Testing Simple Polygons

Esther M. Arkin[(1)*] Patrice Belleville[(2)†] Joseph S.B. Mitchell [(1)‡]
David Mount[(3)§] Kathleen Romanik[(4)¶] Steven Salzberg[(5)||]
Diane Souvaine[(6)**]

[(1)]Department of Applied Mathematics and Statistics
State University of New York, Stony Brook, NY 11794-3600
[(2)]Department of Computer Science
University of British Columbia, Vancouver, BC, Canada V6T 1Z4
[(3)]Department of Computer Science
University of Maryland, College Park, MD 20742
[(4)]Center for Automation Research
University of Maryland, College Park, MD 20742
[(5)]Department of Computer Science
Johns Hopkins University, Baltimore, MD 21218
[(6)]Department of Computer Science
Rutgers University, New Brunswick, NJ 08903

Abstract: We consider the problem of verifying a simple polygon in the plane using "test points". A *test point* is a geometric probe that takes as input a point in Euclidean space, and returns "+" if the point is inside the object being probed or "−" if it is outside. A verification procedure takes as input a description of a target object, including its location and orientation, and it produces a set of test points that are used to verify whether a test object matches the description. We give a procedure for verifying an n-sided, non-degenerate, simple target polygon using $5n$ test points. This testing strategy works even if the test polygon has $n+1$ vertices, and we show a lower bound of $3n+1$ test points for this case. We also give algorithms using $O(n)$ test points for simple polygons that may be degenerate and for test polygons that may have up to $n+2$ vertices. All of these algorithms work for polygons with holes. We also discuss extensions of our results to higher dimensions.

Keywords: Testing, verifying, probing.

*Partially supported by National Science Foundation Grants ECSE-8857642 and CCR-9204585.

†Partially supported by Postgraduate and Postdoctoral fellowships from the National Science and Engineering Research Council of Canada.

‡Partially supported by grants from Hughes Research Laboratories, Boeing Computer Services, Air Force Office of Scientific Research contract AFOSR-91-0328, and by National Science Foundation Grants ECSE-8857642 and CCR-9204585.

§Partially supported by NSF Grants CCR-89-08901 and CCR-93-10705 and by grant JSA 91-5 from the Bureau of the Census.

¶Partially supported by AFOSR Grant F49620-93-1-0039.

||Supported in part by National Science Foundation Grant IRI-9116843

**Supported in part by National Science Foundation Grant CCR-91-04732.

1 Introduction

Geometric probing [CY87, Ski88, Ski92, LR94, BS93, Kar91] is the subarea of computational geometry that investigates how to identify or verify an object using a measuring device called a *probe*. For verification problems we are given a description of a target object, including its location and orientation, and we must use probes to verify whether a given test object correctly matches the description. We call the procedure that produces the set of verification probes a *testing algorithm.*

One parameter that is used to specify a geometric probing problem is the type of probe used. In this paper we use point probes or *test points.* Point probes measure whether a single point is inside or outside the object being probed. That is, a point probe takes as input a point in Euclidean space, and it returns + (*positive*) if the point is inside the object being probed or $-$ (*negative*) if it is outside. This type of probe was developed independently by Romanik [RS94, Rom92] and Mitchell [Mit90]. Here we use the model defined by Romanik, which is defined formally in Section 2 below. This type of probe models the data given by point light detectors or a tactile robot sensor in an automated manufacturing application. It also models pixel data used for automatic target recognition. Note that in this model exact verification cannot be done since we assume infinite precision, and we assume that a point probe does not return information about whether it is on the boundary of an object. Therefore, a testing algorithm can only produce probes that verify an object to within a given error bound. This model was used by Romanik and Salzberg [RS95] for obtaining results on verifying orthogonal simple polygons.

In this paper we give an explicit procedure for verifying a simple n-vertex polygon in the plane by using at most $5n$ test points, provided that the target polygon has no three collinear vertices and the test polygon has at most $n+1$ vertices. We also give a lower bound of $3n+1$ test points in this case. Additionally, we give strategies using $O(n)$ test points for the case in which the target polygon may have collinear vertices and the case in which the test polygon has at most $n+2$ vertices. If the test polygon can have three vertices more than the target polygon, then verification with any finite number of test points is impossible (see Section 4). Previous work by Romanik [Rom92] using this model gave a testing algorithm producing $7n$ points for polygons with no three collinear vertices.

An alternate probing model is that of the finger probe: a *finger probe* takes as input a directed line originating from infinity, and it returns the first point of intersection between the line and the object being probed, if one exists, or it returns ∞ if the line completely misses the object (see [CY87, Ski88]). Previous work by Skiena [Ski88] has shown that even if the number of vertices n is fixed, it is not always possible to distinguish two simple polygons using finger probes (see Figure 1). This is because finger probes originate from infinity and cannot bend, so it is impossible for them to probe edges that are blocked from view by other parts of the polygon. Therefore, simple polygons cannot be verified with finger probes. For convex polygons, an algorithm was given in [CY87] for verifying an n-sided polygon with $2n$ finger probes, and an algorithm was given in [RS94, Rom92] for verifying an n-sided polygon with $2n$ test points.

The remainder of this paper is organized as follows. Section 2 gives definitions, and Section 3 gives results for simple polygons where both the target and test polygon have n vertices. First a result is given for the case where the target polygon has no three collinear vertices, and then the result is generalized to handle collinearities. Section 4 generalizes the results of Section 3 by considering testing strategies for polygons that may have one or two more vertices than the target polygon. Section 5 gives a lower bound on the number of test points needed to verify an n-vertex simple polygon with no three collinear vertices where the test polygon may have $n+1$ vertices. Section 6 discusses extensions of the result for polygons with no three collinear vertices to higher

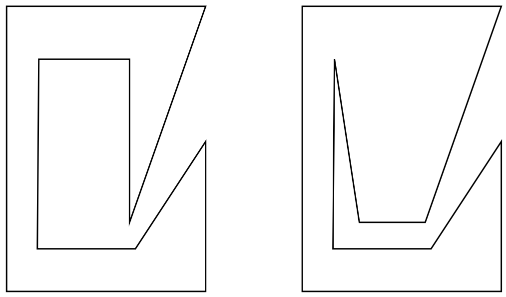

Figure 1: Two simple polygons indistinguishable with finger probes

dimensions. Section 7 summarizes the results and discusses future work.

2 Testing With Point Probes

In this section we give definitions that describe our model for verifying or *testing* with point probes. The model is general enough to be applied to any set of geometric objects. The objects studied in this paper are simple polygons and polyhedra.

An *object* P is a measurable subset $P \subset E^d$ of d-dimensional Euclidean space; in particular, it is a Borel set[1]. An *object class* is a set $\mathcal{Q}$ of objects. Given a *test object* $Q \in \mathcal{Q}$ and a *target object* $P \in \mathcal{Q}$, Q is *consistent* with P on some finite set of test points t if it contains the same subset of t as P, i.e. $t \cap Q = t \cap P$. The *error* of Q, with respect to P, is given by $V(P \triangle Q)$, where $V(P)$ denotes the d-dimensional volume of P and $P \triangle Q$ denotes the symmetric difference of the sets. Let S denote the set of all finite sets of points in E^d, and let I denote the open interval of rationals $(0, 1)$.

Definition. A computable function $T: \mathcal{Q} \times I \rightarrow S$ is a *testing algorithm for $\mathcal{Q}$ with test set size k* if there exists a constant k dependent only on $\mathcal{Q}$ such that for all $\epsilon \in I$ and for all $P \in \mathcal{Q}$, there exists a $t \in S$ such that $T(P, \epsilon) = t$ and $|t| \leq k$ and for all $Q \in \mathcal{Q}$, if Q is consistent with P on t, then $V(P \triangle Q) \leq \epsilon$. $T(P, \epsilon)$ is called a *test set for P with respect to the class $\mathcal{Q}$*. For each $t_i \in T(P, \epsilon)$, if $t_i \in P$ then t_i is a *positive test point*; otherwise, t_i is a *negative test point*.

Thus given a target object $P \in \mathcal{Q}$ and an error bound $\epsilon \in I$, T produces a test set for P such that any test object that is consistent with P on this set has error no more than ϵ. If such a T and k exist, then $\mathcal{Q}$ is *k-testable*. Note that in general a testing algorithm may produce test sets whose sizes are functions of both ϵ and the complexity of the target object, but the testing algorithms we develop in this paper produce constant size test sets.

3 Testing n-Vertex Simple Polygons

In this section we give testing algorithms for simple polygons with n vertices. First we address the case where the target polygon has no collinear edges, and later we handle collinearities.

[1] We use the standard definitions from topology for measurable sets and Borel sets. See [Rud66] for formal definitions of these concepts.

3.1 Testing with no collinear edges

We first develop a testing algorithm for n-sided simple polygons with test set size $6n$ (Theorem 1) subject to the following two assumptions, and then we improve it to achieve size $5n$ (Corollary 2).

Assumption 1 *The target polygon has no two collinear edges.*

Assumption 2 *The target polygon and the test polygon both have n vertices.*

Theorem 1 *The class of simple polygons in E^2 with n vertices and no collinear edges is $6n$-testable.*

Proof: Our goal is to use 6 test points to verify each edge of the target polygon. The 6 points are divided into 3 pairs of positive/negative points, one pair near each end of the edge, and the third pair somewhere in the middle of the edge (see Figure 2). More precisely, each pair of test points is represented by the line segment connecting the two points, which is called a *test-segment.* Assume that these test-segments have some length δ_1, and are placed along an edge. One segment is placed at a distance of at most δ_2 from each of the end vertices of the edge, and the third segment is placed between them, at a position chosen so that we do not create any unnecessary degeneracies. An *unnecessary degeneracy* is three pairs of test points on non-collinear edges that can be split by one line. Think of each pair of test points as being a single point on the boundary of the target polygon (i.e., $\delta_1 = 0$). By Assumption 1, we can place these points such that any line in the plane can stab at most three points, and if it does stab three, these must be points on one of the target polygon's edges. Now as we increase δ_1 from zero, creating pairs of test points, we can ensure, for δ_1 sufficiently small, that any line in the plane will separate at most three test pairs, and if a line does separate three pairs, these pairs are generated by one edge of the target polygon.

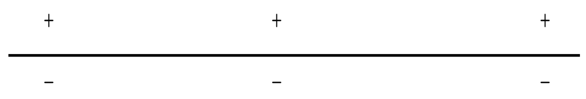

Figure 2: Six points verify an edge of a simple polygon

We use a simple counting argument to show that the edges of a test polygon consistent with the target polygon separate the test pairs in the same circular ordering as does the target polygon. Any n-sided simple polygon consistent on the test points must separate all $3n$ pairs of points with its n edges. By the placement of the test points and Assumption 1, no edge can separate more than two pairs of points except by separating three pairs along an edge. Therefore, in order to separate all $3n$ pairs with n edges, each edge must separate three pairs along some edge of the target polygon.

Finally, we choose parameters δ_1 and δ_2 based on the specific target polygon P and a specific choice of ϵ to guarantee that any n-sided simple polygon Q whose edges separate the $3n$ test pairs in the same circular ordering as the target polygon P satisfies $V(P \triangle Q) \leq \epsilon$. Clearly, $V(P \triangle Q)$ is a continuous function of the parameters δ_i, and decreases to zero as δ_1 and δ_2 approach zero. Thus for any choice of $\epsilon > 0$ one can choose $\delta_1, \delta_2 > 0$ small enough to achieve this error bound. □

Note that if a polygon has collinear vertices, but no collinear edges, our testing method of Theorem 1 applies, since we can always place the test pairs so as not to cause any unnecessary degeneracies.

We can reduce the number of test points to $5n$ by "reusing" either one positive or one negative test point at each vertex, depending on whether the interior angle of the polygon is convex or reflex, and making the three test points collinear (see Figure 3). However, to do this we must strengthen Assumption 1 as follows:

Assumption 3 *The target polygon has no three collinear vertices.*

If Assumption 3 holds, then one can find small disks centered at each vertex, such that no line can stab more than two such disks. If the three test points near a vertex are placed inside this disk, then no line can separate more than two of these pairs (or more than three pairs in total), and any line that separates three pairs must lie close to one of the edges of the target polygon.

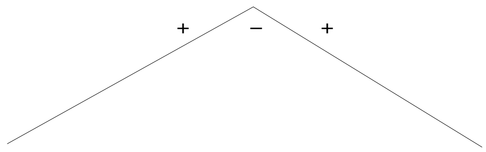

Figure 3: "Reusing" a test point at a vertex

Corollary 2 *The class of simple polygons in E^2 with n vertices, no three of which are collinear, is $5n$-testable.*

Since the proof of Theorem 1 relies only on the general position assumption, it also applies to simple polygons with holes[2].

Corollary 3 *The class of simple polygons with holes in E^2 with n vertices, no three of which are collinear, is $5n$-testable.*

3.2 Testing with collinear edges

The proof of Theorem 1 does not hold for polygons for which two or more edges may be collinear, because a single edge of the test polygon can now separate more than three pairs of test points. As Figure 4 shows, there are polygons that satisfy Assumption 2 but not Assumption 1, where the edges of the test polygon separate different pairs of points than the target. However, we note that these two polygons are very similar by our definition, as the area of the symmetric difference can be made arbitrarily small by shrinking δ_1. Thus, these polygons do not provide a counterexample to the testability of general n-sided polygons.

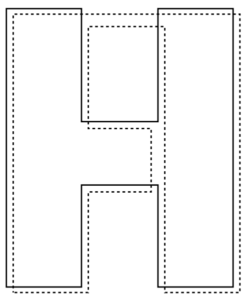

Figure 4: Polygons whose edges separate different pairs of points

We have been unable to identify a target polygon that is not testable with the testing algorithm of Theorem 1, but we have no proof that the testing algorithm is guaranteed to work for polygons

[2]The authors would like to thank Jit Bose for suggesting the problem of testing simple polygons with holes.

with collinear edges. Consequently, we eliminate Assumption 1 at the expense of adding more (but still a linear number of) test points. No testing algorithm can avoid having thin strips by which the target and test polygons may differ, as in Figure 4; that is, no testing algorithm can guarantee that clockwise traversals of the test and target polygons separate the test pairs in the same cyclical order. However, our next testing algorithm, which builds upon the $6n$-size testing scheme of Theorem 1, does guarantee that the two polygons have a small symmetric difference.

The revisions to the testing scheme of Theorem 1 focus on collinear edges. Let $e_1 = (v', v)$ and $e_2 = (w, w')$ be collinear edges of the target polygon such that their endpoints are in the order v', v, w, w' along the line containing both edges. If no other edge of the target polygon lies along the line segment between v and w, these edges are *consecutive*. If a clockwise traversal of the boundary of the polygon moves from v' to v and from w to w', or from w' to w and from v to v', then these edges are *aligned* (see Figure 5). If e_1 and e_2 are consecutive and aligned, then (v, w) will be called a *phantom edge* (depicted by a dashed line in our figures).

Augment the test set of Theorem 1 with three pairs of test points along each such phantom edge, called *phantom edge test pairs* or *phantom pairs*, placed so that they are collinear with the test pairs defined by real edges (*real pairs*), and they avoid any unnecessary degeneracies. Real pairs have one point of each sign; the phantom pairs along one phantom edge are all of the same sign – all positive or all negative (see Figure 5). If an edge of the test polygon separates pairs of test points generated by two or more collinear edges of the target polygon, then this edge will also erroneously separate the phantom pairs between these edges, so in order for the test polygon to be consistent with the target, an additional edge must separate these phantom pairs.

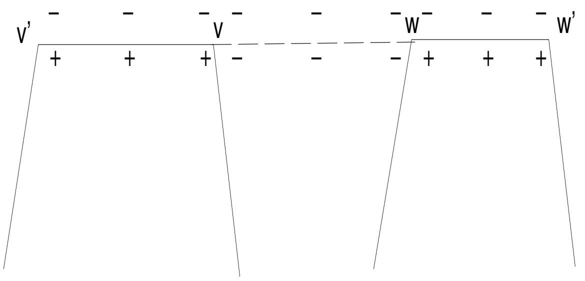

Figure 5: Testing degenerate polygons satisfying Assumption 2

To formalize this argument, we define the *value* of an edge of the test polygon, in a clockwise traversal of the polygon, to be the sum of the values of the real and phantom pairs it separates, as follows. If the directed edge separates a real pair, with the positive test point on its right side, the value is $+1$, whereas if the negative test point is on the right, the value is -1. For a phantom pair, we consider the positive side of the phantom edge generating the pair to be the same as the positive sides of the real edges bounding the phantom edge. The value of a phantom pair is -1 if the pair is separated with the positive side on the right of the separating edge, and $+1$ if it is on the left. See Figure 6 for an illustration, where the edges of the test polygon are denoted as arrows, and the edges of the target polygon by solid lines.

Lemma 4 *The value of each edge of the test polygon is at most three.*

Proof: Assume by contradiction that there exists an edge that has value 4 or more. A value of 4 can only be obtained by separating real pairs along consecutive, aligned edges or by separating phantom pairs along aligned phantom edges. However, an edge that separates real pairs on consecutive,

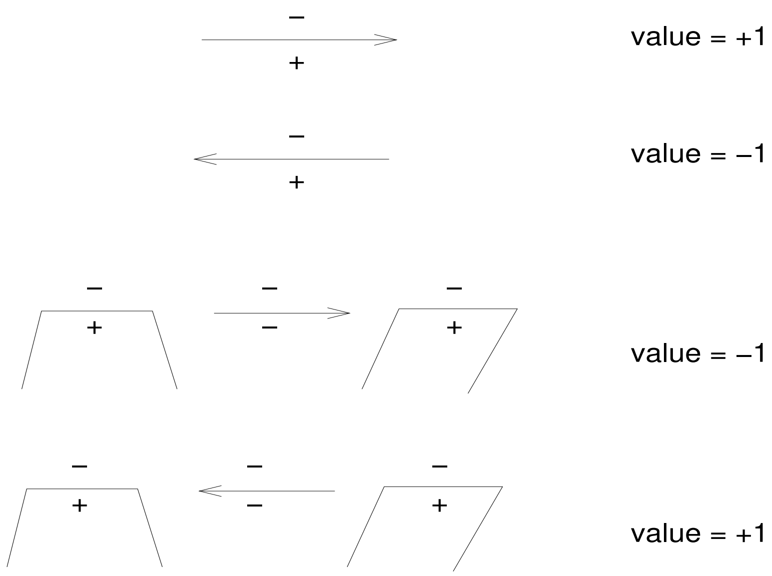

Figure 6: Assignment of value

aligned edges will also separate the pairs on the phantom edge between these edges, and these pairs will yield the opposite sign. Similarly, an edge that separates pairs from more than one phantom edge will separate real pairs giving the opposite sign. □

Theorem 5 *The class of simple polygons in E^2 with n vertices can be tested with $6n + 6l$ test points, where l is the number of phantom edges (pairs of consecutive, aligned edges) of the target polygon.*

Proof: We use the testing algorithm described above. Any test polygon that is consistent with the target polygon on the test points must have value exactly $3n$, since every such polygon must separate phantom pairs an equal number of times in each direction, and real pairs exactly once more with the positive test points on its right. The argument is now reminiscent of the proof of Theorem 1. Since each edge of the test polygon has value at most three (Lemma 4), and the number of edges is n, each edge must have value exactly three. This implies that the three pairs generated by a real or phantom edge of the target polygon are separated by the same edge(s) of the test polygon. As a result, all edges of the test polygon are "close to" edges of the target polygon, and the only place in which the target and test polygons can differ is in the "strips" between the two points of a test pair, as in Figure 4. However, the area of these strips can be made arbitrarily small by decreasing δ_1. Finally, we note that $l \leq n$, so the number of test points remains linear in n. □

The same testing strategy can be used for simple polygons with holes. The edges of a hole are traversed in a counter-clockwise order, so two consecutive collinear edges of the target polygon, $e_1 = (v', v)$ and $e_2 = (w, w')$, where e_2 is an edge of a hole, are *aligned* if a clockwise traversal of the boundary of the polygon moves from v' to v and counter-clockwise traversals of the boundaries of the holes moves from w to w', or visa versa. As before, a phantom edge occurs between every pair of consecutive, aligned edges. Also, the value of an edge of the test polygon is determined by making a clockwise traversal of the boundary of the polygon and counter-clockwise traversals of the boundaries of the holes, with the value of a pair being determined as before. By traversing the holes in counter-clockwise order, the arguments of Theorem 5 remain valid, and the following corollary results.

Corollary 6 *The class of simple polygons with holes in E^2 with n vertices can be tested with $6n+6l$ test points, where l is the number of phantom edges of the target polygon.*

4 Testing Simple Polygons with More Than n Vertices

In this section, we consider the question of whether a linear-size testing scheme can be devised when Assumption 2 is relaxed and the test polygon may have more vertices than the target polygon. We will assume that the vertices of the target polygon are in general position (Assumption 3). We also must make the following assumption:

Assumption 4 *There exists a bound, $\delta > 0$, such that for each test polygon, the minimum distance between a vertex and an edge of which it is not an endpoint is at least δ.*

Without Assumption 4, no finite-sized testing scheme is possible since a consistent test polygon can be created by adding one edge at a convex-hull vertex of the target polygon, forming a long, skinny "needle" (see Figure 7). This needle can be made arbitrarily long, and thus the area of the symmetric difference between the test and target polygons can be made arbitrarily large.

Even with Assumption 4, a test polygon with three more edges than the target polygon can fool any polynomial-sized testing scheme. This test polygon is created by forming a long "needle" somewhere in the middle of one of the externally visible edges of the target polygon, avoiding all test points. This needle replaces one edge with four edges, and the area of the symmetric difference between the two polygons can be made arbitrarily large by increasing the length of the needle. In order to avoid these needles, test pairs must be placed closely together along each externally visible edge, so that a needle placed between two adjacent test pairs would violate Assumption 4. However, this requires a number of test points that is a function of δ and the length of the edges of the target polygon.

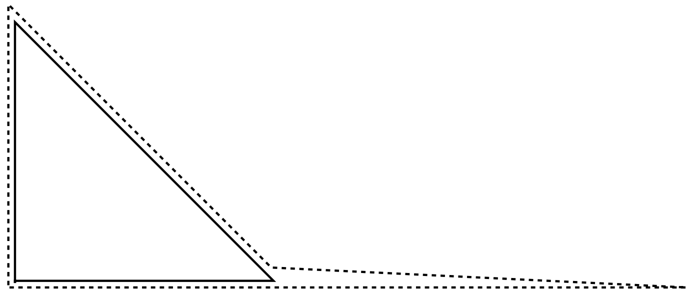

Figure 7: Polygons with $n + 1$ vertices cannot be tested

However, if we use Assumption 4 and we only allow the test polygon to have two more edges than the target polygon, then we do achieve positive results: a testing scheme of size $5n$ when the number of additional edges is at most one; and a testing scheme of size $7n$ when the number of additional edges is at most two. For the proofs of these testing algorithms we make use of the following lemma:

Lemma 7 *Given a line segment f and a point p at distance d from f, any line segment g that contains p but does not cross f must have an endpoint at distance at most d from f, or else f must have an endpoint at distance at most d from g.*

Proof: If f and g are parallel, then since g contains a point p at distance d from f, either f or g will have an endpoint at distance d from the other. If f and g are not parallel, then the proof is an easy case analysis, based on the intersection point of the two lines containing segments f and g, which can lie on exactly one of the segments or on neither segment. $\square$

4.1 Testing polygons with one additional vertex

We first give a testing algorithm for the case where the test polygon may have one more vertex than the target polygon.

Theorem 8 *Under Assumptions 3 and 4, a simple target polygon in E^2 with n vertices (possibly with holes) can be tested using $5n$ test points, provided that the test polygon has at most $n+1$ vertices.*

Proof: We place the test points as in Corollary 2, choosing δ_1 and δ_2 small enough to ensure that the distance from a vertex to any point on either of the two test-segments near it is less than $\frac{\delta}{2}$. The proof proceeds as in Theorem 1. There are $3n$ pairs of test points to be separated, and no one edge of the test polygon can separate more than three pairs. Therefore, any consistent test polygon must have at least n edges. Also, each test pair must be separated by exactly one edge, since two edges separating a pair would cause one of the test points to be inconsistent, and three or more edges separating a pair would, by Lemma 7, cause some vertex to be within δ of a nonadjacent edge, violating Assumption 4. We say that an edge of the target polygon *corresponds to* an edge of the test polygon if the two edges separate the same three pairs of test points. If the test polygon has n edges, then by choosing the parameters δ_i from the proof of Theorem 1 appropriately we can ensure that the error of the test polygon is less than ϵ. If the test polygon has $n+1$ edges, then there are three cases to consider:

1. All n edges of the target polygon have corresponding edges in the test polygon, and so one edge $f = (v_1, v_2)$ of the test polygon separates no pairs. Let f connect two edges f_1 and f_2 of the test polygon. If the corresponding edges e_1 and e_2 of the target polygon do not share a vertex, then the two edges of the target polygon incident on e_1 have corresponding test polygon edges f_3 and f_4, both of which must share an endpoint with f_1 by Lemma 7 and Assumption 4. But f and f_1 also share a vertex, providing a contradiction. If e_1 and e_2 do share a vertex v, then f_1 and f_2 both separate a pair of points near the vertex v and pass at most distance $\frac{\delta}{2}$ from vertex v. Thus either v_1 is within δ of the edge f_2 or v_2 is within δ of f_1, violating Assumption 4.

2. The test polygon Q has $n-1$ edges that separate three pairs each, and hence have corresponding edges in the target polygon P. Of the remaining two edges of the test polygon, one edge f_1 separates one pair and the other f_2 separates two pairs, all of which are separated by the single target edge e. Replace f_1 and f_2 by a single edge f separating all three pairs and adjust the adjacent edges to form a new polygon Q'. By Theorem 1 we know that for every value of $\epsilon_1 > 0$ we can choose the parameters δ_1 and δ_2 so that $V(P \triangle Q') \leq \epsilon_1$. To show that for any $\epsilon_2 > 0$ we can choose the parameters small enough so that $V(Q \triangle Q') \leq \epsilon_2$, we observe

first that as δ_1 is decreased to 0, each test pair becomes a single point on the boundary of the polygon, and $V(Q \triangle Q') = 0$. As the test points within each pair are moved continuously apart from one another, the maximum potential difference in slope between f_2 and f grows continuously. The limits on the slope of f_2 (determined by δ_1) and on the placement of the vertex v_f shared by f_2 and f_1 (dictated by Assumption 4) prevent f_1 from forming a long needle with either of its adjacent edges when δ_1 is increased from zero, so the area $V(Q \triangle Q')$ also grows continuously. Choosing $\epsilon_1 = \epsilon_2 = \epsilon/2$ and the parameters δ_i accordingly, completes the argument.

3. The test polygon has $n-2$ edges that separate three pairs each and hence have corresponding edges in the target polygon. Each of the remaining three edges of the test polygon f_1, f_2, and f_3 separates two pairs.

   The six pairs of test points separated by the f_i edges must be generated by two edges e_1 and e_2 of the target polygon that share a vertex v. Suppose not. Then at least one edge f_i separates a pair from each of e_1 and e_2. Furthermore, since any line segment separating the end pairs along an edge e_i will also separate the middle pair, there is an edge f_3 that separates an end pair from each of e_1 and e_2. Let e_4 and e_5 be the edges of the target polygon adjacent to these two endpoints of e_1 and e_2 respectively, and let f_4 and f_5 be the corresponding edges of the test polygon. Both edges f_4 and f_5 must share an endpoint with f_3 by Lemma 7 and Assumption 4. If a traversal of the edges of the target polygon starting from e_1 and followed by e_4 traverses e_5 before e_2, then this implies that the test polygon will contain a cycle – starting from f_3, followed by f_4, around to f_5 and back to f_3 – that does not contain f_1 and f_2, contradicting simplicity. On the other hand, if a traversal of the target polygon starting from e_1 and followed by e_4 traverses e_2 before e_5, then a traversal of the test polygon starting with f_3 and followed by f_4 will traverse edge f_5 in the opposite direction than the target polygon traverses it, so the test points separated by f_5 will be inconsistent. So edges e_1 and e_2 of the target polygon share a vertex v.

   Since the three test points forming two test pairs near vertex v are collinear, two edges f_1 and f_2 must separate these two pairs. Edges f_1 and f_2 must also each separate a middle pair along the same edge that they separate an end pair, since if either separated a pair on the opposite edge, one of the two pairs separated by the f_i edge would be inconsistent. Thus edge f_3 must separate the two test pairs near the endpoints of e_1 and e_2 other than v, and by Lemma 7 and Assumption 4, f_3 forms a cycle with the other test polygon edges not including f_1 and f_2, which contradicts simplicity. Therefore, by considering all possibilities we have shown that this case is impossible.

□

4.2 Testing polygons with two additional vertices

In order to test polygons that may have two more vertices than the target polygon, our testing algorithm must be modified. Consider a test polygon that looks like the target except that one of the externally visible edges of the target polygon is replaced by three edges in the test polygon, forming a "zig-zag" (see Figure 8). The test polygon is consistent with the target polygon, no matter how small the parameters δ_1 and δ_2 are chosen, but the area of the symmetric difference can be made arbitrarily large compared to the area of the target polygon.

To overcome this difficulty, one pair of test points can be added near each end of each edge of the target polygon at distance δ_3 from the vertex such that $\delta_3 > \delta_2$, every point on the test-segment

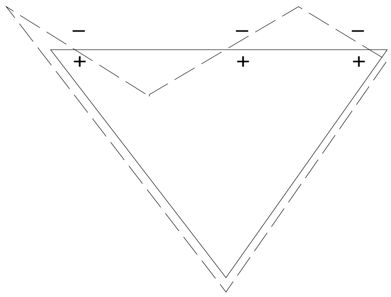

Figure 8: $5n$ test points of Corollary 2 do not suffice for testing a polygon with $n+2$ vertices

is within $\frac{\delta}{2}$ of the vertex, all the positive and all the negative test points along an edge are collinear, and unnecessary degeneracies are avoided. By eliminating the middle pair of test points on each edge the net gain is $2n$ test points. See Figure 9 for an illustration. To prove that this testing algorithm is correct, we employ several lemmas.

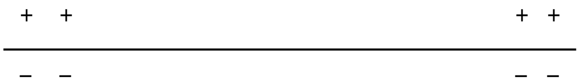

Figure 9: 8 points verify an edge of a polygon with $n+2$ vertices

Lemma 9 *The four pairs of test points near a vertex of the target polygon must be separated by exactly two edges of a successful test polygon.*

Proof: By our placement of test points, no line can separate all four pairs of test points near one vertex. If three or more edges separate the test pairs near a vertex, then there will be at least two edges not sharing an endpoint that pass within $\frac{\delta}{2}$ of the vertex, and by Lemma 7 an endpoint of one of these edges will be within δ of the other edge, violating Assumption 4. Therefore, exactly two edges must be used to separate the four pairs. $\square$

Lemma 10 *The four pairs of test points generated by a single edge of the target polygon can be separated by at most two different edges of a successful test polygon.*

Proof: If four pairs of test points generated by an edge of the target polygon are separated by three different edges of the test polygon, then the two test pairs near one endpoint of the target edge are separated by two different test edges. By Lemma 9 these two edges must also separate the two test pairs on the adjacent target edge near the vertex. The test edge separating the pair closest to the vertex cannot also separate the pair closest to the vertex on the adjacent edge because these two pairs are collinear. If the two edges separate the pairs on the adjacent edge in the opposite way, then they cross, contradicting simplicity. $\square$

As in the proof of Theorem 8, we say that an edge of the target polygon *corresponds to* an edge of the test polygon if the two edges separate the same four pairs of test points.

Lemma 11 *At least $n-2$ edges of the target polygon must have corresponding edges in a successful test polygon.*

Proof: If only $n-3$ edges of the target polygon have corresponding edges in the test polygon, then the test polygon (if it has $n+2$ edges) must use five edges to separate the remaining 12 test pairs. Therefore, some edge f of the test polygon must separate three test pairs, a feat possible only if these three pairs lie along one edge e of the target polygon. Consequently, f separates only one of the test pairs near one endpoint of e, so at least two more edges must be used to separate the remaining three pairs near that vertex, contradicting Lemma 9. □

Theorem 12 *Under Assumptions 3 and 4, a simple target polygon in E^2 with n vertices (possibly with holes) can be tested using $7n$ test points, provided that the test polygon has at most $n+2$ vertices.*

Proof: If the target polygon has n vertices, then there are $4n$ test pairs to be separated, and no one edge of the test polygon can separate more than four pairs. Therefore, any consistent test polygon must have at least n edges. Also, each test pair must be separated by exactly one edge, since two edges separating a pair would cause one of the test points to have an incorrect label, and three or more edges separating a pair would, by Lemma 7, cause some vertex to be within δ of a nonadjacent edge, violating Assumption 4. The proof mimics that of Theorem 8. We assume that the test polygon has $n+2$ edges, although the same arguments hold when the test polygon has $n+1$ edges. We consider the following three cases:

1. All n edges of the target polygon have corresponding edges in the test polygon, so two edges f_1 and f_2 of the test polygon separate no pairs, where f_1 and f_2 may or may not be adjacent in the test polygon. In either case, consideration of the edges of the test polygon adjacent to them will show, by Lemma 7, that there is a vertex that is "too close" to an edge of which it is not an endpoint, contradicting Assumption 4 as in the first case of Theorem 8.

2. There are $n-1$ edges of the test polygon that separate four pairs each, and hence have corresponding edges in the target polygon. The test pairs of the remaining edge of the target polygon are separated by two of the remaining three edges of the test polygon, by Lemma 10. Furthermore, each of these two edges of the test polygon separates exactly two test pairs, since if an edge separated a single test pair, then there would be three edges separating the four pairs near a vertex, which would contradict Lemma 9.

   If the three edges of the test polygon that do not correspond to edges of the target polygon are not consecutive along the boundary of the test polygon, then one of them must connect two edges of the test polygon that do correspond to edges of the target polygon, and as in Case (1) of Theorem 8, we get a contradiction to Assumption 4. If the three edges are consecutive, then the edge that separates no pairs cannot be extreme or the argument of Case (1) in Theorem 8 would produce a contradiction. Therefore, the edge of the test polygon that separates no test pairs must connect the two edges separating two test pairs each.

   An edge that separates two test pairs that are generated by a single edge of the target polygon must be very close in slope to that target edge by our choice of the parameters δ_i. Therefore, although it is possible to form a long needle with two of the three edges, as in Figure 10, such a needle will be close to a vertex that is not on the needle, contradicting Assumption 4. The length of any needle that does not violate Assumption 4 is limited by the length of the edge of the target polygon that the three edges replace, and the width is controlled by the parameters δ_i. Thus, as in Case (2) of Theorem 8, the three edges can be replaced by a single edge, and the resulting polygon will have a small symmetric difference with respect to the test polygon.

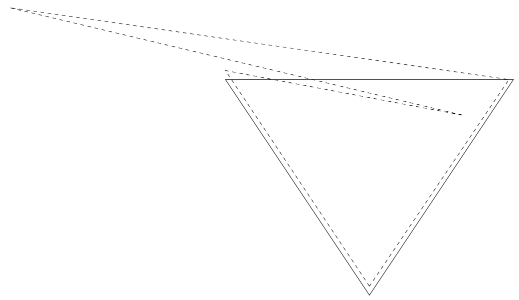

Figure 10: Needle formed with 3 edges is close to vertex

3. There are $n-2$ edges of the test polygon that separate four pairs each, and hence have corresponding edges in the target polygon. The test pairs of each of the remaining edges, e_1 and e_2, of the target polygon are separated by exactly two edges of the test polygon, by Lemma 10. Each edge of the test polygon separating test pairs along e_1 or e_2 must separate exactly two pairs, since if a test polygon edge separated three test pairs, then there would be at least three edges separating the four pairs near a vertex, violating Lemma 9. Therefore, arguments similar to those of Case (2) in Theorem 8 will complete the proof.

□

It is tempting to think that by possibly adding some constant number of test pairs per edge we could test polygons with $n+3$ vertices. However, as we noted before, this is false.

5 Lower bounds

In this section, we prove a lower bound of $3n+1$ on the size of the test set required for the class of simple polygons with n or $n+1$ vertices. This lower bound holds, in fact, for the class of simple polygons with n or $n+1$ vertices and at most one reflex vertex, even when Assumptions 3 and 4 are satisfied. As shown in Theorem 8, this class is testable with $5n$ test points. Romanik and Smith have already proven a lower bound of $2n$ for the class of convex polygons [RS94].

Theorem 13 *Under Assumptions 3 and 4, the class of simple polygons in E^2 with n or $n+1$ vertices and at most one reflex vertex cannot be tested with less than $3n+1$ test points.*

Proof: Let P be a regular convex n-gon whose edges have length 3δ, and let $\mathcal{Q}$ be the set of all simple polygons in E^2 with n or $n+1$ vertices and at most one reflex vertex. We prove that every test set for P with respect to $\mathcal{Q}$ requires at least $2n$ negative test points and $n+1$ positive test points. We will denote the vertices of P by $v_1, \ldots, v_n$, and the point of $\overline{v_i v_{i+1}}$ at distance d from v_i by $x_{i,d}$.

We first prove that at least $2n$ negative test points are required. For each edge e_i of P, let π_i be the ray originating at v_i and oriented from v_{i-1} to v_i, and let T_i^{cw} be the triangle whose vertices are v_i, $x_{i,\delta}$, and the point of π_i at distance δ from v_i, as shown in Figure 11. The area of T_i^{cw} is independent of ϵ, and $P \cup T_i^{cw}$ belongs to $\mathcal{Q}$, so T_i^{cw} must contain a negative test point when ϵ is small enough. Define T_i^{ccw} to be the reflection of T_i^{cw} through the perpendicular bisector of e_i, and consider the triangles $T_1^{cw}, \ldots, T_n^{cw}, T_1^{ccw}, \ldots, T_n^{ccw}$. Each such triangle must contain a negative test point. Since no two of these triangles intersect except at a vertex of P, this means that $2n$ negative test points are needed.