

Chasing Puppies: Mobile Beacon Routing on Closed Curves

Mikkel Abrahamsen ✉ 

Basic Algorithms Research Copenhagen (BARC), University of Copenhagen, Denmark

Jeff Erickson ✉ 

University of Illinois at Urbana-Champaign, USA

Irina Kostitsyna ✉

Eindhoven University of Technology, Netherlands

Maarten Löffler ✉

Utrecht University, Netherlands

Tillmann Miltzow ✉ 

Utrecht University, Netherlands

Jérôme Urhausen ✉

Utrecht University, Netherlands

Jordi Vermeulen ✉

Utrecht University, Netherlands

Giovanni Viglietta ✉ 

Japan Advanced Institute of Science and Technology, Nomi City, Ishikawa, Japan

Abstract

We solve an open problem posed by Michael Biro at CCCG 2013 that was inspired by his and others' work on beacon-based routing. Consider a human and a puppy on a simple closed curve in the plane. The human can walk along the curve at bounded speed and change direction as desired. The puppy runs with unbounded speed along the curve as long as the Euclidean straight-line distance to the human is decreasing, so that it is always at a point on the curve where the distance is locally minimal. Assuming that the curve is smooth (with some mild genericity constraints) or a simple polygon, we prove that the human can always catch the puppy in finite time.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Beacon routing, navigation, generic smooth curves, puppies

Supplementary Material <https://github.com/viglietta/Chasing-Puppies>

Funding *Mikkel Abrahamsen*: Partially supported by the VILLUM Foundation grant 16582.

Maarten Löffler: Partially supported by the Dutch Research Council (NWO) under project number 614.001.504 and 628.011.005.

Tillmann Miltzow: Supported by the Dutch Research Council (NWO) under Veni grant EAGER

Jérôme Urhausen: Supported by the Dutch Research Council (NWO); 612.001.651.


Jordi Vermeulen: Supported by the Dutch Research Council (NWO); 612.001.651.

Acknowledgements Thanks to Ivor van der Hoog, Marc van Kreveld, and Frank Staals for helpful discussions in the early stages of this work, and to Joseph O'Rourke for clarifying the history of the problem. Portions of this work were done while the second author was visiting Utrecht University.

1 Introduction

You have lost your puppy somewhere on a simple closed curve. Both of you are forced to stay on the curve. You can see each other and both want to reunite. The problem is that the puppy runs infinitely faster than you, and it believes naively that it is always a good idea to minimize its straight-line distance to you. What do you do?


To be more precise, let $\gamma: S^1 \hookrightarrow \mathbb{R}^2$ be a simple closed curve in the plane, which we informally call the *track*. Two special points move around the track, called the *puppy* p and the *human* h . The human can walk along the track at bounded speed and change direction as desired. The puppy runs with unbounded speed along the track as long as its Euclidean straight-line distance to the human is decreasing, until it reaches a point on the curve where the distance is locally minimized. As the human moves along the track, the puppy moves to stay at a local distance minimum. The human's goal is to move in such a way that the puppy and the human meet. See Figure 1 for a simple example.



■ **Figure 1** Catching the puppy.

In this paper we show that it is always possible to reunite with the puppy, under the assumption that the curve is well-behaved in a sense to be defined.

The problem was posed in a different guise at the open problem session of the 25th Canadian Conference on Computational Geometry (CCCG 2013) by Michael Biro. In Biro's formulation, the track was a railway, the human a locomotive, and the puppy a train carriage that was attracted to an infinitely strong magnet installed in the locomotive.



■ **Figure 2** If the human keeps walking counterclockwise from h_1 , the human and the puppy will never meet. To the right are closeups of two of the spikes of the star.

Returning to our formulation of catching a puppy, it was also asked if the human will always catch the puppy by choosing an arbitrary direction and walking only in that direction. This turns out not to be the case; consider the star-shaped track in Figure 2. Suppose the human and puppy start at points h_1 and p_1 , respectively, and the human walks counterclockwise around the track. When the human reaches h_2 , the puppy runs from p_2 to p'_2 . When the human reaches h_3 , the puppy runs from p_3 to p'_3 . Then the pattern repeats

indefinitely. Examples of this type, where the human walking in the wrong direction will never catch the puppy, were independently discovered during the conference by some of the authors and by David Eppstein.

1.1 Related work

Biro's problem was inspired by his and others' work on *beacon-based geometric routing*, a generalization of both greedy geometric routing and the art gallery problem introduced at the 2011 Fall Workshop on Computational Geometry [7] and the 2012 Young Researchers Forum [8], and further developed in Biro's PhD thesis [6] and papers [9, 10]. A *beacon* is a stationary point object that can be activated to create a "magnetic pull" towards itself everywhere in a given polygonal domain P . When a beacon at point b is *activated*, a point object p moves greedily to decrease its Euclidean distance to b , alternately moving through the interior of P and sliding along its boundary, until it either reaches b or gets stuck at a "dead point" where Euclidean distance is minimized. By activating different beacons one at a time, one can route a moving point object through the domain. Initial results for this model by Biro and his colleagues [6–10] sparked significant interest and subsequent work in the community [2, 3, 5, 13, 18, 20–22, 26]. More recent works have also studied how to utilize objects that repel points instead of attracting them [11, 24].

Biro's problem can also be viewed as a novel variant of classical *pursuit* problems, which have been an object of intense study for centuries [25]. The oldest pursuit problems ask for a description of the *pursuit curve* traced by a *pursuer* moving at constant speed directly toward a *target* moving along some other curve. Pursuit curves were first systematically studied by Bouguer [12] and de Maupertuis [14] in 1732, who used the metaphor of a pirate overtaking a merchant ship; another notable example is Hathaway's problem [16], which asks for the pursuit curve of a dog swimming at unit speed in a circular lake directly toward a duck swimming at unit speed around its circumference. In more modern *pursuit-evasion* problems, starting with Rado's famous "lion and man" problem [23, pp.114–117], the pursuer and target both move strategically within some geometric domain; the pursuer attempts to *capture* the target by making their positions coincide while the target attempts to evade capture. Countless variants of pursuit-evasion problems have been studied, with multiple pursuers and/or targets, different classes of domains, various constraints on motion or visibility, different capture conditions, and so on. Biro's problem can be naturally described as a *cooperative pursuit* or *pursuit-attraction* problem, in which a strategic target (the human) *wants* to be captured by a greedy pursuer (the puppy).

Kouhestani and Rappaport [19] studied a natural variant of Biro's problem, which we can recast as follows. A *guppy* is restricted to a closed and simply-connected *lake*, while the human is restricted to the boundary of the lake. The guppy swims with unbounded speed to decrease its Euclidean distance to the human as quickly as possible. Kouhestani and Rappaport described a polynomial-time algorithm that finds a strategy for the human to catch the guppy, if such a strategy exists, given a simple polygon as input; they also conjectured that a capturing strategy always exists. Abel, Akitaya, Demaine, Demaine, Hesterberg, Korman, Ku, and Lynch [1] recently proved that for some polygons and starting configurations, the human cannot catch the guppy, even if the human is allowed to walk in the exterior of the polygon, thereby disproving Kouhestani and Rappaport's conjecture. Their simplest counterexample is an orthogonal polygon with about 50 vertices.

1.2 Our results

Before describing our results in detail, we need to carefully define the terms of the problem. The *track* is a simple closed curve $\gamma: S^1 \hookrightarrow \mathbb{R}^2$. We consider the motion of two points on this curve, called the *human* (or *beacon* or *target*) and the *puppy* (or *pursuer*). A *configuration* is a pair $(x, y) \in S^1 \times S^1$ that specifies the locations $h = \gamma(x)$ and $p = \gamma(y)$ for the human and puppy, respectively. Let $D(x, y)$ denote the straight-line Euclidean distance between these two points. When the human is located at $h = \gamma(x)$, the puppy moves from $p = \gamma(y)$ to greedily decrease its distance to the human, as follows.

- If $D(x, y + \varepsilon) < D(x, y)$ for all sufficiently small $\varepsilon > 0$, the puppy runs forward along the track, by increasing the parameter y .
- If $D(x, y - \varepsilon) < D(x, y)$ for all sufficiently small $\varepsilon > 0$, the puppy runs backward along the track, by decreasing the parameter y .

If both of these conditions hold, the puppy runs in an arbitrary direction. While the puppy is running, the human remains stationary. If neither condition holds, the configuration is *stable*; the puppy does not move until the human does. When the configuration is stable, the human can walk in either direction along the track; the puppy walks along the track in response to keep the configuration stable, until it is forced to run again. The human's goal is to *catch* the puppy; that is, to reach a configuration in which the two points coincide.

Our main result is that the human can always catch the puppy in finite time, starting from any initial configuration, provided the track is either a generic simple smooth curve or an arbitrary simple polygon.

The remainder of the paper is structured as follows. We begin in Section 2 by considering some variants and special cases of the problem. In particular, we give a simple self-contained proof of our main result for the special case of orthogonal polygons.

We consider generic smooth tracks in Sections 3 and 4. Specifically, in Section 3 we define two important diagrams, which we call the *attraction diagram* and the *dual attraction diagram*, and prove some useful structural results. At a high level, the attraction diagram is a decomposition of the configuration space $S^1 \times S^1$ according to the puppy's behavior, similar to the *free space diagrams* introduced by Alt and Godau to compute Fréchet distance [4]. We show that for a sufficiently generic smooth track, the attraction diagram consists of a finite number of disjoint simple closed *critical* curves, exactly two of which are topologically nontrivial. Then in Section 4, we argue that the human can catch the puppy on any track whose attraction diagram has this structure.

In Section 5, we describe an extension of our analysis from smooth curves to simple polygonal tracks. Because polygons do not have well-defined tangent directions at their vertices, this extension requires explicitly modeling the puppy's direction of motion in addition to its location. We first prove that the human can catch the puppy on a polygon that has no acute vertex angles and where no three vertices form a right angle; under these conditions, the attraction diagram has exactly the same structure as for generic smooth curves. We then reduce the problem for arbitrary simple polygons to this special case by *chamfering*—cutting off a small triangle at each vertex—and arguing that any strategy for catching the puppy on the chamfered track can be pulled back to the original polygon.


Finally, we close the paper by suggesting several directions for further research.

2 Warmup: other settings and a special case

In this section, we discuss two variants of Biro’s problem and the special case of orthogonal polygons.

In the first variant, both the human h and the puppy p are allowed to move anywhere in the interior and on the boundary of a simple polygon P . Here, as in beacon routing and Kouhestani and Rappaport’s variant [1, 19], the puppy moves greedily to decrease its Euclidean distance to the human, alternately moving through the interior of P and sliding along its boundary.

As we will show in Theorem 1, h has a simple strategy to catch p in this setting, essentially by walking along the dual graph of any triangulation. This is an interesting contrast to the proof by Abel et al. [1] that h and p cannot always meet when h is restricted to the *exterior* of P and p to the interior. Our main result that h and p can meet when both are restricted to the *boundary* of P (even for a much wider class of simple closed curves), somehow sits in between these other two variants.




■ **Figure 3** The proposed strategy when h and p are restricted to the interior of a simple polygon P . The human h will follow the path π . Note that the triangle containing p will change before h reaches d_1 , and π will be updated accordingly.

When both h and p are restricted to the interior of P , we propose the following strategy for h ; see Figure 3. Let \mathcal{T} be a triangulation of P and let t_1, \dots, t_k be the path of pairwise adjacent triangles in \mathcal{T} such that $h \in t_1$ and $p \in t_k$. Let e_i be the common edge of t_i and t_{i+1} and let d_i be the midpoint of e_i . Let $\pi = hd_1d_2 \dots d_{k-1}$ be a path from h to d_{k-1} , which is contained in the triangles t_1, \dots, t_{k-1} . The human starts walking along π . As soon as the puppy enters a new triangle, the human recomputes π as described and follows the new path.

► **Theorem 1.** *The proposed strategy will make h and p meet.*

Proof. First, we observe that if the puppy ever enters the triangle t_1 that is occupied by the human, then the puppy and the human will meet immediately. Assume that the human does not meet the puppy right from the beginning. The region $P \setminus t_1$ consists of one, two, or three polygons, one of which P_p contains p . Thus, whenever the human moves from one triangle to another, the set of triangles that can possibly contain p shrinks. We conclude that the human and the puppy must meet eventually. ◀

In our second variant, the human and the puppy are both restricted to a simple, closed curve γ in \mathbb{R}^3 . Here it is easy to construct curves on which h and p will never meet; the simplest example is a “double loop” that approximately winds twice around a planar circle, as shown in Figure 4.



■ **Figure 4** A double loop in \mathbb{R}^3 ; the human and puppy will never meet.

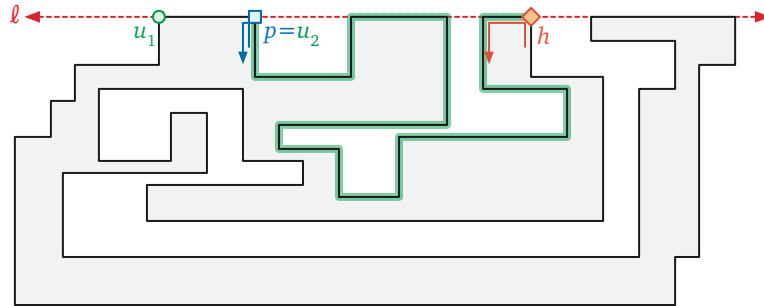
Finally, we consider the special case of Biro’s original problem where the track γ is the boundary of an orthogonal polygon in the plane. This special case of our main results admits a much simpler self-contained proof.

► **Theorem 2.** *The human can catch the puppy on any simple orthogonal polygon, by walking counterclockwise around the polygon at most twice.*

Proof. Let P be an arbitrary simple orthogonal polygon. Let u_1 be its leftmost point with the maximum y -coordinate, and u_2 be the next boundary vertex of P in the clockwise order (see Figure 5). Finally, let ℓ be the horizontal line supporting the segment u_1u_2 .

We break the motion of the human into two phases. In the first phase, the human moves counterclockwise around P from their starting location to u_1 . If the human catches the puppy during this phase, we are done, so assume otherwise. In the second phase, the human walks counterclockwise around P starting from u_1 to u_2 .

We claim that the puppy p is never in the interior of the segment u_1u_2 during the second phase; thus, p always lies on the closed counterclockwise subpath of P from h to u_2 (or less formally, “between h and u_2 ”). This claim implies that the human and the puppy are united during the second phase on u_2 at the latest.



■ **Figure 5** Proof of Theorem 2. During the human’s second trip around P , the puppy lies between u_2 and the human.

The puppy must first be at u_2 if it ever wants to be in the interior of u_1u_2 . So consider any moment during the second phase when p moves upward to the vertex u_2 . At that moment, h must be on the line ℓ to the right of p . (For any point a below ℓ , there is a point b below u_2 that is closer to a than u_2 .) Thus, the puppy will stay on u_2 as long as h is on ℓ . As soon as h leaves ℓ the puppy will leave u_2 downward. Thus the puppy can never go to the interior of the edge u_1u_2 . ◀

The star-shaped track in Figure 2 shows that this simple argument does not extend to arbitrary polygons, even with a constant number of edge directions.

3 Diagrams of smooth tracks

We first formalize both the problem and our solution under the assumption that the track is a generic smooth simple closed curve $\gamma: S^1 \hookrightarrow \mathbb{R}^2$. In particular, for ease of exposition, we assume that γ is regular and C^3 , meaning it has well-defined continuous first, second, and third derivatives, and its first derivative is nowhere zero. We also assume γ satisfies some additional genericity constraints, to be specified later. We consider polygonal tracks in Section 5.

3.1 Configurations and genericity assumptions

We analyze the behavior of the puppy in terms of the *configuration space* $S^1 \times S^1$, which is the standard torus. Each configuration point $(x, y) \in S^1 \times S^1$ corresponds to the human being located at $h = \gamma(x)$ and the puppy being located at $p = \gamma(y)$.

For any configuration (x, y) , recall that $D(x, y)$ denotes the straight-line Euclidean distance between the points $\gamma(x)$ and $\gamma(y)$. We classify all configurations $(x, y) \in S^1 \times S^1$ into three types, according to the sign of the partial derivative of distance with respect to the puppy's position.

- (x, y) is a *forward* configuration if $\frac{\partial}{\partial y} D(x, y) < 0$.
- (x, y) is a *backward* configuration if $\frac{\partial}{\partial y} D(x, y) > 0$.
- (x, y) is a *critical* configuration if $\frac{\partial}{\partial y} D(x, y) = 0$.


Starting in any forward (resp. backward) configuration, the puppy automatically runs forward (resp. backward) along the track γ . Genericity implies that there are a finite number of critical configurations (x, y) with any fixed value of x , or with any fixed value of y . We further classify the critical configurations as follows:

- (x, y) is a *stable* critical configuration if $\frac{\partial^2}{\partial y^2} D(x, y) > 0$.
- (x, y) is an *unstable* critical configuration if $\frac{\partial^2}{\partial y^2} D(x, y) < 0$.
- (x, y) is a *forward pivot* configuration if $\frac{\partial^2}{\partial y^2} D(x, y) = 0$ and $\frac{\partial^3}{\partial y^3} D(x, y) < 0$.
- (x, y) is a *backward pivot* configuration if $\frac{\partial^2}{\partial y^2} D(x, y) = 0$ and $\frac{\partial^3}{\partial y^3} D(x, y) > 0$.

In any stable configuration, the puppy's distance to the human is locally minimized, so the puppy does not move unless the human moves. In any unstable configuration, the puppy can decrease its distance by running in either direction. Finally, in any forward (resp. backward) pivot configuration, the puppy can decrease its distance by moving in one direction but not the other, and thus automatically runs forward (resp. backward) along the track.

Critical points can also be characterized geometrically as follows. Refer to Figure 6. A configuration (x, y) is critical if the human $\gamma(x)$ lies on the line $N(y)$ normal to γ at the puppy's location $\gamma(y)$. Let $C(y)$ denote the center of curvature of the track at $\gamma(y)$. Then (x, y) is a pivot configuration if $\gamma(x) = C(y)$, a stable critical configuration if the open ray from $C(y)$ through the human point $\gamma(x)$ contains the puppy point $\gamma(y)$, and an unstable critical configuration otherwise.

Genericity of the track γ implies that this classification of critical configurations is exhaustive, and moreover, that the set of pivot configurations is finite. In particular, our analysis requires that in any pivot configuration (x, y) , the puppy point $\gamma(y)$ is not a



■ **Figure 6** Three critical configurations: (h_1, p) is unstable; (h_2, p) is a pivot configuration, and (h_3, p) is stable.

local curvature minimum or maximum.¹ Otherwise, we would need higher derivatives to disambiguate the puppy's behavior. In the extreme case where γ contains both an open circular arc α and its center c , all configurations where $h = c$ and $p \in \alpha$ are stable.

3.2 Attraction diagrams

The *attraction diagram* of the track γ is a decomposition of the configuration space $S^1 \times S^1$ by critical configurations. Our genericity assumptions imply that the set of critical points—the common boundary of the forward and backward configurations—is the union of a finite number of disjoint simple closed curves, which we call *critical cycles*. At least one of these critical cycles, the main diagonal $x = y$, consists entirely of stable configurations; critical cycles can also consist entirely of unstable configurations. If a critical cycle is neither entirely stable nor entirely unstable, then its points of vertical tangency are pivot configurations, and these points subdivide the curve into x -monotone paths, which alternately consist of stable and unstable configurations.

Figure 7 shows a sketch of the attraction diagram of a simple closed curve. We visualize the configuration torus $S^1 \times S^1$ as a square with opposite sides identified. Green and red paths indicate stable and unstable configurations, respectively; blue dots indicate pivot configurations; and backward configurations are shaded light gray. Figure 8 shows the attraction diagram for a more complex polygonal track, with slightly different coloring conventions. (Again, we will discuss polygonal tracks in more detail in Section 5.)

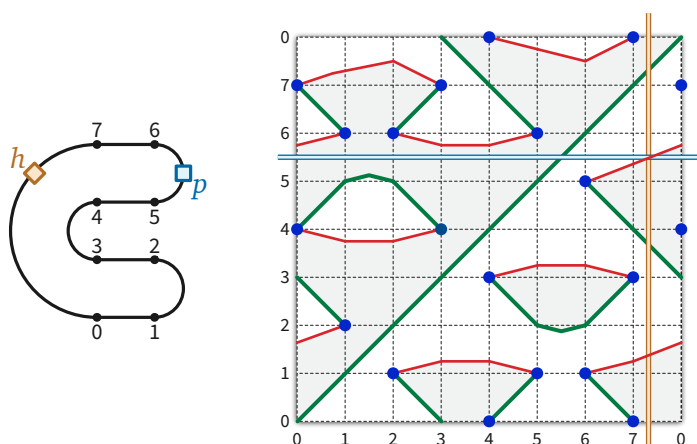
The cycles in any attraction diagram have a simple but important topological structure. A critical cycle in the attraction diagram is *contractible* if it is the boundary of a simply connected subset of the torus $S^1 \times S^1$ and *essential* otherwise. For example, the main diagonal is essential, and the attraction diagram in Figure 7 contains two contractible critical cycles and two essential critical cycles.

► **Lemma 3.** *The attraction diagram of any generic closed curve contains an even number of essential critical cycles.*

Proof. This lemma follows immediately from standard homological arguments, but for the sake of completeness we sketch a self-contained proof.

Fix a generic closed curve γ . Let α be the horizontal cycle $\{(0, y) \mid y \in S^1\}$, and let β be the vertical cycle $\{(x, 0) \mid x \in S^1\}$ in the torus $S^1 \times S^1$. Without loss of generality, assume α and β intersect every critical cycle in the attraction diagram of γ transversely.

¹ More concretely, we assume the track γ intersects its evolute (the locus of centers of curvature) transversely, away from its cusps.



■ **Figure 7** The attraction diagram of a simple closed curve, with one unstable critical configuration emphasized.

A critical cycle C in the attraction diagram is contractible if and only if α and β each cross C an even number of times. (Indeed, this parity condition characterizes all simple contractible closed curves in the torus.) On the other hand, α and β each cross the main diagonal once. It follows that α and β each cross *every* essential critical cycle an odd number of times; otherwise, some pair of essential critical cycles would intersect.

Because the critical cycles are the boundary between the forward and backward configurations, α and β each contain an even number of critical points. The lemma now follows immediately. ◀

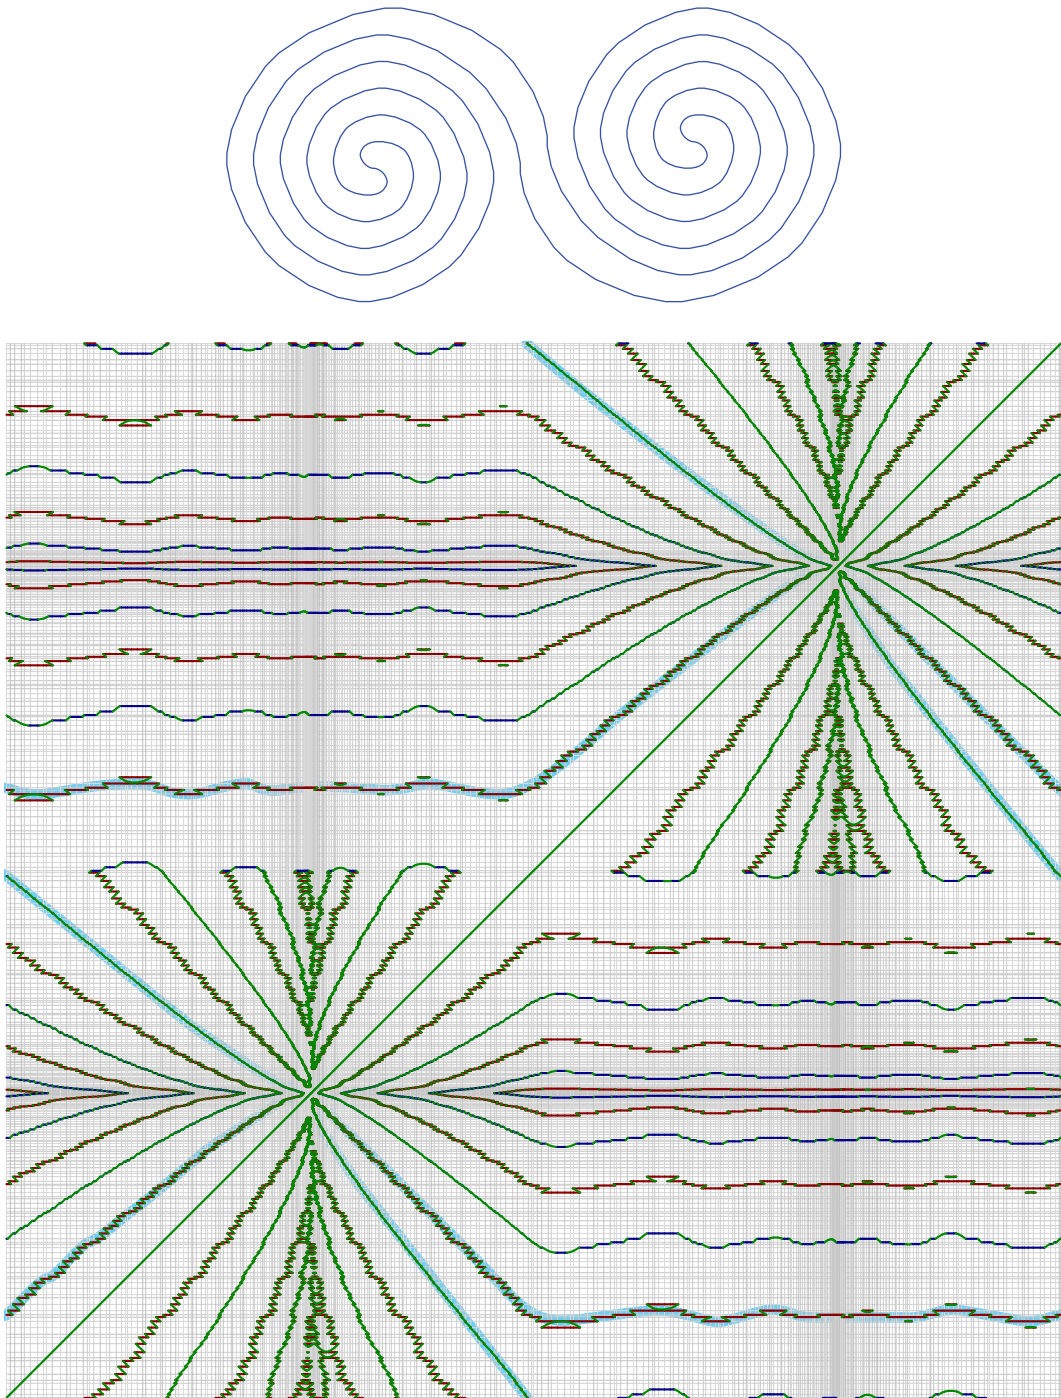
We emphasize that this lemma does *not* actually require the track γ to be simple; the argument relies only on properties of generic functions over the torus that are minimized along the main diagonal.

3.3 Dual attraction diagrams

Our analysis also relies on a second diagram, which we call the *dual attraction diagram* of the track. We hope the following intuition is helpful. While the attraction diagram tells us the possible positions of the puppy depending on the position of the human, the dual attraction diagram gives us the possible positions of the human depending on the position of the puppy. For each puppy configuration $y \in S^1$, we consider the normal line $N(y)$. We are interested in the intersection points of γ with $N(y)$, as those are the possible positions of the human. The idea of the dual attraction diagram is to trace the positions of the human as a function of the position of the puppy, see Figure 10.

Let $T(y)$ denote the directed line tangent to γ at the point $\gamma(y)$. For any configuration (x, y) , let $\ell(x, y)$ denote the distance from $\gamma(x)$ to the tangent line $T(y)$, signed so that $\ell(x, y) > 0$ if the human point $\gamma(x)$ lies to the left of $T(y)$ and $\ell(x, y) < 0$ if $\gamma(y)$ lies to the right of $T(y)$. More concisely, assuming without loss of generality that the track γ is parameterized by arc length, $\ell(x, y)$ is twice the signed area of the triangle with vertices $\gamma(x)$, $\gamma(y)$, and $\gamma(y) + \gamma'(y)$.


Let $L: S^1 \times S^1 \rightarrow S^1 \times \mathbb{R}$ denote the function $L(x, y) = (y, \ell(x, y))$. The dual attraction diagram is the decomposition of the infinite cylinder $S^1 \times \mathbb{R}$ by the points $\{L(x, y) \mid (x, y) \text{ is critical}\}$. At the risk of confusing the reader, we refer to the image $L(x, y) \in S^1 \times \mathbb{R}$



■ **Figure 8** The attraction diagram of a complex simple polygon. Serrations in the diagram are artifacts of the curve being polygonal instead of smooth. The river is highlighted in blue.


of any critical configuration (x, y) as a critical point of the dual attraction diagram.

The dual attraction diagram can also be described as follows. For any $y \in S^1$ and $d \in \mathbb{R}$, let $\Gamma(y, d)$ denote the point on the normal line $N(y)$ at distance d to the left of the tangent vector $\gamma'(y)$. More formally, assuming without loss of generality that γ is parametrized by arc length, we have $\Gamma(y, d) = \gamma(y) + d \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \gamma'(y)$. We emphasize that $\Gamma(y, d)$ does not necessarily lie on the curve γ . The dual attraction diagram is the decomposition of the cylinder $S^1 \times \mathbb{R}$ by the preimage $\Gamma^{-1}(\gamma)$ of γ .



■ **Figure 9** Examples of the functions ℓ and Γ used to define the dual attraction diagram.

Because γ is simple and regular, the dual attraction diagram is the union of simple disjoint closed curves. The function L continuously maps each critical cycle in the attraction diagram to a closed curve in the cylinder $S^1 \times \mathbb{R}$. Thus, the restriction of L to the set of critical configuration is a homeomorphism onto its image in the dual attraction diagram. In particular, L maps the main diagonal $x = y$ to the horizontal axis $\ell(x, y) = 0$ of the dual attraction diagram. We emphasize, however, that the two diagrams are not topologically equivalent. Figure 10 shows the dual attraction diagram of the same track whose attraction diagram is shown in Figure 7; here preimages of points inside the track are shaded.



■ **Figure 10** The dual attraction diagram of a simple closed curve, with one critical configuration emphasized. Compare with Figure 7.

► **Lemma 4.** *For any generic simple closed curve γ , the attraction diagram of γ and the dual attraction diagram of γ contain the same number of essential critical cycles.*

Proof. Let α denote the horizontal cycle $y = 0$ in the torus $S^1 \times S^1$, and let α' be the vertical line $y = 0$ in the infinite cylinder $S^1 \times \mathbb{R}$. Let C be any critical cycle on the attraction diagram, and let $C' = L(C)$ be the corresponding critical cycle in the dual attraction diagram.

Recall from the proof of Lemma 3 that C is contractible on the torus if and only if $|C \cap \alpha|$ is even. Similarly, C' is contractible in the cylinder if and only if $|C' \cap \alpha'|$ is even. The map

$L: S^1 \times S^1 \rightarrow S^1 \times \mathbb{R}$ maps $C \cap \alpha$ bijectively to $C' \cap \alpha'$. We conclude that C is essential if and only if C' is essential. ◀

With this correspondence in hand, we can now more carefully describe the topological structure of the *attraction* diagram when the track is simple.

► **Lemma 5.** *The attraction diagram of a **simple** generic closed curve contains **two** essential critical cycles.*

Proof. Fix a generic closed curve γ . Lemma 3 implies that the attraction diagram of γ contains at least two essential critical cycles, one of which is the main diagonal. Thus, to prove the lemma, it remains to show that there are *at most* two essential critical cycles, in either the attraction diagram or the dual attraction diagram.

Let $\Sigma \subset S^1 \times \mathbb{R}$ denote the set of essential critical cycles in the *dual attraction* diagram. Any two cycles in Σ are homotopic—meaning one can be continuously deformed into the other—because there is only one nontrivial homotopy class of simple cycles on the infinite cylinder $S^1 \times \mathbb{R}$. It follows that the cycles in Σ have a well-defined vertical total order. In particular, the highest and lowest intersection points between any vertical line and Σ always lie on the *same* two essential cycles in Σ .

Without loss of generality, suppose $\gamma(0)$ is a point on the convex hull of γ with a unique tangent line. Let C be any essential critical cycle in the attraction diagram of γ , and let $C' = L(C)$ denote the corresponding essential cycle in the dual attraction diagram. C must pass through all possible puppy positions *and* all possible human positions; thus, C contains a configuration $(0, y)$ for some parameter $y \in S^1$. Recall that $N(y)$ denotes the line normal to γ at $\gamma(y)$. Then $\gamma(0)$ must also lie on the convex hull of $\gamma \cap N(y)$. We conclude that C' must be either the highest or lowest essential critical cycle in the dual attraction diagram. We conclude that there are at most two critical cycles, completing the proof. ◀

In the rest of the paper, we mnemonically refer to the two essential critical cycles in the attraction diagram of a simple track as the *main diagonal* and the *river*.

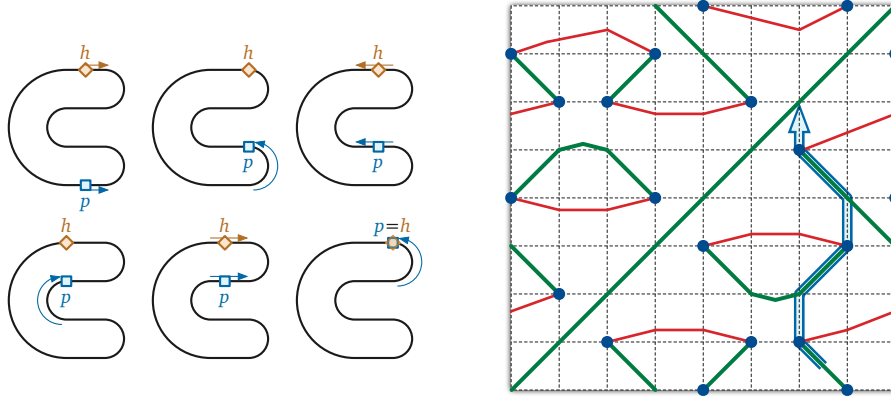
We emphasize that the converse of Lemma 5 is false; there are non-simple tracks whose attraction diagrams have exactly two essential critical cycles. (Consider the figure-eight curve ∞ .) Moreover, we conjecture that Lemma 5 can be generalized to all (smooth) tracks with turning number ± 1 .

4 Dexter and sinister strategies

We can visualize any strategy for the human to catch the puppy as a path through the attraction diagram that consists entirely of segments of stable critical paths and vertical segments, as shown in Figure 11. We refer to the vertical segments as *pivots*. Every pivot (except possibly the first) starts at a pivot configuration, and every pivot ends at a stable configuration.


We call a strategy *dexter* if it ends with a backward pivot—a *downward* segment, approaching the main diagonal to the *right*—and we call a configuration (x, y) *dexter* if there is a dexter strategy for catching the puppy starting at (x, y) . Similarly, a strategy is *sinister* if it ends with a forward pivot—a *skyward* segment, approaching the main diagonal to the *left*—and a configuration is sinister if it is the start of a sinister strategy.² A single

² *Dexter* and *sinister* are Latin for right (or skillful, or fortunate, or proper, from a Proto-Indo-European root meaning “south”) and left (or unlucky, or unfavorable, or malicious), respectively.



■ **Figure 11** A sinister strategy for catching the puppy; compare with Figures 1 and 7.

configuration can be both dexter and sinister; see Figure 12.



■ **Figure 12** Dexter (orange) and sinister (cyan) configurations in the example attraction diagram. Arrows on the stable critical paths describe dexter and sinister strategies for catching the puppy.

► **Theorem 6.** *Let γ be a generic track whose attraction diagram has exactly two essential critical cycles. Every configuration on γ is either dexter or sinister; thus, the human can catch the puppy on γ from any starting configuration.*

Before giving the proof, we emphasize that Theorem 6 does not require the track γ to be simple. Also, it is an open question whether having exactly two essential critical cycle curves is a *necessary* condition for the human to always be able to catch the puppy. (We conjecture that it is not.)

Proof. Fix a generic track γ whose attraction diagram has exactly two essential critical cycles, which we call the *main diagonal* and the *river*. Assume γ has at least one pivot configuration, since otherwise, from any starting configuration, the puppy runs directly to the human.

Let D be the set of all dexter configurations, and let S be the set of all sinister configurations. We claim that D and S are both annuli that contain both the main diagonal and the river. Because S and D meet on opposite sides of the main diagonal, this claim implies that $D \cup S$ is the entire torus, completing the proof of the lemma. We prove our claim explicitly for D ; a symmetric argument establishes the claim for S .




For purposes of argument, we partition the attraction diagram of γ by extending vertical segments from each pivot configuration to the next critical cycles directly above and below. We call the cells in this decomposition *trapezoids*, even though their top and bottom boundaries may not be straight line segments. At each forward pivot configuration p , we color the vertical segment above (x, y) *green* and the vertical segment below p *red*; the colors are reversed for backward vertical segments, see Figure 13.

The first step of any strategy is a (possibly trivial) pivot onto a stable critical path. Because the human and puppy can move freely within any stable critical path σ , either every point in σ is dexter, or no point in σ is dexter. Similarly, for any green pivot segment π , either every point in π is dexter or no point in π is dexter.

Consider any trapezoid τ , and let σ be the stable critical path on its boundary. Starting in any configuration in τ , the puppy immediately moves to a configuration on σ . Thus, if any point in τ is dexter, then σ is dexter, which implies that *every* point in τ is dexter. Thus, we can describe entire trapezoids as dexter or not dexter. It follows that D is the union of trapezoids.

If two trapezoids share a stable critical path *other than the main diagonal*, then either both trapezoids are dexter or neither is dexter. Similarly, if the green pivot segment leaving a pivot configuration p is dexter, then all four trapezoids incident to p are dexter; otherwise, either two or none of these four trapezoids are dexter.

We conclude that aside from the main diagonal, the boundary of D consists entirely of unstable critical paths, pivot configurations, and red vertical segments. Moreover, for every pivot configuration p on the boundary of D , the green pivot segment leaving p is *not* dexter.




■ **Figure 13** Possible arrangements of dexter trapezoids near a forward pivot configuration.

By definition, every point in D is connected by a (dexter) path to the main diagonal, so D is non-empty and connected. On the other hand, D excludes a complete cycle of forward configurations just below the main diagonal. For any $x \in S^1$, let $D(x)$ denote the set of dexter configurations (x, y) ; this set consists of one or more vertical line segments in the attraction diagram.

Suppose for the sake of argument that some set $D(x)$ is disconnected. Because D is connected, the boundary of D must contain a *concave vertical bracket*: A vertical boundary segment π whose adjacent critical boundary segments both lie (without loss of generality) to the right of π , but D lies locally to the left of π . See Figure 14. Let p be the pivot configuration at one end of π . The green vertical segment on the other side of p is dexter, which implies that *all* trapezoids incident to p are dexter, contradicting the assumption that π lies on the boundary of D . We conclude that for all x , the set $D(x)$ is a single vertical line segment; in other words, D is a *monotone* annulus.

The bottom boundary of D is the main diagonal. The monotonicity of D implies that the top boundary of D is a monotone “staircase” alternating between upward red vertical segments and rightward unstable critical paths. Every trapezoid immediately above the top boundary of D contains only forward configurations. Thus, there is a complete essential



■ **Figure 14** A hypothetical concave vertical bracket on the boundary of D .

cycle ϕ of forward configurations just above the upper boundary of D . Because ϕ contains only forward configurations, ϕ must lie entirely above the river. It follows that D contains the entire river.

Symmetrically, S is an annulus bounded above by the main diagonal and bounded below by a non-contractible cycle of backward configurations; in particular, the entire river lies inside S . We conclude that $D \cup S$ is the entire configuration torus. ◀

If the attraction diagram of γ has more than two essential critical cycles curves, then D and S are still monotone annuli, each bounded by the main diagonal and an essential cycle of red vertical segments and unstable paths, and thus S and D each contain at least one essential critical cycle other than the main diagonal. However, $D \cup S$ need not cover the entire torus.

► **Corollary 7.** *The human can catch the puppy on any generic simple closed track, from any starting configuration.*

5 Polygonal tracks

Our previous arguments require, at a minimum, that the track has a continuous derivative that is never equal to zero. We now extend our results to polygonal tracks, which do not have well-defined tangent directions at their vertices.

5.1 Polygonal attraction diagrams

Throughout this section, we fix a simple polygonal track P with n vertices. We regard P as a continuous piecewise-linear function $P: S^1 \rightarrow \mathbb{R}^2$, parametrized by arc length. Without loss of generality $P(0)$ is a vertex of the track. We index the vertices and edges of P in order, starting with $v_0 = P(0)$, where edge e_i connects v_i to v_{i+1} ; all index arithmetic is implicitly performed modulo n .

To properly describe the puppy's behavior, we must also account for the direction that the puppy is facing, even when the puppy lies at a vertex. To that end, we represent the track using both a continuous *position* function $\pi: S^1 \rightarrow \mathbb{R}$ and a continuous *direction* function $\theta: S^1 \rightarrow S^1$, such that for all $y \in S^1$, the derivative vector $\pi'(y)$ is a non-negative scalar multiple of the unit vector $\theta(y)$. Intuitively, as we increase y , the puppy alternately moves at constant speed along edges (when $\pi'(y)$ is a positive multiple of $\theta(y)$) and continuously turns at constant speed at vertices (when $\pi'(y) = 0$).

We classify any human-puppy configuration $(x, y) \in S^1 \times S^1$ as *forward*, *backward*, or *critical*, if the dot product $(P(x) - \pi(y)) \cdot \theta(y)$ is negative, positive, or zero, respectively. In any forward configuration (x, y) , the puppy moves to increase the parameter y ; in any backward configuration, the puppy moves to decrease the parameter y . (The human's direction is


irrelevant.) The *attraction diagram* is the set of all critical configurations $(x, y) \in S^1 \times S^1$. We further classify critical configurations (x, y) as follows:

- *final* if $P(x) = \pi(y)$,
- *stable* if $(x, y - \varepsilon)$ is forward and $(x, y + \varepsilon)$ is backward for all suffic. small $\varepsilon > 0$,
- *unstable* if $(x, y - \varepsilon)$ is backward and $(x, y + \varepsilon)$ is forward for all suffic. small $\varepsilon > 0$,
- *forward pivot* if $(x, y - \varepsilon)$ and $(x, y + \varepsilon)$ are both forward for all suffic. small $\varepsilon > 0$, or
- *backward pivot* if $(x, y - \varepsilon)$ and $(x, y + \varepsilon)$ are both backward for all suffic. small $\varepsilon > 0$.

Straightforward case analysis implies that this classification is exhaustive.


To define the attraction diagram of P , we decompose the torus $S^1 \times S^1$ into a $2n \times n$ grid of rectangular cells, where each column corresponds to an edge e_j containing the human, and each row corresponds to either a vertex v_i or an edge e_i containing the puppy. The *main diagonal* of the attraction diagram is the set of all final configurations. Strictly speaking, in this case the “main diagonal” is not just a straight line, but consists of alternating diagonal and vertical segments. We can characterize the critical points inside each cell as follows:

Each edge-edge cell $e_i \times e_j$ contains at most one boundary-to-boundary path of stable critical configurations (x, y) . Refer to Figure 15.



■ **Figure 15** All edge-edge critical configurations are stable.

Each vertex-edge cell $v_i \times e_j$ contains at most one boundary-to-boundary path of stable critical configurations and at most one boundary-to-boundary path of unstable critical configurations. If the cell contains both paths, they are disjoint. A configuration (x, y) with $\pi(y) = v_i$ is stable if and only if $P(x)$ lies in the outer normal cone at v_i , and unstable if and only if $P(x)$ lies in the inner normal cone at v_i ; see Figure 16.




■ **Figure 16** Stable and unstable vertex-edge critical configurations.

5.2 Polygonal pivot configurations

Unlike the attraction diagrams of generic smooth curves defined in Section 3.2, the attraction diagrams of polygons are not always well-behaved. In particular, a pivot configuration may be incident to more (or fewer) than two critical curves, and in extreme cases, pivot configurations need not even be discrete. We call such a configuration a *degenerate* pivot configuration.


In any pivot configuration (x, y) , the puppy $\pi(y)$ lies at some vertex v_i , the puppy's direction $\theta(y)$ is parallel to either e_i (or e_{i+1}). Generically, each pivot configuration is a shared endpoint of an unstable critical path in cell $v_i \times e_j$ and a stable critical path in cell $e_i \times e_j$ (or $e_{i-1} \times e_j$); see Figure 17.



■ **Figure 17** Near a non-degenerate pivot configuration.

There are three distinct ways in which degenerate pivot configurations can appear.

A **type-1 degeneracy** is caused by an acute angle on P . Specifically, let v_i be a vertex of P . The configuration (x, y) with $P(x) = \pi(y) = v_i$ is degenerate if the angle between e_{i-1} and e_i is strictly acute. In the attraction diagram of a type-1 degeneracy, two stable critical curves and two unstable critical curves end on a single vertical section of the main diagonal (corresponding to the human and the puppy being both at v_i , but the puppy facing in different directions). Refer to Figure 18.




■ **Figure 18** Stable and unstable configurations near an acute vertex angle.

A **type-2 degeneracy** is caused by a more specific configuration. Let e_i be an edge of P , and let ℓ be the line perpendicular to e_i through v_i (or, symmetrically, through v_{i+1}). Let v_j be another vertex of P which lies on ℓ . The configuration (x, y) with $P(x) = v_j$ and $\pi(y) = v_i$ is degenerate if:


- v_{i-1} and v_j lie in the same open halfspace of the supporting line of e_i ; **and**
- v_{j-1} and v_{j+1} lie in the same open halfspace of ℓ .

A type-2 degeneracy corresponds to a vertex (pivot configuration) of degree 4 or 0 in the attraction diagram. We distinguish these further as *type-2a* and *type-2b*. Refer to Figure 19.

Finally, a **type-3 degeneracy** is essentially a limit of both of the previous types of degeneracies. Let e_i be an edge of P , let ℓ be the line perpendicular to e_i through v_i , and let e_j be another edge of P which lies on ℓ . The configuration (x, y) with $P(x) \in e_j$ and $\pi(y) = v_i$ is degenerate if vertices v_{i-1} and v_j lie in the same open halfspace of the supporting line of e_i . When this degeneracy occurs, pivot configurations are not discrete, because the point $P(x) \in e_j$ can be chosen arbitrarily. Moreover, the vertex-vertex configurations (v_j, v_i) and (v_{j-1}, v_i) have odd degree in the attraction diagram. A type-3 degeneracy can be



■ **Figure 19** Type-2a and type-2b degenerate pivot configurations.



■ **Figure 20** Type-3 degenerate pivot configurations on incident and non-incident edge pairs.

connected to (2 or more) other critical curves, or be isolated. We distinguish these further as *type-3a* and *type-3b*. See Figure 20.

In Section 5.3 we first consider polygonal tracks which do not have any degeneracies of these three types. To simplify exposition, we forbid degeneracies by assuming that no vertex angle in P is acute and that no three vertices of P define a right angle. In Section 5.5 we lift these assumptions by *chamfering* the polygon, cutting off a small triangle at each vertex.

5.3 Catching puppies on generic obtuse polygons

Generic obtuse polygonal tracks behave almost identically to smooth tracks, once we properly define the attraction diagram and dual attraction diagram.

► **Lemma 8.** *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. The attraction diagram of P is the union of disjoint simple critical cycles.*


Proof. Each edge-edge cell $e_i \times e_j$ contains at most one section of stable critical configurations (x, y) (Figure 15). For each such configuration, the points $\pi(y) \in e_i$ and $P(x) \in e_j$ are connected by a line perpendicular to e_i . Because no three vertices of P define a right angle, these points cannot both be vertices of P ; thus, any critical path inside the cell $e_i \times e_j$ avoids the corners of that cell.

Each vertex-edge cell $v_i \times e_j$ contains at most one section of a stable and one section of an unstable path (Figure 16). Again, because no three vertices of P define a right angle, these paths avoid the corners of the cell $v_i \times e_j$.

For every pivot configuration (x, y) , the puppy $\pi(y)$ lies at a vertex v_i , the puppy's direction $\theta(y)$ is parallel to either e_i (or e_{i+1}), and because no three vertices of P form a right angle, the human $P(x)$ lies in the interior of some edge e_j . Since we avoid degenerate pivot configurations, each pivot configuration is a shared endpoint of an unstable critical path in cell $v_i \times e_j$ and a stable critical path in cell $e_i \times e_j$ (or $e_{i-1} \times e_j$).

It now follows that the set of unstable critical configurations is the union of monotone paths whose endpoints are pivot configurations. Similarly, the set of stable critical configurations is also the union of monotone paths whose endpoints are pivot configurations. Each unstable critical path lies in a single vertex strip.

Because every vertex angle in P is obtuse, every configuration (x, y) where the human $P(x)$ lies on an edge e_i and the puppy $\pi(y)$ lies on the previous edge e_{i-1} is either stable or final. In particular, the only non-pivot critical configurations (x, y) with $\pi(y) = v_i$ are the final configurations with $P(x) = v_i$. Thus, the main diagonal is disjoint from all other critical cycles; in fact, no other critical cycle intersects any grid cell that touches the main diagonal.



■ **Figure 21** Near the main diagonal.

This completes the classification of all critical configurations. We conclude that the attraction diagram consists of the (simple, closed) main diagonal and possibly other simple closed curves composed of stable and unstable critical paths meeting at pivot configurations. All these critical cycles are disjoint. ◀

► **Lemma 9.** *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. If the attraction diagram of P has exactly two essential critical cycles, then the human can catch the puppy on P , starting from any initial configuration.*

The remainder of the proof is essentially unchanged from the smooth case. For any configuration (x, y) , let $T(y)$ denote the directed “tangent” line through $\pi(y)$ in direction $\theta(y)$, and let $L(x, y)$ denote the signed distance from $P(x)$ to $T(y)$, signed positively if $P(x)$ lies to the left of $T(y)$ and negatively if $P(x)$ lies to the right of $T(y)$. The *dual attraction diagram* of P consists of all points $(y, L(x, y)) \in S^1 \times \mathbb{R}$ where (x, y) is a critical configuration. As in the smooth case, the map $(x, y) \mapsto (y, L(x, y))$ is a homeomorphism from the critical cycles in the attraction diagram to the curves in the dual attraction diagram; moreover, this map preserves the contractibility of each critical cycle.

► **Lemma 10.** *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. The attraction diagram of P contains exactly two essential critical cycles.*


► **Theorem 11.** *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. The human can catch the puppy on P , starting from any initial configuration.*

5.4 Chamfering

We now extend our analysis to arbitrary simple polygons. We define a *chamfering* operation which transforms a polygon P into a new polygon \bar{P} . First we show that \bar{P} has no more degeneracies of type 1, 2a, or 3a. The polygon \bar{P} may still have degenerate pivot configurations

of type 2b and type 3b; however, since these correspond to isolated forward or backward pivot configurations in the attraction diagram, they do not impact the existence of a strategy to catch the puppy on \bar{P} . The puppy will just move over them as if they were normal forward or backward configurations. Thereafter we show that such a strategy can be correctly translated back to a strategy on P .

Let P be an arbitrary simple polygon, and let $\varepsilon > 0$ be smaller than half of any distance between two non-incident features of P . Then the ε -*chamfered* polygon \bar{P} is another simple polygon with twice as many vertices as P , defined as follows. Refer to Figure 22. For each vertex v_i of P , we create two new vertices v'_i and v''_i , where v'_i is placed on e_{i-1} at distance ε from v_i , and v''_i is placed on e_i at distance ε from v_i . Edge e'_i in \bar{P} connects v''_i to v'_{i+1} , and a new *short edge* s_i connects v'_i to v''_i .



■ **Figure 22** The chamfering operation.

The chamfering operation alters the local structure of the attraction diagram near every vertex. The idea is that for non-degenerate configurations, the change will not influence the behaviour of the puppy in such configurations, and as such will not influence the existence of any catching strategies. However, at degenerate configurations, the change in the structure is significant. We will argue in Section 5.5 that the changes are such that every strategy in the chamfered polygon translates to a strategy in the original polygon.

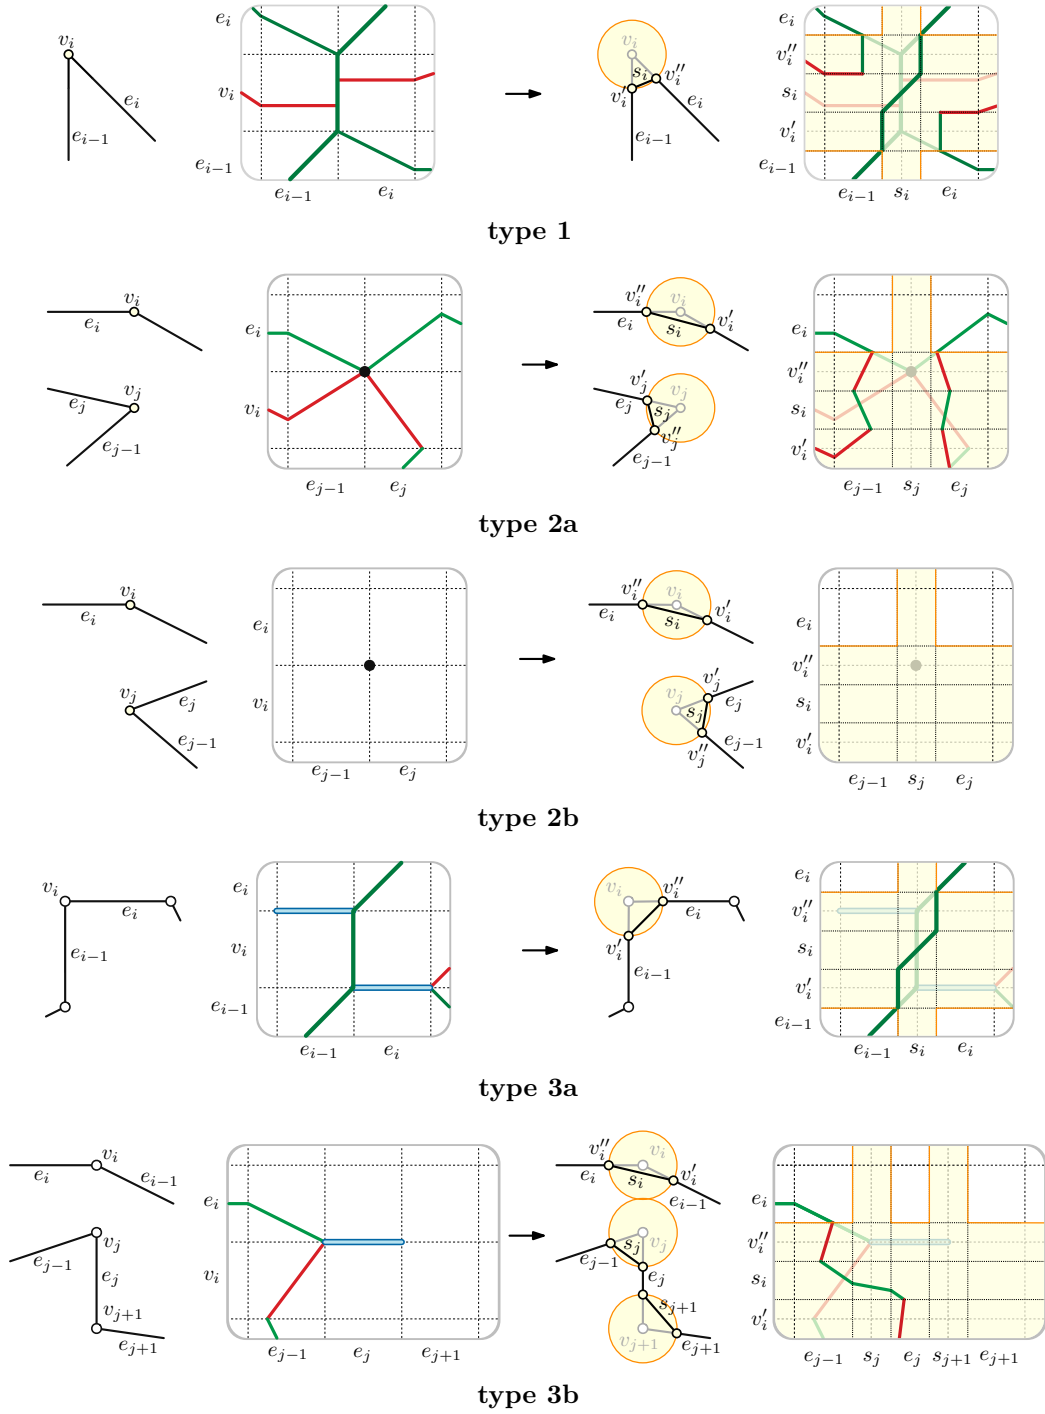
Here we review again the different types of degenerate pivot configurations, and how the chamfering operation affects the structure of the attraction diagram in each case. Refer to Figure 23.

- Near type 1 degeneracies, the higher-degree vertices on the main diagonal disappear. Instead, two separate critical curves almost touch the main diagonal: one from above and one from below.
- Near type 2a degeneracies, the degree-4 vertex disappears. Instead, the two incident critical curves coming from the left are connected, and the two incident curves coming from the right are connected.
- Near type-2b degeneracies, the isolated pivot vertex simply disappears.
- Near type-3 degeneracies, the degenerate pivot “vertex” disappears. Any connected critical curve is locally rerouted away from the degenerate location.

5.5 Catching puppies on arbitrary simple polygons

► **Lemma 12.** *Let P be an arbitrary simple polygon. There exists an ε such that the ε -chamfered polygon \bar{P} has no degenerate pivot configurations of type 1, type 2a, or type 3a.*

Proof. First, note that P has no type 1 degeneracies: we replace each vertex v_i with angle α_i by two new vertices v'_i and v''_i with angles $\alpha'_i = \alpha''_i = \pi - \frac{1}{2}(\pi - \alpha_i) = \frac{1}{2}\pi + \frac{1}{2}\alpha_i > \frac{1}{2}\pi$.




■ **Figure 23** Effect of the chamfering operation on the attraction diagram near degenerate pivot configurations. The size of ε is exaggerated; the figures show the combinatorial structure of the chamfered diagram for a much smaller value of ε . Only the effect of chamfering vertices relevant for the degeneracy is shown.

Next, we argue about type 2 degeneracies, which may occur for some values of ε . We argue that each potential type 2a degeneracy only occurs for a specific value of ε ; since there are finitely many potential degeneracies the lemma then follows.

Note that, as we vary ε , all vertices of P' move linearly and with equal speed. Suppose such a configuration is *not* unique for a specific choice of ε ; that is, sustained for all values of ε . Then two vertices involved in the degenerate configuration must move in the same direction; that is, two edges of P , say e_i and e_j , must be parallel. There are two configurations in P' which could potentially give rise to a type 2 degeneracy. We argue that then, in fact, it cannot satisfy all requirements of a type 2 degeneracy.

- An edge e'_i has endpoint v'_i (or symmetrically, v''_{i-1}) such that the line ℓ through v'_i and perpendicular to e'_i contains another vertex v'_j (or v''_{j-1}). Refer to Figure 24. Then, as v'_i move along e'_i , ℓ move at the same speed as v'_i , so v'_j moves in the same direction at the same speed along e'_{j-1} . So e'_{j-1} is parallel to e'_i . But now consider v''_j . Since v_j lies on the supporting line of e'_{j-1} at distance ε from v'_j , and v''_j lies at distance ε from v_j , we conclude that v''_{j-1} and v''_j lie on the opposite side of ℓ ; thus, this is not a type 2 degeneracy.




■ **Figure 24** Potential new degenerate pivot configurations based on a (shortened) original edge e'_i . For ε small enough, there can be no degeneracy.

- A short edge s_i of P' has an endpoint v'_i (or symmetrically, v''_i) such that the line ℓ through v'_i and perpendicular to s_i contains another vertex v'_j (or v''_{j-1}). Refer to Figure 25. Then, as ε increases, v'_i move along e'_i , ℓ moves at a slower speed; hence, we cannot conclude anything about the orientation of e_j . However, in this case, note that e'_i and s_i lie on opposite sides of ℓ ; therefore, e'_j and thus v''_{j-1} lies on the opposite side of ℓ . It is possible that v''_j lies on the same side, in which case we have a degenerate pivot configuration of type 2b (Figure 25 (left)), or that v''_j lies on ℓ , in which case we have a degenerate pivot configuration of type 3b (Figure 25 (middle)). If v''_j lies on the opposite side of ℓ , there is no degeneracy (Figure 25 (right)).



Let P be an arbitrary simple polygon and \bar{P} a chamfered copy without degeneracies. We say a parameter value x is *verty* whenever $P(x)$ is within distance ε from a vertex. We say a parameter value x is *edgy* whenever $P(x)$ is not within distance ε from a vertex. We reparameterize \bar{P} such that $P(x) = \bar{P}(x)$ whenever x is edgy; the parameterization of \bar{P} is uniformly scaled for verty parameters. We say a configuration (x, y) is edgy when x and y are both edgy.

We say a path σ is a *valid* path in the attraction diagram if it describes a human and puppy behavior that obeys the rules imposed on the puppy and the human, as explained in Section 1.



■ **Figure 25** Potential new degenerate pivot configurations based on a short edge s_i . For any ε we may still have a new degeneracy of type 2b (left), 3b (middle), or no degeneracy (right).

► **Lemma 13.** *Assuming ε is sufficiently small. For any valid path σ between two edgy configurations (x_1, y_1) and (x_2, y_2) in the attraction diagram of \bar{P} , there is an equivalent path between (x_1, y_1) and (x_2, y_2) in the attraction diagram of P .*

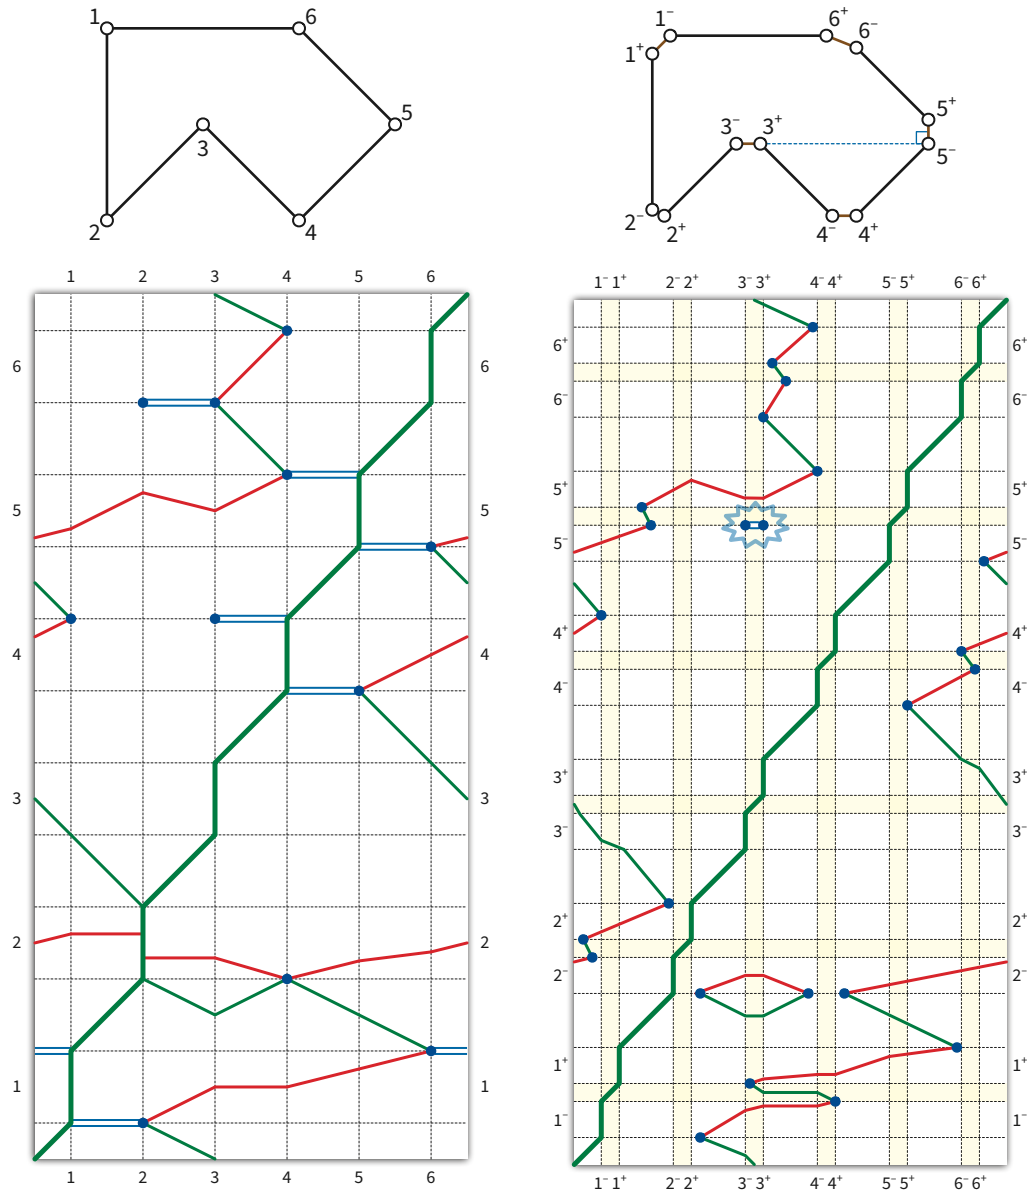
Proof. The intuition is that the behavior of the puppy around vertices is the same as in the original polygon. Note that when the human reaches a vertex in P then it can happen that the puppy suddenly moves infinitely fast until it reaches a stable position. However, if the human is *leaving* a vertex then the puppy may move, but only at a finite speed. (At least for some short distance.) Thus if the puppy is at an edgy position, we can always move the human such that we get into an edgy configuration, given that ε is small enough. Recall that there is a one to one correspondence between edgy configurations of P and \bar{P} . Also once the puppy is in an edgy configuration, we can also move the human into an edgy configuration, assuming we have chosen ε small enough. Thus, we only care about the situation that the puppy is in a vertex configuration. We need to make sure that the same way the puppy can leave the vertex position in \bar{P} can be simulated in P .

We first handle the situation that the puppy traverses over a generic vertex v_i . To be precise, let v_i be a vertex of P and e_{i-1} and e_i its incident edges. Let ℓ and ℓ' be the lines through v_i orthogonal to e_{i-1} and e_i respectively. We say that v_i is generic if there is no vertex that lies on either ℓ or ℓ' . We denote the infinite strip bounded by ℓ and v_i' by the letter S . Similarly, we denote the infinite strip bounded by ℓ and v_i'' by the symbol S' . As no vertex of P lies on one of the two lines ℓ, ℓ' then, we can choose ε small enough such that also there is no vertex of \bar{P} in one of the strips S, S' , other than v_i' and v_i'' .


Thus whenever the human crosses over one of the strips S or S' , the human is on an edge of \bar{P} . Thus, when the puppy will leave the vertex position defined by the segment s_i , then the human will be in an edgy configuration, as is the puppy. Thus we can simulate the same valid path moving the puppy out of the vertex configuration at s_i .

We are now handling the case that the puppy moves towards a degenerate configuration. Recall all degenerate positions as listed in Figure 23. Interestingly, while those configurations are degenerate for the attraction diagram the behavior of the puppy and the human is easily understood. For instance in a type 1 degeneracy the human walks towards the acute angle and the puppy will move to the same position. Note that type 2b configurations are not changing the behavior of the puppy at all, as the isolated pivot point will never be reached in the attraction diagram. Similarly, type 3a and type 3b degeneracies are not altering the puppy behavior in an interesting way. Given a movement of the human in \bar{P} it is straightforward how the movement of the human should be in P . See to the right of Figure 27, for the most difficult example of a type 2a degeneracy. ◀

We are now ready to prove our main result.



■ **Figure 26** The attraction diagram of a degenerate polygon, before and after chamfering. All existing degeneracies disappeared in the chamfered polygon, which does have one new but harmless degeneracy.



■ **Figure 27** Left: When the puppy is at the vertex v_i , then its behavior is determined on the region that the human lies. We have four regions, called A, B, C, D . Right: Once, the human gets into a verty configuration, the puppy makes an infinitely fast jump forward. This can be simulated in P , as the lower vertex stems from a type 2a degeneracy.

▶ **Theorem 14.** *Let P be a simple polygon. The human can catch the puppy on P , starting from any initial configuration.*

Proof. By Lemma 12, there is a value ε such that the ε -chamfered polygon \bar{P} has no degeneracies of type 1 or type 2a or type 3. There may still be some type 2b degeneracies, but they result in isolated pivots in the attraction diagram, which can be safely ignored.

Consider an arbitrary starting configuration on P . If the starting configuration is not stable, we let p move until it is. If the resulting configuration is not edgy, we walk h along P until we reach an edge configuration (x, y) . (This must be possible, except if the puppy stays on a vertex for the entire time, in that case, we can catch the puppy trivially, by going to that vertex.)

By Theorem 11, there exists a strategy for h to catch p on \bar{P} . If the end configuration of this strategy is not edgy, note that we may now simply move h and p together to an edgy final configuration (f, f) . By Lemma 13, there is an equivalent strategy to reach (f, f) from (x, y) on P . Combined with the initial path to (x, y) this gives us a path from an arbitrary starting configuration to a final configuration on P . ◀

6 Further questions

For simple curves, we have only proved that a catching strategy exists. At least for polygonal tracks, it is straightforward to compute such a strategy in $O(n^2)$ time by searching the attraction diagram. In fact, we can compute a strategy that minimizes the total distance traveled by either the human or the puppy in $O(n^2)$ time, using fast algorithms for shortest paths in toroidal graphs [15, 17]. Unfortunately, this quadratic bound is tight in the worst case if the output strategy must be represented as an explicit path through the attraction diagram. We conjecture that an optimal strategy can be described in only $O(n)$ space by listing only the human's initial direction and the sequence of points where the human reverses direction. On the other hand, an algorithm to compute such an optimal strategy in subquadratic time seems unlikely.

If the track is a *smooth curve* of length ℓ whose attraction diagram has k pivot configurations, a trivial upper bound on the distance the human must walk to catch the puppy is $\ell \cdot k/2$. In any optimal strategy, the human walks straight to the point on the curve corresponding to a pivot located at one of the two endpoints of the current “stable sub-curve” of a critical curve (walking less than ℓ). Then the configuration moves to another stable

sub-curve, and so on, never visiting the same stable sub-curve twice. Our question is whether a better upper bound can be proved.

In fact, if minimizing distance is not a concern, we conjecture that *no* reversals are necessary. That is, on *any* simple track, starting from *any* configuration, we conjecture that the human can catch the puppy *either* by walking only forward along the track *or* by walking only backward along the track. Figure 2 and its reflection show examples where each of these naïve strategies fails, but we have no examples where both fail. (Our proof of Theorem 2 implies that the human can always catch the puppy on an *orthogonal* polygon by walking *at most once* around the track in some direction, depending on the starting configuration.)

More ambitiously, we conjecture that the following *oblivious* strategy is always successful: walk twice around the track in one (arbitrary) direction, then walk twice around the track in the opposite direction.

Another interesting question is to what extent our result applies to self-intersecting curves in the plane, when we consider the two strands of the curve at an intersection point to be distinct. It is easy to see that the human cannot catch the puppy on a curve that traverses a circle twice; see Figure 4. Indeed, we know how to construct examples of bad curves with any rotation number *except* -1 and $+1$. We conjecture that Lemma 5, and therefore our main result, extends to all non-simple tracks with rotation number ± 1 . Similarly, are there interesting families of curves in \mathbb{R}^3 where the human and puppy can always meet?

Finally, it is natural to consider similar pursuit-attraction problems in more general domains. Theorem 1 shows that the human can always catch the puppy in the interior of a simple polygon, by walking along the dual tree of any triangulation. Can the human always catch the puppy in any planar straight-line graph? Inside any polygon with holes?

References

- 1 Zachary Abel, Hugo A. Akitaya, Erik D. Demaine, Martin L. Demaine, Adam Hesterberg, Matias Korman, Jason S. Ku, and Jayson Lynch. Negative instance for the edge patrolling beacon problem. *Preprint*, <https://arxiv.org/abs/2006.01202>, 2020.
- 2 Israel Aldana-Galván, Jose L. Alvarez-Rebollar, Juan Carlos Cataa-Salazar, Nestaly Marin-Nevárez, Erick Solís-Villareal, Jorge Urrutia, and Carlos Velarde. Beacon coverage in orthogonal polyhedra. In *Proceedings of the 29th Canadian Conference on Computational Geometry (CCCG 2017)*, pages 156–161, 2017.
- 3 Israel Aldana-Galván, Jose L. Alvarez-Rebollar, Juan Carlos Cataa-Salazar, Nestaly Marin-Nevárez, Erick Solís-Villareal, Jorge Urrutia, and Carlos Velarde. Tight bounds for illuminating and covering of orthotrees with vertex lights and vertex beacons. *Graphs and Combinatorics*, 36:617–630, 2020.
- 4 Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geom. Appl.*, 5:75–91, 1995. doi:10.1142/S0218195995000064.
- 5 Sang Won Bae, Chan-Su Shin, and Antoine Vigneron. Tight bounds for beacon-based coverage in simple rectilinear polygons. *Computational Geometry*, 80:40–52, 2019. doi:10.1016/j.comgeo.2019.02.002.
- 6 Michael Biro. *Beacon-based routing and guarding*. PhD thesis, State University of New York at Stony Brook, 2013.
- 7 Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon-based routing and coverage. In *Proceedings of the 21st Fall Workshop on Computational Geometry*, 2011.
- 8 Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon based structures in polygonal domains. In *Abstracts of the 1st Computational Geometry: Young Researchers Forum*, 2012.

- 9 Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Combinatorics of beacon routing and coverage. In *Proceedings of the 25th Canadian Conference on Computational Geometry (CCCG 2013)*, 2013.
- 10 Michael Biro, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon-based algorithms for geometric routing. In *13th International Symposium on Algorithms and Data Structures (WADS 2013)*, pages 158–169, 2013. doi:10.1007/978-3-642-40104-6_14.
- 11 Prosenjit Bose and Thomas C. Shermer. Gathering by repulsion. *Computational Geometry*, 90:101627, 2020. doi:10.1016/j.comgeo.2020.101627.
- 12 Pierre Bouguer. Sur de nouvelles courbes ausquelle on peut donner le nom de linges de poursuite. *Mémoires de mathématique et de physique tirés des registres de l'Académie royale des sciences*, pages 1–14, 1732. URL: <https://gallica.bnf.fr/ark:/12148/bpt6k35294>.
- 13 Johans Cleve and Wolfgang Mulzer. Combinatorics of beacon-based routing in three dimensions. *Computational Geometry*, 91:101667, 2020. doi:10.1016/j.comgeo.2020.101667.
- 14 Pierre de Maupertuis. Sure les courbes de poursuite. *Mémoires de mathématique et de physique tirés des registres de l'Académie royale des sciences*, pages 15–17, 1732. URL: <https://gallica.bnf.fr/ark:/12148/bpt6k35294>.
- 15 John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *J. Algorithms*, 5(3):391–407, 1984. doi:10.1016/0196-6774(84)90019-1.
- 16 Arthur S. Hathaway. Problems and solutions: Problem 2801. *American Mathematical Monthly*, 27(1):31, 1920. doi:10.2307/2973244.
- 17 Monika R. Henzinger, Philip Klein, Satish Rao, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. *J. Comput. Syst. Sci.*, 55(1):3–23, 1997. doi:10.1006/jcss.1997.1493.
- 18 Irina Kostitsyna, Bahram Kouhestani, Stefan Langerman, and David Rappaport. An Optimal Algorithm to Compute the Inverse Beacon Attraction Region. In *34th International Symposium on Computational Geometry (SoCG 2018)*, 2018. doi:10.4230/LIPIcs.SoCG.2018.55.
- 19 Bahram Kouhestani and David Rappaport. Edge patrolling beacon. In *Abstracts from the 20th Japan Conference on Discrete and Computational Geometry, Graphs, and Games (JCDCGGG 2017)*, pages 101–102, 2017.
- 20 Bahram Kouhestani, David Rappaport, and Kai Salomaa. On the inverse beacon attraction region of a point. In *Proceedings of the 27th Canadian Conference on Computational Geometry (CCCG 2015)*, 2015.
- 21 Bahram Kouhestani, David Rappaport, and Kai Salomaa. The length of the beacon attraction trajectory. In *Proceedings of the 28th Canadian Conference on Computational Geometry (CCCG 2016)*, pages 69–74, 2016.
- 22 Bahram Kouhestani, David Rappaport, and Kai Salomaa. Routing in a polygonal terrain with the shortest beacon watchtower. *Computational Geometry*, 68:34–47, 2018. doi:10.1016/j.comgeo.2017.05.005.
- 23 John E. Littlewood. *Littlewood's Miscellany: edited by Béla Bollobás*. Cambridge University Press, 1986.
- 24 Amirhossein Mozafari and Thomas C. Shermer. Transmitting particles in a polygonal domain by repulsion. In *12th International Conference on Combinatorial Optimization and Applications (COCOA 2018)*, pages 495–508, 2018. doi:10.1007/978-3-030-04651-4_33.
- 25 Paul J. Nahin. *Chases and Escapes: The Mathematics of Pursuit and Evasion*. Princeton University Press, 2007.
- 26 Thomas C. Shermer. A combinatorial bound for beacon-based routing in orthogonal polygons. In *Proceedings of the 27th Canadian Conference on Computational Geometry (CCCG 2015)*, 2015.