Algorithmica (2024) 86:2845–2884
https://doi.org/10.1007/s00453-024-01245-0

Approximation Algorithms for the Two-Watchman Route in a Simple Polygon

Bengt J. Nilsson[1] · Eli Packer[2]

Received: 17 September 2023 / Accepted: 3 June 2024 / Published online: 19 June 2024

Abstract

The *two-watchman route problem* is that of computing a pair of closed tours in an environment so that the two tours together see the whole environment and some length measure on the two tours is minimized. Two standard measures are: the minmax measure, where we want the tours where the longest of them has smallest length, and the minsum measure, where we want the tours for which the sum of their lengths is the smallest. It is known that computing a minmax two-watchman route is NP-hard for simple rectilinear polygons and thus also for simple polygons. Also, any c-approximation algorithm for the minmax two-watchman route is automatically a $2c$-approximation algorithm for the minsum two-watchman route. We exhibit two constant factor approximation algorithms for computing minmax two-watchman routes in simple polygons with approximation factors 5.969 and 11.939, having running times $O(n^8)$ and $O(n^4)$ respectively, where n is the number of vertices of the polygon. We also use the same techniques to obtain a 6.922-approximation for the *fixed two-watchman route problem* running in $O(n^2)$ time, i.e., when two starting points of the two tours are given as input.

Keywords Art gallery problems · Visibility · Watchman routes · Approximation algorithms

1 Introduction

Some of the most intriguing problems in computational geometry concern visibility and motion planning in polygonal environments. A classical problem is that of computing a *shortest watchman route* in an environment, i.e., the shortest closed tour that sees the complete free-space of the environment. Watchman routes can either be *fixed*, requiring the tour to pass a given boundary point or *floating*, with no requirement to

✉ Bengt J. Nilsson
bengt.nilsson.TS@mau.se

Eli Packer
eli@yoom.com

1 Malmö University, 205 06 Malmö, Sweden

2 Yoom, Ramat Gan, Israel

pass any specific point. These problems have been shown NP-hard [7, 10] and even $\Omega(\log n)$-inapproximable [23] for polygons with holes having a total of n segments.

After a sequence of false starts [8, 14, 31, 32], Tan *et al.* [33] prove an $O(n^4)$ time dynamic programming algorithm for computing a shortest fixed watchman route through a given boundary point in a simple polygon. This is later improved to $O(n^3 \log n)$ time by Dror *et al.* [9] and to $O(n^3)$ time by Tan and Jiang [34]. Carlsson *et al.* [6] show how to generalize an algorithm that computes a shortest fixed watchman route to compute a shortest floating watchman route in a simple polygon with a quadratic factor overhead. Tan [29] improves this to a linear factor overhead. Hence, the currently best algorithm for a shortest floating watchman route in a simple polygon uses $O(n^4)$ time.

Given the relatively high polynomial time complexity for computing watchman routes in simple polygons, efficient approximation algorithms are also of interest. Nilsson [25] and Tan [30] have independently developed linear time approximation algorithms for a shortest floating watchman route in a simple polygon.

The more general problem of computing multiple watchman routes that together see the environment has received much less attention. Mitchell and Wynters [24] show that already computing the pair of tours that together see a simple rectilinear polygon is NP-hard, if we want to minimize the length of the longest of the two tours, the *minmax* measure. It is still an open problem whether it is possible to compute a pair of tours for which the sum of the lengths of the two tours is minimal, the *minsum* measure, in polynomial time. Packer [26] gives some experimental results for multiple watchman routes in simple polygons. For point sized watchmen, so-called *static guards*, Belleville [3, 4] shows an efficiently computable characterization of all simple polygons that are two-guardable with point guards.

Our Results. We present a polynomial time constant factor approximation algorithm to compute a minmax or minsum pair of tours that together see a simple polygon. We first consider the floating version of the problem and obtain a 5.969-approximation algorithm for the minmax pair of tours and 11.939-approximation for the minsum pair of tours that runs in $O(n^8)$ time, where n is the number of vertices of the polygon.

In the next three sections, we provide some preliminary results and prove some crucial properties that we use continuously in the sequel. In Sect. 5, we give the algorithm for the minmax two-watchman route prove its correctness and analyze its running time. In Sect. 6, we show how to modify the previous algorithm to run in $O(n^4)$ time while maintaining constant approximation factor, albeit only guaranteeing a factor twice as large as the previous algorithm. In Sect. 7 we modify the algorithm to handle the fixed two-watchman route, the case when we have fixed starting points for the tours that they have to pass through, arriving at an $O(n^2)$ time algorithm with approximation factor 6.922. We conclude the presentation in Sect. 8.

2 Preliminaries

Let $\mathbf{P}$ be a simple polygon having n vertices and let $\partial\mathbf{P}$ denote the boundary of $\mathbf{P}$. We say that two points in $\mathbf{P}$ *see* each other, if the line segment connecting the points does not intersect the exterior of $\mathbf{P}$. For any arbitrary connected object X inside $\mathbf{P}$, we

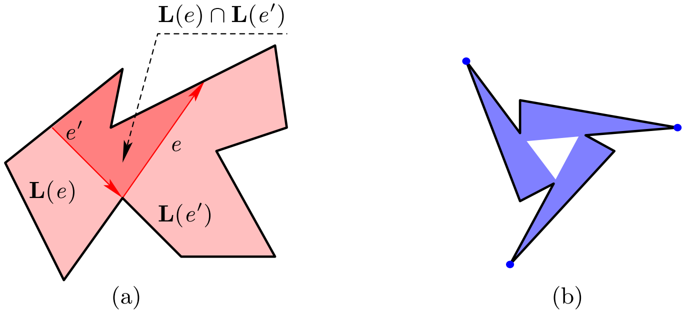

Fig. 1 Illustrating the definition of extensions (**a**) and an example where a disconnected set has points to the left of every extension (the marked blue convex vertices) but does not see the whole polygon (**b**) (Color figure online)

denote by $\mathbf{VP}(X)$ the *weak visibility polygon* of X in $\mathbf{P}$, i.e., the set of points in $\mathbf{P}$ that see some point of X. The boundary of a visibility polygon $\mathbf{VP}(X)$ consists of edges that are either (sub)edges of $\mathbf{P}$ or edges that have their end points on $\partial\mathbf{P}$ but their interior points in the interior of $\mathbf{P}$. These latter edges are denoted the *windows* of $\mathbf{VP}(X)$ and they have at least one end point on a reflex vertex of $\mathbf{P}$. We henceforth assume the existence of linear time algorithms to compute $\mathbf{VP}(X)$ when X is a point or a segment inside $\mathbf{P}$. Such algorithms have previously been presented in the literature [12, 16, 21, 28].

A *cut* is a directed line segment in $\mathbf{P}$ with both end points on $\partial\mathbf{P}$ and each interior point is an interior point of $\mathbf{P}$. Hence, a directed segment incident to a polygon edge or a directed segment intersecting more than two vertices is not a cut. A cut always separates $\mathbf{P}$ into exactly two sub-polygons of nonzero area. If a cut is represented by the segment $[p, q]$ we say that the cut is directed from p to q. For a cut c in $\mathbf{P}$, we define the *left polygon*, $\mathbf{L}(c)$, to be the set of points in $\mathbf{P}$ locally to the left of c according to c's direction.

Assume a counterclockwise walk of $\partial\mathbf{P}$. Such a walk imposes a direction on each of the edges of $\mathbf{P}$ in the direction of the walk. Consider a reflex vertex of $\mathbf{P}$. The two edges incident to the vertex can each be extended inside $\mathbf{P}$ until the extensions reach a boundary point. These extended segments form cuts given the same direction as the edge they are collinear to. We call these cuts *extensions*; see Fig. 1a and denote the set of extensions in $\mathbf{P}$ by $\mathcal{E}$.

We define a *guard set* to be any set of points $\mathcal{G}$ that together see all of $\mathbf{P}$, i.e., $\bigcup_{g\in\mathcal{G}} \mathbf{VP}(g) = \mathbf{P}$. It is clear that any guard set must have points intersecting $\mathbf{L}(e)$ for every extension e of $\mathbf{P}$, since otherwise the edge collinear to e will not be seen by the guard set; see Fig. 1a. Chin and Ntafos [8] prove that this is indeed also a sufficient requirement when the guard set is connected, as it is for a shortest watchman route. For disconnected guard sets, it is easy to construct examples where this requirement is not sufficient; see Fig. 1b where the guard set consisting of the three marked convex vertices has points to the left of each essential extension but it does not see the complete polygon.

Let c be a cut. If a guard set $\mathcal{G}$ intersects $\mathbf{L}(c)$, i.e., $\mathcal{G} \cap \mathbf{L}(c)$ is non-empty, we say that c is *covered* by $\mathcal{G}$. Furthermore, if $\mathcal{G}$ intersects the interior of $\mathbf{L}(c)$, then $\mathcal{G}$ *properly*

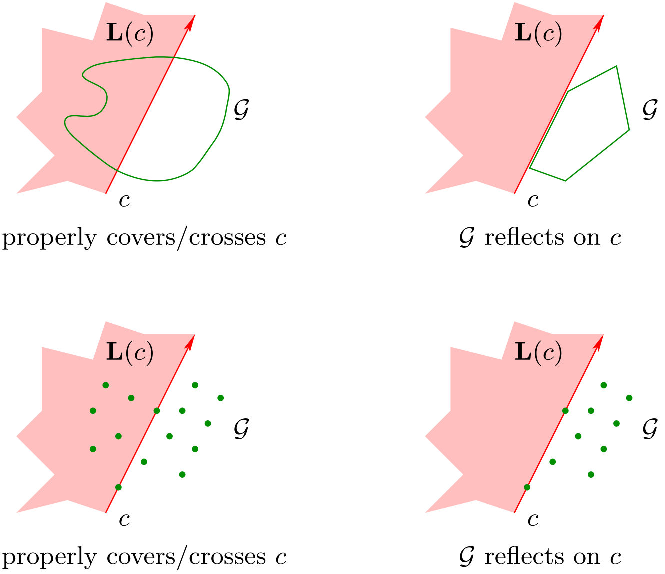

Fig. 2 Illustrating covering, crossing, and reflecting guard sets. The cut c is covered in all four examples

covers c. If $\mathcal{G}$ properly covers c and intersects c, we say that $\mathcal{G}$ *crosses* c. Finally, if $\mathcal{G}$ covers c, but does not properly cover c, then $\mathcal{G}$ *reflects* on c; see Fig. 2.

We also make use of the fact that shortest paths in $\mathbf{P}$ between combinations of segments and points can be computed efficiently [12, 13, 21]. We denote the shortest path between two objects X and Y in $\mathbf{P}$ by $SP(X, Y)$.

Let X_1 and X_2 be two closed polygonal cycles contained in a simple polygon $\mathbf{P}$, such that each point in $\mathbf{P}$ sees some point on X_1 or X_2. We call the pair $\mathcal{X} = (X_1, X_2)$, a *two-watchman route*. The length of a cycle X in $\mathbf{P}$ is denoted $\|X\|$ and we let $\|\mathcal{X}\|_{\mathsf{sum}} \stackrel{\text{def}}{=} \|X_1\| + \|X_2\|$ be the *sum length* of $\mathcal{X}$ and $\|\mathcal{X}\|_{\mathsf{max}} \stackrel{\text{def}}{=} \max\left\{\|X_1\|, \|X_2\|\right\}$ be the *max length* of $\mathcal{X}$.

Let $\mathcal{S} = (S_1, S_2)$ and $\mathcal{T} = (T_1, T_2)$ be two two-watchman routes such that $\|\mathcal{S}\|_{\mathsf{sum}} \leq \|\mathcal{X}\|_{\mathsf{sum}}$ and $\|\mathcal{T}\|_{\mathsf{max}} \leq \|\mathcal{X}\|_{\mathsf{max}}$ for any two-watchman route $\mathcal{X}$ in $\mathbf{P}$. We say that $\mathcal{S}$ is a *minsum* two-watchman route and $\mathcal{T}$ is a *minmax* two-watchman route. The following inequalities are immediate from the definitions, for any two-watchman route $\mathcal{X}$,

$$\|\mathcal{X}\|_{\mathsf{max}} = \max\left\{\|X_1\|, \|X_2\|\right\} \leq \|X_1\| + \|X_2\| = \|\mathcal{X}\|_{\mathsf{sum}}, \tag{1}$$

$$\|\mathcal{X}\|_{\mathsf{sum}} = \|X_1\| + \|X_2\| \leq 2\max\left\{\|X_1\|, \|X_2\|\right\} = 2\|\mathcal{X}\|_{\mathsf{max}}, \tag{2}$$

and therefore

$$\|\mathcal{T}\|_{\mathsf{max}} \leq \|\mathcal{S}\|_{\mathsf{sum}} \leq 2\|\mathcal{T}\|_{\mathsf{max}}. \tag{3}$$

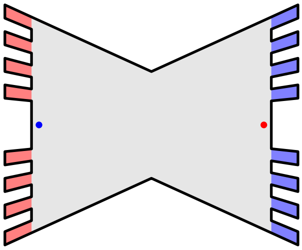

Fig. 3 A counterexample, showing that no cut through the polygon separates the two-watchman route problem into two independent single watchman route problems having reasonable approximation ratio

Hence, computing a c-approximation for one measure also gives at most a $2c$-approximation for the other measure.

One could imagine that there always is a cut in the polygon such that the two tours that are shortest watchman tours for the two sub-polygons formed are almost as short as the optimum two-watchman tour. Trying out all possible cuts, and solving the two single watchman tour subproblems in each case would give a simple and efficient algorithm with good approximation ratio. One would then imagine wrong, as the counterexample in Fig. 3c shows. The red and blue point sized tours, each see the grey region and the regions of their own color. Here, any partition of the polygon into two pieces by a cut will make the watchman tour solution of each piece infinitely longer than the optimal solution. The example can easily be modified for non-point sized tours.

3 General Properties of Two-Watchman Routes

We make the following assumption about the polygons considered in this section.

Assumption 3.1 The polygon **P** considered in this section is simple and not guardable by one or two point guards. Hence, **P** is not convex or starshaped.

The following lemma is at the heart of our construction of an algorithm to compute a two-watchman route.

Lemma 3.1 *If two tours in* ***P*** *see all of* ∂***P****, then they see all of* ***P****.*

Proof We do a proof by contradiction. Let X_1 and X_2 be two tours in **P** and assume that p is an interior point of **P** not seen by any of them. We show that there must be some boundary point that is also not seen by the tours contradicting that X_1 and X_2 together see ∂**P**. Let **VP**(X_1) and **VP**(X_2) be the two visibility polygons of the tours. Each of them must have boundary segments w_1 and w_2 separating the subpolygon containing p from the subpolygon containing the tour. The segment w_1 cuts **P** into two subpolygons. Let $\mathbf{P}_{X_1}$ be the subpolygon containing X_1 and let $\mathbf{P}_p$ be the subpolygon containing p. If X_2 intersects w_1, we interchange the roles of X_1 and X_2, thus we can assume that X_2 is either completely in $\mathbf{P}_{X_1}$ or completely in $\mathbf{P}_p$. We have three cases:

X_2 **lies in** $\mathbf{P}_{X_1}$. Let $SP(p, X_2)$ be the shortest path from p to X_2 and let s be the first segment of this path. It connects p with a reflex vertex v on the boundary of **P**. Extend s away from v until it hits the boundary at p'; see Fig. 4a. The point

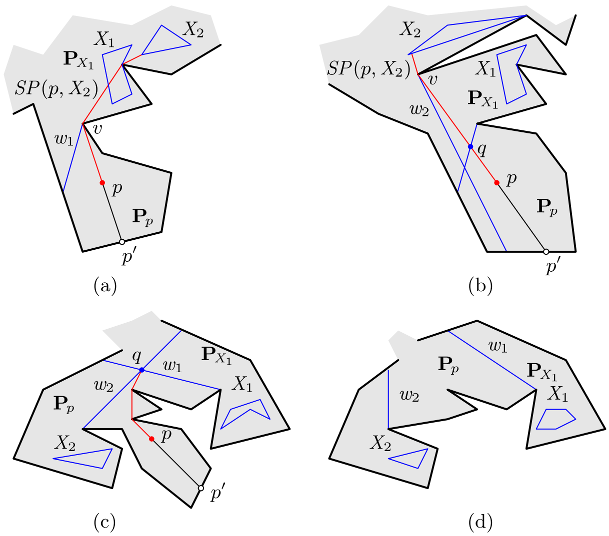

Fig. 4 Illustrating the three cases of the proof of Lemma 3.1. Examples **a**, **b** illustrate the case when X_2 lies in $\mathbf{P}_{X_1}$, **c** illustrates the case when X_2 lies in $\mathbf{P}_p$ and w_1 and w_2 intersect, and **d** the case when X_2 lies in $\mathbf{P}_p$ and w_1 and w_2 do not intersect

p' is clearly not seen by X_2 since if it were, then p would also bee seen by X_2. The point p' is not seen by X_1 either since $SP(p, X_2)$ crosses w_1 at some point q with $SP(p, q)$ in $\mathbf{P}_p$ and since $SP(p, q)$ is a shortest path, the extension from p to p' cannot cross w_1; see Fig. 4b.

X_2 lies in $\mathbf{P}_p$ and w_1 and w_2 intersect. Let q be the intersection point between w_1 and w_2 and let s be the first segment of $SP(p, q)$. It connects p either with a reflex vertex v on the boundary of $\mathbf{P}$ or with q. Extend s away from v (or q) until it hits the boundary at p'; see Fig. 4c. The point p' cannot be seen by X_1 or X_2 since $SP(p, q)$ lies outside both $\mathbf{VP}(X_1)$ and $\mathbf{VP}(X_2)$ (except for the point q), and hence, since $SP(p, q)$ is a shortest path, the extension from p to p' cannot cross any of w_1 or w_2.

X_2 lies in $\mathbf{P}_p$ and w_1 and w_2 do not intersect. It is clear that there exist points on the boundary not seen by the two tours since the boundary points close to the end points of w_1 not in $\mathbf{P}_{X_1}$ and the boundary points close to the end points of w_2 not in $\mathbf{P}_{X_2}$ are not seen by any of the tours; see Fig. 4d.

This completes the proof. □

The lemma implies that it is sufficient that our algorithm constructs two tours that together see the whole boundary of $\mathbf{P}$ to guarantee that all of $\mathbf{P}$ is guarded. The example in Fig. 1b shows that the claim does not hold for three guard sets.

Consider two tours X_1 and X_2 and a polygon boundary edge b.

Lemma 3.2 *For any two tours X_1 and X_2 and a polygon boundary edge b, the sets $\mathbf{VP}(X_1) \cap b$, $\mathbf{VP}(X_2) \cap b$ and $\mathbf{VP}(X_1) \cap \mathbf{VP}(X_2) \cap b$ are each connected.*

Proof If the set under consideration is empty, the lemma follows by definition so assume in each case that it is nonempty. If any of $\mathbf{VP}(X_i) \cap b$, for $i \in \{1, 2\}$, is connected, it is a subsegment of b and therefore convex. Since the intersection of two connected and convex sets is also a connected and convex set, it follows that the connectedness of $\mathbf{VP}(X_1) \cap \mathbf{VP}(X_2) \cap b = \bigcap_{i\in\{1,2\}} \mathbf{VP}(X_i) \cap b$ is an immediate consequence, if we can show that the two sets $\mathbf{VP}(X_i) \cap b$, for $i \in \{1, 2\}$, are connected.

We make a proof by contradiction and assume that $\mathbf{VP}(X_i) \cap b$ is disconnected, for each $i \in \{1, 2\}$. By following b from one end point to the other we pass each of the connected components giving us an ordering of them. Let r_1 be a point in the first component and let r_2 be a point in the last component. Each point r_1 and r_2 is seen by X_i. Hence, for X_i there is a point r' on b between r_1 and r_2 that is not seen by X_i. Let p_j, for $j \in \{1, 2\}$, be two points on X_i that see r_j respectively. Without loss of generality, we can assume that $[p_j, r_j]$ does not intersect X_i except at p_j, for $i \in \{1, 2\}$. If the segments $[p_1, r_1]$ and $[p_2, r_2]$ do not intersect, we construct a closed simple polygon $\mathbf{R}_i$ in $\mathbf{P}$ as follows: follow b from r_1 to r_2, from r_2 to p_2, from p_2 to p_1 along a simple path in X_i, and finally from p_1 to r_1. The only edges of the constructed polygon $\mathbf{R}_i$ that are not part of X_i are the three edges adjacent to r_1 and r_2 so $\mathbf{R}_i$ is completely seen by the part of the tour X_i connecting p_2 to p_1. Hence, the point r' is also seen since it lies on b between r_1 and r_2, giving us a contradiction.

On the other hand if the segments $[p_1, r_1]$ and $[p_2, r_2]$ intersect, let q be an intersection point between $[p_1, r_1]$ and $[p_2, r_2]$. The three points p_1, p_2, and q form a triangle interior to $\mathbf{P}$ and we construct a closed simple polygon $\mathbf{R}'_i$ in $\mathbf{P}$ having q as a vertex as follows: follow a simple path of X_i from p_1 to p_2, a straight edge from p_2 to q, and a straight edge from q to p_1. If we extend the segment $[p', q]$ inside $\mathbf{R}'_i$, it must intersect X_i, since p_1 connects to p_2 in $\mathbf{R}'_i$ using a simple path of X_i. Hence, the point p' is also seen from X_i, again giving us a contradiction. □

From Lemma 3.2 we have that T_1 and T_2 both see connected components of any boundary edge of $\mathbf{P}$. Hence, a boundary edge b can be partitioned into at most three subsegments, at most one of which is seen by both T_1 and T_2 and the remaining at most two are seen by one of T_1 and T_2.

Our next lemma shows that $\|(T_1, T_2)\|_{\mathsf{max}} < \|W_{\mathsf{opt}}\|$ and this strict inequality is used in the proof of Lemma 3.5.

Lemma 3.3 *Let* (T_1, T_2) *be the shortest minmax two-watchman route and* W_{opt} *the shortest watchman route. If* $\|W_{\mathsf{opt}}\| > 0$,

$$\|(T_1, T_2)\|_{\mathsf{max}} < \|W_{\mathsf{opt}}\|.$$

Proof Consider a shortest watchman route W_{opt} and let p and q be two points on W_{opt} such that if you follow the tour W_{opt} from p a distance of $\|W_{\mathsf{opt}}\|/2$ in counterclockwise order, you reach the point q. It is clear that following the tour from p to q in clockwise order also gives a path of length $\|W_{\mathsf{opt}}\|/2$.

Let $SP(p, q)$ be the shortest path from p to q in $\mathbf{P}$. If $\|SP(p, q)\| < \|W_{\mathsf{opt}}\|/2$, we construct a two-watchman route (X_1, X_2) such that $\|(X_1, X_2)\|_{\mathsf{max}} < \|W_{\mathsf{opt}}\|$ as follows. Let X_1 be the tour obtained by following $SP(p, q)$ from p to q and W_{opt} from q to p in counterclockwise order. Since $\|SP(p, q)\| < \|W_{\mathsf{opt}}\|/2$, the length of X_1 is strictly smaller than $\|W_{\text{opt}}\|$. The tour X_2 is obtained by following W_{opt} from p to q in counterclockwise order and then $SP(p, q)$ from q to p. Again, the length of X_2 is strictly smaller than $\|W_{\text{opt}}\|$.

If $\|SP(p, q)\| = \|W_{\text{opt}}\|/2$, on the other hand, then both the clockwise and counterclockwise paths from p to q along W_{opt} follow $SP(p, q)$, since shortest paths are unique. We construct a two-watchman route (X_1, X_2) as follows. Let r be the midpoint on $SP(p, q)$ and let X_1 be the tour obtained by following $SP(p, q)$ from p to r and back to p. Similarly, let X_2 be the tour obtained by following $SP(p, q)$ from q to r and back to q. We have that $\|(X_1, X_2)\|_{\max} = \|W_{\text{opt}}\|/2 < \|W_{\text{opt}}\|$ also in this case.

Since (T_1, T_2) is the shortest minmax two-watchman route, it has max-length bounded by that of (X_1, X_2) and is hence also strictly smaller than $\|W_{\text{opt}}\|$, concluding the proof. □

In the previous section, we defined $\mathcal{E}$ to be the set of extensions of the edges in $\mathbf{P}$. There is a subdivision of $\mathcal{E}$ into nonempty subsets $\mathcal{E}_1$ and $\mathcal{E}_2$, such that each tour T_1 and T_2 of a minmax two-watchman route covers the extensions in $\mathcal{E}_1$ and $\mathcal{E}_2$ respectively. It is clear that neither $\mathcal{E}_1$ nor $\mathcal{E}_2$ can be empty since if one of them is empty, the other contains all the extensions of $\mathcal{E}$. This means that one of T_1 or T_2 covers all the extensions in $\mathcal{E}$, but the shortest tour that covers all extensions in $\mathcal{E}$ is W_{opt}, the shortest watchman route, contradicting Lemma 3.3. Note also that the subdivision is not necessarily a partition since an extension in $\mathcal{E}$ can be covered by both T_1 and T_2.

Assume from now on that each tour T_1 and T_2 is as short as possible and consider the points of the boundary $\partial\mathbf{P}$ of $\mathbf{P}$ that are seen by T_1. The visibility from T_1 subdivides $\partial\mathbf{P}$ into disjoint (maximal) subpaths and we color the interior points of each subpath *white* if T_1 sees these points and *black* if T_1 does not see the points. The end points of each (maximal) subpath is colored *grey*. For a color $c \in \{black, grey, white\}$, we say that a point on $\partial\mathbf{P}$ has color *c for* T_1. We can similarly color the boundary for T_2 in which case we say that a point has color c' for T_2, $c' \in \{black, grey, white\}$.

We refine the coloring of the boundary somewhat by considering the convex vertices of T_1. Let $\|T_1\| > 0$ and let u be a convex vertex of T_1. There exists at least one grey boundary point $p(u)$ that is seen from u but not from any other point of T_1. The point $p(u)$ must exist, otherwise T_1 can be made shorter, contradicting that the tours T_1 and T_2 are as short as possible. In fact, u can have many such points. We therefore consider a convex vertex u of T_1 to have *multiplicity* k, if there are k different points $p(u)$ associated to u. We differentiate between $p(u)$ and $p(u')$ even though $u = u'$ and has multiplicity at least two. We let the color of each point $p(u)$ be *dark grey*. The remaining grey points are considered to be *light grey*.

Let $l(u)$ be the maximal line segment in $\mathbf{P}$ passing through the points u on T_1 and $p(u)$ on $\partial\mathbf{P}$. The segment $l(u)$ must intersect at least one reflex vertex of $\partial\mathbf{P}$ between u and $p(u)$ that hides $p(u)$ from $T_1\backslash u$. Let $v(u)$ be the reflex vertex of $\partial\mathbf{P}$ on $l(u)$ closest to $p(u)$. The line segment $[p(u), v(u)]$ partitions $\mathbf{P}$ into two subpolygons $\mathbf{P}_1$ containing T_1 and $\mathbf{P}_0$. Let $b(u)$ be the polygon boundary edge adjacent to the reflex vertex $v(u)$ in $\mathbf{P}_0$ and let $e(u)$ be the extension collinear to $b(u)$ inside $\mathbf{P}_1$. We associate a direction to $l(u)$ so that $e(u)$ lies locally to the left of $l(u)$. Henceforth, we denote $\mathbf{P}_0$ by $\mathbf{P}_{e(u)}$ and $\mathbf{P}_1$ by $\mathbf{P}_{l(u)}$ to let them depend on the vertex u of T_1. We refer to Fig. 5a for an illustration of the given definitions and prove the following lemma.

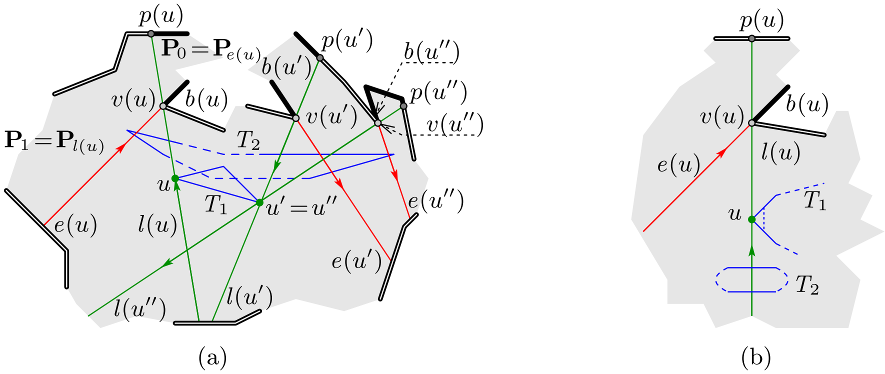

Fig. 5 **a** Illustrates the coloring definitions of the points on the boundary of **P**. **b** Illustrates the proof of Lemma 3.4, if p is white for T_2, then T_1 can be shortened

Lemma 3.4 *For* $\|T_1\| > 0$*, any boundary point that is dark grey for* T_1 *must also be grey for* T_2*, and for* $\|T_2\| > 0$*, any boundary point that is dark grey for* T_2 *must also be grey for* T_1.

Proof We show the lemma for a boundary point p that is dark grey for T_1. The reverse claim follows by symmetry. We have two cases.

If p is white for T_2, we have an immediate contradiction since the single convex vertex u of T_1 that sees p can be cut away from T_1 arbitrarily close to u, thus shortening the length of T_1; see Fig. 6b.

If p is black for T_2, this means that there is open interval on the boundary $\partial\mathbf{P}$ centered at p that is not seen by any point of T_2. Since p is one end point of a maximal subpath that is seen by T_1, there exist boundary points not seen by either T_1 or T_2, a contradiction.

This concludes the proof. □

We can now show the following claim.

Lemma 3.5 *There exists a minmax two-watchman route* $\mathcal{T} = (T_1, T_2)$*, such that each tour* T_i intersects *some extension in* $\mathcal{E}_i$*,* $i \in \{1, 2\}$.

Proof We will use the assumption that both tours T_1 and T_2 are as short as possible. At least one of them will have the length $\|(T_1, T_2)\|_{\mathsf{max}}$ and the other the shortest possible length given that the first tour achieves the length $\|(T_1, T_2)\|_{\mathsf{max}}$. We will show the result only for tour T_1, but the same argument holds also for T_2. The proof is by case analysis, first subdividing into the cases when $\|T_1\| = 0$ and thus the region $\mathbf{P} \setminus \mathbf{VP}(T_2)$ is starshaped, and the case $\|T_1\| > 0$. In this latter case we make a proof by contradiction assuming that T_1 does not intersect any extension in $\mathcal{E}$. The argument in this case becomes slightly involved but relies on showing that under this assumption, Lemma 3.4 does not hold, giving us a contradiction.

Assume first that $\|T_1\| = 0$, i.e., T_1 is a point sized tour. This implies that $\mathbf{P} \setminus \mathbf{VP}(T_2)$ is starshaped even though it may be a disconnected set. The set of points in $\mathbf{P}$ that see all of $\mathbf{P}\setminus\mathbf{VP}(T_2)$, also known as the *kernel*, is the intersection of the left halfplanes collinear to the boundary edges from $\partial\mathbf{P}\cap(\mathbf{P}\setminus\mathbf{VP}(T_2))$ [20]. Left means here locally to the left of the associated edge in the order of the counterclockwise traversal. The kernel boundary cannot be made up only of boundary edges of **P**, since that would make **P** a convex polygon and thus guardable with one point guard contradicting Assumption 3.1. Hence, there is a kernel boundary edge collinear with some extension of **P** and we can let T_1 be a point on this kernel boundary edge.

Now, assume that $\|T_1\| > 0$. We make a proof by contradiction and assume further that T_1 does not intersect any extension in $\mathcal{E}$. If this is the case, T_1 cannot have any reflex vertices, since if T_1 does, then any such vertex coincides with a reflex vertex of **P** and thus T_1 intersects the two extensions adjacent to this vertex. Hence, T_1 has only convex vertices. If T_1 is a line segment, we consider the two end points of the segment to be the convex vertices of the tour that goes back and forth between them.

The rest of the proof now shows that, if T_1 does not intersect some extension, then there exists a dark grey boundary point for T_1 that is either black or white for T_2, contradicting Lemma 3.4. We prove this by a case analysis. The case analysis first considers the two cases when T_1 has two convex vertices u and u', such that $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(e(u')\big)=\emptyset$, i.e., T_1 lies in a region of **P** bounded by two extensions on either side of T_1 and the case when it does not. This second case is then subdivided into two further cases, when the left hand sides of all the segments $l(u)$ of convex vertices u of T_1, all have non-empty intersection and when they have empty intersection. In the first of these cases, T_1 can be slid in one direction shortening it and in the second of these cases, there must exist exist a convex vertex u of T_1 and three extensions that T_2 covers ensuring that $p(u)$ is white for T_2, contradicting Lemma 3.4. We proceed with the formal details.

If T_1 has two convex vertices u and u', such that $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(e(u')\big)=\emptyset$, then the intersection $\mathbf{L}\big(e(u)\big)\cap\mathbf{L}\big(e(u')\big)$ is also empty, since $\Big(\mathbf{L}\big(e(u)\big)\setminus\mathbf{P}_{e(u)}\Big)\subset\mathbf{L}\big(l(u)\big)$. This means that T_1 lies properly in the region $\mathbf{P}\setminus\Big(\mathbf{L}\big(e(u)\big)\cup\mathbf{L}\big(e(u')\big)\Big)$ without touching $e(u)$ or $e(u$') whereby T_2 must intersect the two disjoint regions $\mathbf{L}\big(e(u)\big)$ and $\mathbf{L}\big(e(u')\big)$. Since $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(e(u')\big)=\emptyset$, T_2 must cross $l(u)$ and hence $p(u)$ is white for T_2 giving us a contradiction to Lemma 3.4; see Fig. 5a illustrating this case.

If all pairs of vertices u and u' of T_1 (with multiplicity) have the property that $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(e(u')\big)\neq\emptyset$, then since $\Big(\mathbf{L}\big(e(u')\big)\setminus\mathbf{P}_{e(u')}\Big)\subset\mathbf{L}\big(l(u')\big)$ it also holds that $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(l(u')\big)\neq\emptyset$ and we have further subcases.

If $\bigcap_{u\in\mathcal{U}}\mathbf{L}\big(l(u)\big)\neq\emptyset$, where $\mathcal{U}$ is the set of vertices of T_1 (with multiplicity), then let q be a point in $\bigcap_{u\in\mathcal{U}}\mathbf{L}\big(l(u)\big)$ and let u_1 and u_2 be two distinct vertices of T_1 such that T_1 is completely contained in the cone $\angle u_1, q, u_2$. We note that T_1 is the shortest tour that visits the regions $\mathbf{L}\big(l(u)\big)$, for $u\in\mathcal{U}$. Let $\triangle u_1, q, u_2$ denote the triangle with corners at u_1, q, and u_2. The interior region of $\angle u_1, q, u_2\setminus\triangle u_1, q, u_2$ cannot contain convex vertices of T_1, since any such vertex u' would have $q\notin\mathbf{L}\big(l(u')\big)$ contradicting that $\bigcap_{u\in\mathcal{U}}\mathbf{L}\big(l(u)\big)\neq\emptyset$; see Fig. 6a. T_1 thus connects u_1 and u_2 by a line segment. Vertex u_1 (with multiplicity) intersects $\bigcap_{u\in\mathcal{U}_{u_1}}\mathbf{L}\big(l(u)\big)$ for some subset $\mathcal{U}_{u_1}$ of $\mathcal{U}$. Similarly, the vertex u_2 (with multiplicity) intersects $\bigcap_{u\in\mathcal{U}_{u_2}}\mathbf{L}\big(l(u)\big)$ for some subset $\mathcal{U}_{u_2}$ of $\mathcal{U}$, with $\mathcal{U}_{u_1}\cap\mathcal{U}_{u_2}=\emptyset$. Since the intersection of locally convex regions is also locally convex, any point on the line segment $[u_1, q]$ intersects $\bigcap_{u\in\mathcal{U}_{u_1}}\mathbf{L}(l(u))$ and similarly, any point on the line segment $[u_2, q]$ intersects $\bigcap_{u\in\mathcal{U}_{u_2}}\mathbf{L}\big(l(u)\big)$. Thus we can move u_1 and u_2 (with multiplicity) along their corresponding line segments towards q, thus shortening T_1, giving us a contradiction; see Fig. 6a. If $\bigcap_{u\in\mathcal{U}}\mathbf{L}\big(l(u)\big)=\emptyset$ and for all pairs u and u', $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(l(u')\big)\neq\emptyset$, then T_1 must have at least three vertices. We first show that there exists a subset of three vertices u_1, u_2, and u_3 such that $\mathbf{L}\big(l(u_1)\big)\cap\mathbf{L}\big(l(u_2)\big)\cap\mathbf{L}\big(l(u_3)\big)=\emptyset$. Since $T_1\cap\mathbf{L}\big(l(u)\big)=u$, for every vertex u of T_1, not only do $\mathbf{L}\big(l(u)\big)\cap\mathbf{L}\big(l(u')\big)\neq\emptyset$ but the segments $l(u)$ and $l(u')$ intersect in a point, for every vertex pair u and u'. Pick u_1 to be any vertex of T_1 and assume without loss of generality that $l(u_1)$ is horizontal and directed towards the right; see Fig. 6b. Initialize the set $\mathcal{L}$ to be $\mathcal{L}:=\{u_1\}$ and sort the remaining vertices u on the angle the corresponding cut $l(u)$ makes with $l(u_1)$, from 0 to 2π. Now, add vertices u (with multiplicity), one by one, according to the sorted order to $\mathcal{L}$, for $u\in\mathcal{U}$ until $\bigcap_{u\in\mathcal{L}}\mathbf{L}\big(l(u)\big)=\emptyset$. Let u_2 denote the last vertex added to $\mathcal{L}$ during the process above, and let u_3 be a vertex in $\mathcal{L}$ such that the intersection $q_1=l(u_3)\cap l(u_1)$ lies after $q_2=l(u_3)\cap l(u_2)$ on the directed cut $l(u_3)$. To see that $\mathbf{L}\big(l(u_1)\big)\cap\mathbf{L}\big(l(u_2)\big)\cap\mathbf{L}\big(l(u_3)\big)=\emptyset$, any point in the intersection $\mathbf{L}\big(l(u_1)\big)\cap\mathbf{L}\big(l(u_3)\big)$ lies above (or on) $l(u_1)$ and any point in the intersection $\mathbf{L}\big(l(u_2)\big)\cap\mathbf{L}\big(l(u_3)\big)$ lies below $l(u_1)$, since q_2 lies below $l(u_1)$; see Fig. 6b. This also shows that u_3 must exist, since if, for all vertices $u\in\mathcal{L}\setminus\{u_1,u_2\}$, the point $l(u)\cap l(u_1)$ lies before $l(u)\cap l(u_2)$ on $l(u)$, the first such intersection point on $l(u_2)$ along $l(u_2)$ lies to the left (or on) each cut $l(u)$, for $u\in\mathcal{L}$, contradicting that $\bigcap_{u\in\mathcal{L}}\mathbf{L}\big(l(u)\big)=\emptyset$; see Fig. 6b. Since $\mathbf{L}\big(l(u_1)\big)\cap\mathbf{L}\big(l(u_2)\big)\cap\mathbf{L}\big(l(u_3)\big)=\emptyset$, it also follows that $\mathbf{L}\big(e(u_1)\big)\cap\mathbf{L}\big(e(u_2)\big)\cap\mathbf{L}\big(e(u_3)\big)=\emptyset$, as $\Big(\mathbf{L}\big(e(u)\big)\setminus\mathbf{P}_{e(u)}\Big)\subset\mathbf{L}\big(l(u)\big)$, for every vertex u of T_1. The tour T_1 has no points in any of the regions $\mathbf{L}\big(e(u_1)\big)$, $\mathbf{L}\big(e(u_2)\big)$, or $\mathbf{L}\big(e(u_3)\big)$, so T_2 must intersect each of them. Assume that T_2 does not intersect $l(u_2)$ or $l(u_3)$, otherwise at least one of $p(u_2)$ or $p(u_3)$ is white for T_2, contradicting Lemma 3.4. However, then T_2 must have interior points in $\mathbf{L}\big(l(u_2)\big)\cap\mathbf{L}\big(l(u_3)\big)$ since T_2 intersects $\mathbf{L}\big(e(u_2)\big)$ and $\mathbf{L}\big(e(u_3)\big)$ but does not intersect $l(u_2)$ or $l(u_3)$. Hence, T_2 has points below $l(u_1)$ and since $\mathbf{L}\big(e(u_1)\big)$ lies above (or on) $l(u_1)$, T_2 must intersect $l(u_1)$ and thus $p(u_1)$ is white for T_2, again contradicting Lemma 3.4; see Fig. 6b.

This concludes the proof. □

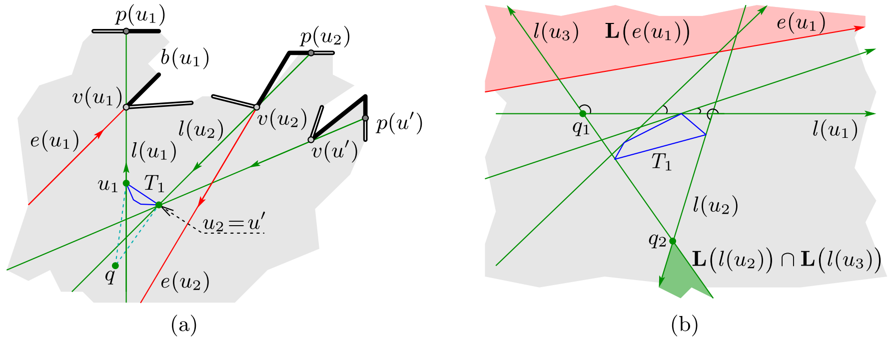

Fig. 6 Illustrating the proof of Lemma 3.5. **a** Illustrates the case when $\bigcap_{u\in\mathcal{U}} \mathbf{L}(l(u)) \neq \emptyset$ where $\mathcal{U}$ is the set of vertices of T_1. **b** Illustrates the case when $\bigcap_{u\in\mathcal{U}} \mathbf{L}(l(u)) = \emptyset$, showing how to select u_1, u_2, and u_3, so that $\mathbf{L}(l(u_1)) \cap \mathbf{L}(l(u_2)) \cap \mathbf{L}(l(u_3)) = \emptyset$

4 Tentacles and Jellyfish

We will use the following definitions extensively in the sequel.

Definition 4.1 For a point q in $\mathbf{P}$ and a point r on the boundary $\partial\mathbf{P}$, we call the shortest path from q to some point in $\mathbf{P}$ that sees r, a *tentacle* from q to r, denoted Z_q^r. We say that q is the *head* of the tentacle and that a tentacle is *attached* to its head. The other end point of the tentacle is called the *tip*.

The idea of tentacles is not new, as similar geometric objects have, for example, previously been defined to compute the quickest visibility path, the visibility center, and the geodesic center in simple polygons [2, 22, 27].

With a tentacle Z_q^r (where q does not see r) we also associate a *tentacle cut* $c(Z_q^r)$. Consider the maximal line segment l in $\mathbf{P}$ passing through the tip p of Z_q^r and the boundary point r. If the segment l has end points r and r' and (possibly) subdivides into connected pieces $l_1, l_2, \ldots$ intersecting the boundary only at the end points and where l_i partitions $\mathbf{P}$ into two subpolygons, one containing q and one containing r. The cut $c(Z_q^r)$ is the segment l_i directed so that $q \in \mathbf{P} \setminus \mathbf{L}(c(Z_q^r))$; see Fig. 7a. The first boundary vertex u encountered as you move from r along l towards p is the *hiding vertex* of the tentacle Z_q^r.

We can prove the following technical lemma.

Lemma 4.1 *If q is a point in $\mathbf{P}$ and s_b is a subsegment of a boundary edge b of $\partial\mathbf{P}$ having end points r_1 and r_2, then the two tentacles $Z_q^{r_1}$ and $Z_q^{r_2}$ together see the whole subsegment s_b.*

Proof The proof is a simple modification of the proof of Lemma 3.2. Let p_1 and p_2 be the two tips of $Z_q^{r_1}$ and $Z_q^{r_2}$, respectively. Let r be some arbitrary point in s_b. If the segments $[p_1, r_1]$ and $[p_2, r_2]$ do not intersect, we construct a closed simple polygon $\mathbf{R}$ in $\mathbf{P}$ as follows: follow s_b from r_1 to r_2, from r_2 to p_2, from p_2 to q along the tentacle $Z_q^{r_2}$, from q to p_1 along $Z_q^{r_1}$, and finally from p_1 to r_1. The only edges of the constructed polygon $\mathbf{R}$ that are not part of the two tentacles are the three edges adjacent to r_1 and

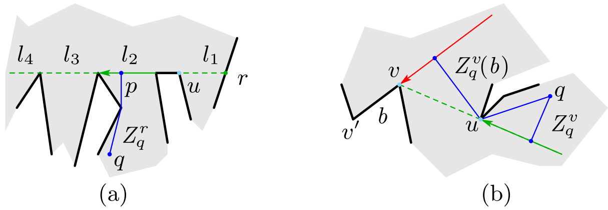

Fig. 7 **a** Illustrating the tentacle Z_{q}^{r}, its hiding vertex u, and its associated cut. **b** The difference between tentacle and edge restricted tentacle. Vertex u is the hiding vertex for Z_q^v and v is the hiding vertex for $Z_q^v(b)$

r_2 so **R** is completely seen by the two tentacles. Hence, the point r is also seen since it lies on s_b between r_1 and r_2.

On the other hand if the segments $[p_1, r_1]$ and $[p_2, r_2]$ intersect, let p be an intersection point between $[p_1, r_1]$ and $[p_2, r_2]$. The three points r_1, r_2, and p form a triangle interior to **P** and we construct a closed simple polygon $\mathbf{R}'$ in **P** having p as a vertex as follows: follow the tentacles from r_1 to r_2 via q, a straight edge from r_2 to p, and a straight edge from p to r_1. If we extend the segment $[r, p]$ inside $\mathbf{R}'$, it must intersect a tentacle, since r_1 connects to r_2 in $\mathbf{R}'$ using the tentacles. Hence, the point r is also seen from the tentacles. □

Let s be a line segment in **P** and let $b = [v, v']$ be some boundary edge of $\partial\mathbf{P}$. The next lemma establishes that if we move the head of the tentacle Z_q^r, from $q \in s$, a small distance to the point $q' \in s$ and the point $r \in]v, v'[$ to the point $r' \in]v, v'[$ also a small distance, the difference in length between the two tentacles $\|Z_q^r\| - \|Z_{q'}^{r'}\|$ is a smooth function.

Lemma 4.2 *Let q move a distance δ to q' on a line segment s and let r move a distance ϵ to r', where both r and r' lie in the open interval $]v, v'[$ of a boundary edge $b = [v, v']$, in such a way that the first segment of the tentacles from q and q' intersect the same reflex vertex, if the tentacle consists of multiple segments, and $c(Z_q^r)$ and $c(Z_{q'}^{r'})$ have the same hiding vertex, then*

$$\|Z_{q'}^{r'}\| = \|Z_q^r\| + \mathcal{F}(\delta, \epsilon),$$

where $\mathcal{F}(\delta, \epsilon)$ is continuous and differentiable.

The proof of the lemma is a lengthy case analysis where all arguments are based on similarity and the cosine theorem and is therefore deferred to "Appendix A".

To alleviate the fact that Lemma 4.2 only holds for points r in the interior of boundary edges we define the *edge restricted tentacle* $Z_q^r(b)$ to be

$$Z_q^r(b) = \lim_{]v,v'[\ni p \to r} Z_q^p \quad \text{where } b = [v, v']. \tag{4}$$

An edge restricted tentacle $Z_q^r(b)$ can differ from Z_q^r only when r is a vertex of b. If v is a reflex vertex and q lies to the right of extension e collinear to boundary edge b, then $Z_q^v(b)$ is the shortest path from q that sees b, not just v; see Fig. 7b. In all other cases, $Z^r_q(b) = Z^r_q$. The proof of Lemma 4.1 does not make use of the edge restriction of a tentacle $Z^r_q(b)$ so it still holds for edge restricted tentacles. We can generalize Lemma 4.2 to also hold for vertices using edge restricted tentacles. We claim this as a corollary.

Corollary 4.3 *Let q move a distance δ to q' on a line segment s and let r move a distance ϵ to r' along a boundary edge $b = [v, v']$, in such a way that the first segment of the tentacles from q and q' intersect the same reflex vertex, if the tentacle consists of multiple segments, and $c\big(Z^r_q(b)\big)$ and $c\big(Z^{r'}_{q'}(b)\big)$ have the same hiding vertex, then*

$$\|Z^{r'}_{q'}(b)\| = \|Z^r_q(b)\| + \mathcal{F}(\delta, \epsilon),$$

where $\mathcal{F}(\delta, \epsilon)$ is continuous and differentiable.

We will henceforth only work with edge restricted tentacles and just call them tentacles.

Given two points q and q' in **P**, consider the tentacles $Z^v_q(b)$ and $Z^v_{q'}(b)$, for each boundary edge b and each vertex v of b. We define two sets $\mathcal{J}_q$ and $\mathcal{J}_{q'}$ of tentacles such that $Z^v_q(b) \in \mathcal{J}_q$ iff $\|Z^v_q(b)\| \leq \|Z^v_{q'}(b)\|$ and $Z^v_{q'}(b) \in \mathcal{J}_{q'}$ iff $\|Z^v_q(b)\| > \|Z^v_{q'}(b)\|$. In this way, each end point v of each boundary edge b has exactly one tentacle in one of $\mathcal{J}_q$ or $\mathcal{J}_{q'}$.

From Lemma 4.1 it is clear that if both end points of a boundary edge $b = [v, v']$ have tentacles in the same set, either $\mathcal{J}_q$ or $\mathcal{J}_{q'}$, then the tentacles in the set sees the whole edge b. However, assume that $Z^v_q(b) \in \mathcal{J}_q$ and $Z^{v'}_{q'}(b) \in \mathcal{J}_{q'}$, then $\|Z^v_q(b)\| \leq \|Z^v_{q'}(b)\|$ and $\|Z^{v'}_q(b)\| > \|Z^{v'}_{q'}(b)\|$. Since the length of a tentacle $Z^r_q(b)$ changes smoothly as r moves along b from v to v'; see Corollary 4.3, there is some point r^* on b such that $\|Z^{r^*}_q(b)\| = \|Z^{r^*}_{q'}(b)\|$. If we include $Z^{r^*}_q(b)$ into $\mathcal{J}_q$ and $Z^{r^*}_{q'}(b)$ into $\mathcal{J}_{q'}$, this guarantees that the tentacles in $\mathcal{J}_q$ and $\mathcal{J}_{q'}$ together see b.

Thus, if each edge b either has the two tentacles of its end points in one set or there is a point r^* on b such that $Z^{r^*}_q(b) \in \mathcal{J}_q$ and $Z^{r^*}_{q'}(b) \in \mathcal{J}_{q'}$, then the whole boundary is seen by the tentacles in the set. By considering the tour constructed by following each tentacle from the head to the tip and back in some order for each set $\mathcal{J}_q$ and $\mathcal{J}_{q'}$, by Lemma 3.1, the polygon **P** is guarded by the tentacles in the two sets. We call each of the two sets $\mathcal{J}_q$ and $\mathcal{J}_{q'}$ a *jellyfish* with *head* q and q', respectively, and $\mathcal{J}_{q,q'} = \mathcal{J}_q \cup \mathcal{J}_{q'}$ the *jellyfish pair* with heads q and q'. We define $\|\mathcal{J}_{q,q'}\|$, the *length* of a jellyfish pair, to be the length of its longest tentacle; see Fig. 8a, b.

Let e_1 and e_2 be two extensions intersected by T_1 and T_2 respectively. These extensions exist by Lemma 3.5 and we have the following lemma.

Lemma 4.4 *If u_1 and u_2 are intersection points of T_1 and T_2 with extensions e_1 and e_2 respectively, then $\|\mathcal{J}_{u_1,u_2,}\| \leq \|(T_1, T_2)\|_{\max}/2$.*

Proof Without loss of generality, assume that a longest tentacle in $\mathcal{J}_{u_1,u_2,}$ is $Z^r_{u_1}(b)$, for some point r on boundary edge b. We distinguish five cases.

1. If r is an interior point of b seen by T_1, then follow a path from u_1 along T_1 until we reach the first point u_1' on T_1 that sees r. Denote the subpath of T_1 thus constructed $T_1[u_1, u_1']$ and let the other subpath $(T_1 \setminus T_1[u_1, u_1']) \cup \{u_1, u_1'\}$ (the closure of the

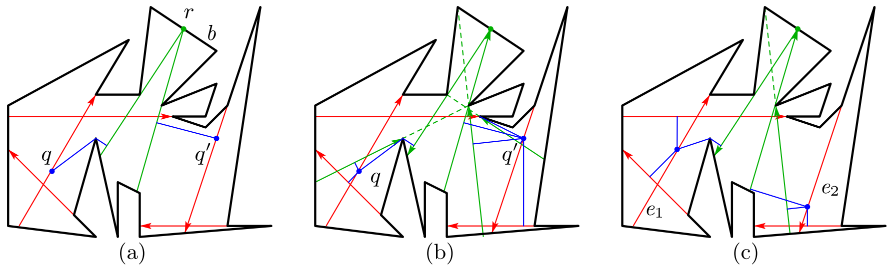

Fig. 8 **a** A pair of tentacles $Z^r_q(b)$ and $Z^r_{q'}(b)$ in blue, **b** a jellyfish pair $\mathcal{J}_{q,q'}$ in blue with jellyfish $\mathcal{J}_q$ and $\mathcal{J}_{q'}$, **c** a minimum jellyfish pair $\mathcal{J}^{\min}_{e_1,e_2}$ with minimum jellyfish $\mathcal{J}^{\min}_{e_1}$ and $\mathcal{J}^{\min}_{e_2}$ in blue (Color figure online)

path) be denoted $T_1[u_1', u_1]$. The path $T_1[u_1', u_1]$ also sees r since u_1' sees r. The tentacle $Z^r_{u_1}(b)$ is at most as long as the shorter of the subpaths $T_1[u_1, u_1']$ and $T_1[u_1', u_1]$, whereby $\|Z^r_{u_1}(b)\| \le \|T_1\|/2$ in this case.

2. If r is an interior point of of b but is not seen by T_1, then it is seen by T_2, and since $\|Z^r_{u_1}(b)\| \le \|Z^r_{u_2}(b)\|$ and by the argument in Case 1. using u_2 and T_2 instead of u_1 and T_1, giving us $\|Z^r_{u_2}(b)\| \le \|T_2\|/2$, we have $\|Z^r_{u_1}(b)\| \le \|T_2\|/2$ in this case.
3. If $r = v$ is a vertex of b seen by T_1 and T_1 also sees other points of b, then by Lemma 3.2, T_1 sees some connected set of b that includes v. Hence $\lim_{b \ni p \to v} \|Z^p_{u_1}\| = \|Z^v_{u_1}(b)\|$ by the definition of edge restricted tentacle and using our argument from Case 1 (including the fact that T_1 sees all points p on b in a connected neighbourhood of v), we have $\|Z^v_{u_1}(b)\| \le \|T_1\|/2$ in this case.
4. If $r = v$ is a vertex of b seen by T_1 but T_1 sees no other points of b, then T_2 must see v and other points of b, otherwise there are points on b not seen by any of T_1 and T_2. Using the same argument as in Case 3 for u_2 and T_2, we have $\|Z^v_{u_2}(b)\| \le \|T_2\|/2$. Since $\|Z^v_{u_1}(b)\| \le \|Z^v_{u_2}(b)\|$ we have $\|Z^v_{u_1}(b)\| \le \|T_2\|/2$ in this case.
5. If $r = v$ is a vertex of b not seen by T_1, then b contains some connected set that includes v and that is not seen by T_1. This set must be seen by T_2 and we can use the argument in the previous step to establish that $\|Z^v_{u_1}(b)\| \le \|T_2\|/2$ also in this case.

This concludes the proof. □

Given two segment s and s' in **P**, we define the *bases* along s and s' to be a pair of points having the property $q_*, q_*' = \arg\min_{q \in s, q' \in s'} \left\{\|\mathcal{J}_{q,q'}\|\right\}$, i.e., two points, q_* on s and q_*' on s', where $\|\mathcal{J}_{q_*,q_*'}\|$ is minimal. We denote the jellyfish pair $\mathcal{J}_{q_*,q_*'}$ by $\mathcal{J}^{\min}_{s,s'}$, the *minimum jellyfish pair*; see Fig. 8b, c.

From this definition and Lemma 4.4, we have that

$$\|\mathcal{J}^{\min}_{e_1,e_2}\| \le \|\mathcal{J}_{u_1,u_2}\| \le \|(T_1, T_2)\|_{\max}/2, \tag{5}$$

where again, e_1 and e_2 are two extensions intersected by T_1 and T_2, respectively.

Let q_1 on e_1 and q_2 on e_2 be the two bases of $\mathcal{J}^{\min}_{e_1,e_2}$; see Fig. 8c. We denote by $\mathcal{J}^{\min}_{e_1}$ the subset of tentacles having their heads at q_1 and by $\mathcal{J}^{\min}_{e_2}$ the subset of tentacles having their heads at q_2. Each set of tentacles $\mathcal{J}^{\min}_{e_1}$ and $\mathcal{J}^{\min}_{e_2}$ is a *minimum jellyfish*. We call q_1 the base of $\mathcal{J}^{\min}_{e_1}$ and q_2 the base of $\mathcal{J}^{\min}_{e_2}$. Thus, $\mathcal{J}^{\min}_{e_1,e_2} = \mathcal{J}^{\min}_{e_1} \cup \mathcal{J}^{\min}_{e_2}$.

Algorithm *Two-Watchman-Route*
Input: A simple polygon $\mathbf{P}$
Output: A two-watchman route $\mathcal{W}_{\mathsf{T}}$ that sees $\mathbf{P}$

1 Run Belleville's algorithm to establish if the polygon is guardable by two point guards. If this is the case, return the two point guards computed by the algorithm
2 Compute the set of extensions $\mathcal{E}$ in $\mathbf{P}$
3 Compute a shortest watchman route W_{opt} in $\mathbf{P}$
4 Let $\mathcal{W}^*_{\mathsf{T}} \leftarrow (W_{\mathsf{opt}}, \text{arbitrary point in } \mathbf{P})$
5 **for** every pair of extensions $e_1, e_2 \in \mathcal{E}$, $e_1 \neq e_2$ **do**
5.1 Establish the bases q_1 and q_2, and compute $\mathcal{J}^{\min}_{e_1,e_2} = \mathcal{J}^{\min}_{q_1,q_2}$
5.2 Sort tentacles in $\mathcal{J}^{\min}_{e_1,e_2}$ in decreasing order and remove covered tentacles $\rightarrow \mathcal{J}^{\text{red}}_{e_1,e_2}$
5.3 Compute the two tours $\mathcal{W}_{\mathsf{T}} = \big(t(\mathcal{J}^{\text{red}}_{e_1}), t(\mathcal{J}^{\text{red}}_{e_2})\big)$
5.4 **if** $\|\mathcal{W}_{\mathsf{T}}\|_{\max} < \|\mathcal{W}^*_{\mathsf{T}}\|_{\max}$ **then** $\mathcal{W}^*_{\mathsf{T}} \leftarrow \mathcal{W}_{\mathsf{T}}$ **endif**
endfor
6 **return** $\mathcal{W}^*_{\mathsf{T}}$
End *Two-Watchman-Route*

Fig. 9 The two-watchman-route algorithm

The main part of our algorithm, presented in the next section, computes, for every pair of extensions and every pair of boundary edges, the heads of the shortest tentacle pairs that see these edges. The heads are then used as potential bases and the length of the jellyfish pair with these heads as bases is computed. We keep the jellyfish pair with minimum length and its bases through the iteration and from the previous discussion we know that at the end of the iteration the final bases are q_1 and q_2.

5 The Algorithm

Our algorithm is illustrated in pseudo-code in Fig. 9 and we show that it approximates a minmax two-watchman route by a factor of $\big(7\pi/6 + 3 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\big)$ and therefore by Inequality (2) also a minsum two-watchman route by a factor twice as large.

The algorithm begins by running Belleville's algorithm [3, 4] to establish if the polygon is guardable by two point guards. If this is the case, it returns the two point guards computed by the algorithm. Note that if $\mathbf{P}$ is two-guardable by point guards, our algorithm must obtain two such point guards to satisfy the approximation guarantee. Otherwise, it computes the set of extensions $\mathcal{E}$ in $O(n \log n)$ time using a dynamic ray-shooting data structure [17], and initializes the solution to be a single shortest watchman route W_{opt}[29, 34] together with some arbitrary point in $\mathbf{P}$.

The rest of this section is devoted to showing how to implement Step 5 of the algorithm. It consists of an iteration over all pairs of extensions. For each pair, we assume that each tour in a minmax two-watchman route intersects one of the extensions in the pair.

5.1 Computing Tentacles and Bases

The algorithm needs to find the two bases q_1 on e_1 and q_2 on e_2. To this end, let s_{e_i} be the maximal line segments through q_i, orthogonal to e_i inside $\mathbf{P}$, for $i \in \{1, 2\}$. The segment s_{e_i} partitions $\mathbf{P}$ into two subpolygons $\mathbf{P}_L$ and $\mathbf{P}_R$. The minimum jellyfish $\mathcal{J}^{\min}_{e_i}$

either has one tentacle that attains the length $\|\mathcal{J}_{e_i}^{\min}\|$ with its first segment orthogonal to e_i or at an endpoint of e_i, or it has two tentacles, one in $\mathbf{P}_L$ and the other in $\mathbf{P}_R$, that attain this length. To prove this, note that if the single tentacle attaining the length $\|\mathcal{J}_{e_i}^{\min}\|$ does not have a first segment orthogonal to e_i, then by moving the head slightly along e_i, we can reduce the length of the tentacle and thereby the jellyfish. Similarly, if all tentacles attaining the maximal length are in the same subpolygon, say $\mathbf{P}_L$, then we can move the head along e_i into $\mathbf{P}_L$, again reducing the length of the jellyfish. In both cases, this contradicts that $\mathcal{J}_{e_i}^{\min}$ is a minimum jellyfish from a minimum jellyfish pair on e_1 and e_2. Thus, there are at most two longest tentacle pairs of $\mathcal{J}_{e_1,e_2}^{\min}$, at least one pair of which attains the length $\|\mathcal{J}_{e_1,e_2}^{\min}\|$.

We construct the data structures needed to compute the bases. These data structures are, for each reflex vertex in $\mathbf{P}$ and each extension endpoint, the shortest path tree to every vertex in $\mathbf{P}$[13]. The shortest path trees are augmented with the additional edges and leaves obtained by extending each tree edge until it intersects a boundary edge of $\mathbf{P}$ without crossing any other tree edge. Also, for each augmented tree, a data structure is built, enabling us to find the common ancestor of any pair of nodes in the tree (vertices of $\mathbf{P}$) in constant time [15]. These can be precomputed in linear time for each root vertex, i.e., in quadratic time in total.

Next, we do a case analysis based on the number of longest tentacles occurring in a jellyfish pair. There are six cases to deal with.

Case 1

If $\mathcal{J}_{e_1,e_2}^{\min}$ has one unique longest tentacle, then we know from the previous discussion that it is a tentacle $Z_{q_{e_i}}^{v}(b)$ for some vertex endpoint v of some boundary edge b and $i = 1$ or $i = 2$. The point q_{e_i} is the point on e_i that minimizes this distance and furthermore $\|Z_{q_{e_i}}^{v}(b)\| \leq \|Z_{q_{e_{3-i}}}^{v}(b)\|$. If v is a reflex vertex, let e be the extension associated to it and we have that $Z_{q_{e_i}}^{v}(b) = SP\big(e_i, \mathbf{VP}(v) \cap \mathbf{L}(e)\big)$ and if v is a convex vertex, $Z_{q_{e_i}}^{v}(b) = SP\big(e_i, \mathbf{VP}(v)\big)$. In both cases, q_{e_i} is the point of intersection between the shortest path and e_i. Given e_1 and e_2 we can, for each boundary edge b and for each vertex end point v, verify if v is reflex or convex, compute $SP\big(e_1, \mathbf{VP}(v) \cap \mathbf{L}(e)\big)$ and $SP\big(e_2, \mathbf{VP}(v) \cap \mathbf{L}(e)\big)$ or $SP\big(e_1, \mathbf{VP}(v)\big)$ and $SP\big(e_2, \mathbf{VP}(v)\big)$, depending on the case, in linear time [11, 12, 18, 19, 21, 28]. We denote these tentacles $Z_{e_1}^{v}(b)$ and $Z_{e_2}^{v}(b)$ to indicate that the intersection points with e_1 and e_2 are not fixed given points. We select the shorter of the two, $Z_{e_i}^{v}(b)$, identifying the intersection with the corresponding extension e_i, $i = 1$ or 2, keeping the q_{e_i} for which the tentacle is the longest. After iterating over all boundary edges, one potential base q_{e_i} remains on e_i and one $q_{e_{3-i}}$ remains on e_{3-i}. The whole process takes quadratic time and gives us the potential bases $q_1^{(1)}$ and $q_2^{(1)}$ on e_1 and e_2, respectively.

Case 2

If $\mathcal{J}_{e_1,e_2}^{\min}$ has two longest tentacles, this can occur in three different ways. One possibility is that two tentacles $Z_{e_i}^{v}(b)$ and $Z_{e_j}^{v'}(b')$, $i \in \{1, 2\}$, $j \in \{1, 2\}$ and different boundary edges b and b', have exactly the same length. If $i = j$, the heads of $Z_{e_i}^{v}(b)$ and $Z_{e_j}^{v'}(b')$

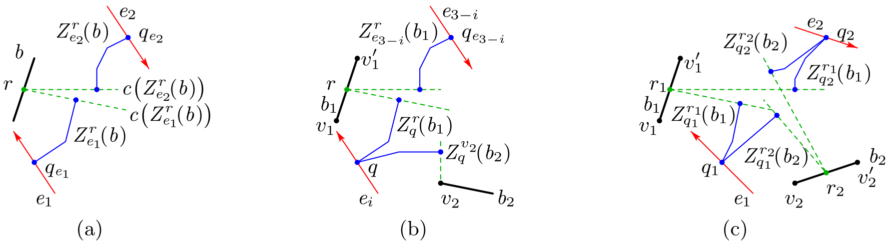

Fig. 10 Computing the bases of the minimum jellyfish pair

also coincide in this case, otherwise the base is found in Case 3. Such tentacles are discovered using the method described in the previous case.

The second way is that two tentacles in $\mathcal{J}^{\min}_{e_1,e_2}$ are $Z^{r^*}_{e_1}(b)$ and $Z^{r^*}_{e_2}(b)$, for some boundary edge b, where r^* is a point on b such that $\|Z^{r^*}_{q_{e_1}}(b)\| = \|Z^{r^*}_{q_{e_2}}(b)\|$, q_{e_1} and q_{e_2} being the points on e_1 and e_2 that make these tentacles as short as possible. We can, for each boundary edge $b = [v, v']$, compute $Z^{v}_{e_1}(b)$, $Z^{v}_{e_2}(b)$, $Z^{v'}_{e_1}(b)$, and $Z^{v'}_{e_2}(b)$ as in Case 1. Now, if $\|Z^{v}_{e_1}(b)\| < \|Z^{v}_{e_2}(b)\|$ and $\|Z^{v'}_{e_1}(b)\| > \|Z^{v'}_{e_2}(b)\|$, or the inequalities are reversed, we let a point r slide along b from v to the other end point v'; see Fig. 10a. By Corollary 4.3, the lengths of $Z^r_{e_1}(b)$ and $Z^r_{e_2}(b)$ as r moves along b are smooth functions (continuous and differentiable) except at positions denoted *event points* where any of the participating tentacles has one of the following properties:

1. an interior point of the first edge of a tentacle intersects a vertex of **P**, or the first and second edge of the tentacle become collinear,
2. an interior point of the last edge of a tentacle intersects a vertex of **P**, or the last and penultimate edge of the tentacle become collinear,
3. the cut $c\big(Z^r_{e_i}(b)\big)$, $i \in \{1, 2\}$, becomes intersects more than one vertex of **P**,
4. the head or the tip of a tentacle either hits or leaves the boundary of **P**, and
5. one tentacle goes from being shorter than the longer tentacle to becoming the longer tentacle, at which point the two tentacles have equal length.

The positions on b for the first three event point types can be obtained in linear time, if $Z^r_{e_i}(b)$ consists of at least two segments having as common end point, the reflex boundary vertex u_i, by a traversal of the augmented shortest path tree to each vertex in **P** from the root u_i, giving us a total of $O(n)$ such event points. If $Z^r_{e_i}(b)$ consists of just one segment, the positions for the first three event point types on b can be obtained in linear time by a traversal of the augmented shortest path trees rooted at the two end points of e_i.

The fourth type is obtained by finding the point at which the first edge of the tentacle hits either end point of e_i and where the last edge of the tentacle intersects $c\big(Z^r_{e_i}(b)\big)$ orthogonally on the boundary which can happen only once for each tentacle in each interval between event points of the first three types. Finally, the last type of event points can occur only a constant number of times in each interval between event points of the first four types, since here each function is the sum of at most three square roots of rational polynomials of constant degree and two such functions can be equal in at most a constant number of points. The number of event points between v and v'

on b is at most linear and iterating over all possible boundary edges, takes quadratic time to solve this case. This gives us the potential bases $q_1^{(2)}$ and $q_2^{(2)}$ on e_1 and e_2, respectively.

Case 3

In the third possibility that $\mathcal{J}_{e_1,e_2}^{\min}$ has two longest tentacles, the two tentacles are $Z_q^{v_1}(b_1)$ and $Z_q^{v_2}(b_2)$ for two boundary edges b_1 and b_2 with vertex end points v_1 and v_2 respectively and q on e_i such that $\|Z_q^{v_1}(b_1)\| = \|Z_q^{v_2}(b_2)\|$. We compute $Z_{e_i}^{v_1}(b_1)$ and $Z_{e_i}^{v_2}(b_2)$ and let q_{e_i} and q_{e_i}' be the heads of $Z_{e_i}^{v_1}(b_1)$ and $Z_{e_i}^{v_2}(b_2)$, respectively, on e_i. We let a point q slide along e_i from q_{e_i} to q_{e_i}'. Again, by Corollary 4.3, the lengths of $Z_q^{v_1}(b_1)$ and $Z_q^{v_2}(b_2)$ as q moves along e_i are smooth functions (continuous and differentiable) except at the event points established in Case 2. At these event points, the structures of the tentacles are updated and we can test whether the two tentacles have equal length for some head point on e_i before the next event point. The number of event points between q_{e_i} and q_{e_i}' on e_i is at most linear. We perform these steps also on e_{3-i} and take the shortest pair as the representative base for the pair (v_1, v_2) of vertices. Iterating over all possible pairs of boundary edges and their end points, the process takes cubic time in total and gives us the potential bases $q_1^{(3)}$ and $q_2^{(3)}$ on e_1 and e_2, respectively.

Case 4

If $\mathcal{J}_{e_1,e_2}^{\min}$ has three longest tentacles, this can occur in two ways. The first possibility if some combination of three tentacles with common heads from the three previous cases exist, in which case these can be established using the previously described methods.

The second possibility that $\mathcal{J}_{e_1,e_2}^{\min}$ has three longest tentacles occurs if there are two boundary edges $b_1 = [v_1, v_1']$ and $b_2 = [v_2, v_2']$ such that $\|Z_{q^*}^{r^*}(b_1)\| = \|Z_{e_{3-i}}^{r^*}(b_1)\| = \|Z_{q^*}^{v_2}(b_2)\|$, with q^* on e_i; see Fig. 10b. We can establish the event points by a combination of the methods in Cases 2 and 3. Begin by finding the interval $[q_{e_i}, q'_{e_i}]$ on e_i by computing the tentacles $Z_{e_i}^{v_1}(b_1)$ and $Z_{e_i}^{v_2}(b_2)$ having the heads q_{e_i} and q_{e_i}', and then the event points generated by $Z_q^{v}(b)$ and $Z_q^{v'}(b')$ in order as q moves along e_i from q_{e_i} to q_{e_i}'. Subsequently, establish the event points generated by $Z_{e_i}^{r}(b)$ and $Z_{e_{3-i}}^{r}(b)$ in order as r moves along b. For each pair of intervals between event points along e_i and along b_1, we verify if we can establish a point q in the interval on e_i and a point r in the interval on b_1 such that $Z_q^{r}(b_1)$, $Z_q^{v_2}(b_2)$, and $Z_{e_{3-i}}^{r}(b_1)$ have equal length by solving the system of equalities given by the length functions of the three tentacles. Within each interval the length functions are smooth according to Corollary 4.3 so this takes constant time. In the worst case, we have to consider a linear number of intervals each on e_i and b_1 taking quadratic time. Iterating over all possible pairs of boundary edges and their end points, the computation takes quartic time in total and gives us the potential bases $q_1^{(4)}$ and $q_2^{(4)}$ on e_1 and e_2, respectively.

Case 5

If $\mathcal{J}^{\min}_{e_1,e_2}$ has four longest tentacles, this can occur in two ways. Again, the first possibility is if some combination of four tentacles with common heads from the four previous cases exist, in which case these can be established using the previously described methods.

The second possibility that $\mathcal{J}^{\min}_{e_1,e_2}$ has four longest tentacles occurs if there are two boundary edges b_1 and b_2 such that $\|Z^{r_1^*}_{q_1^*}(b_1)\| = \|Z^{r_1^*}_{q_2^*}(b_1)\| = \|Z^{r_2^*}_{q_1^*}(b_2)\| = \|Z^{r_2^*}_{q_2^*}(b_2)\|$, with q_1^* on e_1, q_2^* on e_2, r_1^* on b_1, and r_2^* on b_2; see Fig. 10c. We can establish the event points by extending the method in Case 4, giving a linear number of intervals on each segment e_1, e_2, b_1, and b_2, where the length functions of the tentacles are smooth according to Corollary 4.3. We verify if we can establish point q_1, q_2, r_1, and r_2 in their respective interval such that $Z^{r_1}_{q_1}(b_1)$, $Z^{r_1}_{q_2}(b_1)$, $Z^{r_2}_{q_1}(b_2)$, and $Z^{r_2}_{q_2}(b_2)$ have equal length by solving the system of equalities given by the length functions of the four tentacles which takes constant time in each interval. In the worst case, we have to consider a linear number of intervals each on e_1, e_2, b_1, and b_2, thus taking quartic time. Iterating over all possible pairs of boundary edges and their end points, the computation takes $O(n^6)$ time in total and gives us the potential bases $q_1^{(5)}$ and $q_2^{(5)}$ on e_1 and e_2, respectively.

Case 6

If $\mathcal{J}^{\min}_{e_1,e_2}$ has five or more longest tentacles, they must occur as some combination of tentacles structured as in the previous cases and they can therefore be obtained using the methods described above.

Analysis of the Preprocessing Step

The case analysis above gives us five pairs of potential bases, $(q_1^{(1)}, q_2^{(1)}), (q_1^{(2)}, q_2^{(2)}), \dots, (q_1^{(5)}, q_2^{(5)})$, for which we can compute the jellyfish pairs $\mathcal{J}_{q_1^{(1)},q_2^{(1)}}, \mathcal{J}_{q_1^{(2)},q_2^{(2)}}, \dots, \mathcal{J}_{q_1^{(5)},q_2^{(5)}}$, each in quadratic time, and from these we select the minimum one. By Corollary 4.3 and the previous discussion, it follows that this jellyfish pair is a minimum jellyfish pair $\mathcal{J}^{\min}_{e_1,e_2}$ with bases q_1 and q_2 on e_1 and e_2, respectively.

In this way, we have an $O(n^6)$ time subroutine to find the bases q_1 and q_2 on e_1 and e_2, respectively. Given the bases, the computation of $\mathcal{J}^{\min}_{e_1,e_2}$ takes an additional quadratic time. The whole process of Step 5.1 of the algorithm thus takes $O(n^6)$ time.

5.2 Establishing the Tours

Given a minimum jellyfish pair $\mathcal{J}^{\min}_{e_1,e_2}$, we sort the tentacles in decreasing order of length and for each tentacle $Z^r_{q_i}(b)$ in decreasing order check if $c\big(Z^r_{q_i}(b)\big)$ is already covered by a longer tentacle. If so, $Z^r_{q_i}(b)$ is removed from the jellyfish pair. We call the resulting jellyfish pair *reduced* and denote it by $\mathcal{J}^{\text{red}}_{e_1,e_2}$ with individual jellyfish $\mathcal{J}^{\text{red}}_{e_1}$ and $\mathcal{J}^{\text{red}}_{e_2}$.

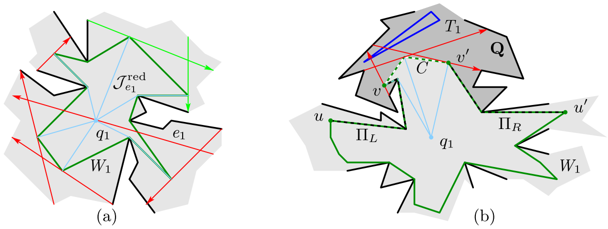

Fig. 11 Illustrating the construction of the tour and the partition into subpaths used in the proof of Lemma 5.1. **a** The green tour W_1 is the relative convex hull of the jellyfish $\mathcal{J}^{\rm red}_{e_1}$ (light blue). **b** Partitioning the tour W_1 into subpaths C, Π_R, and Π_L, when W_1 does not intersect or contain T_1 (Color figure online)

Let $t(\mathcal{J})$ denote the two relative convex hulls of each of the jellyfish in the jellyfish pair $\mathcal{J}$ inside **P**. Without loss of generality, if $\mathcal{J} = \mathcal{J}_1 \cup \mathcal{J}_2$ with $\mathcal{J}_1$ and $\mathcal{J}_2$ being the two jellyfish, $t(\mathcal{J}_1)$ and $t(\mathcal{J}_2)$ also denote the individual relative convex hulls of $\mathcal{J}_1$ and $\mathcal{J}_2$ in **P**.

Let W_1 and W_2 be the relative convex hulls $t(\mathcal{J}^{\rm red}_{e_1})$ and $t(\mathcal{J}^{\rm red}_{e_2})$ in **P**, respectively. These tours can be computed in linear time by first following the shortest path from each tentacle tip to the next, cyclically around the corresponding head of each jellyfish and then applying the algorithm by Toussaint [35]; see Fig. 11a.

Consider a polygonal tour that has reflex vertices only at reflex vertices of **P**. A maximal consecutive subsequence of edges of the tour is called a *reflex chain* of the tour, if each interior vertex of the chain is reflex in the tour. The end vertices of a reflex chain must therefore be convex vertices of the tour. In contrast, a *convex chain* of the tour is a maximal consecutive subsequence of edges such that each interior and end vertex of the chain is convex in the tour. We note that a single convex vertex of a tour is a convex chain if the preceding and subsequent vertices are reflex.

Lemma 5.1 *The tours* (W_1, W_2) *obtained by algorithm Two-Watchman-Route form a two-watchman route and*

$$\|(W_1, W_2)\|_{\max} \le (7\pi/6 + 3 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5})\|(T_1, T_2)\|_{\max} \approx 5.969\|(T_1, T_2)\|_{\max}.$$

Proof The correctness of the algorithm follows from Lemmas 3.1, 4.1, and the fact that since the two tours together see every boundary edge this ensures that they form a two-watchman route.

To prove the approximation bound, we assume that T_1 and T_2 do not intersect, otherwise $\|W_{\rm opt}\| \le \|T_1\| + \|T_2\| \le 2\|(T_1, T_2)\|_{\max}$ immediately proving the lemma since $\|(W_1, W_2)\|_{\max} \le \|W_{\rm opt}\|$ and the algorithm initializes the two-watchman route with a shortest watchman tour and a point in Step 3 and then only updates its two-watchman route in Step 5.4 if its max length is strictly smaller than that of the current route pair.

The algorithm computes the reduced minimum jellyfish pair $\mathcal{J}^{\rm red}_{e_1,e_2}$ in Step 5.2. By trying all pairs of extensions in Step 5, the algorithm must necessarily consider a pair, e_1 and e_2, intersected by the tours T_1 and T_2; see Lemma 3.5. Consider the tentacles in

$\mathcal{J}_{e_1}^{\text{red}}$ and $\mathcal{J}_{e_2}^{\text{red}}$ centered on the bases q_1 on e_1 and q_2 on e_2, respectively. Every tentacle has length at most $R = \|(T_1, T_2)\|_{\max}/2$ by Inequality (5), since we can assume that T_1 intersects e_1 and T_2 intersects e_2, whereby the geodesic radii of W_1 and W_2 are both at most R.

Each convex chain of W_1 has length at most $2\pi R$, where R is an upper bound on the length of each geodesic shortest path from q_1 to any point on W_1, since the circle is the longest convex curve of radius R. This follows from Archimedes' axioms for arc-length [1, 5].

We make a case analysis and bound the length of W_1 for each case separately. The first case when W_1 has one convex chain can be easily bounded. When W_1 has two or more convex chains, we separate that into two further cases, when W_1 intersects or contains both T_1 and T_2 and when it does not. The first of these cases can be bounded by estimating the total length of a tour including both T_1 and T_2. For the last case, we assume that W_1 does not intersect T_1 and partition W_1 into the subpath C_1 and the subtour W_1' and bound the length of these separately. The path C_1 passes those tentacle cuts for tentacles in $\mathcal{J}_{e_1}^{\text{red}}$ that are covered exclusively by T_1 and the subtour W_1' is the relative convex hull of the tentacles in in $\mathcal{J}_{e_1}^{\text{red}}$ for which the tentacle cuts are covered by T_2. The length of the subpath C_1 is bounded in Lemma 5.2 and the length of the subtour W_1' is bounded in Lemma 5.3. We provide the details below.

W_1 has one convex chain. The convex chain of W_1 has length at most $2\pi R$ as we noted above. The length of the possible reflex chain of W_1 is bounded by at most two radii from q_1 to the circle perimeter since its length is bounded by that of the two tentacles of $\mathcal{J}_{e_1}^{\min}$ connecting to the end points of the reflex chain. Hence,

$$\|W_1\| \leq 2\pi R + 2R = (\pi + 1)\|(T_1, T_2)\|_{\max}. \tag{6}$$

W_1 has at least two convex chains. This case is further subdivided into the cases:

W_1 or its interior intersects both T_1 and T_2. If W_1 intersects T_1, let p_1 be an intersection point of W_1 with T_1. From p_1 move counterclockwise along W_1 until the end point of a tentacle from $\mathcal{J}_{e_1}^{\text{red}}$ is reached at p_{L}. Similarly, move clockwise along W_1 until the end point of a tentacle from $\mathcal{J}_{e_1}^{\text{red}}$ is reached at p_{R}. Since the region bounded by moving counterclockwise from p_{R} along W_1 to p_{L}, from p_{L} along the tentacle of $\mathcal{J}_{e_1}^{\text{red}}$ to q_1, and from q_1 along a tentacle of $\mathcal{J}_{e_1}^{\text{red}}$ to p_{R} forms a pseudotriangle, relatively convex inside $\mathbf{P}$, the shortest path from p_1 to q_1 has length at most that of one of the tentacles of $\mathcal{J}_{e_1}^{\text{red}}$, which is bounded by R. If W_1 does not intersect T_1, then T_1 lies interior to W_1 and we let p_1 be a point closest to q_1. Again, the length of shortest path from p_1 to q_1 is bounded by R, since extending the first edge until it hits the tour W_1 at some point p_1', the shortest path from p_1' to q_1 contains p_1 and since p_1' lies on W_1, we can use the previous argument to show that the length of the shortest path from p_1' to q_1 is bounded by R.
If W_1 intersects T_2, let p_2 be intersection point of W_1 with T_2, otherwise let p_2 be a point on T_2 closest to q_1. In the same way as above, we can bound the length of the shortest path from p_2 to q_1 by R.

Now, construct a tour W_{bnd} by following T_1 from p_1 around in counterclockwise order, following the shortest path from p_1 to q_1, the shortest path from q_1 to p_2, from p_2 around T_2 in counterclockwise order, and then back along the shortest paths via q_1 to p_1. It is clear that W_{bnd} is a watchman tour and therefore has length no shorter than W_{opt} and since the algorithm initializes the two-watchman route with W_{opt} and an arbitrary point in Step 3 and then only updates its two-watchman route in Step 5.4 if it has max-length strictly smaller than the current route pair, the length of W_1 is bounded by

$$\|W_1\| < \|W_{\mathsf{opt}}\| \le \|W_{\mathsf{bnd}}\| \le \|T_1\| + \|T_2\| + 4R \le 4\|(T_1, T_2)\|_{\mathsf{max}}. \quad (7)$$

W_1 or its interior intersects at most one of T_1 or T_2. Without loss of generality, we assume that T_1 does not intersect W_1 or the interior of W_1, otherwise we interchange the roles of T_1 and T_2 in the argument below.

Cut **P** along the segments of W_1 into disjoint pieces, thus partitioning **P** into separate components; see Fig. 11b. Let **Q** be the component containing T_1. The subpath Π of W_1 bounding **Q** is either a single segment connecting two reflex vertices of W_1 or it consists of a single convex chain C between two segments having reflex vertices v and v' at the end points. Extend Π at both ends along W_1 until convex vertices are reached at both ends, these must exist since W_1 has at least two convex chains. Let the path thus obtained be Π'. If Π includes a convex chain, consider the two paths $\Pi' \setminus C$ that we denote by Π_L and Π_R, where Π_L appears before Π_R during a counterclockwise traversal of W_1 starting at a point in C; see Fig. 11b.

Walk along C from each end point v and v' until convex vertices $\bar{v}$ and $\bar{v}$' are reached for which the associated tentacle cuts are not covered by T_2, if there are any. These cuts must be covered by T_1. Let C_1 be the subpath of C from $\bar{v}$ to $\bar{v}$'. If $C_1 = C$, then already the end points of C have associated tentacle cuts that are not covered by T_2 and if all associated tentacle cuts to the vertices of C are covered by T_2, then we let $C_1 = \emptyset$.

If C_1 is not empty, Lemma 5.2 below gives us that $\|C_1\| \le 4\pi R/3$ and if it is empty or a single point $\|C_1\| = 0$. Thus, we need to bound the length of the remainder of W_1. Let p_x be some intersection point between W_1 and e_1. From p_x, walk counterclockwise along W_1 back to p_x, shortcutting those convex vertices associated to tentacle cuts only covered by T_1 and let W_1' be the tour thus obtained. Let $\mathcal{J}_{e_1}^2$ be the subset of tentacles of $\mathcal{J}_{e_1}^{\text{red}}$ for which each associated tentacle cut is covered by T_2. By construction, the tour W_1' is the relative convex hull of the tentacles of $\mathcal{J}_{e_1}^2$ in **P**. All convex vertices on W_1 associated to tentacle cuts covered exclusively by T_1 lie on C_1 by our observation above, giving us

$$\|W_1\| \le \|W_1'\| + 2R + \|C_1\|, \quad (8)$$

since the union of the two tentacles connecting q_1 with the end points of C_1 intersects W_1', they together cover the tentacle cuts covered by W_1 at the same points, and W_1 is relatively convex.

From Inequality (8), Lemmas 5.2 and 5.3 below, bounding the lengths of C_1 and W_1', we have for this case

$$\begin{aligned}\|W_1\| &\leq \|W_1'\| + 2R + \|C_1\| \\ &\leq \left(\pi/2 + 2 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\right)\|(T_1, T_2)\|_{\mathsf{max}} + 2R + 4\pi R/3 \\ &= \left(7\pi/6 + 3 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\right)\|(T_1, T_2)\|_{\mathsf{max}}. \qquad (9)\end{aligned}$$

From the three cases that we have dealt with, Inequalities (6), (7), and (9), we have

$$\|W_1\| \leq \max\begin{cases}\left(\pi + 1\right)\|(T_1, T_2)\|_{\mathsf{max}}, \\ 4\|(T_1, T_2)\|_{\mathsf{max}}, \\ \left(7\pi/6 + 3 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\right)\|(T_1, T_2)\|_{\mathsf{max}}.\end{cases} \qquad (10)$$

The last case has the worst upper bound
$\|W_1\| \leq \left(7\pi/6 + 3 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\right)\|(T_1, T_2)\|_{\mathsf{max}} \approx 5.969\|(T_1, T_2)\|_{\mathsf{max}}$.
We can make the same argument for the second tour W_2, obtaining the same bound, whereby the lemma follows. □

We bound the length of the subpath C_1 next.

Lemma 5.2 *Given a jellyfish $\mathcal{J}_{e_1}^{\text{red}}$ with relative convex hull W_1 that does not intersect T_1, let C_1 be the maximal subpath of W_1 connecting convex vertices of W_1 such that the corresponding tentacle cuts are covered by T_1. We have*

$$\|C_1\| \leq \frac{4\pi \cdot R}{3}.$$

Proof Consider the two end points $\bar{v}$ and $\bar{v}$' of C_1 (if the end points coincide, C_1 is a single point and has length 0, so we assume that the end points do not coincide); see Fig. 11b. Let e and e' be the two tentacle cuts covered by the two end points $\bar{v}$ and $\bar{v}$' of C_1. Since T_2 does not cover any of e or e', T_1 must have points to the left of these cuts. Without loss of generality, we assume that $\bar{v}$ is passed before $\bar{v}$' during a counterclockwise traversal of W_1 starting at a point interior to C_1. We make the additional observation that any tentacle cut associated to convex vertices of $W_1 \setminus C_1$ must be covered by T_2, otherwise T_1 would intersect W_1 or have points in its interior.

Let l and l' be the lines through e and e' and let p be the intersection point between l and l', assuming that it exists. If l and l' are parallel or p lies after $\bar{v}$ on l in the direction of e, then the angle between the tentacle of $\mathcal{J}_{e_1}^{\text{red}}$ connecting q_1 and $\bar{v}$ and the tentacle of $\mathcal{J}_{e_1}^{\text{red}}$ connecting q_1 and $\bar{v}$' containing C_1 is at most π, since q_1, $\bar{v}$', p, and $\bar{v}$ form a convex quadrilateral (p can be taken to be a point on l implicitly at infinity, if l and l' are parallel); see Fig. 12a. Each segment of C_1 can be projected onto a half circle centered at q_1 having radius R without overlap so the length of C_1 is therefore bounded by πR in this case.

Now, if p lies before $\bar{v}$ on l in the direction of e; see Fig. 12b, then since T_1 lies in **Q** without intersecting C_1, T_1 must intersect the two cuts e and e'. Assume these

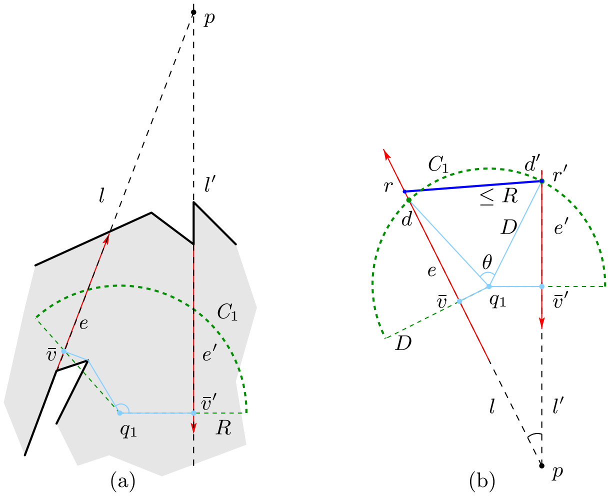

Fig. 12 Illustrating the proof of Lemma 5.2. Bounding the length of C_1

intersection points are r and r'. The distance between r and r' is at most R since $\|T_1\| \leq \|(T_1, T_2)\|_{\mathsf{max}} = 2R$ by Inequality 5. Assume that the longest tentacle of $\mathcal{J}_{e_1}^{\mathrm{red}}$ connecting q_1 with a convex vertex of C_1 has length $D \leq R$ and let d and d' be the points at distance exactly D from q_1 to l and l', respectively. The point d must lie between $\bar{v}$ and r on e, otherwise T_1 intersects the longest tentacle of $\mathcal{J}_{e_1}^{\mathrm{red}}$ and therefore also C_1. For the same reason, the point d' must lie between $\bar{v}$' and r' on e'. The distance $\|d, d'\| \leq R$ since T_1 does not intersect W_1. Given that the angle at d, q_1, $\bar{v}$, and the angle at d', q_1, $\bar{v}$', are each at most $\pi/2$, that the angle at d, q_1, d' is θ, the segments of C_1 can be projected onto a semicircle centered at q_1 having radius D without overlap. The length of this semicircle is $(\pi + \theta)D$, which is maximized for $\theta = \pi/3$ and $D = R$ as obtained by standard analytic methods. Thus, $\|C_1\| \leq 4\pi R/3$ is the maximum bound in all cases. □

It only remains to show the bound for the length of the subtour W_1'.

Lemma 5.3 *Given a jellyfish $\mathcal{J}_{e_1}^{\mathrm{red}}$ with base q_1 and relative convex hull W_1 that does not intersect T_1, let W_1' be the relative convex hull of $\mathcal{J}_{e_1}^{2}$, the subset of tentacles Z_{q_1} in $\mathcal{J}_{e_1}^{\mathrm{red}}$ for which each tentacle cut $c(Z_{q_1})$ is covered by T_2. We have*

$$\|W_1'\| \leq \left(\frac{\pi}{2} + 2 - \sqrt{3} + \sqrt{5}\arcsin\frac{1}{\sqrt{5}}\right)\|(T_1, T_2)\|_{\mathsf{max}}.$$

Proof Let e be an extension covered by T_1 such that T_1 makes a reflection on e. Such an extension must exist by Lemma 3.5 since, if T_1 properly covers e, then it must cover some other extension in $\mathbf{L}(e)$. Without loss of generality, we can therefore assume that $e_1 = e$ and thus that all but one point of T_1 lie in $\mathbf{P} \setminus \mathbf{L}(e_1)$. It also follows that T_2

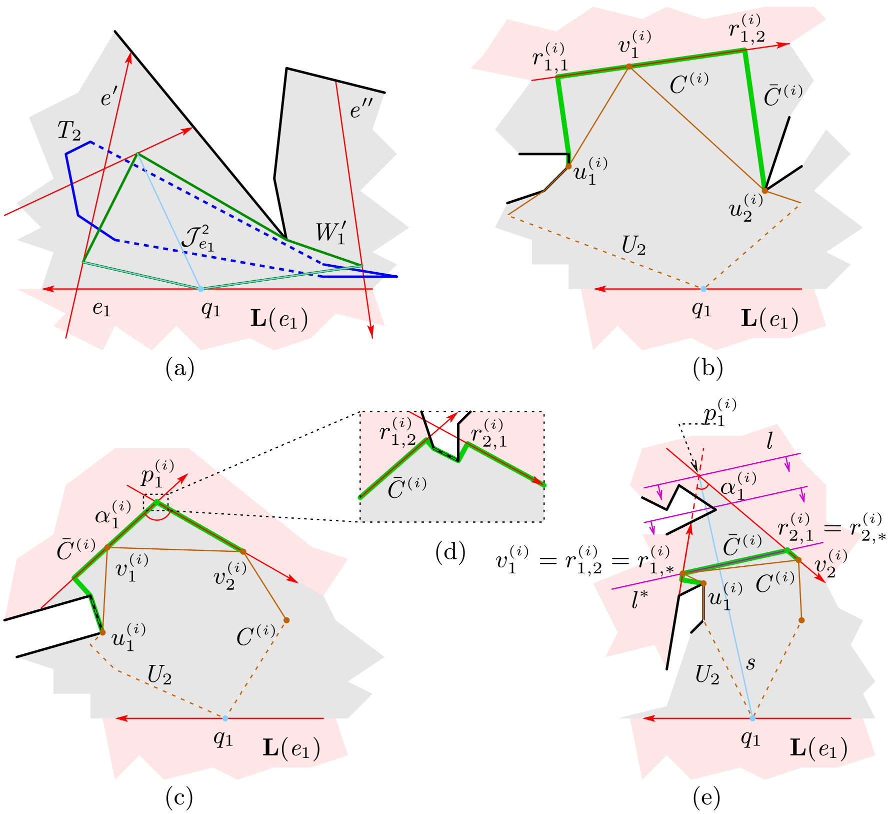

Fig. 13 Illustrating the proof of Lemma 5.3. **a** Tours T_2 and W_1' intersect if W_1' has reflex vertices. **b**–**e** Constructing the chain $\bar{C}^{(i)}$ from $C^{(i)}$. **b** The base case of one extension, **c** the inductive case with obtuse angle between two extensions, **d** the inductive case with acute angle between two extensions

cannot cover e_1, since this would mean that T_1 could be made shorter contradicting that T_1 and T_2 are as short as possible. Hence, all points of T_2 lie in $\mathbf{P} \setminus \mathbf{L}(e_1)$.

We identify two different cases. In the first case, W_1' has at least two convex chains, we bound the length of W_1' by considering a tour V_2 that contains W_1' and show inductively on the number number of convex chains it contains that $\|V_2\| \leq 2\sqrt{2}\|(T_1, T_2)\|_{\max}$. In the second case, W_1' has one convex chain and we bound its length by a short sequence of circular arcs. We provide the details below.

W_1' has at least two reflex chains or one reflex chain with both end points in $\mathbf{P} \setminus \mathbf{L}(e_1)$ We prove that T_2 must intersect W_1'. Let e' and e" be the two tentacle cuts associated to convex vertices of W_1' on either side of a reflex chain with both end points in $\mathbf{P} \setminus \mathbf{L}(e_1)$. The cuts e' and e" must exist and do not intersect in $\mathbf{P} \setminus \mathbf{L}(e_1)$, otherwise W_1' does not have the stated reflex chain. Now, T_2 must cover e' and e" so it either covers e' exactly at the reflection point of W_1' and e' or T_2 covers e" exactly at the reflection point of W_1' and e", giving us an intersection point between W_1' and T_2 in either case, or T_2 contains a path that connects a point in

$\big(\mathbf{P}\setminus\mathbf{L}(e_1)\big)\cap\mathbf{L}(e')$ with a point in $\big(\mathbf{P}\setminus\mathbf{L}(e_1)\big)\cap\mathbf{L}(e'')$ which must intersect W_1' since e' and e" are separated by reflex vertices on W_1' touching the boundary of $\mathbf{P}$ and T_2 has no point in the interior of $\mathbf{L}(e_1)$; see Fig. 13a. We can construct a tour that connects q_1 with the closest point on T_2, follows T_2 around in counterclockwise order and then connects to q_1. This tour has the following properties, that we call *bounding properties*,

1. it covers all the tentacle cuts covered by W_1' (and thus by $\mathcal{J}^2_{e_1}$),
2. it lies completely in $e_1 \cup (\mathbf{P} \setminus \mathbf{L}(e_1))$,
3. it has a convex vertex at q_1, and
4. it has length at most $2R + \|T_2\|$.

The last bounding property follows since T_2 intersects W_1' (or its interior). Now, let U_2 be the shortest tour that obeys the bounding properties. We show how to extend U_2 to a new tour V_2 so that W_1' is contained in V_2 and $\|V_2\| \le \sqrt{2}\|U_2\|$. Since W_1' is relatively convex and contained in V_2, we immediately have

$$\|W_1'\| \le \|V_2\| \le \sqrt{2}\|U_2\| \le 2\sqrt{2}R + \sqrt{2}\|T_2\|. \tag{11}$$

To prove Inequality (11), follow U_2 counterclockwise from q_1, subdividing the tour into convex chains $C^{(i)}$, $1 \le i \le k$, such that $u_1^{(i)}$ is the reflex vertex before $C^{(i)}$ on U_2, and $u_2^{(i)}$ is the reflex vertex after $C^{(i)}$ on U_2. In the cases where the first vertex after q_1 or before q_1 is convex, we consider q_1 to be the vertex $u_1^{(1)}$ and $u_2^{(k)}$ according to the case. For each convex chain $C^{(i)}$, let $m^{(i)}$ be the number of convex vertices on the chain, denoted by $v_1^{(i)}, \ldots, v_{m^{(i)}}^{(i)}$, these are the tips of the tentacles of $\mathcal{J}^2_{e_1}$, and let $e_1^{(i)}, \ldots, e_{m^{(i)}}^{(i)}$ be the corresponding associated tentacle cuts; see Fig. 13b–e. It is clear that no tentacle of $\mathcal{J}^2_{e_1}$ intersects the interior of $\mathbf{L}(e_j^{(i)})$, for any tentacle cut $e_j^{(i)}$, since $\mathcal{J}^2_{e_1} \subseteq \mathcal{J}^{\text{red}}_{e_1}$ and if it did, it would cover $e_j^{(i)}$ whereby $e_j^{(i)}$ and its associated tentacle would be discarded as $\mathcal{J}^{\text{red}}_{e_1}$ is constructed. (To guarantee this property is in fact the reason why the algorithm constructs $\mathcal{J}^{\text{red}}_{e_1,e_2}$ in Step 5.2.) To obtain V_2 from U_2, we replace each convex chain $C^{(i)}$, from $u_1^{(i)}$ to $u_2^{(i)}$ by a different path $\bar{C}^{(i)}$. If $m^{(i)} = 1$, i.e., $C^{(i)}$ has only one convex vertex $v_1^{(i)}$, then let $r_{1,1}^{(i)}$ be the end point of $SP(u_1^{(i)}, e_1^{(i)})$ on the tentacle cut $e_1^{(i)}$. Similarly, let $r_{1,2}^{(i)}$ be the end point of $SP(u_2^{(i)}, e_1^{(i)})$ on $e_1^{(i)}$. We let $\bar{C}^{(i)}$ be the shortest path from $u_1^{(i)}$ to $e_1^{(i)}$, the path from $r_{1,1}^{(i)}$ to $r_{1,2}^{(i)}$ on $e_1^{(i)}$, followed by the shortest path from $e_1^{(i)}$ to $u_2^{(i)}$; see Fig. 13b. The convex vertex $v_1^{(i)}$ lies on the segment from $r_{1,1}^{(i)}$ to $r_{1,2}^{(i)}$ on $e_1^{(i)}$, otherwise U_2 can be made shorter contradicting that U_2 is the shortest tour obeying the bounding properties. The length of the subpath of $\bar{C}^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$ is upper bounded by the length of the two catheti connecting $u_1^{(i)}$ with $v_1^{(i)}$, one being parallel to $e_1^{(i)}$ and the other orthogonal to $e_1^{(i)}$. Similarly, the length of the subpath of $C^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$ is lower bounded by the length of the line segment from $u_1^{(i)}$ to $v_1^{(i)}$ forming the hypotenuse of a right angled triangle with $u_1^{(i)}$ and $v_1^{(i)}$ as two corners. From elementary geometry it follows that the sum of the lengths of the catheti is bounded by $\sqrt{2}$ times the length of the hypotenuse in a right angled triangle. Thus, the length of the subpath of $\bar{C}^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$ is at most $\sqrt{2}$ times the length of the subpath of $C^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$. The same argument ensures that the length of the subpath of $\bar{C}^{(i)}$ from $v_1^{(i)}$ to $u_2^{(i)}$ is also at most $\sqrt{2}$ times the length of the subpath of $C^{(i)}$ from $v_1^{(i)}$ to $u_2^{(i)}$. We proceed inductively on the parameter $m^{(i)}$, the number of tentacle cuts covered along the path $C^{(i)}$, as follows. Let $p_1^{(i)}$ be the intersection of $e_1^{(i)}$ and $e_2^{(i)}$ (or the directed lines along $e_1^{(i)}$ and $e_2^{(i)}$). The point $p_1^{(i)}$ must exist, otherwise $C^{(i)}$ is not a convex chain but contains reflex vertices. Let $\alpha_1^{(i)}$ be the angle of the wedge $\mathbf{P} \setminus \left(\mathbf{L}(e_1^{(i)}) \cap \mathbf{L}(e_2^{(i)})\right)$ at $p_1^{(i)}$; see Fig. 13c–e. We let $\bar{C}^{(i)} \overset{\text{def}}{=} \bar{C}^{(i)}\big(u_1^{(i)}, u_{m^{(i)}}^{(i)}\big) = \bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big) \cup \bar{C}^{(i)}\big(v_2^{(i)}, u_{m^{(i)}}^{(i)}\big)$ be the path from $u_1^{(i)}$ to $u_2^{(i)}$, where $\bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big)$ is a path from $u_1^{(i)}$ to $v_2^{(i)}$ that covers $e_1^{(i)}$, contains $v_1^{(i)}$, and is at most $\sqrt{2}$ times longer than the path from $u_1^{(i)}$ to $v_2^{(i)}$ along $C^{(i)}$. Since $C^{(i)}$ from $v_2^{(i)}$ to $u_2^{(i)}$ covers fewer than $m^{(i)}$ tentacle cuts, we can assume that $\bar{C}^{(i)}\big(v_2^{(i)}, u_{m^{(i)}}^{(i)}\big)$ from $v_2^{(i)}$ to $u_2^{(i)}$ constructed inductively covers the remaining tentacle cuts, contains the vertices $v_2^{(i)}, \ldots, v_{m^{(i)}}^{(i)}$, and is at most $\sqrt{2}$ times longer than the path from $v_2^{(i)}$ to $u_2^{(i)}$ along $C^{(i)}$. It remains to describe the construction of $\bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big)$ from $u_1^{(i)}$ to $v_2^{(i)}$, that covers one tentacle cut $e_1^{(i)}$, passes one vertex $v_1^{(i)}$, and is at most $\sqrt{2}$ times longer than the path from $u_1^{(i)}$ to $v_2^{(i)}$ along $C^{(i)}$.

If $\alpha_1^{(i)}$ is obtuse or a right angle, let $SP(e_1^{(i)}, e_2^{(i)})$ be the shortest path between $e_1^{(i)}$ end $e_2^{(i)}$. The path either degenerates into the single point $p_1^{(i)}$, if $e_1^{(i)}$ end $e_2^{(i)}$ intersect, or it is a path connecting a point $r_{1,2}^{(i)}$ on $e_1^{(i)}$ to a point $r_{2,1}^{(i)}$ on $e_2^{(i)}$. In this case, $\bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big)$ is the shortest path from $u_1^{(i)}$ to $e_1^{(i)}$, the path from $r_{1,1}^{(i)}$ to $r_{1,2}^{(i)}$ (or $p_1^{(i)}$) on $e_1^{(i)}$, followed by the shortest path from $e_1^{(i)}$ to $e_2^{(i)}$, the path from $r_{2,1}^{(i)}$ (or $p_1^{(i)}$) to $v_2^{(i)}$; see Fig. 13c, d. As previously noticed, the length of the subpath of $\bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big)$ from $u_1^{(i)}$ to $v_1^{(i)}$ is upper bounded by the length of the two catheti connecting $u_1^{(i)}$ with $v_1^{(i)}$, one being parallel to $e_1^{(i)}$ and the other orthogonal to $e_1^{(i)}$. Similarly, the length of the subpath of $C^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$ is lower bounded by the length of the line segment from $u_1^{(i)}$ to $v_1^{(i)}$ forming the hypotenuse of a right angled triangle with $u_1^{(i)}$ and $v_1^{(i)}$ as two corners. From elementary geometry it follows that the sum of the lengths of the catheti is bounded by $\sqrt{2}$ times the length of the hypotenuse in a right angled triangle. Thus, the length of the subpath of $\bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big)$ from $u_1^{(i)}$ to $v_1^{(i)}$ is at most $\sqrt{2}$ times the length of the subpath of $C^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$. The same argument ensures that the length of the subpath of $\bar{C}^{(i)}\big(u_1^{(i)}, v_2^{(i)}\big)$ from

$v_1^{(i)}$ to $v_2^{(i)}$ is also at most $\sqrt{2}$ times the length of the subpath of $C^{(i)}$ from $v_1^{(i)}$ to $v_2^{(i)}$, since the angle $\alpha_1^{(i)}$ is obtuse or a right angle.

If $\alpha_1^{(i)}$ is acute, consider the segment $s \stackrel{\text{def}}{=} [q_1, p_1^{(i)}]$ allowed to pass through the boundary of $\mathbf{P}$, if $e_1^{(i)}$ and $e_2^{(i)}$ do not intersect. Let l be the line orthogonal to s passing through $p_1^{(i)}$. We slide l along the segment s from $p_1^{(i)}$ towards q_1 until l touches a point of $C^{(i)}$ between $v_1^{(i)}$ and $v_2^{(i)}$ inclusive, giving the line l^*. In fact, this intersection point is one of $v_1^{(i)}$ or $v_2^{(i)}$. Let $r_{1,2}^{(i)}$ be the intersection point of $SP(e_1^{(i)}, l^*)$ on $e_1^{(i)}$ and $r_{1,*}^{(i)}$ the intersection point of $SP(e_1^{(i)}, l^*)$ on l^*. Similarly, let $r_{2,1}^{(i)}$ be the intersection point of $SP(e_2^{(i)}, l^*)$ on $e_2^{(i)}$ and $r_{2,*}^{(i)}$ the intersection point of $SP(e_2^{(i)}, l^*)$ on l^*. As before, $r_{1,1}^{(i)}$ is the intersection point of $SP(u_1^{(i)}, e_1^{(i)})$ with $e_1^{(i)}$. Now, $\bar{C}^{(i)}(u_1^{(i)}, v_2^{(i)})$ is the shortest path from $u_1^{(i)}$ to $e_1^{(i)}$, the path from $r_{1,1}^{(i)}$ to $r_{1,2}^{(i)}$, followed by the shortest path from $e_1^{(i)}$ to l^* at $r_{1,*}^{(i)}$, the path from $r_{1,*}^{(i)}$ to $r_{2,*}^{(i)}$ on l^*, the shortest path from l^* to $e_2^{(i)}$ ending at $r_{2,1}^{(i)}$, followed by the path from $r_{2,1}^{(i)}$ to $v_2^{(i)}$; see Fig. 13e. Since $\alpha_1^{(i)}$ is acute, the interior angles between $e_1^{(i)}$ and l^* and between $e_2^{(i)}$ and l^* are both obtuse or right angles. Hence, the lengths of the subpaths of $\bar{C}^{(i)}$ from $u_1^{(i)}$ to $v_1^{(i)}$ and from $v_1^{(i)}$ to $v_2^{(i)}$ are upper bounded by the lengths of the corresponding catheti of the right angled triangles, one connecting $u_1^{(i)}$ and $v_1^{(i)}$, with one catheter being parallel to $e_1^{(i)}$, and one connecting $v_1^{(i)}$ to $v_2^{(i)}$, with one catheter parallel to l^*. As before, the length of this subpath of $\bar{C}^{(i)}(u_1^{(i)}, v_2^{(i)})$ is at most $\sqrt{2}$ times the length of the subpath of $C^{(i)}$ from $u_1^{(i)}$ to $v_2^{(i)}$.

Thus, we have proved Inequality (11) for both cases.

W_1' has no reflex chain or one reflex chain with at most one end point in $\mathbf{P} \setminus \mathbf{L}(e_1)$ If there is a point of T_2 at distance at most R from q_1, then there exists a tour having the bounding properties as defined above and we can bound the length of W_1' in exactly the same way as in the previous case. Hence, from now on, we assume that all points of T_2 have distance greater than R from q_1. Since each tentacle of $\mathcal{J}_{e_1}^2$ has length at most R, every point of the tour W_1' has distance at most R to q_1 and each tentacle is perpendicular to the tentacle cut which in turn intersects T_2, except where W_1' has a reflex chain. Let $\mathbf{D}_1$ be the possible set of points that can be points of W_1', given T_2 but for any possible set of tentacle cuts. Thus, W_1' is contained in $\mathbf{D}_1$. The perimeter of $\mathbf{D}_1$ consists of connected circular arcs with possible line segments in between. If W_1' has a reflex chain, the perimeter alternates between circular arcs centered at q_1 and straight line segments connecting those reflex boundary vertices where W_1' also has reflex vertices. Any other polygon boundary part intersecting the interior of the region is disregarded as we want to obtain the maximum possible perimeter length in relation to R, given that W_1' has at most one reflex chain. $\mathbf{D}_1$ is the region painted light green in Fig. 14a. Since W_1' passes q_1 and the remaining perimeter of $\mathbf{D}_1$ is convex (except where W_1' follows the reflex chain), it follows that $\|W_1'\| \leq \sup \|\partial \mathbf{D}_1\|$ over the set of all possible tentacle cuts. We bound the length of the perimeter of $\mathbf{D}_1$ as follows. Consider the smallest angle cone $\angle T_2$ with apex q_1 that contains T_2. Let t_1 and t_2 be two points of T_2 that touch the sides of this cone; see Fig. 14a. We assume a coordinate system where q_1 is at the origin, t_1 is at distance $R + y_1$, with $y_1 > 0$, directly above q_1 and t_2 lies in the first quadrant in this coordinate system at distance $R + y_2$, with $y_2 > 0$, from q_1. We can also assume that any reflex chain of W'_1 lies in the first quadrant in the coordinate system. If the reflex chain has points in the fourth quadrant, the length of $\mathbf{D}_1$ can never attain its maximum possible ratio to R since all tentacle cuts then lie in quadrants three and four and have positive slope. Thus, all tentacles of $\mathcal{J}^2_{e_1}$ lie in the fourth quadrant and $\|W'_1\| \leq \sup \|\partial \mathbf{D}_1\| \leq (\pi/2+2)R$ in this case, which is not maximal as we shall see below.

Denote the part of $\mathbf{D}_1$ in the fourth quadrant (to the left of $\angle T_2$) by $\mathbf{D}_1^{\mathsf{L}}$. We can bound the length of the two circular arcs that form part of the perimeter of $\mathbf{D}_1^{\mathsf{L}}$ by

$$\|\partial \mathbf{D}_1^{\mathsf{L}}\| = \pi R/2 - \alpha_1 R/2 + \alpha_1(R + y_1) = (\pi + \alpha_1)R/2 + \alpha_1 \cdot y_1, \qquad (12)$$

where $\alpha_1 = 2 \arcsin \big(R \big/ (2R + 2y_1)\big)$; see Fig. 14c. If we denote the part of $\mathbf{D}_1$ inside the cone $\angle T_2$ by $\mathbf{D}_1^{\mathsf{C}}$, the light green region in the cone $\angle T_2$ in Fig. 14a, we can bound the length of the circular arc perimeter of $\mathbf{D}_1^{\mathsf{C}}$ by $\|\partial \mathbf{D}_1^{\mathsf{C}}\| \leq R \cdot \beta$, where $\|[t_1, t_2]\|^2 = (R+y_1)^2 + (R+y_2)^2 - 2(R+y_1)(R+y_2)\cos\beta$ and β is the angle of the cone $\angle T_2$, by the cosine theorem; see Fig. 14a. Again, we can assume that any reflex chain of W'_1 does not intersect the cone $\angle T_2$ and therefore also not inside $\mathbf{D}_1^{\mathsf{C}}$, otherwise the length of the perimeter of $\mathbf{D}_1$ can never attain its maximum possible ratio to R since this would give a coarse bound of $\|W'_1\| \leq \sup \|\partial \mathbf{D}_1\| \leq (\pi + 2)R$ in this case, which again is not maximal.

If the remaining part of $\mathbf{D}_1$ intersects the single reflex chain of W'_1, the length of the sequence of alternating circle arcs and line segments is bounded be the length of two line segments along two radii from q_1 of length R and t_2 of length $R + y_2$, from their intersection point to their respective ends at p_1 and p_2; see Fig. 14a, b. We denote these segments by s_1 and s_2 as in Fig. 14b. Denote the angle between s_1 and s_2 by γ. Let p_3 be the intersection point of the two circular arcs with center at q_1 and length R and center at t_2 and length $R + y_2$ having positive x-coordinate in our coordinate system and let l_1 and l_2 be the subarcs of these two circular arcs from p_3 to p_1 and p_2, respectively; see Fig. 14b. Let the radius $[q_1, p_1]$ have angle θ_1 to the radius $[q_1, p_3]$ and let the radius $[t_2, p_2]$ have angle θ_2 to the radius $[t_2, p_3]$. Thus, $\gamma = \pi/2 - \alpha_2/2 + \theta_1 + \theta_2$. Let s'_1 be the subsegment of s_1 outside the triangle $\triangle q_1, t_2, p_3$ and let s'_2 be the subsegment of s_2 outside $\triangle q_1, t_2, p_3$. The angle of the triangle at p_3 is $\pi/2 - \alpha_2/2 \leq \gamma$ so $\|s_1\| - \|s'_1\| + \|s_2\| - \|s'_2\| \leq \|l_1\| + \|s'_1\| + \|l_2\| + \|s'_2\|$ giving $\|s_1\| + \|s_2\| \leq \|l_1\| + 2\|s'_1\| + \|l_2\| + 2\|s'_2\|$. Let the two circle arcs symmetric to $\mathbf{D}_1^{\mathsf{L}}$ in the first quadrant, to the right of $\angle T_2$ bound the region $\mathbf{D}_1^{\mathsf{R}}$ and we have

$$\|\partial \mathbf{D}_1^{\mathsf{R}}\| = (\pi + \alpha_2)R/2 + \alpha_2 \cdot y_2, \qquad (13)$$

where $\alpha_2 = 2\arcsin\left(R\big/(2R+2y_2)\right)$; see Fig. 14a, c. Define $\mathbf{D}_1^{\mathsf{X}}$ to be the union of $\mathbf{D}_1^{\mathsf{L}}$, $\mathbf{D}_1^{\mathsf{C}}$, and $\mathbf{D}_1^{\mathsf{R}}$, i.e., the locus $\mathbf{D}_1$ when W_1' does not have a reflex chain in the interior of $\mathbf{D}_1$; see Fig. 14c. Since $\|s_1\| + \|s_2\| \leq \|l_1\| + 2\|s_1'\| + \|l_2\| + 2\|s_2'\|$, we have that

$$\|\partial \mathbf{D}_1\| \leq \|\partial \mathbf{D}_1^{\mathsf{X}}\| + 2\|s_1'\| + 2\|s_2'\| \tag{14}$$

and we bound the length of these three parts separately.

To maximize the length of the perimeter of $\mathbf{D}_1^{\mathsf{X}}$ we observe that the angle β is as large as possible when the segment $[t_1, t_2]$ is as large as possible and the values y_1 and y_2 are as small as possible. The length of $[t_1, t_2]$ is bounded above by R and y_1 and y_2 are minimal if the segment $[t_1, t_2]$ is almost tangent to the circle of radius R centered at q_1. With this setup we can write the length of the perimeter of $\mathbf{D}_1^{\mathsf{X}}$ as a function of y_1, where $0 \leq y_1 \leq (\sqrt{2}-1)R$, and using standard variational calculus we obtain that

$$\sup_{0\leq y_1\leq(\sqrt{2}-1)R} \|\partial \mathbf{D}_1^{\mathsf{X}}(y_1)\| < \left(\pi + 2\sqrt{5}\arcsin\frac{1}{\sqrt{5}}\right)R, \tag{15}$$

which occurs when $y_1 = y_2 = (\sqrt{5}/2 - 1)R$.

To maximize the lengths of s_1' and s_2', we realize that they are the longest when the radius $[q_1, p_1]$ intersects the segment $[t_2, p_3]$ perpendicularly and when the radius $[t_2, p_2]$ intersects the segment $[q_1, p_3]$ perpendicularly, i.e., when $\theta_1 = \theta_2 = \alpha_2/2$. Hence, the two segments are the longest when α_2 is as large as possible, which occurs when the triangle $\triangle q_1, t_2, p_3$ is equilateral, whereby

$$\sup_{0\leq\alpha_2\leq\pi/6} \|s_1'\| = \sup_{0\leq\alpha_2\leq\pi/6} \|s_2'\| = (1-\sqrt{3}/2)R. \tag{16}$$

This gives us a final bound of

$$\begin{aligned}\|W_1'\| &< \left(\pi + 4 - 2\sqrt{3} + 2\sqrt{5}\arcsin 1/\sqrt{5}\right)R \\ &= \left(\pi/2 + 2 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\right)\|(T_1, T_2)\|_{\mathsf{max}}.\end{aligned} \tag{17}$$

We have from the two cases above, Inequalities (11) and (17), that

$$\|W_1'\| \leq \max\begin{cases} 2\sqrt{2}R + \sqrt{2}\|T_2\| = 2\sqrt{2}\|(T_1, T_2)\|_{\mathsf{max}} \approx 2.8284\|(T_1, T_2)\|_{\mathsf{max}}, \\ \left(\pi/2 + 2 - \sqrt{3} + \sqrt{5}\arcsin 1/\sqrt{5}\right)\|(T_1, T_2)\|_{\mathsf{max}} \approx 2.8754\|(T_1, T_2)\|_{\mathsf{max}}. \end{cases} \tag{18}$$

The larger of the two values is our bound for the length of W_1'. □

We note that we have made a slight overestimation when bounding the length of the perimeter of $\mathbf{D}_1$ as we are maximizing the length of the perimeter of $\mathbf{D}_1^{\mathsf{X}}$ and the lengths

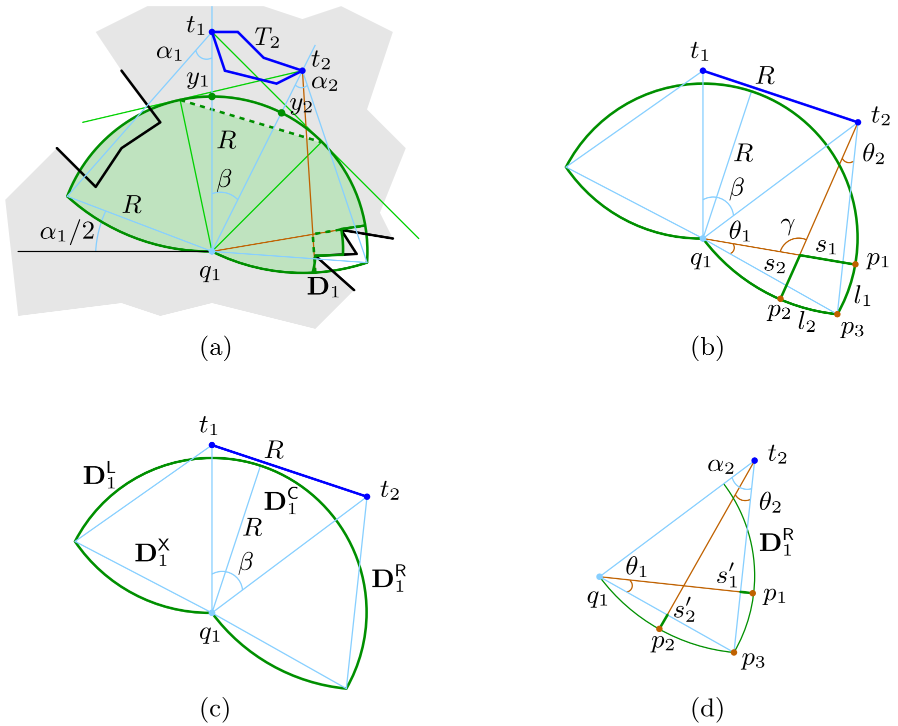

Fig. 14 Illustrating the proof of Lemma 5.3. Bounding the length of W_1' when it has at most one reflex chain with at most one end point in $\mathbf{P} \setminus \mathbf{L}(e_1)$ and the distance from T_2 to q_1 is greater than R. **a** $\mathbf{D}_1$ in green is the locus of possible points for W_1'. **b** The worst case placement of T_2 relative to $\mathbf{D}_1$ maximizing its boundary length. **c** Partitioning $\mathbf{D}_1$ into $\mathbf{D}_1^{\mathsf{L}}$, $\mathbf{D}_1^{\mathsf{C}}$, and $\mathbf{D}_1^{\mathsf{R}}$. **d** Detail of $\mathbf{D}_1^{\mathsf{R}}$ indicating the subsegments s_1' and s_2' outside the triangle $\triangle q_1, t_2, p_3$ (Color figure online)

of s_1' and s_2' separately; see Inequalities (15) and (16). Some numerical experimentation indicates that our overestimation only affects the approximation bound in the second decimal, leading us to not pursue any improvement further.

5.3 Complexity Analysis of the Algorithm

The complexity analysis of the algorithm is straightforward. The for-loop in Step 5 considers $O(n^2)$ pairs of extensions. Computing the tentacles and bases in Step 5.1 takes $O(n^6)$ time as established in Section 5.1. The work in Step 5.2 takes $O(n \log n)$ time since it is dominated by sorting the tentacles by length. Step 5.3 requires linear time using the algorithm by Toussaint [35] and the test in Step 5.4 takes constant time. Hence, the total time complexity for the algorithm is $O(n^8)$. The running time analysis together with Lemma 5.1 give us our main result.

Theorem 1 *The Two-Watchman-Route algorithm computes a* 5.969*-approximation of the minmax two-watchman route and a* 11.939*-approximation of the minsum two-watchman route in* $O(n^8)$ *time.*

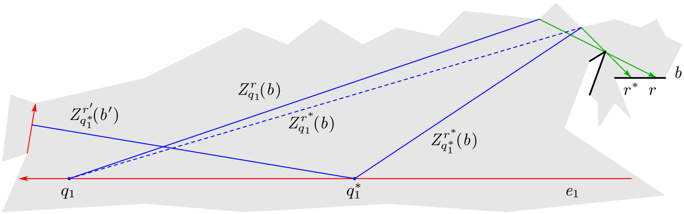

Fig. 15 Illustrating the trade-off between the optimal base placement and the approximate base placement

6 A Trade-off Between Computation Time and Accuracy

We can trade the approximation bound for computational efficiency by realizing that relaxing the computation as described in Sect. 5.1 still provides base positions q_1 and q_2 sufficiently close to their optimal positions on e_1 and e_2. To do this, we reduce the computation in Sect. 5.1 to only use the first two cases, taking $O(n^2)$ time rather than $O(n^6)$ time. Let q_1 be the base placement on e_1 using the reduced computation and let q_1^* be its optimal placement on e_1. Let $Z^r_{q_1}(b)$ be a longest tentacle in the jellyfish $\mathcal{J}^{\min}_{q_1}$. Now, consider the tentacle $Z^{r^*}_{q_1^*}(b)$, taking into account that r moves to r^* on b when the head moves from q_1 to q_1^* on e_1. There must exist a boundary edge b' such that the tentacle $Z^{r'}_{q_1^*}(b')$ intersects the line through q_1 orthogonal to e_1, otherwise q_1 would be closer to q_1^*; see Fig. 15. Hence, $\|q_1, q_1^*\| \le \|Z^{r'}_{q_1^*}(b')\|$ and we have

$$\|Z^r_{q_1}(b)\| \le \|Z^{r^*}_{q_1}(b)\| \le \|Z^{r^*}_{q_1^*}(b)\| + \|q_1, q_1^*\| \le \|Z^{r^*}_{q_1^*}(b)\| + \|Z^{r'}_{q_1^*}(b')\| \le 2R, \tag{19}$$

since R is the length of the longest tentacle in $\mathcal{J}^{\min}_{e_1}$. We can argue similarly for $\mathcal{J}^{\min}_{q_2}$ giving us the jellyfish pair $\mathcal{J}^{\min}_{q_1,q_2}$ for which we can compute the reduced jellyfish pair and then the relative convex hulls.

Using $R' = 2R$ in the proofs of Lemmas 5.1, 5.2 and 5.3 instead of R, all the arguments hold, giving us a tour with an approximation ratio at worst twice that of Theorem 1. We state this as a corollary.

Corollary 2 *The simplified Two-Watchman-Route algorithm computes a* 11.939-*approximation of the minmax two-watchman route and a* 23.879-*approximation of the minsum two-watchman route in* $O(n^4)$ *time.*

7 Computing the Fixed Two-Watchman Route

Since the heads q_1 and q_2 are given as input in this simpler case of the problem, it suffices to compute the Jellyfish pairs $\mathcal{J}_{q_1,q_2}$ with q_1 and q_2 as heads which takes $O(n^2)$ time as explained in Sect. 5.1 and since W_1 intersects T_1 and W_2 intersects T_2, the second case in the proof of Lemma 5.3 cannot occur, whereby the approximation bound for W_1 becomes

$$\|W_1\| \leq \|W_1'\| + 2R + \|C_1\| \leq 2\sqrt{2}R + \sqrt{2}\|T_2\| + 2R + 4\pi R/3$$
$$= (2\sqrt{2} + 1 + 2\pi/3)\|(T_1, T_2)\|_{\mathsf{max}}. \quad (20)$$

However, we note that this tour does not necessarily pass through q_1, unless q_1 lies on the polygon boundary, hence we have to account for an extra reflex chain of length at most $2R$ to guarantee this. Making the same analysis for W_2, the approximation factor becomes $2\sqrt{2} + 2 + 2\pi/3 \approx 6.922$ in this case. The time complexity is dominated by computing the jellyfish, giving us the following theorem.

Theorem 3 *The algorithm computes a* 6.922*-approximation of the fixed minmax two-watchman route and a* 13.845*-approximation of the fixed minsum two-watchman route in* $O(n^2)$ *time given two starting points for the two tours. If both starting points lie on the boundary, the approximation factors are* 5.922 *and* 11.845*, respectively.*

8 Conclusions

We have shown a polynomial time algorithms for computing constant factor approximations for the minmax and minsum two-watchman route in a simple polygon.

Our algorithms rely heavily on the fact that for two tours it is sufficient to guarantee that the boundary is seen to ensure that the complete polygon is seen. This does not hold for three or more tours. Thus, our method for the two-watchman tours does not generalize to the problem for three or more watchmen. Solving these problems within a constant factor bound remains elusive.

Establishing the complexity for the minsum two-watchman route is still open although our algorithm provides a polynomial time 11.939-approximation.

The authors would like to thank Doc. Åse Jevinger and Prof. Paweł Żyliński for fruitful discussions that have improved the text immensely.

Funding Open access funding provided by Malmö University.

Declarations

Conflict of interest We, the authors, declare that we have no Conflict of interest as defined by Springer, or other interests that might be perceived to influence the results and/or discussion reported in this manuscript. The results/data/figures in this manuscript have not been published elsewhere, nor are they under consideration by another publisher. All of the material is owned by the authors and no permissions of use are required. We furthermore declare that we have no conflict of interest concerning any of the material in this manuscript, that we have developed the algorithms and written the manuscript together, that author B.J.N. did the final editing of the manuscript, and that we submit this manuscript in accordance with the Nature Portfolio journal policies.

A Proof of Lemma 4.2

Lemma *Let q move a distance δ to q' on a line segment s and let r move a distance ϵ to r', where both r and r' lie in the open interval $]v, v'[$ of a boundary edge $b = [v, v']$, in such a way that the first segment of the tentacles from q and q' intersect the same reflex vertex, if the tentacle consists of multiple segments, and $c(Z_q^r)$ and $c(Z_{q'}^{r'})$ have the same hiding vertex, then*

$$\|Z_{q'}^{r'}\| = \|Z_q^r\| + \mathcal{F}(\delta, \epsilon),$$

such that

$$\mathcal{F}(\delta, \epsilon) = -F_0 + \sqrt{F_0^2 + F_1\delta + F_2\delta^2} - F_3 + \frac{F_3 + F_4\epsilon + F_5\delta + F_6\epsilon\delta + F_7\epsilon^2 + F_8\epsilon^2\delta}{\sqrt{1 + F_9\epsilon + F_{10}\epsilon^2 + F_{11}\epsilon^3 + F_{12}\epsilon^4}}$$
$$- F_{13} + \sqrt{\frac{F_{13}^2 + F_{14}\epsilon + F_{15}\delta + F_{16}\epsilon\delta + F_{17}\epsilon^2 + F_{18}\delta^2 + F_{19}\epsilon^2\delta + F_{20}\epsilon\delta^2 + F_{21}\epsilon^2\delta^2}{1 + F_{22}\epsilon + F_{23}\epsilon^2}},$$

where $F_0, \ldots, F_{23}$ are constants.

Proof We begin with the assumption that Z_q^r consists of at least two segments whose end points touch reflex boundary vertices of the polygon. Since this is the case, we can separate the motion from r to r' on b and the motion from q to q' on s and handle those two cases separately. We first look at the motion from r to r' on b and since a small change of q on s does not affect the motion on b.

Given a tentacle Z_q^r, we emulate a point sliding from r to r' along b. If the tentacle tip intersects the boundary of the polygon at a point p, where the tentacle sees r, then we let the segment $[r, r'']$ be parallel to $[p, p']$ and intersecting the line collinear to $[u_h, r']$ at r", where u_h is the hiding vertex; see Fig. 16a. The lines through $[p, p']$ and $[v, v']$ intersect at a point t (unless they are parallel, in which case we assume t to be a point at infinity). Since the triangles $\triangle r', r, r''$ and $\triangle r', t, p'$ are similar, we have

$$\frac{\epsilon}{\|r, r''\|} = \frac{\|r, t\| - \epsilon}{\|p, t\| + \epsilon'}$$

and since the triangles $\triangle r, u_h, r''$ and $\triangle p, u_h, p'$ are also similar, we have

$$\frac{\|r, r''\|}{\|u_h, r\|} = \frac{\epsilon'}{\|u_h, p\|},$$

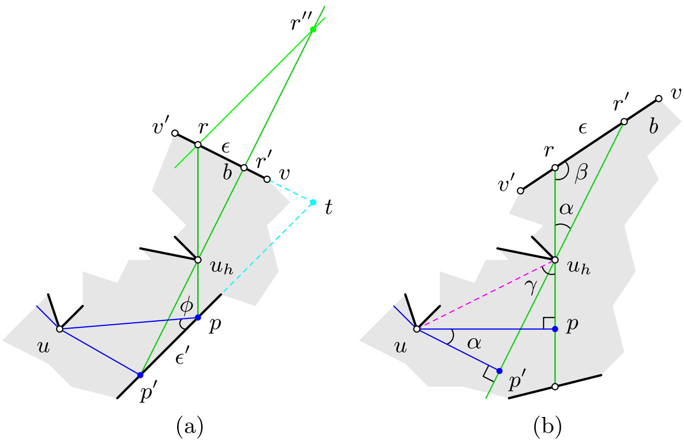

Fig. 16 Showing the motions of a tentacle along the interior of b

giving us that

$$\epsilon' = \frac{\|p,t\| \cdot \|u_h, p\| \cdot \epsilon}{\|r,t\| \cdot \|u_h, r\| - (\|u_h, p\| + \|u_h, r\|) \cdot \epsilon} = \frac{a \cdot \epsilon}{1 - c \cdot \epsilon}, \tag{21}$$

for constants a and c depending on the points p, r, u_h, and t.

Thus, the length of the tentacle as the point moves locally from r to r' along b is

$$\begin{aligned} \|Z_q^{r'}\| &= \|Z_q^r\| - \|u,p\| + \|u,p'\| \\ &= \|Z_q^r\| - \|u,p\| + \sqrt{\|u,p\|^2 + \epsilon'^2 - 2\|u,p\|\epsilon' \cos\phi} \\ &= \|Z_q^r\| - \|u,p\| + \frac{\sqrt{\|u,p\|^2(1-c\epsilon)^2 + a^2\epsilon^2 - 2\|u,p\|a\epsilon\cos\phi(1-c\epsilon)}}{1-c\epsilon} \\ &= \|Z_q^r\| - A_0 + \frac{\sqrt{A_0^2 + A_1\epsilon + A_2\epsilon^2}}{1 + A_3\epsilon}, \end{aligned} \tag{22}$$

for constants $A_0, \ldots, A_3$ that only depend on the points u, p, r, u_h, t, and the angle ϕ; see Fig. 16a.

If the tentacle tip does not touch the boundary of the polygon; see Fig. 16b, then we realize that as the point moves locally from r to r', the length of the tentacle changes

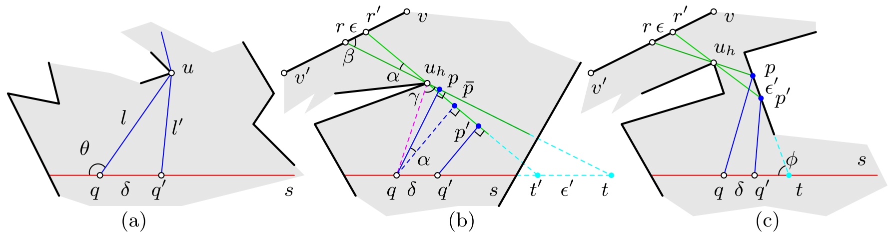

Fig. 17 Showing the motion of a tentacle along line segment e

as

$$\begin{aligned}
\|Z_q^{r'}\| &= \|Z_q^r\| - \|u,p\| + \|u,p'\| \\
&= \|Z_q^r\| - \|u,p\| + \|u,u_h\|\sin(\gamma-\alpha) \\
&= \|Z_q^r\| - \|u,p\| + \|u,u_h\|(\sin\gamma\cos\alpha - \cos\gamma\sin\alpha) \\
&= \|Z_q^r\| - \|u,p\| + \|u,u_h\|\sin\gamma\frac{\|u_h,r\| - \epsilon\cos\beta}{\sqrt{\|u_h,r\|^2 - 2\epsilon\|u_h,r\|\cos\beta + \epsilon^2}} \\
&\qquad - \|u,u_h\|\cos\gamma\frac{\epsilon\sin\beta}{\sqrt{\|u_h,r\|^2 - 2\epsilon\|\,r\|\cos\beta + \epsilon^2}} \\
&= \|Z_q^r\| - \|u,p\| + \frac{\|u,u_h\|\|u_h,r\|\sin\gamma - \epsilon\|u,u_h\|(\cos\beta\sin\gamma + \cos\gamma\sin\beta)}{\sqrt{\|u_h,r\|^2 - 2\epsilon\|u_h,r\|\cos\beta + \epsilon^2}} \\
&= \|Z_q^r\| - \|u,p\| + \frac{\|u,p\|\|u_h,r\| - \epsilon\|u,u_h\|\sin(\beta+\gamma)}{\sqrt{\|u_h,r\|^2 - 2\epsilon\|u_h,r\|\cos\beta + \epsilon^2}} \\
&= \|Z_q^r\| - B_0 + \frac{B_0 + B_1\epsilon}{\sqrt{1 + B_2\epsilon + B_3\epsilon^2}} \qquad (23)
\end{aligned}$$

for constants $B_0, \ldots, B_3$ that only depend on the points u, p, r, u_h and the angles β and γ; see Fig. 16b.

Let us now consider the change that happens as the head moves a distance δ from q to q' on s, still under the assumption that Z_q^r and $Z_{q'}^{r'}$ consists of at least two segments where the end points different from q (q') and r (r') touch reflex vertices of the polygon thus making turns at the boundary. The length of the tentacle changes as

$$\begin{aligned}
\|Z_{q'}^r\| &= \|Z_q^r\| - \|q,u\| + \|q',u\| \\
&= \|Z_q^r\| - \|q,u\| + \sqrt{\|q,u\|^2 + \delta^2 - 2\|q,u\|\delta\cos(\pi-\theta)} \\
&= \|Z_q^r\| - C_0 + \sqrt{C_0^2 + C_1\delta + C_2\delta^2}, \qquad (24)
\end{aligned}$$

for constants C_0, C_1 and C_2 that only depend on the points u, q, and the angle θ; see Fig. 17a.

If the Z_q^r and $Z_{q'}^{r'}$ each consists of a single segment, then either the tip touches the boundary or it does not. If the tentacle Z_q^r is a single segment that intersects

$c(Z^r_q)$ without the tip touching the boundary, then $\|q', p'\|$ depends on both ϵ and δ. The length $\|q', p'\|$ is established from $\|q, p\|$ by first computing $\|q, \bar{p}\|$, a case we have already solved; see Equality 23 and Fig. 16b; and then computing $\|q', p'\|$ from $\|q, \bar{p}\|$. $\|q', p'\| = \|q, \bar{p}\|(1 - \delta/\|q, t'\|)$ since $\|q', p'\|/(\|q, t'\| - \delta) = \|q, \bar{p}\|/\|q, t'\|$ by similarity. Hence, the length of the tentacle changes as

$$
\begin{aligned}
\|Z^{r'}_{q'}\| &= \|Z^r_q\| - \|q, p\| + \|q', p'\| \\
&= \|Z^r_q\| - \|q, p\| + \left(1 - \frac{\delta}{\|q, t'\|}\right) \|q, \bar{p}\| \\
&= \|Z^r_q\| - \|q, p\| + \left(1 - \frac{\delta}{\|q, t\| - \epsilon'}\right) \|q, u_h\| \sin(\gamma - \alpha) \\
&= \|Z^r_q\| - \|q, p\| + \left(1 - \frac{\delta}{\|q, t\| - \frac{a\epsilon}{1-c\epsilon}}\right) \|q, u_h\| (\sin\gamma \cos\alpha - \sin\gamma \cos\alpha) \\
&= \|Z^r_q\| - \|q, p\| + \\
&\qquad + \left(1 - \frac{\delta}{\|q, t\| - \frac{a\epsilon}{1-c\epsilon}}\right) \|q, u_h\| \sin\gamma \frac{\|u_h, r\| - \epsilon \cos\beta}{\sqrt{\|u_h, r\|^2 - 2\epsilon \|u_h, r\| \cos\beta + \epsilon^2}} \\
&\qquad - \left(1 - \frac{\delta}{\|q, t\| - \frac{a\epsilon}{1-c\epsilon}}\right) \|q, u_h\| \cos\gamma \frac{\epsilon \sin\beta}{\sqrt{\|u_h, r\|^2 - 2\epsilon \|u_h, r\| \cos\beta + \epsilon^2}} \\
&= \|Z^r_q\| - \|q, p\| + \\
&\qquad + \left(1 - \frac{\delta(1 - c\epsilon)}{\|q, t\| - (c\|q, t\| + a)\epsilon}\right) \frac{\|q, p\| \|u_h, r\| - \epsilon \|u, u_h\| \sin(\beta + \gamma)}{\sqrt{\|u_h, r\|^2 - 2\epsilon \|u_h, r\| \cos\beta + \epsilon^2}} \\
&= \|Z^r_q\| - D_0 + \frac{D_0 + D_1\epsilon + D_2\delta + D_3\epsilon\delta + D_4\epsilon^2 + D_5\epsilon^2\delta}{\sqrt{1 + D_6\epsilon + D_7\epsilon^2 + D_8\epsilon^3 + D_9\epsilon^4}},
\end{aligned}
\tag{25}
$$

for constants $D_0, \ldots, D_9$ that only depend on the points p, q, r, and u_h, together with the angles β and γ; see Fig. 17b.

If the tentacle Z^r_q is a single segment that intersects $c(Z^r_q)$ with the tip touching the boundary, then the length $\|q', p'\|$ is established from $\|q, p\|$ by the cosine theorem. Hence, the length of the tentacle changes as

$$
\begin{aligned}
\|Z^{r'}_{q'}\| &= \|Z^r_q\| - \|q, p\| + \|q', p'\| \\
&= \|Z^r_q\| - \|q, p\| + \sqrt{(\|p, t\| - \epsilon')^2 + (\|q, t\| - \delta)^2 - 2(\|p, t\| - \epsilon')(\|q, t\| - \delta)\cos\phi} \\
&= \|Z^r_q\| - \|q, p\| + \\
&\qquad + \sqrt{\left(\|p, t\| - \frac{a\epsilon}{1 - c\epsilon}\right)^2 + (\|q, t\| - \delta)^2 - 2\left(\|p, t\| - \frac{a\epsilon}{1 - c\epsilon}\right)(\|q, t\| - \delta)\cos\phi} \\
&= \|Z^r_q\| - E_0 + \sqrt{\frac{E_0^2 + E_1\epsilon + E_2\delta + E_3\epsilon\delta + E_4\epsilon^2 + E_5\delta^2 + E_6\epsilon^2\delta + E_7\epsilon\delta^2 + E_8\epsilon^2\delta^2}{1 + E_9\epsilon + E_{10}\epsilon^2}},
\end{aligned}
\tag{26}
$$

for constants $E_0, \ldots, E_{10}$ that only depend on the points p, q, r, u_h and the angle ϕ; see Fig. 17c.

Combining Equalities (22)–(26), we obtain the equality

$$\|Z_{q'}^{r'}\| = \|Z_q^r\| - F_0 + \sqrt{F_0^2 + F_1\delta + F_2\delta^2} - F_3 + \frac{F_3 + F_4\epsilon + F_5\delta + F_6\epsilon\delta + F_7\epsilon^2 + F_8\epsilon^2\delta}{\sqrt{1 + F_9\epsilon + F_{10}\epsilon^2 + F_{11}\epsilon^3 + F_{12}\epsilon^4}} - F_{13} + \sqrt{\frac{F_{13}^2 + F_{14}\epsilon + F_{15}\delta + F_{16}\epsilon\delta + F_{17}\epsilon^2 + F_{18}\delta^2 + F_{19}\epsilon^2\delta + F_{20}\epsilon\delta^2 + F_{21}\epsilon^2\delta^2}{1 + F_{22}\epsilon + F_{23}\epsilon^2}}, \tag{27}$$

for constants $F_0, \ldots, F_{23}$, as claimed. We note further that since F_0, F_3, and F_{13} represent actual distances, these must be non-negative. □

References

1. Archimedes.: On the Sphere and Cylinder, Book I. The Works of Archimedes. T.L. Heath, editor, Dover, New York, pp. 1–55 (1953)
2. Arkin, E.M., Efrat, A., Knauer, C., Mitchell, J.S.B., Polishchuk, V., Rote, G., Schlipf, L., Talvitie, T.: Shortest path to a segment and quickest visibility queries. J. Comput. Geom. **7**, 77–100 (2016)
3. Belleville, P.: Computing two-covers of simple polygons. Master's thesis, School of Computer Science, McGill University (1991)
4. Belleville, P.: Two-Guarding simple polygons. In: Proceedings of 4th Canadian Conference on Computational Geometry, pp. 103–108 (1992)
5. Brodie, S.E.: Archimedes' axioms for arc-length and area. Math. Mag. **53**(1), 36–39 (1980)
6. Carlsson, S., Jonsson, H., Nilsson, B.J.: Finding the shortest watchman route in a simple polygon. Discrete Comput. Geom. **22**, 377–402 (1999)
7. Chin, W., Ntafos, S.: Optimum watchman routes. Inf. Process. Lett. **28**, 39–44 (1988)
8. Chin, W., Ntafos, S.: Shortest watchman routes in simple polygons. Discrete Comput. Geom. **6**(1), 9–31 (1991)
9. Dror, M., Efrat, A., Lubiw, A., Mitchell, J.S.B.: Touring a sequence of polygons. In: Proceedings of 35th ACM Symposium on Theory of Computing, STOC'03, pp. 473–482. ACM (2003)
10. Dumitrescu, A., Tóth, C.D.: Watchman tours for polygons with holes. Comput. Geom. Theory Appl. **45**(7), 326–333 (2012)
11. ElGindy, H., Avis, D.: A linear algorithm for computing the visibility polygon from a point. J. Algorithms **2**, 186–197 (1981)
12. Goodman, J.E., O'Rourke, J., (eds.): Handbook of Discrete and Computational Geometry. CRC Press (1997)
13. Guibas, L., Hershberger, J., Leven, D., Sharir, M., Tarjan, R.: Linear time algorithms for visibility and shortest path problems inside triangulated simple polygons. Algorithmica **2**, 209–233 (1987)
14. Hammar, M., Nilsson, B.J.: Concerning the time bounds of existing shortest watchman route algorithms. In: Proceedings of 11th International Symposium on Fundamentals in Computation Theory, FCT'1997, volume 1279 of Lecture Notes in Computer Science, pp. 210–221. Springer Verlag (1997)
15. Harel, D., Tarjan, R.E.: Fast algorithms for finding nearest common ancestors. SIAM J. Comput. **13**(2), 338–355 (1984)
16. Heron of Alexandria. Catoptrica (On Reflection). ∼62
17. Hershberger, J., Suri, S.: A pedestrian approach to ray shooting: shoot a ray. Take a walk. J. Algorithms **18**(3), 403–431 (1995)
18. Joe, B., Simpson, R.B.: Correction to Lee's visibility polygon algorithm. BIT **27**, 458–473 (1987)
19. Lee, D.T.: Visibility of a simple polygon. Comput. Vis. Graph. Image Process. **22**, 207–221 (1983)
20. Lee, D.T., Preparata, F.P.: An optimal algorithm for finding the kernel of a polygon. J. ACM **26**, 415–421 (1979)
21. Lee, D.T., Preparata, F.P.: Euclidean shortest paths in the presence of rectilinear barriers. Networks **14**, 393–410 (1984)
22. Lubiw, A., Naredla, A.M.: The visibility center of a simple polygon. In: Proceedings of 29th Annual European Symposium on Algorithms, ESA'2021. Leibniz International Proceedings in Informatics (LIPIcs), Volume 204, pp. 65:1–65:14, Schloss Dagstuhl — Leibniz-Zentrum für Informatik (2021)

23. Mitchell, J.S.B.: Approximating watchman routes. In: Proceedings of 24th ACM-SIAM Symposium on Discrete Algorithms, SODA'13, pp. 844–855 (2013)
24. Mitchell, J.S.B., Wynters, E.L.: Watchman routes for multiple guards. In: Proceedings of 3rd Canadian Conference on Computational Geometry, pp. 126–129 (1991)
25. Nilsson, B.J.: Approximating a shortest watchman route. Fund. Inform. **45**(3), 253–281 (2001)
26. Packer, E.: Computing multiple watchman routes. In: Proceedings of 7th International Workshop, WEA 2008, pp. 114–128. Springer Verlag, Lecture Notes in Computer Science 5038 (2008)
27. Pollack, R., Sharir, M., Rote, G.: Computing the geodesic center of a simple polygon. Discrete Comput. Geom. **4**, 611–626 (1989)
28. Sack, J.R., Urrutia, J., (eds).: Handbook on Computational Geometry. Elsevier Science Publishers (1999)
29. Tan, X.-H.: Fast computation of shortest watchman routes in simple polygons. Inf. Process. Lett. **77**(1), 27–33 (2001)
30. Tan, X.-H.: A linear-time 2-approximation algorithm for the watchman route problem for simple polygons. Theoret. Comput. Sci. **384**(1), 92–102 (2007)
31. Tan, X.-H., Hirata, T.: Constructing Shortest Watchman Routes by Divide and Conquer. In: Proceedings of 4th International Symposium on Algorithms and Computation, pp. 68–77. Springer Verlag, Lecture Notes in Computer Science 762 (1993)
32. Tan, X.-H., Hirata, T., Inagaki, Y.: An incremental algorithm for constructing shortest watchman routes. Int. J. Comput. Geom. Appl. **3**, 351–365 (1993)
33. Tan, X.-H., Hirata, T., Inagaki, Y.: Corrigendum to "An incremental algorithm for constructing shortest watchman routes". Int. J. Comput. Geom. Appl. **9**(3), 319–324 (1999)
34. Tan, X.-H., Jiang, B.: Efficient algorithms for touring a sequence of convex polygons and related problems. In: Proceedings of Theory and Applications of Models of Computation, TAMC 2017, pp. 614–627. Springer Verlag, Lecture Notes in Computer Science 10185 (2017)
35. Toussaint, G.T.: An optimal algorithm for computing the relative convex hull of a set of points in a polygon. In: Signal Processing III: Theories and Applications, European Signal Processing Conference, pp. 853–856. North-Holland (1986)