CERTIFIED SIMULTANEOUS ISOTOPIC APPROXIMATION OF PAIRS OF CURVES VIA SUBDIVISION

MICHAEL BURR AND MICHAEL BYRD

Abstract. We present a certified algorithm based on subdivision for computing an isotopic approximation to a pair of curves in the plane. Our algorithm is based on the certified curve approximation algorithm of Plantinga and Vegter. The main challenge in this computation is to correctly and efficiently compute the intersections of the curves. To address this issue, we introduce a new, but simple test that guarantees the global correctness of our output.

1. Introduction

In [10, 11], Plantinga and Vegter introduced an algorithm to construct topologically correct piecewise-linear approximations to smooth and bounded real hypersurfaces in two and three dimensions. Their algorithm is particularly interesting as it is a symbolic-numeric algorithm based on subdivision whose predicates are simple and easy to implement. On singular input, however, the Plantinga and Vegter algorithm does not terminate as both of their predicates fail on regions containing singular points. The current paper presents an algorithm in the spirit of the original Plantinga and Vegter algorithm for correctly approximating the union of two smooth curves in the plane with simple transverse crossings.

Main question. *Suppose that $f, g \in \mathbb{Z}[x, y]$ define two smooth curves in the real plane and their corresponding varieties $\mathcal{V}(f)$ and $\mathcal{V}(g)$ intersect transversely in simple crossings. Our goal is to construct a pair of approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$ to $\mathcal{V}(f)$ and $\mathcal{V}(g)$, respectively, such that $\mathcal{A}(f) \cup \mathcal{A}(g)$ is a topologically correct piecewise-linear approximation to $\mathcal{V}(f) \cup \mathcal{V}(g)$, see Figure 1(b).*

In our setting, topologically correct means that there is an ambient isotopy that deforms space while taking both $\mathcal{A}(f)$ to $\mathcal{V}(f)$ and $\mathcal{A}(g)$ to $\mathcal{V}(g)$. In particular, the crossings of the approximations form a topologically correct approximation to the intersection points of $\mathcal{V}(f)$ and $\mathcal{V}(g)$. The approximation and the varieties can also be made as close as desired in Hausdorff distance by further subdivision.

The main challenge is that while the Plantinga and Vegter algorithm computes the individual approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$, the algorithm only guarantees the existence of two ambient isotopies. One of the ambient isotopies takes $\mathcal{A}(f)$ to $\mathcal{V}(f)$ and the other takes $\mathcal{A}(g)$ to $\mathcal{V}(g)$, but there is no guarantee that these isotopies are compatible in any sense, see Figure 1(a). This can cause $\mathcal{A}(f) \cap \mathcal{A}(g)$ to fail to include intersections between $\mathcal{V}(f)$ and $\mathcal{V}(g)$ as well as for $\mathcal{A}(f) \cap \mathcal{A}(g)$ to include extraneous intersections, which do correspond to intersections between $\mathcal{V}(f)$ and $\mathcal{V}(g)$.

An extension of the Plantinga and Vegter algorithm was introduced in [1] to handle unbounded and singular input. This approach can solve the current problem by computing an approximation to the variety $\mathcal{V}(fg)$, but this approach relies on separation bounds between singular points. These separation bounds are typically so pessimistic that it is questionable whether this algorithm is practical. On the other hand, the algorithm introduced in [6] also studies the problem considered here, but their algorithm uses more restrictive tests than what we propose, and they may require a significant number of subdivisions to characterize the local behavior of curves within a region. For instance, their algorithm has more topological requirements on boxes that contain intersections of curves than our approach. This makes our correctness statement a little bit weaker than the correctness statement in Lien *et al.*, but they are still quite strong and more in line with the statement appearing in the original work of Plantinga and Vegter.

Our main contribution is the design and correctness for Algorithm 2. This algorithm is a certified symbolic-numeric subdivision-based algorithm for solving the Main question, with the following correctness statement:

Theorem 1. *Suppose that $f, g \in \mathbb{Z}[x, y]$ and $R = [a, b] \times [c, d]$ is a rectangular subset of $\mathbb{R}^2$ such that $a, b, c, d \in \mathbb{Z}$. In addition, suppose that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ define smooth curves in R which intersect simply and*

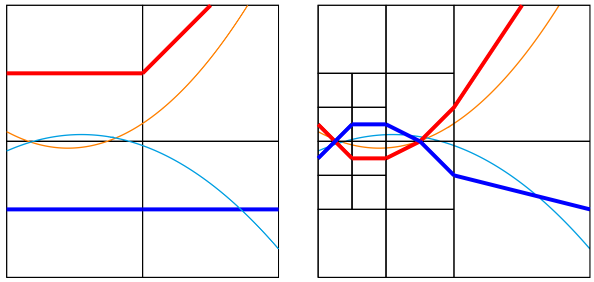

(a) (b)

FIGURE 1. Two approximations (drawn with thick lines) of a pair of curves (drawn with thin lines). The approximation and the curve are paired by color. A naïve approach (a) misses a pair of intersections due to excursions while our approach (b) correctly approximates the curves simultaneously.

transversely. Moreover, suppose that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ do not intersect on the boundary of R. Let $\mathcal{A}(f)$ and $\mathcal{A}(g)$ be the output of Algorithm 2. There exists a neighborhood U of R and an ambient isotopy defined on U that simultaneously takes $\mathcal{A}(f)$ to $\mathcal{V}(f)$ and $\mathcal{A}(g)$ to $\mathcal{V}(g)$, respectively.

Outline. In Section 2, we recall the details of the Plantinga and Vegter algorithm for curve approximation. In Section 3, we provide the details of the subdivision step of our algorithm. Finally, in Section 4, we describe how the curve approximations are constructed.

2. Background

The Plantinga and Vegter algorithm for curve approximation is an adaptive subdivision-based algorithm based on the marching cube algorithm [8]. The input to the algorithm is a bivariate polynomial $f \in \mathbb{Z}[x, y]$ and an input region $R = [a, b] \times [c, d]$ in $\mathbb{R}^2$ such that $a, b, c, d \in \mathbb{Z}$. The output is a piecewise linear approximation $\mathcal{A}(f)$ which is guaranteed to be ambient isotopic to the real variety $\mathcal{V}(f)$. The algorithm consists of three steps: (1) the subdivision step, (2) the balancing step, and (3) the approximation step.

In the subdivision step, square regions are considered by the algorithm and either accepted and further processed or rejected and split into four equal-sized subboxes for further consideration. The acceptance and rejection is based on two predicates called C_0 and C_1. These predicates take, as input, a square region $B \subseteq R$ and produce TRUE or FALSE. A square is accepted when either of the predicates is TRUE and is rejected if both predicates return FALSE. The predicates have the following properties: When $C_0(B) = $ TRUE, we conclude that $\mathcal{V}(f)$ does not intersect B. On the other hand, when $C_1(B) = $ TRUE, there do not exist any pair of points $(x_1, y_1), (x_2, y_2) \in B$ such that $\nabla f(x_1, y_1)$ and $\nabla f(x_2, y_2)$ are perpendicular.

In the balancing step, additional subdivisions are performed until the side length of neighboring boxes differ in length by at most a factor of two. Finally, in the approximation step, for every box which satisfies $C_1(B)$ but not $C_0(B)$, f is evaluated on each vertex. For each edge in the subdivision where f changes signs, the algorithm adds a vertex on this edge. Finally, in each box, the vertices are connected in a way so that the edges do not intersect, and, if there are four vertices on the sides of a box, the two vertices on the same side of the box are not connected.

2.1. **Correctness and complexity.** In the original presentation, [10, 11], the correctness of the curve approximation algorithm was only proved for smooth and bounded curves. In [1], by slightly weakening the correctness statement, the algorithm was extended to unbounded curves. In addition, in [1], the algorithm and correctness statement were extended to nonsingular curves, but the practicality of this approach remains in question. In [7], the authors extend the algorithm to non-square regions within the subdivision.

In each of these algorithms, an important feature of their correctness statements is that they guarantee global correctness, not local correctness. For instance, the correctness statement does not guarantee that the approximation and the variety are isotopic when restricted to a box, only that they are isotopic in the input region R (or an open set containing the input region). The main difference between the variety and its approximations are *excursions* An excursion occurs when the variety briefly enters a neighboring box, but this behavior does not appear in the approximation, see Figure 1 and Definition 2. Instead, the ambient isotopy stretches space so that the approximation is moved into the neighboring box.

This global correctness without local correctness is a key feature of this family of algorithms. In many cases, it leads to many fewer boxes created since these algorithms do not need to resolve the behavior of small excursions. This key feature makes the problem of approximating a pair of curves given by $f, g \in \mathbb{Z}[x, y]$ more challenging since an excursion may involve an intersection between the varieties $\mathcal{V}(f)$ and $\mathcal{V}(g)$, but the ambient isotopies separate the curves and remove the intersection from the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$.

The complexity of the Plantinga and Vegter algorithm was first studied in [2] using continuous amortization, see, e.g., [4, 3], there, the authors found both adaptive and worst-case complexity bounds for the number of regions formed by subdivision as well as the bit-complexity of the algorithm. In addition, the authors found examples which were guaranteed to exhibit the worst-case exponential complexity bounds. In [5], a smoothed-analysis based approach was used to show that the average complexity of the algorithm is polynomial. In [12], a condition number-based approach also showed that the average complexity of the algorithm is polynomial, but for a larger class of random polynomials including some sparse families.

2.2. **Predicate details.** The two predicates in the Plantinga and Vegter algorithm are typically implemented using interval arithmetic, see, e.g., [9] for more details. Interval arithmetic extends the standard arithmetic operations to intervals. For instance,

$$\begin{aligned}[a, b] + [c, d] &= [a + c, b + d],\\ [a, b] - [c, d] &= [a - d, b - c], \text{ and}\\ [a, b][c, d] &= [\min\{ac, ad, bc, bd\}, \max\{ac, ad, bc, bd\}].\end{aligned}$$

These interval operations can be extended to the evaluation of functions, and we use the symbol $\square$ to denote any such extension. In particular, for a polynomial $f \in \mathbb{Z}[x, y]$ and a region B, $\square f(B)$ is an interval containing the image $f(B)$. The interval $\square f(B)$ is often larger than $f(B)$, but significantly easier to compute.

The C_0 test is implemented as $C_0(B) = \textsc{True}$ if and only if $0 \notin \square f(B)$. Since $\square f(B)$ is an over-approximation to $f(B)$, if $0 \notin \square f(B)$, then $0 \notin f(B)$, so the variety $\mathcal{V}(f)$ cannot intersect B. The C_1 test is slightly more complicated, as $C_1(B) = \textsc{True}$ if and only if $0 \notin \square\langle\nabla f, \nabla f\rangle(B \times B)$. In this formulation, each of the factors of $B \times B$ is the argument to one ∇f. If $0 \notin \square\langle\nabla f, \nabla f\rangle(B \times B)$, then there cannot be a pair of points $(x_1, y_1), (x_2, y_2) \in B$ such that $\langle\nabla f(x_1, y_1), \nabla f(x_2, y_2)\rangle = 0$, i.e., the gradient vectors cannot be perpendicular.

We call this family of algorithms symbolic-numeric algorithms for two reasons: The predicates perform exact computations using the coefficients of f, i.e., not merely treating f as a function. The computations themselves are performed using arbitrary-precision floating point computations on dyadic points. In other words, the evaluations are exact, but leverage the speed of floating point calculations.

2.3. **Topological details.** We collect some key facts from [10, 11] and provide a description of the ambient isotopy in the Plantinga and Vegter algorithm. These facts are used throughout our correctness proofs. Given a subdivision of R into boxes, we use the word *side* to denote one of the four sides of a box in the subdivision. We define an *edge* of that subdivision to be a side of a box that is not composed of a union of sides of smaller neighboring boxes. In particular, if B is a box of the final subdivision, then the edges of B are either the sides of B or a half-side of B when B's neighbor in that direction is smaller than B.

Definition 2. *Let B be a box of a subdivision and $f \in \mathbb{Z}[x, y]$. An* excursion *of $\mathcal{V}(f)$ is a component of $\mathcal{V}(f) \cap B$ whose two endpoints are on the same edge of the subdivision, see Figure 1(a).*

We note that excursions do not appear in the piecewise-linear approximation $\mathcal{A}(f)$ as they are deformed into neighboring boxes.

In [10, 11], the authors use several topological lemmas to show that the predicates C_0 and C_1 exert control over the behavior of $\mathcal{V}(f)$ within a box.

Lemma 3 ([10, 11])**.** *Suppose that B is a box of the subdivision, and suppose that there are two segments s_1 and s_2 in B such that*

(1) the lines formed from extending s_1 and s_2 are perpendicular,
(2) the value of f on both endpoints of s_1 is the same, and

(3) the value of f on both endpoints of s_2 is the same.
Then, $C_1(B) = \text{FALSE}$.

This lemma follows from applying the intermediate value theorem on segments s_1 and s_2 to show that each segment contains a point such that the gradient at that point is perpendicular to the segment. This lemma leads to several corollaries, three of which we list here:

Corollary 4 ([10, 11])**.** *Suppose that B is a box of a subdivision such that $C_1(B) = \text{TRUE}$. In addition, let γ_f be a component of $\mathcal{V}(f) \cap B$ which is an excursion on edge e of B. Then, γ_f is entirely contained within the semicircle in B whose diameter is e.*

Corollary 5 ([10, 11])**.** *Suppose that B is a box of a subdivision. If $\mathcal{V}(f)$ intersects two adjacent sides of B twice on each of these sides, then $C_1(B) = \text{FALSE}$.*

Corollary 6 ([10, 11])**.** *Suppose that B is a box of a subdivision such that $C_1(B) = \text{TRUE}$. There is at most one component of $\mathcal{V}(f) \cap B$ that extends from the northern to southern edges of B.*

Finally, we briefly describe how the ambient isotopy deforms the variety $\mathcal{V}(f)$ to $\mathcal{A}(f)$ as a two-step procedure. The first step of the ambient isotopy is to remove all excursions by deforming space so that the excursions are moved into neighboring boxes. Briefly, we let $\widetilde{\mathcal{V}}(f)$ be the result of applying the first step of the ambient isotopy to $\mathcal{V}(f)$. For every box B of the subdivision, $\widetilde{\mathcal{V}}(f) \cap B$ and $\mathcal{A}(f) \cap B$ are ambient isotopic *within* B. In particular, this means that they have the same number of components within B. The second step of the ambient isotopy simultaneously deforms $\widetilde{\mathcal{V}}(f) \cap B$ to $\mathcal{A}(f) \cap B$ within each box B by sliding the points on the boundary of B to their appropriate places and straightening the curves within each box B. By carefully considering these steps, we observe that the ambient isotopies derived from the Plantinga and Vegter algorithm move points at most one box away as the only points that move between boxes are those near excursions, but by Corollary 4, these points are never further than one box away.

Definition 7. *Let S be a union of boxes from a subdivision and $\gamma_f : [0,1] \to \mathcal{V}(f) \cap S$ a curve in the variety of f. The* extension of γ_f without excursions *is denoted $\overline{\gamma}_f$ and is component of $\mathcal{V}(f) \cap S$ containing γ_f. The* extension of γ_f with excursions *is the curve $\widetilde{\gamma}_f$ which is formed by following γ_f forward and backwards until either (1) the curve becomes a closed loop or (2) the curve reaches the first and last intersections of the curve with the boundary of S before passing through a box not in S.*

The extension can be constructed by iteratively adding excursions and curve components at the ends of the path until the curve leaves S and passes through other boxes of the subdivision.

Lemma 8. *Let S be a union of boxes from a subdivision and $\gamma_f : [0,1] \to \mathcal{V}(f) \cap S$ a curve in the variety of f. Suppose that γ_f deforms to $\alpha_f \subseteq \mathcal{A}(f)$ within S. Let $\widetilde{\gamma}_f$ be the extension of this path with excursions. Let $\overline{\alpha}_f$ be the component of $\mathcal{A}(f) \cap S$ containing α_f. Either $\widetilde{\gamma}_f$ and $\overline{\alpha}_f$ are topological circles within S or the endpoints of $\widetilde{\gamma}_f$ and $\overline{\alpha}_f$ are on the same edges of the subdivision.*

3. Subdivision step

Given $f, g \in \mathbb{Z}[x, y]$, a first attempt to solve the main question may be to simultaneously run the standard Plantinga and Vegter algorithm on f and g, but to use a common refinement of the region R. This approach leads to three different types of potential errors in the approximations, see Figures 1, 2, and 3, respectively:

1. Missing intersections: intersections of $\mathcal{V}(f)$ and $\mathcal{V}(g)$ which do not correspond to intersections of $\mathcal{A}(f)$ and $\mathcal{A}(g)$.
2. Extra intersections: intersections of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ which do not correspond to intersections of $\mathcal{V}(f)$ and $\mathcal{V}(g)$.
3. Shared edges: some of the edges of the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$ may be shared between the two approximations.

Our main tool to avoid all three of these errors is a new predicate, called $C_1^\times$. On a square B, if $C_1^\times(B) = \text{TRUE}$, then there do not exist any pair of points $(x_1, y_1), (x_2, y_2) \in B$ such that $\nabla f(x_1, y_1)$ and $\nabla g(x_1, x_2)$ are parallel. In the plane, we may implement this test using the cross product, i.e., $C_1^\times(B) = \text{TRUE}$ if and only if $0 \notin \square(\nabla f \times \nabla g)(B \times B)$, where each of the factors of $B \times B$ is an argument to one of the gradients. As an initial illustration of the utility of this new $C_1^\times$ predicate, we provide the following motivating result:

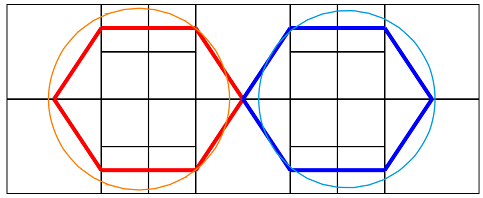

FIGURE 2. Approximations from the standard Plantinga and Vegter algorithm which intersect at a point even though the curves do not.

Lemma 9. *Let B be a rectangle and s line segment in B. Suppose that f attains the same value twice on s and g also attains the same value twice on s. Then $C_1^{\times}(B) = \text{FALSE}$.*

Proof. Suppose that $(x_1, y_1), (x_2, y_2) \in s$ such that $f(x_1, y_1) = f(x_2, y_2)$. By applying Rolle's theorem to $f(tx_1 + (1-t)x_2, ty_1 + (1-t)y_2)$, there is some $t_f \in (0, 1)$ such that

$$\frac{d}{dt} f(tx_1 + (1-t)x_2, ty_1 + (1-t)y_2)\Big|_{t=t_f} = 0.$$

This, however, can be rewritten as the following dot product:

$$\nabla f(t_f x_1 + (1-t_f)x_2, t_f y_1 + (1-t_f)y_2) \cdot \langle x_2 - x_1, y_2 - y_1 \rangle = 0.$$

In other words, there is some point in B where ∇f is perpendicular to s. Repeating the argument for g also gives that there is some point in B where ∇g is perpendicular to s. Therefore, the gradients of f and g are parallel for some pair of points in the box B and $C_1^{\times}(B) = \text{FALSE}$. □

This leads to the following special case:

Corollary 10. *Let B be a rectangle and assume that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ intersect more than once in B. Then $C_1^{\times}(B) = \text{FALSE}$.*

We now present the algorithm for the subdivision step of the pairwise curve approximation algorithm. For simplicity, we focus on the case where the input region R is a square and leave the details for the general rectangular case to Section 4. Since there are two polynomials $f, g \in \mathbb{Z}[x, y]$, we write C_0^f and C_1^f for the standard tests from the Plantinga and Vegter algorithm for the function f. Similarly, we define C_0^g and C_1^g for g. In addition, we need the notion of a neighborhood of a box:

Definition 11. *Let $R = [a, b] \times [c, d]$ be a square region and $\mathcal{S}$ a partition of R into squares. For any square $B \in \mathcal{S}$, the* neighborhood *of B in $\mathcal{S}$ is denoted by $\mathcal{N}(B)$ and consists of B along with all of the other squares in $\mathcal{S}$ that have a positive-length intersection with B, i.e., squares that only meet B at its corners are not in $\mathcal{N}(B)$. More generally, we define $\mathcal{N}_1(B) := \mathcal{N}(B)$ and $\mathcal{N}_i(B)$ to be the union of all the neighborhoods of boxes in $\mathcal{N}_{i-1}(B)$.*

In other words, $\mathcal{N}_i(B)$ consists of all the boxes which are at most i boxes away from B. We write $C_1^{\times}(\mathcal{N}_i(B)) = \text{TRUE}$ to denote that $C_1^{\times}$ holds in the smallest rectangle containing $\mathcal{N}_i(B)$.

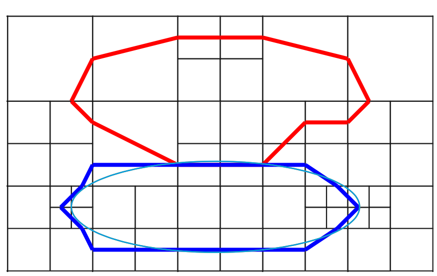

FIGURE 3. The approximations from the standard Plantinga and Vegter algorithm which intersect tangentially even though the curves do not.

Algorithm 1 Subdivision step

Input: polynomials $f, g \in \mathbb{Z}[x, y]$ and a square region R with integral corners
Output: a partition of R for further processing

```
initialize queue Q to contain R
while Q is not empty do
    pop square B from Q
    accept B if any of the following holds:
        • C_0^f(B) = C_0^g(B) = TRUE
        • C_0^f(B) = C_1^g(B) = TRUE
        • C_1^f(B) = C_0^g(B) = TRUE
        • C_1^f(B) = C_1^g(B) = C_1^×(N_2(B)) = TRUE
    reject B if none of the previous hold:
        subdivide B into four equal-sized boxes B_1, ..., B_4
        push B_1, ..., B_4 into Q
end while
return accepted boxes
```

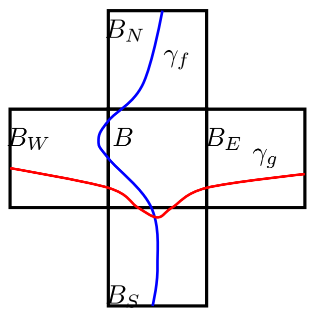

FIGURE 4. The neighborhood of the box B with two crossing components.

Suppose that after this new subdivision step, the balancing and approximation steps of the standard Plantinga and Vegter algorithm are performed, resulting in approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$. We now consider the crossing properties of these approximations.

3.1. **Transversal crossing of approximations.** For the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$, we call a crossing *transversal* if the approximations cross within the interior of a box. This only happens when one approximation has an edge from the north side of a box to the south side, while the other approximation extends from the east side of the box to the west side. We show that every transversal crossing of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ corresponds to a unique crossing of the varieties $\mathcal{V}(f)$ and $\mathcal{V}(g)$ as follows:

Proposition 12. *Suppose that B is a box such that $C_1^f(B) = C_1^g(B) = C_1^\times(\mathcal{N}(B)) =$ TRUE. If $\mathcal{A}(f)$ and $\mathcal{A}(g)$ intersect transversely in B, then $\mathcal{V}(f)$ and $\mathcal{V}(g)$ intersect exactly once and transversely in $\mathcal{N}(B)$.*

Proof. By Corollary 10, the number of intersections of $\mathcal{V}(f)$ and $\mathcal{V}(g)$ is at most one in $\mathcal{N}(B)$. Moreover, any such intersection must be transversal, since otherwise the gradients agree at the intersection and $C_1^\times(\mathcal{N}(B))$ would be FALSE.

Without loss of generality, we assume that $\mathcal{A}(f)$ extends from the northern to the southern edges of B and $\mathcal{A}(g)$ extends from the eastern to the western edges of B. Moreover, since the Plantinga and Vegter algorithm implies the existence of two homotopies which do not deform space further than one box away, there are unique components, which we call *crossing components* of $\mathcal{V}(f) \cap \mathcal{N}(B)$ and $\mathcal{V}(g) \cap \mathcal{N}(B)$ which cross B from its north side to its south side and from its east side to its west side. Let γ_f and γ_g denote these two crossing components of $\mathcal{V}(f) \cap \mathcal{N}(B)$ and $\mathcal{V}(g) \cap \mathcal{N}(B)$. We note that these components are not required to stay within B as there may be excursions to the neighboring boxes. See Figure 4 for additional details.

While there is considerable flexibility in the local shapes of γ_f and γ_g, their larger structures are well-constrained. We restrict our attention to γ_f since the behavior of γ_g is analogous. The endpoints of γ_f must be on the boundary of $\mathcal{N}(B)$ since, otherwise, $\mathcal{N}(B)$ would contain a closed loop. This is not possible because inside any closed loop, there is an extremal point where ∇f vanishes, but this point would force $C_1^{\times}(\mathcal{N}(B))$ to fail.

Moreover, we let B_N, B_E, B_S, and B_W denote the unions of the boxes of $\mathcal{N}(B)$ lying to the north, east, south, and west of B, respectively. We note that there may be at most two boxes in any cardinal direction due to the balancing step in the Plantinga and Vegter algorithm. In addition, we show that the endpoints of γ_f are on the *external boundaries* of B_N and B_S, i.e., the boundaries of B_N and B_S that are also boundaries of $\mathcal{N}(B)$. In particular, suppose that there is no endpoint of γ_f in the external boundary of B_N. By construction, γ_f crosses from the northern edge of B to its southern edge, so it intersects the northern and southern edges of B at least once. Moreover, since γ_f does not have an endpoint on the external boundary of B_N, it must cross the northern edge of B a second time. Therefore, γ_f must cross at least one of eastern, southern, or western edges of B an additional time.

First, we show that γ_f cannot cross the eastern or western edges of B. Without loss of generality, we assume that γ_f crosses the eastern edge of B. Since there is no vertex of $\mathcal{A}(f)$ placed on the eastern edge of B, it must be that there are an even number of crossing of the eastern edge of B. Since we have assumed that there is at least one crossing, there must be at least two crossings. Then Corollary 5 shows that this configuration would violate $C_1^f(B) = \text{TRUE}$.

Therefore, the only remaining possibility is for γ_f to have both of its endpoints in the boundary of B_S. This, however, is also impossible as it would imply the existence of two crossing components of $\mathcal{V}(f)$ from the north to the south of B. This is also impossible by Corollary 6 as it would violate $C_1^f(B) = \text{TRUE}$.

Finally, since γ_f extends from the external boundary of B_N to the external boundary of B_S, it separates the external boundary of B_E from the external boundary of B_W. On the other hand, since γ_g extends from the external boundary of B_E to the external boundary of B_W and does not intersect the boundary of $\mathcal{N}(B)$ except at its endpoints, γ_g must intersect γ_f, which implies the desired existence of a crossing in $\mathcal{N}(B)$. $\square$

By investigating the proof of Proposition 12 in more detail, we find that each intersection of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ corresponds to an intersection between γ_f and γ_g within $\mathcal{N}(B)$. Therefore, no two crossings of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ can correspond to the same intersection of $\mathcal{V}(f)$ and $\mathcal{V}(g)$ because at least one of γ_f and γ_g change when considering a different crossing. Therefore, there is an injective map between transversal intersections of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ and transversal intersections of $\mathcal{V}(f)$ and $\mathcal{V}(g)$.

3.2. **Missing intersections.** We begin by noting that excursions are the only reason that the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$ can miss an intersection of $\mathcal{V}(f)$ and $\mathcal{V}(g)$. In particular, we have the following result:

Lemma 13. *Suppose that B is a box such that $C_1^f(B) = C_1^g(B) = C_1^{\times}(B) = \text{TRUE}$. Suppose, in addition, that there are no excursions either entering or exiting B. If the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$ do not intersect in B, including on the boundary of B, then $\mathcal{V}(f)$ and $\mathcal{V}(g)$ do not intersect in B.*

Proof. By the properties of the first step of the Plantinga and Vegter algorithm, see Section 2.3, since there are no excursions, in each box B, the approximations $\mathcal{A}(f) \cap B$ and $\mathcal{A}(g) \cap B$ are each ambient isotopic to $\mathcal{V}(f) \cap B$ and $\mathcal{V}(g) \cap B$ *within the box* B, respectively. We recall, however, that these isotopies are not necessarily the same. In particular, this implies that the number of components of $\mathcal{A}(f)$ and $\mathcal{V}(f)$ agree within B, and, similarly for $\mathcal{A}(g)$ and $\mathcal{V}(g)$.

Suppose that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ intersect in B. Let γ_f be a component of $\mathcal{V}(f) \cap B$ that intersects $\mathcal{V}(g) \cap B$. Let α_f be the corresponding component of $\mathcal{A}(f)$ to which γ_f deforms under the ambient isotopy. Let e_1 and e_2 be the edges of the subdivision that contain the endpoints of α_f. By Lemma 8, γ_f must also begin and end on these edges. On the other hand, since $\mathcal{A}(g)$ does not have vertices on these edges, it follows that on each edge e_i, the value of g is the same at both endpoints of e_i. Since there are no excursions, it follows that $\mathcal{V}(g)$ does not intersect this edge as $\mathcal{V}(g)$ would need to intersect this edge twice to maintain the sign properties of the endpoints and this would be an excursion.

Since $\mathcal{A}(f)$ and $\mathcal{A}(g)$ do not intersect, the signs of g at both endpoints of α_f must be the same. Since we showed that the sign of g is constant on the edges containing the endpoints of f, the signs of g on both endpoints of γ_f must be the same. This implies $\mathcal{V}(g)$ must intersect γ_f an even number of times as each intersection changes the sign of the restriction $g|_{\gamma_f}$. Since $\mathcal{V}(g)$ and γ_f intersect once, they must intersect at least twice, but this is impossible by Corollary 10. $\square$

Lemma 13 implies that missing intersections must involve at least one excursion. Our plan is to show that any missing intersection must induce a pair of intersections in the neighborhood $\mathcal{N}_2(B)$, which is not possible since $C(\mathcal{N}_2(B)) = \textsc{True}$.

Proposition 14. *Suppose that B is a box such that $C_1^f(B) = C_1^g(B) = C_1^\times(\mathcal{N}_2(B)) = \textsc{True}$. Suppose that $\mathcal{A}(f)$ and $\mathcal{A}(g)$ do not intersect in $\mathcal{N}(B)$, including on the boundary of $\mathcal{N}(B)$, then $\mathcal{V}(f)$ and $\mathcal{V}(g)$ do not intersect in B.*

Proof. Suppose that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ intersect in B. Let γ_f be the component of $\mathcal{V}(f)$ in $\mathcal{V}(f) \cap B$ that includes this intersection. Let $\widetilde{\gamma}_f$ be the extension of γ_f into $\mathcal{N}(B)$ including all excursions into $\mathcal{N}_2(B)$.

Let α_f be the component of $\mathcal{A}(f) \cap \mathcal{N}(B)$ to which γ_f is deformed to under the ambient isotopy. We note that α_f is in $\mathcal{N}(B)$ since the Plantinga and Vegter algorithm does not deform the curve further than one box away. Let e_1 and e_2 be the edges of the subdivision that contain the endpoints of α_f. By Lemma 8, the endpoints of $\widetilde{\gamma}_f$ are also on the edges e_1 and e_2. Let p_1 be the endpoint of $\widetilde{\gamma}_f$ on e_1 and p_2 be the endpoint of $\widetilde{\gamma}_f$ on e_2.

Consider the restriction $g|_{\widetilde{\gamma}_f}$. The signs of this function at the two endpoints must be opposite because each intersection between $\mathcal{V}(g)$ and $\widetilde{\gamma}_f$ changes the sign of $g|_{\widetilde{\gamma}_f}$. Since there is one intersection, having the same sign at both endpoints would require two intersections for the two sign changes, but this is impossible by Corollary 10 and that $C_1^\times(\mathcal{N}_2(B)) = \textsc{True}$.

Now, we prove that the sign of g at p_1 agrees with its signs at the endpoints of e_1. Since $\mathcal{A}(g)$ does not intersect $\mathcal{A}(f)$ and $\mathcal{A}(f)$ has a vertex on e_1, this implies that the signs of g on the endpoints of e_1 are the same. Hence, any intersection of $\mathcal{V}(g)$ with e_1 must be an excursion. Suppose, for contradiction, that the sign of g at p_1 does not match the sign of g at the endpoints of e_1. Then, there are an odd number of intersections from $\mathcal{V}(g) \cap e_1$ on either side of p_1. Therefore, there is at least one pair of points of $\mathcal{V}(g) \cap e_1$ on either side of p_1 which are connected by an excursion.

We show that this excursion implies that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ intersect once more in $\mathcal{N}_2(B)$, but such an intersection is not possible since $C_1^\times(\mathcal{N}_2(B)) = \textsc{True}$ and Corollary 10. See Figure 5 for reference in this argument. Suppose first that this excursion is internal to $\mathcal{N}(B)$. By Corollary 9, since $\mathcal{V}(g)$ has an excursion on e_1, $\mathcal{V}(f)$ does not have an excursion on this edge. By Corollary 4, this excursion is contained within a semicircle along the edge e_1. The curve $\widetilde{\gamma}_f$, however, cannot stay within this semicircle. More precisely, let $\widetilde{\gamma}_g$ be an excursion of g in $\mathcal{N}(B)$ with endpoints on either side of p_1 on e_1. Then, $\widetilde{\gamma}_f$ is within the region bounded by $\widetilde{\gamma}_g$ and e_1, but it must leave this region to reach the intersection in B and it can only intersect e_1 once. Therefore, $\widetilde{\gamma}_f$ and $\widetilde{\gamma}_g$ must intersect in a box other than B.

The argument in the case where the excursion is external to $\mathcal{N}(B)$ is similar. In this case, $\widetilde{\gamma}_g$ is the excursion of g with endpoints on either side of p_1, but is outside $\mathcal{N}(B)$. Then, $\widetilde{\gamma}_f$ must pass through a box in $\mathcal{N}_2(B) \setminus \mathcal{N}(B)$, so $\widetilde{\gamma}_f$ cannot remain within the region bounded by the excursion $\widetilde{\gamma}_g$ and e_1, but it must leave this region and it can only intersect e_1 once. Therefore, $\widetilde{\gamma}_f$ and $\widetilde{\gamma}_g$ must intersect in a box other than B.

Now, we have shown that the sign of g at p_1 agrees with the sign of g at the endpoints of e_1. Similarly, the sign of p_2 agrees with the sign of g at the endpoints of e_2. Moreover, since the signs of g at p_1 and p_2 differ, the sign of g at the endpoints of e_1 differ from the sign of g at the endpoints of e_2. We also observe that removing edges e_1 and e_2 from the boundary $\partial\mathcal{N}(B)$ results in two components. Since the signs of g are different at the endpoints of the two components, there must be an odd number of vertices of $\mathcal{A}(g)$ on each of these two components. However, since $\mathcal{A}(g)$ forms a perfect matching on its vertices in $\mathcal{N}(B)$, one of the edges of $\mathcal{A}(g)$ must intersect α_f, but this is not possible. $\square$

Therefore, by Proposition 12, every transversal intersection of the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$ corresponds to a unique intersection of the varieties $\mathcal{V}(f)$ and $\mathcal{V}(g)$. On the other hand, Proposition 14 implies that every intersection of the varieties $\mathcal{V}(f)$ and $\mathcal{V}(g)$ corresponds to an intersection of the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$. However, this intersection does not need to be transversal, see Figure 6 which is the one remaining case.

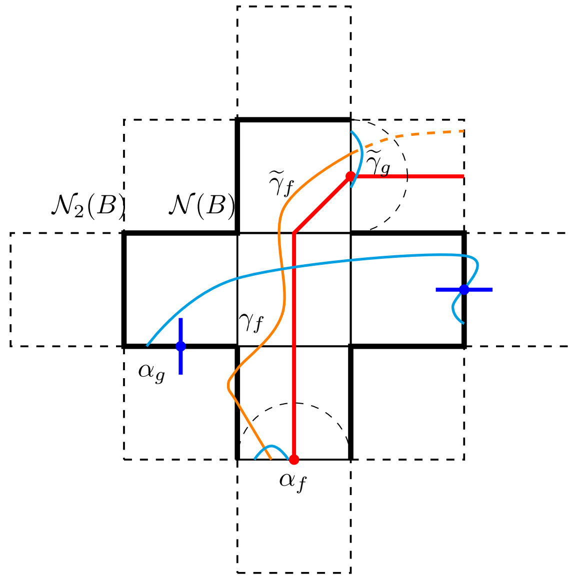

FIGURE 5. Accompanying diagram for the proof of Proposition 14. It illustrates all the impossible features whose nonexistence forces a crossing of the approximations.

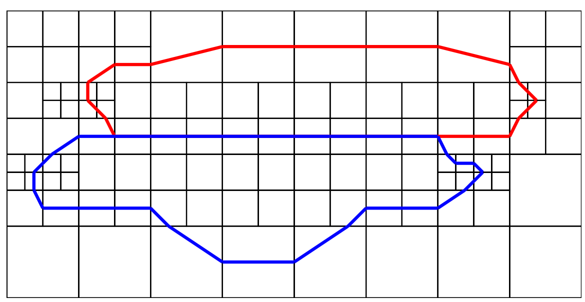

FIGURE 6. Approximations from the standard Plantinga algorithm which share a segment while the curves intersect in two places.

3.3. **Shared edges in approximations.**

For the approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$, we call an intersection *nontransversal* if the approximations meet on the boundary of a box or the approximations coincide along shared segments.

Definition 15. *A contiguous sequence of boxes with shared segments of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ is called a* snake*, see Figure 7. The boxes where the approximations separate are called the* heads *of the snake. A* neighborhood *$\mathcal{N}(S)$ of a snake is the union of all the neighborhoods of boxes in the snake along with the neighborhoods of the heads of the snake. The neighborhood $\mathcal{N}_i(S)$ is defined similarly.*

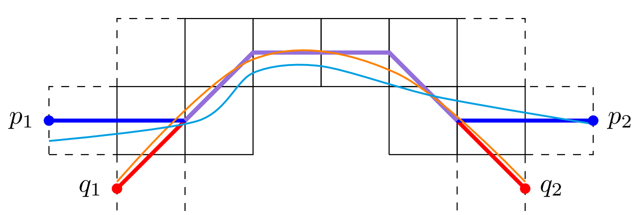

FIGURE 7. Example of a snake where the approximations share segments.

Proposition 16. *Suppose that S is a snake and for every box $B \in S$, $C_1^\times(\mathcal{N}(B))$ holds. There is at most one crossing in $\mathcal{N}(S)$ corresponding to the snake.*

Proof. Let α_f and α_g be the two components of the approximation in S which share vertices or edges. Let γ_f and γ_g be the components of $\mathcal{V}(f) \cap \mathcal{N}(S)$ and $\mathcal{V}(g) \cap \mathcal{N}(S)$, respectively, that α_f and α_g deform to, respectively, under the ambient isotopies. Let $\overline{\gamma}_f$ and $\overline{\gamma}_g$ be the extensions of γ_f and γ_g, respectively, in $\mathcal{N}_2(S)$, but without including excursions. We consider $\overline{\gamma}_f$ as a path and we choose an orientation to this path so that f is positive in a tubular neighborhood of the left of the path and negative in a tubular neighborhood to the right of the path. Since gradients point in the direction of greatest increase, $\nabla f(\overline{\gamma}_f)$ points to the left of $\overline{\gamma}'_f$.

Suppose that $\overline{\gamma}_f$ and $\overline{\gamma}_g$ intersect multiple times in $\mathcal{N}(S)$. We show that this violates $C_1^\times(\mathcal{N}(B))$ for some $B \in S$. See Figure 8 for reference in this argument. Let r_1 and r_2 be two intersection points of $\overline{\gamma}_f$ and $\overline{\gamma}_g$ in $\mathcal{N}(S)$. In addition, we assume that r_1 occurs before r_2 along the path $\overline{\gamma}_f$. We restrict our attention to the portion of $\overline{\gamma}_g$ between r_1 and r_2. We choose the orientation on $\overline{\gamma}_g$ so that r_1 comes before r_2 along this path. As above, the gradient $\nabla g(\overline{\gamma}_g)$ is perpendicular to $\overline{\gamma}'_g$, but, since the orientation of $\overline{\gamma}_g$ is fixed, it points to one side of $\overline{\gamma}'_g$, i.e., the gradient points to either the right or the left of the tangent vector.

Since $\overline{\gamma}_f$ divides $\mathcal{N}(S)$ into two pieces, every time $\overline{\gamma}_g$ crosses $\overline{\gamma}_f$, it crosses from the negative side to the positive side or from the positive side to the negative side. Moreover, since $\overline{\gamma}_g$ stays within $\mathcal{N}(S)$, these types of crossings alternate. Since $\overline{\gamma}_f$ and $\overline{\gamma}_g$ intersect at least twice, there is at least one of each type of crossing.

Since the positive side of f is to the left of $\overline{\gamma}'_f$, a crossing from the positive side to the negative side of $\overline{\gamma}_f$ corresponds to a clockwise turn from $\overline{\gamma}'_f$ to $\overline{\gamma}'_g$. Similarly, a crossing from the negative side to the positive side of $\overline{\gamma}_f$ corresponds to a counter-clockwise turn from $\overline{\gamma}'_f$ to $\overline{\gamma}'_g$. By redefining r_2, if necessary, we may assume that $\nabla f(r_1) \times \nabla g(r_1)$ and $\nabla f(r_2) \times \nabla g(r_2)$ correspond to two different turn directions, so they are different signs.

Finally, we apply the function $\nabla f \times \nabla g$ to the path $\overline{\gamma}_g$. In other words, we look at $\nabla f(\overline{\gamma}_g) \times \nabla g(\overline{\gamma}_g)$. We know that the value of this continuous function has different signs at r_1 and r_2. Hence, by the intermediate value theorem, there is some point on this curve where the gradients of f and g are parallel, but this is not possible as it would contradict $C_1^\times(\mathcal{N}(B)) = \text{TRUE}$ for all $B \in S$. $\square$

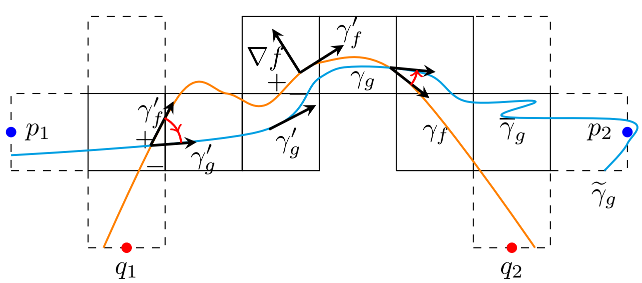

FIGURE 8. Accompanying diagram for the proof of Proposition 16. It illustrates how at two consecutive intersections, the turn directions of the gradients differ.

We have shown that every snake corresponds to at most one intersection of $\mathcal{V}(f)$ and $\mathcal{V}(g)$. It remains to discuss how to decide if a snake corresponds to an intersection. We begin by defining the orientation of points with respect to the snake. Let S be a snake and let B be a head of the snake. Suppose that p and q are two points on the external boundary of $\mathcal{N}(B)$, i.e., boundaries of $\mathcal{N}(B)$ that are also boundaries of $\mathcal{N}(S)$. The external boundary of $\mathcal{N}(B)$ is, therefore, a piecewise-linear path around $\mathcal{N}(B)$. We say that q is *clockwise* from p with respect to the snake if walking around $\partial\mathcal{N}(B) \setminus S$ in a clockwise direction starting at the snake reaches p before q. In this case, p is *counterclockwise* from q with respect to the snake, see Figures 7 and 8.

Lemma 17. *Let S be a snake and let α_f and α_g be the two components of $\mathcal{A}(f) \cap \mathcal{N}(S)$ and $\mathcal{A}(g) \cap \mathcal{N}(S)$, respectively, that include the shared edges of the snake. Let B_1 and B_2 be the two heads of S. Let p_1 and q_1*

be the ends of α_f and α_g, respectively, in B_1, and define p_2 and q_2 similarly. The snake corresponds to an intersection if and only if the orientations from p_1 to q_1 and p_2 to q_2 are the same.

Proof. Let γ_f and γ_g be the two components of $\mathcal{V}(f)$ and $\mathcal{V}(g)$, respectively, in $\mathcal{N}(S)$ to which α_f and α_g deform. Let $\overline{\gamma}_f$ and $\overline{\gamma}_g$ be the extensions *without excursions* of γ_f and γ_g, respectively. Similarly, let $\widetilde{\gamma}_f$ and $\widetilde{\gamma}_g$ be the extensions *with excursions* of γ_f and γ_g, respectively. By the definition of the heads of a snake, the paths $\widetilde{\gamma}_f$ and $\widetilde{\gamma}_g$ pass into different neighbors of B_1, since, otherwise, the snake would be longer. Moreover, the endpoints of $\widetilde{\gamma}_f$ and $\overline{\gamma}_f$ must be in the same neighbors of B_1, even though the endpoints might be on different sides of those boxes. This is because differences between the endpoints of $\widetilde{\gamma}_f$ and $\overline{\gamma}_f$ are due to excursions which, with Corollaries 5 and 6, prevent $\widetilde{\gamma}_f$ from re-entering B. Corresponding statements hold for g as well as for B_2.

Thus, the endpoints of α_f and α_g have the same clockwise or counterclockwise relationship as the endpoints of $\widetilde{\gamma}_f$ and $\widetilde{\gamma}_g$ as well as the endpoints of $\overline{\gamma}_f$ and $\overline{\gamma}_g$. Therefore, we study the relationship between the endpoints of $\overline{\gamma}_f$ and $\overline{\gamma}_g$. Topologically, the boundary $\partial\mathcal{N}(S)$ is a circle and $\overline{\gamma}_f$ and $\overline{\gamma}_g$ form cords of this circle. The paths $\overline{\gamma}_f$ and $\overline{\gamma}_g$ intersect if and only if their endpoints interweave along the boundary of the circle. Interweaving is equivalent to having the same clockwise or counterclockwise order around the circle. □

4. Simultaneous approximation

We now prove Theorem 1 and provide the main algorithm of the paper. Suppose that $f, g \in \mathbb{Z}[x, y]$, $R = [a, b] \times [c, d] \subseteq \mathbb{R}^2$ a rectangle such that $a, b, c, d \in \mathbb{Z}$. Suppose also that $\mathcal{V}(f)$ and $\mathcal{V}(g)$ are nonsingular within R, do not have a common intersection on the boundary of R, and intersect transversely within R. Through a rescaling, we reduce to the case where R is a square. Second, we use the techniques for unbounded curves from [1] for the boundary boxes as long as in each boundary box, either $C_0^f(B) = \textsc{True}$ or $C_0^g(B) = \textsc{True}$. Therefore, we focus on the topological correctness.

By Proposition 14, every intersection between $\mathcal{V}(f)$ and $\mathcal{V}(g)$ corresponds to either a transversal crossing of $\mathcal{A}(f)$ and $\mathcal{A}(g)$ or a snake. By investigating the proofs of Propositions 12 and 16, we see that the these two features correspond to different types of crossings and cannot identify the same crossing twice. Thus, by Lemma 17, we can identify exactly when a crossing occurs. This gives a bijection between crossings in the approximation and crossings in the varieties. Once the intersections are identified, the remaining isotopy steps, such as removing excursions, of the Plantinga and Vegter algorithm can be applied to both $\mathcal{V}(f)$ and $\mathcal{V}(g)$ simultaneously, by Lemma 13, giving topologically correct approximations.

Algorithm 2 Simultaneous approximation algorithm

Input: polynomials $f, g \in \mathbb{Z}[x, y]$ and a square region R with integral corners
Output: approximations $\mathcal{A}(f)$ and $\mathcal{A}(g)$ such that $\mathcal{A}(f) \cap \mathcal{A}(g)$ approximates $\mathcal{V}(f) \cap \mathcal{V}(g)$.

```
Subdivide R using Algorithm 1.
Further subdivide boxes until the side lengths of neighboring boxes differ by at most a factor of two.
Compute the Plantinga and Vegter curve approximation.
For any snakes, apply Lemma 17.
if there is no crossing then
    slightly separate the edges of the snake so that the common edges do not overlap.
else
    slightly separate the ends of the snake so that the approximations don't overlap at the ends of the
  snake, then add an explicit crossing in the middle of the snake.
end if
return the approximations
```

We note that a small Hausdorff distance can also be achieved by making sure that the boxes containing the approximations are sufficiently small and that any snakes are also small. We end with corrected examples of images that appeared earlier in the text, see Figure 9.

Acknowledgements

Burr was supported by NSF grant DMS-1913119 Simons Foundation collaboration grant #964285.

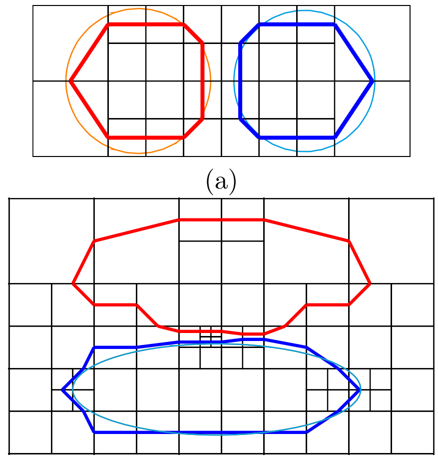

(a)

(b)

FIGURE 9. Topologically correct versions of (a) Figure 2 and (b) Figure 3 using our algorithm.

REFERENCES

[1] Michael Burr, Sung Woo Choi, Ben Galehouse, and Chee K. Yap. Complete subdivision algorithms, II: Isotopic meshing of singular algebraic curves. *Journal of Symbolic Computation*, 47(2):131–152, 2012.

[2] Michael Burr, Shuhong Gao, and Elias Tsigaridas. The complexity of subdivision for diameter-distance tests. *Journal of Symbolic Computation*, 101:1–27, 2020.

[3] Michael A. Burr. Continuous amortization and extensions: with applications to bisection-based root isolation. *Journal of Symbolic Computation*, 77:78–126, 2016.

[4] Michael A. Burr and Felix Krahmer. `SqFreeEVAL`: an (almost) optimal real-root isolation algorithm. *Journal of Symbolic Computation*, 47(2):153–166, 2012.

[5] Felipe Cucker, Alperen A. Ergür, and Josué Tonelli-Cueto. On the complexity of the Plantinga-Vegter algorithm. *Discrete & Computational Geometry. An International Journal of Mathematics and Computer Science*, 68(3):664–708, 2022.

[6] Jyh-Ming Lien, Vikram Sharma, Gert Vegter, and Chee Yap. Isotopic arrangement of simple curves: an exact numerical approach based on subdivision. In *Mathematical software—ICMS 2014*, volume 8592 of *Lecture Notes in Computer Science*, pages 277–282. Springer, Heidelberg, 2014.

[7] Long Lin and Chee Yap. Adaptive isotopic approximation of nonsingular curves: the parameterizability and nonlocal isotopy approach. *Discrete & Computational Geometry. An International Journal of Mathematics and Computer Science*, 45(4):760–795, 2011.

[8] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '87, pages 163–169, 1987.

[9] Ramon E. Moore, R. Baker Kearfott, and Michael J. Cloud. *Introduction to interval analysis*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2009.

[10] Simon Plantinga and Gert Vegter. Isotopic implicit surface meshing. In *Proceeding Symposium on Geometry Processing*, pages 251–260, 2004.

[11] Simon Plantinga and Gert Vegter. Isotopic meshing of implicit surfaces. *The Visual Computer*, 23:45–58, 2007.

[12] Josué Tonelli-Cueto and Elias Tsigaridas. Condition numbers for the cube. i: Univariate polynomials and hypersurfaces. In *Proceedings of the 45th International Symposium on Symbolic and Algebraic Computation*, pages 434–441, 2020.

220 PARKWAY DRIVE, CLEMSON UNIVERSITY, CLEMSON, SC, 29634
Email address: `burr2@clemson.edu`

220 PARKWAY DRIVE, CLEMSON UNIVERSITY, CLEMSON, SC, 29634
Email address: `mbyrd6@clemson.edu`