, for example, reducing the recognition of segment intersection graphs to ETR requires exactly one clause for each pair of vertices, stating either that the corresponding segments cross or that they do not [42, 45].) We exploit the nondeterminism of ETR by guessing which edges of the isotopic polygon contain each self-crossing; using this information essentially requires us to implement integer arrays inside our formula.

Interestingly, given a closed curve γ with n self-intersection points, we can compute the minimum-complexity *orthogonal* polygon isotopic to γ in $O(n^{4/3} \operatorname{polylog} n)$ time, by combining Tamassia's reduction from minimum-bend grid embeddings to minimum-cost flows [65] with a recent planar minimum-cost flow algorithm of Karczmarz and Sankowski [41]; see also Cornelson and Karrenbauer [23].

2 Background

2.1 Closed Curves

A *parametrized closed curve* is a continuous function $\vec{\gamma}\colon S^1 \to \mathbb{R}^2$ from the circle $S^1 = \mathbb{R}/\mathbb{Z} = [0,1]/\{0,1\}$ to the plane. In this paper, we consider only *generic* closed curves, which are injective except at a finite number of transverse double points. The *basepoint* of a parametrized closed curve $\vec{\gamma}$ is the point $\vec{\gamma}(0)$.

Two parametrized closed curves $\vec{\gamma}$ and $\vec{\gamma}'$ are *equivalent* if they differ only by a reparametrization, meaning there is a homeomorphism $\theta\colon S^1 \to S^1$ such that $\vec{\gamma}' = \vec{\gamma} \circ \theta$. Equivalent parametrized curves have exactly the same images, but possibly different basepoints and orientations. A *closed curve* is an equivalence class of parametrized closed curves under this equivalence relation.

An *ambient isotopy* of the plane is a continuous family $H\colon [0,1] \to (\mathbb{R}^2 \to \mathbb{R}^2)$ of homeomorphisms of the plane, where without loss of generality $H(0)$ is the identity map. Two parametrized closed curves $\vec{\gamma}$ and $\vec{\gamma}'$ are **isotopic** if there is an ambient isotopy H such that $\vec{\gamma}' = H(1) \circ \vec{\gamma}$, and two closed curves are isotopic is they have isotopic parametrizations. Less formally, two curves are isotopic if there is a continuous deformation of the plane taking (the image of) one to (the image of) the other.

The *image graph* of a closed curve γ is the plane graph whose vertices are the curve's self-intersection points, and whose edges are maximal subpaths of the curve between vertices. Two closed curves are isotopic if and only if their image maps are combinatorially isomorphic (forbidding reflection). Thus, up to isotopy, closed curves can be represented either by any reasonable data structure for topological planar maps [21, 37, 49, 66, 68], or by more specialized structures encoding the order and directions of self-intersections along the curve [19, 25, 30, 31, 64]. We describe one such encoding in Section 4.1.

We regard any polygon P with vertices $(p_1, p_2, \ldots, p_m)$ as a piecewise-linear parametrized closed curve with basepoint p_1 that visits the vertices p_i in increasing index order. A polygon is *generic* if all vertices are distinct, no vertex lies in the interior of an edge, and no three edges contain a common point.

Lemma 2.1. *Every generic closed curve with n self-crossings is isotopic to a generic polygon with $O(n)$ vertices.*

Proof: Let γ be a generic closed curve with n self-crossings. Subdivide the edges of the image graph of γ as necessary to to obtain a simple plane graph G. Let $\bar{G}$ be a straight-line plane graph that is combinatorially equivalent (and therefore isotopic) to G. Let P be the unique Euler tour of $\bar{G}$ that crosses itself at every vertex; P is a polygon with (crudely) at most $6n$ vertices. Finally, for sufficiently small ε, uniformly and independently perturbing each vertex of P inside a ball of radius ε yields an isotopic polygon $\tilde{P}$, which is generic with probability 1. $\qquad\qquad\square$

2.2 Pseudoline Arrangements

An **pseudoline arrangement** is a cellular decomposition of the (projective) plane by a finite set of infinite curves (called *pseudolines*), each of which is isotopic to a straight line, and any pair of which intersect transversely, exactly once. In particular, every arrangement of (pairwise non-parallel) straight lines is a pseudoline arrangement. A pseudoline arrangement is **simple** if each point in the plane lies on at most two pseudolines. Two pseudoline arrangements are **isomorphic** if they define combinatorially isomorphic decompositions, or equivalently [9] if there is a homeomorphism of the plane to itself that maps one arrangement to the other.

A pseudoline arrangement is **stretchable** if it is isomorphic to an arrangement of (pairwise non-parallel) lines. The first non-stretchable pseudoline arrangements were constructed by Levi [44], based on classical theorems of projective geometry; Ringel later refined Levi's construction to obtain the first

simple non-stretchable pseudoline arrangements [56]. Figure 2(a) shows one of Ringel's arrangements, based on Pappus's hexagon theorem.[1]

For further background on pseudoline arrangements and stretchability, see Felsner and Goodman [28], Björner *et al.* [9], Bokowski and Sturmfels [11], or Grünbaum [36].

2.3 The Existential Theory of the Reals

The *existential theory of the reals* (ETR) is the set of all true existentially quantified boolean combinations of multivariate polynomial inequalities. More formally, we can recursively define a (fully parenthesized) *sentence* in the existential theory of the reals (in *prenex form*) as follows:

- A *polynomial* is either the constant 0, the constant 1, a variable x_i, or an expression of the form $(Q + Q')$ or $(Q - Q')$ or $(Q \cdot Q')$ for some polynomials Q and Q'.

- An *atomic predicate* is a comparison of the form $(Q = 0)$ or $(Q < 0)$ for some polynomial Q.

- A *predicate* is either an atomic predicate or an expression of the form $(\neg\Phi)$ or $(\Phi \vee \Phi')$ or $(\Phi \wedge \Phi')$ for some predicates Φ and Φ'. The set of all real vectors $(x_1, x_2, \ldots, x_N) \in \mathbb{R}^N$ that satisfy such a predicate is called a *semi-algebraic set*.

- Finally, a *sentence* is an existentially quantified formula $\exists x_1, x_2, \ldots, x_N \colon \Phi(x_1, x_2, \ldots, x_N)$, where $\Phi(x_1, x_2, \ldots, x_N)$ is any predicate in which the only variables are $x_1, x_2, \ldots, x_N$.

ETR is the set of all such sentences that are true when all variables are quantified over the set of real numbers. Equivalently, ETR is the problem of deciding whether a given predicate describes a non-empty semi-algebraic set.

The precise computational complexity of ETR is unknown. The problem is easily seen to be NP-hard (see [4, p. 513], for example), and Canny proved that ETR $\in$ PSPACE [14]. The fastest algorithm known to solve ETR, due to Renegar [52], runs in time $(nd)^{O(N)}$, where n is the number of atomic predicates in the input sentence, d is the maximum degree of any polynomial, and N is the number of distinct real variables; see also Basu *et al.* [4].

The definition of ETR is actually quite flexible. We can freely include arbitrary bounded-degree algebraic real constants, subtraction, division, square roots, absolute values, more general comparisons ($\leq, \neq, >, \geq$), and more general boolean operators ($\Rightarrow, \Leftrightarrow, \oplus$, etc.); these can all be eliminated with only a small increase in the length of the sentence. On the other hand, ETR remains $\exists\mathbb{R}$-complete (see below) if we allow only strict inequalities, or only a single equation $\exists \vec{x} : Q(\vec{x}) = 0$, or only conjunctions of atomic predicates of the form $x_i < x_j$, $x_i + x_j = x_k$, and $x_i \cdot x_j = x_k$ [62].

2.4 $\exists\mathbb{R}$-hardness and Universality

Schaefer [57, 58] and Schaefer and Štefankovič [59] defined the complexity class $\exists\mathbb{R}$ as the set of all decision problems that can be (many-one) reduced to ETR in polynomial time. A problem is $\exists\mathbb{R}$-*hard* if there is a polynomial-time (many-one) reduction from ETR to that problem, and a decision problem is $\exists\mathbb{R}$-*complete* if it is both in $\exists\mathbb{R}$ and $\exists\mathbb{R}$-hard.[2]

A seminal result of Mnëv [47, 48] implies that deciding whether a given simple pseudoline arrangement is stretchable is $\exists\mathbb{R}$-complete; indeed, this was the first problem (other than ETR itself) to be

---

[1]This is in fact the smallest simple non-stretchable pseudoline arrangement. Every arrangement of at most eight pseudolines is stretchable [33], and up to isomorphism, there is exactly one simple non-stretchable arrangement of nine pseudolines [55].

[2]$\exists\mathbb{R}$ is also equal to the Blum-Shub-Smale complexity class $\{0, 1\}^* \cap NP^0_{\mathbb{R}}$, the set of boolean vectors recognizable in polynomial time by a non-deterministic BSS machine over the reals [10] whose only constants are 0 and 1 [13, 57].

identified as ∃ℝ-complete. In fact, Mnëv proved a significantly stronger *universality* result. The *realization space* of an arrangement of n pseudolines is the set of all isomorphic line arrangements, each encoded as a vector in $\mathbb{R}^{2n}$. Mnëv proved that for any (open) semialgebraic set V, there is a (simple) pseudoline arrangement whose realization space is homotopy equivalent to V. Mnëv's proof relies on von Staudt's *algebra of throws* [63], which implements arithmetic using projective geometric constructions. More detailed proofs of Mnëv's results were later given by Shor [62] and Richter-Gebert [53, 54].

Mnëv's universality theorem has since been used to prove the ∃ℝ-hardness of several other geometric and topological problems. Examples include realizability of abstract 4-polytopes [54], cross-product systems [38], Delaunay triangulations [2, 50], and planar linkages [40, 58]; recognition of several types of graphs [12, 16, 17, 42, 45, 46, 58]; several problems in graph drawing [6, 16, 18, 20, 27, 39, 57]; several problems related to fixed points and Nash equilibria [7, 8, 29, 59]; problems in social choice theory [51]; optimal polytope nesting [24, 61]; nonnegative matrix factorization [22, 24, 61]; minimum rank of a matrix with a given sign pattern [3, 5]; real tensor rank [60]; and the art gallery problem [1]. (Some of these papers claim only NP-hardness but describe reductions from pseudoline stretchability that imply ∃ℝ-hardness; others make no claims about computational complexity but derive appropriate universality theorems that imply ∃ℝ-hardness.[3])

For more examples and further background on ∃ℝ-hardness, we refer the reader to surveys by Schaefer [57], Matoušek [45], and Cardinal [15].

3 Straightening Curves is Hard

We consider the decision problem CURVETOPOLYGON, which asks, given a nonnegative integer m and (any reasonable representation of) a generic closed curve γ in the plane, whether any polygon with m vertices is isotopic to γ.

To prove that CURVETOPOLYGON is ∃ℝ-hard, we describe a straightforward polynomial-time reduction from deciding stretchability of simple pseudoline arrangements. Given a simple arrangement of n pseudolines $\Psi = \{\psi_1, \psi_2, \ldots, \psi_n\}$, we construct a corresponding generic closed curve γ_Ψ that is isotopic to a $4n$-gon if and only if Ψ is stretchable. Figure 1 gives a pictorial summary of our reduction.

We assume without loss of generality that the curves in Ψ constitute a *wiring diagram*, as defined by Goodman [32]. That is, the curves are piecewise-linear, and each curve is horizontal except in a small neighborhood of each crossing. We assume the pseudolines are indexed from top to bottom to the left of all crossings, and therefore from bottom to top to the right of all crossings, as shown in Figure 3.
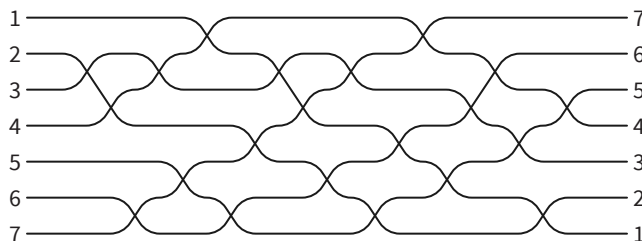

Figure 3. A "wiring diagram" pseudoline arrangement.

To avoid parity issues, we also assume without loss of generality that n is odd. Otherwise, consider the set $\Psi' = \Psi \cup \{\psi_0\}$, where ψ_0 is a new "wire" that passes above all $\binom{n}{2}$ intersection points of Ψ and

---

[3]Mnëv's universality theorem also implies several other universality results in algebraic geometry, collectively dubbed "Murphy's Law" by Vakim [43, 67].

then downward across all n wires in Ψ to the right of those intersection points. We easily observe that the arrangement of Ψ' is simple, and that Ψ' is stretchable if and only if Ψ is stretchable.

We derive our closed curve γ_Ψ from the wiring diagram Ψ as follows. Truncate the pseudolines to an axis-aligned rectangle R that contains all crossings in its interior. For each index i, let u_i and v_i denote the left and right endpoints of the curves $\Psi_i \cap R$. We connect the truncated pseudolines using n pairwise-disjoint paths outside R, each containing two simple outward-facing loops and no other self-intersections, and each connecting two adjacent endpoints on the boundary of R, as shown in Figure 4. Specifically, for each index i, we connect the left endpoints u_{2i} and u_{2i+1} to the left of R, and we connect the right endpoints v_{2i-1} and v_{2i} to the right of R; finally, we connect the top endpoints u_1 and v_{2n} above R. Concatenating these n exterior paths and the n truncated pseudolines yields the closed curve γ_Ψ. We collectively refer to the $2n$ exterior loops as the *fringe* of γ_Ψ.
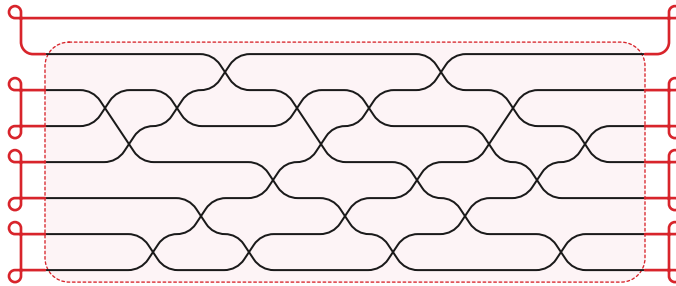
Figure 4. The closed curve γ_Ψ derived from the pseudoline arrangement Ψ in Figure 3.

Lemma 3.1. *Let Ψ be any set of n pseudolines, for any odd integer n, whose arrangement is simple. Any polygon isotopic to the closed curve γ_Ψ has at least $4n$ vertices.*

Proof: Let P be any polygon isotopic to γ_Ψ, and let $H: \mathbb{R}^2 \to \mathbb{R}^2$ be any homeomorphism such that $P = H \circ \gamma_\Psi$. Consider any simple loop ℓ in the fringe of γ_Ψ. The corresponding path $L = H \circ \ell$ is a simple loop of P, which starts and ends at the unique intersection point of two edges e and e' of P. These two edges are distinct, and because they cross, they do not share an endpoint. Thus, L contains at least two vertices of P, specifically, one endpoint of e and one endpoint of e'.
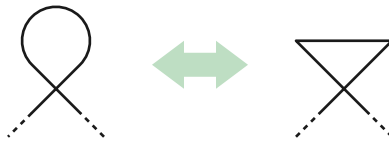
Figure 5. Polygonalizing a simple loop requires at least two vertices

It follows immediately that P contains at least two vertices for each of the $2n$ loops in the fringe of γ_Ψ, and thus at least $4n$ vertices in total. □

Lemma 3.2. *Let Ψ be any set of n pseudolines, for any odd integer n, whose arrangement is simple. The curve γ_Ψ is isotopic to a $4n$-gon if and only if the arrangement of Ψ is stretchable.*

Proof: First, suppose P is a $4n$-gon isotopic to γ_Ψ. Recall that R is an axis-aligned rectangle containing all the vertices of the arrangement of Ψ in its interior. Let $H: \mathbb{R}^2 \to \mathbb{R}^2$ be any homeomorphism such that $P = H \circ \gamma_\Psi$, and let $D = H(R)$ be the image of R under this homeomorphism. The proof of Lemma 3.1 implies that D does not contain any vertex of P. Thus, the intersection $P \cap D$ is an arrangement of n line

segments, each pair of which crosses exactly once. By construction, the homeomorphism H maps the arrangement of arcs $\Psi \cap B^2$ to the arrangement of segments $P \cap D$. Extending these segments to lines does not introduce any new intersection points or change the order of intersections along any curve. Thus, the resulting line arrangement is isomorphic to the arrangement of Ψ; in short, Ψ is stretchable.

On the other hand, suppose L is a set of lines in the plane whose arrangement is isomorphic to the arrangement of Ψ. Let $H\colon \mathbb{R}^2 \to \mathbb{R}^2$ be any homeomorphism that maps the arrangement of Ψ to the arrangement of L, and for each index i, let l_i denote the line $H \circ \psi_i$. Without loss of generality, the lines in L are indexed by increasing slope, each line in L has slope strictly between -1 and 1, each pair of lines in L intersect strictly between the vertical lines $x = 1$ and $x = 1$.

For each index i, let p_i and q_i denote the intersection of l_i with the lines $x = -5/4$ and $x = 5/4$. We can construct a $4n$-gon P_L isotopic to γ_Ψ by connecting the segments $p_i q_i$ with gadgets we call *staples*, each consisting of a segment of length ε and slope 1, followed by a horizontal or vertical segment, followed by a segment of length ε and slope -1, for some suitable real $\varepsilon > 0$ to be chosen later. Specifically, we connect p_1 to q_n with a horizontal staple, and for each index i, we connect the left endpoints p_{2i-1} and p_i and the right endpoints q_{2i} and q_{2i+1} with vertical staples, as shown in Figure 6. Each staple crosses the two lines that it connects, creating two simple loops. For sufficiently small ε, the staples are pairwise disjoint, and so the resulting $4n$-gon P_L is isotopic to γ_Ψ. □
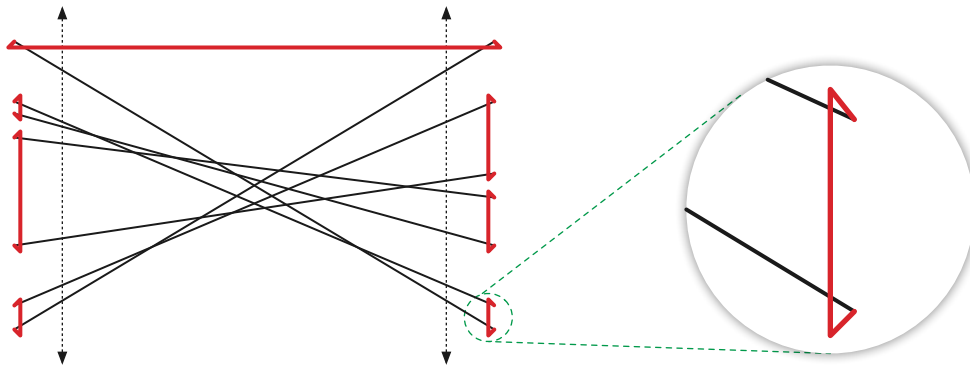

Figure 6. Stapling a line arrangement into a polygon; compare with Figure 4.

Given any set Ψ of pseudolines, we can easily construct the curve γ_Ψ in polynomial time. Thus, Lemma 3.2 immediately implies the following:

Theorem 3.3. CurveToPolygon *is* ∃ℝ-*hard.*

In fact, our reduction implies much stronger result, as a natural corollary of Mnëv's universality theorem for pseudoline arrangements [47, 48, 62]. The *coordinate vector* of a polygon P is a vector $(x_1, y_1, x_2, y_2, \ldots, x_m, y_m) \in \mathbb{R}^{2m}$, where $(x_1, y_1), (x_2, y_2), \ldots, (x_m, y_m)$ are the coordinates of P's vertices in cyclic order. Two polygons (or their coordinate vectors) are *equivalent* if they are equivalent as closed curves, or equivalently, if they differ only by a rotation or reflection of vertex indices. Thus, every generic m-gon is equivalent to exactly $2m - 1$ other m-gons. Finally, for any closed curve γ and any integer m, we define the *isotopy realization space* $\Sigma(\gamma, m) \subseteq \mathbb{R}^{2m}/D_m$ to be the set of all equivalence classes of coordinate vectors of m-gons that are isotopic to γ. (Here D_m is the dihedral group of permutations acting on the m vertex indices.)

Theorem 3.4. *Every semialgebraic set is homotopy equivalent to the isotopy realization space $\Sigma(\gamma, m)$ of some closed curve γ and some integer m. Moreover, every **open** semialgebraic set is homotopy equivalent to the isotopy realization space $\Sigma(\gamma, m)$ of some **generic** closed curve γ and some integer m.*

Every generic m-gon is isotopic to a generic m-gon with integer coordinates. In light of our reduction, results of Goodman, Pollack, and Sturmfels [34, 35] on the intrinsic spread of order types imply that some isotopy classes of polygons require an exponential number of bits to describe. (McDiarmid and Müller [46] derive a similar lower bound for intersection graphs of segments of unit disks in the plane.)

Corollary 3.5. *For every integer m, there is an m-gon P such that every integer m-gon isotopic to P has diameter $2^{2^{\Omega(m)}}$.*

4 ...But Not Too Hard

Reducing our curve-straightening problem to ETR is somewhat more subtle. Intuitively, the reduction appears straightforward, because there is a simple efficient algorithm to verify that an n-crossing curve γ and an m-vertex polygon P are isotopic. First, compute all self-intersection points of P, and their order along each edge of P. If P has either more or less than n self-intersections, or if any self-intersection of P is degenerate (at a vertex, or on more than two edges), fail immediately. Otherwise, fix an arbitrary basepoint $\gamma(0)$ on γ, and then for each vertex p_i of P, check whether corresponding crossings appear in the same order and orientation along both curves, starting at p_i on P and starting at $\gamma(0)$ on γ.

But an efficient verification algorithm is insufficient; we require an efficient formula in the existential theory of the reals. That is, we need to argue that the realization space of all m-gons isotopic to a given curve γ is (the projection of) a semialgebraic set with description complexity polynomial in n.

4.1 Crossing Codes

Our reduction assumes that the closed curve γ is represented using a variant of encodings proposed by Gauss [30, 31], Tait [64], and Dowker and Thistlethwaite [25].[4] Fix an arbitrary parametrization $\vec{\gamma}$ of γ, such that the basepoint $\gamma(0)$ is not a self-intersection point. Call $t \in S^1$ a *crossing value* for $\vec{\gamma}$ if $\vec{\gamma}(t) = \vec{\gamma}(t')$ for some $t' \neq t$ (which must also be a crossing value). Let $0 < t_1 < t_2 < \cdots < t_{2n} < 1$ denote the sorted sequence of crossing values. The **signed crossing code** of $\vec{\gamma}$ consists of a pair of vectors:

- The first vector **twin** $\in [2n]^{2n}$ is an involution of $[2n]$ that describes the pairing of crossing values; $\vec{\gamma}(t_i) = \vec{\gamma}(t_j)$ and $twin_j = i$ if and only if $j = twin_i$.

- The second vector **sign** $\in \{-1, +1\}^{2n}$ indicates the direction of each crossing; $sign_i = +1$ if and only if, for arbitrarily small ε, the subpath $\vec{\gamma}[t_i - \varepsilon, t_i + \varepsilon]$ crosses the subpath $\vec{\gamma}[t_{twin_i} - \varepsilon, t_{twin_i} + \varepsilon]$ from right to left.

Each (unparametrized) closed curve γ has up to $4n$ different signed crossing codes, depending on the choice of basepoint and direction.

A signed crossing code of a curve determines its embedding *on the sphere* [19], but not necessarily *in the plane*, because the code does not specify an outer face. To remove this ambiguity, we declare that the outer face lies immediately to the right of the curve near the basepoint. Thus, to encode a planar curve, we must first choose a basepoint on the outer face and a direction that puts the outer face to the right of that basepoint. Figure 7 shows a closed curve whose image graph has three arcs on the outer face, each leading to a different signed crossing code.

---

[4]Specifically, Gauss [30, 31] encoded curves by assigning each crossing point a unique label, and then listing the labels in order along the curve, either with or without the corresponding sequence of signs. Tait [64] observed that each label in any Gauss code appears once at an even position and once in an odd position, the labels at (say) even positions can be discarded with no loss of information. Dowker and Thistlethwaite observed that the crossing involution we are calling *twin* is completely determined by the subsequence $(twin_1, twin_3, \ldots, twin_{2n-1})$, because $twin_i$ is even if and only if i is odd [25].
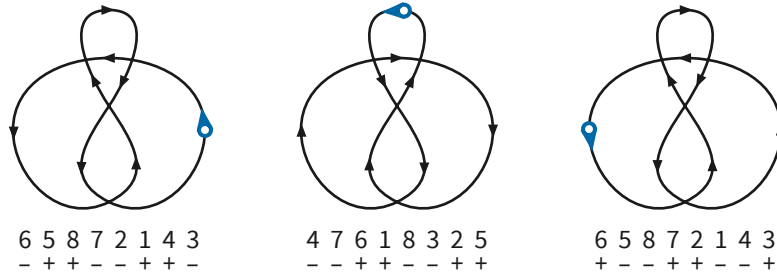
Figure 7. A closed curve with exactly three valid signed crossing codes; compare with Figure 8.

In particular, to obtain a valid signed crossing sequence for a generic polygon P, it suffices to choose any vertex on the convex hull of P as the basepoint vertex p_1, and direct the polygon so that the triple (p_m, p_1, p_2) is oriented counterclockwise. (Different convex hull vertices may yield the same crossing code, and some valid crossing codes for P may require a basepoint that is not a convex hull vertex, or not a vertex at all.)

We also associate with every generic polygon P an **edge code**, which records which edges of P contain each self-crossing. The edge code is a vector $\textbf{\textit{edge}} \in [m]^{2n}$, where for each index i, the ith crossing along P occurs at the point $p_e p_{e+1} \cap p_{e'} p_{e'+1}$, where $e = edge_i$ and $e' = edge_{twin_i}$. The vector $\textit{edge}$ is (weakly) sorted; for each index $1 \le i < 2n$, we have $edge_i \le edge_{i+1}$.

Figure 8 shows a generic polygon P whose convex hull has four vertices; along with the signed crossing and edge codes that result from choosing each convex hull vertex as the basepoint. Each signed crossing code is also a valid signed crossing code for the curve γ in Figure 7 (but not vice versa); it follows that γ and P are isotopic.
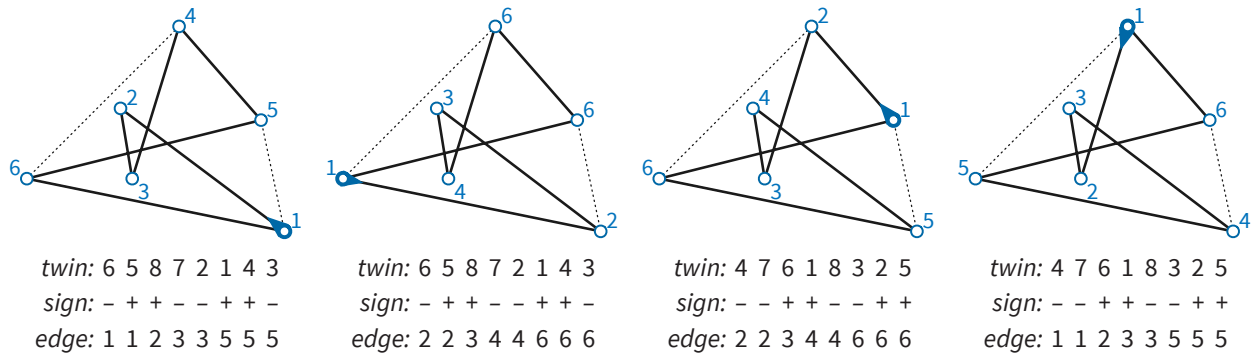
Figure 8. All valid signed crossing and edge codes for the same generic polygon; compare with Figure 7.

4.2 Overview

For the remainder of this section, we use modular index arithmetic implicitly; in particular, $p_{m+1} = p_1$. All indexed conjunctions ($\bigwedge$), disjunctions ($\bigvee$), and summations ($\sum$) are notational shorthand; each term in these expressions appears explicitly in the actual formula. For example, the indexed disjunction

$$\bigwedge_{i=1}^{2n-1} (edge_i \le edge_{i+1})$$

is notational shorthand for the explicit disjunction

$$(edge_1 \le edge_2) \wedge (edge_2 \le edge_3) \wedge \cdots \wedge (edge_{2n-1} \le edge_{2n})$$

For any closed curve γ and any integer m, we describe an ETR sentence ISOTOPICTOPOLYGON(γ, m) that is true if and only if some generic m-gon is isotopic to γ. Lemma 2.1 implies that we can assume without loss of generality that $m = O(n)$. Our sentence has the following high-level structure:

$$\text{ISOTOPICTOPOLYGON}(\gamma, m) \equiv \exists P \in \mathbb{R}^{2m} : \exists \, edge \in \mathbb{R}^{2n} :$$
$$\bigvee_{(twin, sign) \sim \gamma} \text{CODEDPOLYGON}(P, twin, sign, edge)$$

Here $(twin, sign)$ ranges over all valid signed crossing codes of the input curve γ; in each term of this disjunction, the vectors $twin$ and $sign$ are hard-coded constants. The polygon P is represented by its vector $(x_1, y_1, x_2, y_2, \ldots, x_m, y_m)$ of vertex coordinates; for notational convenience, we write $p_i = (x_i, y_i)$. Crudely, the predicate CODEDPOLYGON($P, twin, sign, edge$) states that P is a generic polygon with valid signed crossing code $(twin, sign)$ and edge code $edge$. Although in principle we could compute the edge code of P from its vertex coordinates, we use the inherent nondeterminism in the ETR model to "guess" the edge code, which simplifies our formula.

In more detail, the predicate CODEDPOLYGON($P, twin, sign, edge$) certifies the following conditions:

- P is a generic polygon; p_1 is the rightmost vertex of P, and the triple (p_m, p_1, p_2) is oriented counterclockwise.

- The vertex coordinates of P describe a polygon with exactly n self-crossings.

- The vector $edge$ is a well-formed edge code.

- The crossings occur on the correct edges, with the correct signs; specifically, for each index i, the edges with index $edge_i$ and $edge_{twin_i}$ cross with sign $sign_i$.

- Finally, the crossings along each edge of P occur in the correct order.

We describe subformulas for each of these conditions in the following subsections.

$$\text{CODEDPOLYGON}(P, twin, sign, edge) \equiv$$
$$GoodPolygon(P) \wedge (NumCrossings(P) = n) \wedge WellFormed(edge) \wedge$$
$$CrossingSigns(P, twin, sign, edge) \wedge CrossingOrder(P, twin, edge)$$

4.3 Good Polygon

To avoid boundary cases elsewhere in our formula, we require that (1) no two vertices of P have the same x-coordinate, (2) no three vertices of P are collinear, (3) no two edges of P are parallel, and (4) no three edges of P lie on concurrent lines. These conditions, which can be enforced if necessary by arbitrarily small vertex perturbations, imply that all m vertices of P are distinct; no edge of P is vertical; no two edges of P are collinear; no vertex of P lies in the interior of an edge; and at most two edges intersect at any point. To establish a valid signed crossing code, we also require that p_1 is the rightmost vertex of P, and the triple (p_m, p_1, p_2) is oriented counterclockwise.

The signed area of the triangle (p_i, p_j, p_k) is given by the following determinant:

$$\Delta(i, j, k) \equiv \begin{vmatrix} 1 & x_i & y_i \\ 1 & x_j & y_j \\ 1 & x_k & y_k \end{vmatrix} \equiv (x_i - x_k)(y_j - y_k) - (y_i - y_k)(x_j - x_k)$$

This signed area is positive if the triple (p_i, p_j, p_k) is oriented counterclockwise, negative if the triple is oriented clockwise, and zero if the three points are collinear. Edges $p_i p_{i+1}$ and $p_j p_{j+1}$ lie on parallel

lines (or one of the edges has length 0) if and only if

$$parallel(i,j) \equiv \begin{vmatrix} x_{i+1}-x_i & y_{i+1}-y_i \\ x_{j+1}-x_j & y_{j+1}-y_j \end{vmatrix} = 0.$$

Similarly, edges $p_i p_{i+1}$ and $p_j p_{j+1}$ and $p_k p_{k+1}$ lie on concurrent lines (or some pair is parallel) if and only if

$$concurrent(i,j,k) \equiv \begin{vmatrix} x_{i+1}-x_i & y_{i+1}-y_i & x_{i+1}y_i-x_iy_{i+1} \\ x_{j+1}-x_j & y_{j+1}-y_j & x_{j+1}y_j-x_jy_{j+1} \\ x_{k+1}-x_j & y_{k+1}-y_j & x_{k+1}y_j-x_ky_{k+1} \end{vmatrix} = 0.$$

Our overall general position expression has length $\Theta(m^3)$:

$$GoodPolygon(P) \equiv \left(\bigwedge_{i=1}^{m}(x_i \neq x_{i+1})\right) \wedge \left(\bigwedge_{i=2}^{m}(x_1 \leq x_i)\right) \wedge \left(\Delta(m,1,2) > 0\right)$$

$$\wedge \left(\bigwedge_{i=1}^{m-1}\bigwedge_{j=i+1}^{m} \neg parallel(i,j)\right)$$

$$\wedge \left(\bigwedge_{i=1}^{m-2}\bigwedge_{j=i+1}^{m-1}\bigwedge_{k=j+1}^{m} \left((\Delta(i,j,k) \neq 0) \wedge \neg concurrent(i,j,k)\right)\right)$$

4.4 Number of Crossings

We can determine whether two edges $p_i p_{i+1}$ and $p_j p_{j+1}$ cross in their interiors by checking the orientations of all four triples of endpoints:

$$cross(i,j) \equiv \left(\Delta(p_i,p_j,p_{j+1}) \cdot \Delta(p_{i+1},p_j,p_{j+1}) < 0\right)$$
$$\wedge \left(\Delta(p_i,p_{i+1},p_j) \cdot \Delta(p_i,p_{i+1},p_{j+1}) < 0\right)$$

We count crossing edge pairs by defining new indicator variables X_{ij} that record which pairs of edges cross, and then summing these indicator variables by brute force. (These new variables are existentially quantified at the start of the top-level sentence.)

$$\left(NumCrossings(P) = n\right) \equiv \left(\bigwedge_{i=1}^{m-2}\bigwedge_{j=i+2}^{m} \left((X_{ij}=1) \wedge cross(i,j)\right) \vee \left((X_{ij}=0) \wedge \neg cross(i,j)\right)\right)$$

$$\wedge \left(\sum_{i=1}^{m-2}\sum_{j=i+2}^{m} X_{ij} = n\right)$$

4.5 Well-Formed Edge Code

The vector *edge* is a well-formed edge code if and only if its is monotonically non-decreasing and each coordinate $edge_i$ is an integer between 1 and m.

$$WellFormed(edge) \equiv \left(\bigwedge_{i=1}^{2n-1}(edge_i \leq edge_{i+1})\right) \wedge \left(\bigwedge_{i=1}^{2n}\bigvee_{e=1}^{m}(edge_i = e)\right)$$
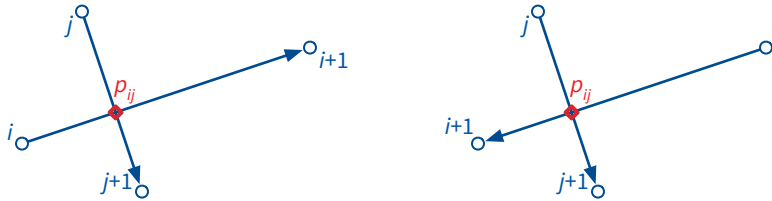
4.6   Crossing Signs

The following refinement to our earlier crossing primitive states that the directed edge $p_i \rightarrow p_{i+1}$ crosses the directed edge $p_j \rightarrow p_{j+1}$ from right to left, defining a *positive* crossing:

$$cross^+(i,j) \equiv \big(\Delta(p_i,p_j,p_{j+1}) < 0\big) \wedge \big(\Delta(p_{i+1},p_j,p_{j+1}) > 0\big)$$
$$\wedge \big(\Delta(p_i,p_{i+1},p_j) > 0\big) \wedge \big(\Delta(p_i,p_{i+1},p_{j+1}) < 0\big)$$

We can indicate negative crossings similarly:

$$cross^-(i,j) \equiv \big(\Delta(p_i,p_j,p_{j+1}) > 0\big) \wedge \big(\Delta(p_{i+1},p_j,p_{j+1}) < 0\big)$$
$$\wedge \big(\Delta(p_i,p_{i+1},p_j) < 0\big) \wedge \big(\Delta(p_i,p_{i+1},p_{j+1}) > 0\big)$$

We easily verify that $cross^-(i,j) = cross^+(j,i)$. The identity $cross(i,j) = cross^+(i,j) \vee cross^-(i,j)$ follows from the fact that if two segments cross, their endpoints alternate around their convex hull.

Figure 9. Left: $cross^+(i,j) = cross^-(j,i)$. Right: $cross^-(i,j) = cross^+(j,i)$

At this point we would like to define

$$CrossingSigns(P, twin, sign, edge) \equiv \bigwedge_{i=1}^{2n} \Big(\big((sign_i = +1) \wedge cross^+(edge_i, edge_{twin_i})\big)$$
$$\vee \big((sign_i = -1) \wedge cross^-(edge_i, edge_{twin_i})\big)\Big).$$

Unfortunately, even after expanding this expression into an explicit conjunction of $2n$ clauses, we are using the *variables* $edge_i$ and $edge_{twin_i}$ in the ith clause as *indices* inside the functions $cross^+$ and $cross^-$. To simulate this indirection, we expand the offending subexpressions into disjunctions as follows:

$$cross^\pm(edge_i, edge_{twin_i}) \equiv \bigvee_{e=1}^{m} \bigvee_{et=1}^{m} \big((e = edge_i) \wedge (et = edge_{twin_i}) \wedge cross^\pm(e, et)\big)$$

Again, this indexed disjunction is shorthand for an explicit disjunction of $\Theta(m^2)$ terms; within each term

$$\big((e = edge_i) \wedge (et = edge_{twin_i}) \wedge cross^\pm(e, et)\big)$$

the crossing indices i and $twin_i$ and the edge indices e and et are explicit constants, so all variables, including the coordinate variables in the subexpression $cross^+(e, et)$, are identifiable at "compile time".

Our final expression $CrossingSigns(P, twin, sign, edge)$ has length $\Theta(m^2 n)$. As a side effect, this expression guarantees that every pair of polygon edges that is supposed to cross actually does. Moreover, because we have already verified that the correct *number* of edge-pairs cross, it follows that *only* the edge pairs that are supposed to cross actually do.

4.7 Crossing Order

Finally, we need to ensure that the crossings along each edge of P appear in the correct order. For all indices $i \neq j$, let $p_{ij} = (x_{ij}, y_{ij})$ denote the intersection point of the lines $\overleftrightarrow{p_i p_{i+1}}$ and $\overleftrightarrow{p_j p_{j+1}}$; our general position assumptions imply that these points are well-defined. Of course we could compute these points explicitly, but it is simpler to define them implicitly by the clauses

$$(\Delta(p_{ij}, p_i, p_{i+1}) = 0) \wedge (\Delta(p_{ij}, p_j, p_{j+1}) = 0).$$

For any three indices i, j, and k, the following expression states that the endpoint p_i and the intersection points p_{ij} and p_{ik} appear in order along the line $\overleftrightarrow{p_i p_{i+1}}$:

$$OrderedP(i, j, k) \equiv \begin{cases} \text{FALSE} & \text{if } i = j \text{ or } i = k \\ (x_i - x_{ij})(x_{ij} - x_{ik}) > 0 & \text{otherwise} \end{cases}$$

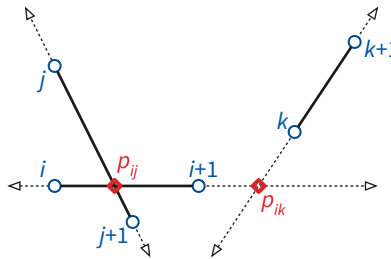(Recall that our general position assumption implies that $x_i \neq x_{i+1}$.)

Figure 10. *OrderedP(i, j, k)*

We need to express this relation in terms of *crossing* indices, instead of *edge* indices. The following expression states that the ith and $(i+1)$st crossings are correctly ordered, unless they appear on different edges.

$$OrderedX(i) \equiv (edge_i \neq edge_{i+1}) \vee OrderedP(edge_i, edge_{twin_i}, edge_{twin_{i+1}})$$

Our earlier crossing-sign expressions guarantee that the relevant *edges* actually intersect, and therefore intersect in the correct order. As before, we expand this expression to simulate indirection, as follows.

$$OrderedP(edge_i, edge_{twin_i}, edge_{twin_{i+1}})$$
$$\equiv \bigvee_{e=1}^{m} \bigvee_{et=1}^{m} \bigvee_{et'=1}^{m} \left((e = edge_i) \wedge (et = edge_{twin_i}) \wedge (et' = edge_{twin_{i+1}}) \wedge OrderedP(e, et, et') \right)$$

Our final ordering expression

$$CrossingOrder(P, twin, edge) \equiv \bigvee_{i=1}^{2n-1} OrderedX(i)$$

has length $\Theta(nm^3)$.

Theorem 4.1. *For any generic curve γ with n self-intersection points and any integer $m = O(n)$, there is a sentence of length $\Theta(nm^3) = O(n^4)$ in the existential theory of the reals that is true if and only if γ is isotopic to a generic polygon with m vertices.*

Corollary 4.2. CurveToPolygon *is $\exists\mathbb{R}$-complete.*

References

[1] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is ∃ℝ-complete. *Proc. 50th Ann. ACM Symp. Theory Comput.*, 65–73, 2018. arXiv:1704.06969.

[2] Karim A. Adiprasito, Arnau Padrol, and Louis Theran. Universality theorems for inscribed polytopes and Delaunay triangulations. *Discrete Comput. Geom.* 54(2):412–431, 2015. arXiv:406.7831.

[3] Ronen Basri, Pedro F. Felzenszwalb, Ross B. Girshick, David W. Jacobs, and Caroline J. Klivans. Visibility constraints on features of 3D objects. *Proc. 2009 IEEE Conf. Comput. Vision Pattern Recog.*, 1231–1238, 209.

[4] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*, 2nd edition. Algorithms and Computation in Mathematics 10. Springer-Verlag, 2006.

[5] Amey Bhangale and Swastik Kopparty. The complexity of computing the minimum rank of a sign pattern matrix. Preprint, May 2015. arXiv:1503.04486v2.

[6] Daniel Bienstock. Some provably hard crossing number problems. *Discrete Comput. Geom.* 6(3):443–459, 1991.

[7] Vittorio Bilò and Marios Mavronicolas. A catalog of ∃ℝ-complete decision problems about Nash equilibria in multi-player games. *Proc. 33rd Symp. Theor. Aspects Comput. Sci.*, 17:1–17:13, 2016. Leibniz Int. Proc. Informatics 47, Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[8] Vittorio Bilò and Marios Mavronicolas. ∃ℝ-complete decision problems about symmetric Nash equilibria in symmetric multi-player games. *Proc. 34rd Symp. Theor. Aspects Comput. Sci.*, 13:1–13:14, 2017. Leibniz Int. Proc. Informatics 66, Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[9] Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler. *Oriented Matroids*, second edition. Encyclopedia of Mathematics and its Applications 46. Cambridge Univ. Press, 2000.

[10] Lenore Blum, Felipe Cucker, Michael Shub, and Steve Smale. *Complexity and Real Computation*. Springer, 1998.

[11] Jürgen Bokowski and Bernd Sturmfels. *Computational Synthetic Geometry*. Lecture Notes Math. 1355. Springer-Verlag, 1989.

[12] Prosenjit Bose, Hazel Everett, and Stephen Wismath. Properties of arrangement graphs. *Int. J. Comput. Geom. Appl.* 13(6):447–462, 2003.

[13] Peter Bürgisser and Felipe Cucker. Counting complexity classes for numeric computations ii: Algebraic and semialgebraic sets. *J. Complexity* 22(2):147–191, 2006.

[14] John Canny. Some algebraic and geometric computations in PSPACE. *Proc. 20th ACM Sympos. Theory. Comput.*, 460–467, 1988. ⟨https://www2.eecs.berkeley.edu/Pubs/TechRpts/1988/6041.html⟩.

[15] Jean Cardinal. Computational geometry column 62. *SIGACT News* 46(4):69–78, 2015.

[16] Jean Cardinal, Stefan Felsner, Tillmann Miltzow, Casey Tompkins, and Birgit Vogtenhuber. Intersection graphs of rays and grounded segments. *J. Graph Algorithms Appl.* 22(2):273–295, 2018. arXiv:1612.03638.

[17] Jean Cardinal and Udo Hoffmann. Recognition and complexity of point visibility graphs. *Discrete Comput. Geom.* 57(1):164–178, 2017. arXiv:1503.07082.

[18] Jean Cardinal and Vincent Kusters. The complexity of simultaneous geometric graph embedding. *J. Graph Algorithms Appl.* 19(1):259–272, 2015. arXiv:1302.7127.

[19] J. Scott Carter. Classifying immersed curves. *Proc. Amer. Math. Soc.* 111(1):281–287, 1991.

[20] Steven Chaplick, Krzysztof Fleszar, Fabian Lipp, Alexander Ravsky, Oleg Verbitsky, and Alexander Wolff. The complexity of drawing graphs on few lines and few planes. *Proc. 15th Algorithms Data Struct. Symp. (WADS)*, 265–276, 2017. Lecture Notes Comput. Sci. 10389, Springer-Verlag. arXiv:1607.06444.

[21] Jianer Chen. Algorithmic graph embeddings. *Theoret. Comput. Sci.* 161(2):247–266, 1997.

[22] Joel E. Cohen and Uriel G. Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Alg. Appl.* 190:149–168, 1993.

[23] Sabine Cornelson and Andreas Karrenbauer. Accelerated bend minimization. *J. Graph Algorithms Appl.* 16(3):635–650, 2012.

[24] Michael G. Dobbins, Andreas Holmsen, and Tillman Miltzow. A universality theorem for nested polytopes. Preprint, August 2019. arXiv:1908.02213.

[25] Clifford H. Dowker and Morwen B. Thistlethwaite. Classification of knot projections. *Topology Appl.* 16(1):19–31, 1983.

[26] Vida Dujmović, David Eppstein, Matthew Suderman, and David R. Wood. Drawings of planar graphs with few slopes and segments. *Comput. Geom. Theory Appl.* 38(3):194–212, 2007. arXiv:math.CO/0606450.

[27] Stephane Durocher, Debajyoti Mondal, Rahnuma Islam Nishat, and Sue Whitesides. A note on minimum-segment drawings of planar graphs. *J. Graph Algorithms Appl.* 17(3):301–328, 2013.

[28] Stefan Felsner and Jacob E. Goodman. Psuedoline arrangements. *CRC Handbook of Discrete and Computational Geometry*, chapter 5, 125–157, 2017. CRC Press. ⟨https://www.csun.edu/~ctoth/Handbook/HDCG3.html⟩.

[29] Jugal Garg, Ruta Mehta, Vijay V. Vazirani, and Sadra Yazdanbod. $\exists\mathbb{R}$-completness for decision version of multi-player (symmetric) Nash equilibria. *ACM Trans. Econ. Comput.* 6(1):1:1–1:23, 2018.

[30] Carl Friedrich Gauß. Nachlass. I. Zur Geometria situs. *Werke*, vol. 8, 271–281, 1900. Teubner. Originally written between 1823 and 1840.

[31] Carl Friedrich Gauß. Nachlass. II. Zur Geometrie der Lage, für zwei Raumdimensionen. *Werke*, vol. 8, 282–285, 1900. Teubner.

[32] Jacob E. Goodman. Proof of a conjecture of Burr, Grünbaum, and Sloane. *Discrete Math.* 32:27–35, 1980.

[33] Jacob E. Goodman and Richard Pollack. Proof of Grünbaum's conjecture on the stretchability of certain arrangements of pseudolines. *J. Comb. Theory Ser. A* 29(3):385–390, 1980.

[34] Jacob E. Goodman, Richard Pollack, and Bernd Sturmfels. Coordinate representation of order types requires exponential storage. *Proc. 21st Ann. ACM Symp. Theory Comput.*, 405–410, 1989.

[35] Jacob E. Goodman, Richard Pollack, and Bernd Sturmfels. The intrinsic spread of a configuration in $\mathbf{R}^d$. *J. Amer. Math. Soc.* 3(3):639–651, 1990.

[36] Branko Grünbaum. *Arrangements and Spreads*. Regional Conf. Ser. Math. 10. American Mathematical Society, 1972.

[37] Leonidas J. Guibas and Jorge Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graphics* 4(2):75–123, 1985.

[38] Christian Herrmann, Johanna Sokoli, and Martin Ziegler. Satisfiability of cross product terms is complete for real nondeterministic polytime Blum-Shub-Smale machines. *Proc. 6th Conf. Mach. Comput. Universality*, 85–92, 2013. Electron. Proc. Theoret. Comput. Sci. 128. arXiv:1309.1270.

[39] Udo Hoffmann. On the complexity of the planar slope number problem. *J. Graph Algorithms Appl.* 21(2):183–193, 2017.

[40] Michael Kapovich and John J. Millson. Universality theorems for configuration spaces of planar linkages. *Topology* 41(6):1051–107, 2002.

[41] Adam Karczmarz and Piotr Sankowski. Min-cost flow in unit-capacity planar graphs. Preprint, July 2019. arXiv:1907.02274.

[42] Jan Kratochvíl and Jiří Matoušek. Intersection graphs of segments. *J. Comb. Theory Ser. B* 62(2):289–315, 1994.

[43] Seok Heyong Lee and Ravi Vakil. Mnëv-Sturmfels universality for schemes. *A Celebration of Algebraic Geometry*, 457–468, 2013. Clay Math. Proc. 18, Amer. Math. Soc. arXiv:1202.3934.

[44] Friedrich Levi. Die teilung der projektiven ebene durch gerade und pseudogerade. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.* 78:256–267, 1926.

[45] Jiří Matoušek. Intersection graphs of segments and ∃ℝ. Preprint, June 2014. arXiv:1406.2636.

[46] Colin McDiarmid and Tobias Müller. Integer realizations of disk and segment graphs. *J. Comb. Theory Ser. B* 103(1):114–143, 2013. arXiv:1111.2931.

[47] Nikolai E. Mnëv. Varieties of combinatorial types of projective configurations and convex polyhedra. *Dokl. Akad. Nauk SSSR* 283(6):1312–1314, 1985. in Russian.

[48] Nikolai E. Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. *Topology and geometry—Rohlin Seminar*, 527–543, 1988. Lecture Notes Math. 1346, Springer.

[49] David E. Muller and Franco P. Preparata. Finding the intersection of two convex polyhedra. *Theoret. Comput. Sci.* 7:217–236, 1978.

[50] Arnau Padrol and Louis Theran. Delaunay triangulations with disconnected realization spaces. *Proc. 30th Ann. Symp. Comput. Geom.*, 163–170, 2014.

[51] Dominik Peters. Recognising multidimensional euclidean preferences. *Proc. 31st AAAI Conf. Artif. Intell.*, 642–648, 2017. ⟨https://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14947⟩.

[52] James Renegar. On the computational complexity and geometry of the first-order theory of the reals, Part I. *J. Symb. Comput.* 13(3):255–299, 1992.

[53] Jürgen Richter-Gebert. Mnëv's universality theorem revisited. *Séminaire Lotharingien de Combinatoire* B34h:15 pp., 1995. ⟨https://www.mat.univie.ac.at/~slc/wpapers/s34berlin.html⟩.

[54] Jürgen Richter-Gebert. *Realization spaces of polytopes*. Lecture Notes Math. 1643. Springer, 1996.

[55] Jürgen Richter. Kombinatorische Realisierbarkeitskriterien für orientierte Matroide. *Mitt. Math. Sem. Gießen* 194:1–112, 1989.

[56] Gerhard Ringel. Teilungen der Ebene durch Geraden oder topologische Geraden. *Math. Z.* 64(1):79–102, 1956.

[57] Marcus Schaefer. Complexity of some geometric and topological problems. *Proc. 17th Int. Symp. Graph Drawing*, 334–344, 2009. Lecture Notes Comput. Sci. 5849, Springer.

[58] Marcus Schaefer. Realizability of graphs and linkages. *Thirty Essays on Geometric Graph Theory*, 461–482, 2013. Springer-Verlag.

[59] Marcus Schaefer and Daniel Štefankovič. Fixed points, Nash equilibria, and the existential theory of the reals. *Theory Comput. Systems* 60(2):172–193, 2017.

[60] Marcus Schaefer and Daniel Štefankovič. The complexity of tensor rank. *Theory Comput. Systems* 62(5):1161–1174, 2018. arXiv:1612.04338.

[61] Yaroslav Shitov. The nonnegative rank of a matrix: Hard problems, easy solutions. *SIAM Review* 59(4):794–800, 2017. arXiv:1605.04000.

[62] Peter W. Shor. Stretchability of pseudolines is NP-hard. *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, 531–554, 1991. DIMACS Series Discrete Math. Theor. Comput. Sci. 4, Amer. Math. Soc. and Assoc. Comput. Mach.

[63] Karl Georg Christian von Staudt. *Geometrie der Lage*. Verlag von Bauer and Rapse (Julius Merz), Nürnberg, 1847.

[64] Peter Guthrie Tait. On knots I. *Trans. Royal Soc. Edinburgh* 28(1):145–190, 1876–7. ⟨https://babel.hathitrust.org/cgi/pt?id=njp.32101074834365&seq=199⟩.

[65] Roberto Tamassia. On embedding a graph in the grid with the minimum numberof bends. *SIAM J. Comput.* 16(3):421–444, 1987.

[66] William T. Tutte. A census of planar maps. *Canad. J. Math.* 15:249–271, 1963.

[67] Ravi Vakil. Murphy's law in algebraic geometry: Badly-behaved deformation spaces. *Invent. Math.* 164(3):569–590, 2006. arXiv:math.AG/0411469.

[68] Kevin Weiler. Edge-based data structures for solid modeling in curved-surface environments. *IEEE Comput. Graph. Appl.* 5(1):21–40, 1985.