


The Reverse Makeya Problem

Sang Won Bae¹


Kyonggi University
Suwon, Korea
swbae@kgu.ac.kr

 <https://orcid.org/0000-0002-8802-4247>

Sergio Cabello²

University of Ljubljana and IMFM
Slovenia

sergio.cabello@fmf.uni-lj.si


 <https://orcid.org/0000-0002-3183-4126>

Otfried Cheong³

KAIST

Daejeon, Korea

otfried@kaist.airpost.net

 <https://orcid.org/0000-0003-4467-7075>

Yoonsung Choi

KAIST

Daejeon, Korea

giantsol2@kaist.ac.kr

Fabian Stehn

Bayreuth University

Germany

fabian.stehn@uni-bayreuth.de

Sang Duk Yoon⁴

Postech

Pohang, Korea

sangduk.yoon@gmail.com

Abstract

We prove a generalization of Pál's 1921 conjecture that if a convex shape P can be placed in any orientation inside a convex shape Q in the plane, then P can also be turned continuously through 360° inside Q . We also prove a lower bound of $\Omega(mn^2)$ on the number of combinatorially distinct maximal placements of a convex m -gon P in a convex n -gon Q . This matches the upper bound proven by Agarwal et al.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Makeya problem, convex, isodynamic point, turning

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.6

¹ Supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2015R1D1A1A01057220).

² Supported by Slovenian Research Agency, program P1-0297.

³ Supported by ICT R&D program of MSIP/IITP [R0126-15-1108].

⁴ Supported by NRF grant 2011-0030044 (SRC-GAIA) funded by the government of Korea.



© Sang Won Bae, Sergio Cabello, Otfried Cheong, Yoonsung Choi, Fabian Stehn, and Sang Duk Yoon;

licensed under Creative Commons License CC-BY


34th International Symposium on Computational Geometry (SoCG 2018).

Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 6; pp. 6:1–6:13

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Acknowledgements This research was started during the Korean Workshop on Computational Geometry 2017, organized by Tohoku University in Zao Onsen. We thank all workshop participants for the helpful discussions. We also thank Xavier Goaoc and Helmut Alt for helpful discussions.

1 Introduction

In 1917, Soichi Kakeya posed the following problem: What is the minimum area region in the plane in which a needle of length 1 can be turned through 360° continuously, and return to its initial position [6]? For convex regions, the problem was solved by Pál [8], who showed that the solution is the equilateral triangle of height one, having area $1/\sqrt{3}$. For the general case, Besicovitch gave the surprising answer that one could rotate a needle using an arbitrary small area [2, 3]. Kakeya-type problems have received considerable attention in the literature, as there are strong connections to problems in number theory [4], geometric combinatorics [10], arithmetic combinatorics [7], oscillatory integrals, and the analysis of dispersive and wave equations [9].

If one generalizes the problem for convex regions slightly, and asks for the smallest convex region in which a given convex shape P can be turned through 360° , the problem seems to be still wide open: the answer is not even known when P is an equilateral triangle or a square.

In this paper, we consider a “reverse” version of the problem, where the convex compact shapes P and Q are already given, and we ask: how large can we make P such that it can turn through 360° inside Q ?

Let’s assume that the origin is in the interior of P , and denote P rotated by θ around the origin by P_θ . Being able to turn P inside Q obviously implies that P_θ can be translated into Q for any orientation θ . Is this condition also sufficient?

In his 1921 paper solving the convex case of the Kakeya problem, Pál [8] conjectured that this is the case. Intriguingly, the paper contains a footnote added during the proof stage, stating that Harald Bohr had proven this conjecture. Unfortunately, this proof seems to have never been published, and we have not been able to find another proof in the literature. We also do not know how exactly Pál defined “turning” in this context: Is the angle changing in a strictly monotone way, or merely monotonically?

We therefore prove a stronger version of Pál’s conjecture: For a given angle θ , let $\lambda(\theta)$ be the largest scaling factor such that $\lambda(\theta)P_\theta$ can be translated into Q . We prove that the function $\theta \mapsto \lambda(\theta)$ is continuous, and show that there is a continuous function $\tau: [0, 2\pi] \rightarrow \mathbb{R}^2$ such that $\lambda(\theta)P_\theta + \tau(\theta) \subseteq Q$. In other words, P can be rotated, while continuously scaling and translating it to maintain the largest possible size that will fit inside Q at that orientation.

Our result implies Pál’s conjecture: If P_θ can be translated into Q for any θ , then $\lambda(\theta)$ is always at least one, and so $P_\theta + \tau(\theta)$ is a continuous motion that turns P inside Q .

When P and Q are convex polygons, then our problem is closely related to work by Agarwal et al. [1]. They showed that the set of all *similar copies* of a convex m -gon P that lie in a given convex n -gon Q can be represented as a convex polytope \mathcal{F} of complexity $O(mn^2)$ in \mathbb{R}^4 . (A *similar copy* of a shape P is the image of P under scaling and translation.) One can project \mathcal{F} onto a plane representing the scaling and rotation of P , and obtains a convex polygon \mathcal{P} , again of complexity $O(mn^2)$, whose boundary corresponds to points (θ, λ) where $\lambda P_\theta \subseteq Q$ and $\lambda = \lambda(\theta)$. In other words, the boundary of \mathcal{P} immediately gives a description of the function $\theta \mapsto \lambda(\theta)$.

Agarwal et al. give a construction of a convex m -gon P and a convex n -gon Q such that there are $\Theta(mn^2)$ placements of similar copies of P inscribed into Q and realizing distinct

sets of vertex-edge contacts. This implies that the complexity of \mathcal{F} is $\Theta(mn^2)$. However, their construction does not give a lower bound on the complexity of the projection \mathcal{P} (which is equal to the complexity of the function $\theta \mapsto \lambda(\theta)$), since not all the placements in their construction are maximal.

We construct a convex m -gon P and a convex n -gon Q such that there are $\Theta(mn^2)$ maximal similar placements of P inscribed to Q . This implies that the complexity of \mathcal{P} , and therefore the complexity of the function $\theta \mapsto \lambda(\theta)$, is $\Theta(mn^2)$. For the special case of P being an equilateral triangle, our lower bound construction gives a convex n -gon Q that has $\Theta(n^2)$ combinatorially distinct inscribed maximal equilateral triangles. This implies that it may be difficult to improve on the quadratic-time algorithm for finding the largest equilateral triangle inscribed to a convex polygon by DePano et al. [5].

On the algorithmic side, Agarwal et al. give an $O(mn^2 \log n)$ time algorithm that computes \mathcal{F} and its projection \mathcal{P} . From \mathcal{F} and \mathcal{P} , we can construct the functions $\theta \mapsto \lambda(\theta)$ and $\theta \mapsto \tau(\theta)$. Given these functions, we can then answer questions such as:

1. *What is the largest similar copy of P inscribed into Q ?* The answer to this question is given by the *maximum* of $\lambda(\theta)$, and its computation was the goal of Agarwal et al.
2. *What is the largest similar copy of P that can be turned through 360° inside Q ?* This question is answered by the *minimum* of $\lambda(\theta)$.

To better understand the geometry of the problem, we give a purely geometric characterization of the local minima of $\theta \mapsto \lambda(\theta)$, leading to a necessary condition for the solution to question 2.

Finally, we consider the following problem: Given a triangle Q , how can we place one point on each of its edges such that the diameter of the resulting three-point set is minimized? We prove that if Q has no angle larger than 120° , then the answer is given by the corners of the largest equilateral triangle that can be turned inside Q (that is, the solution to question 2 above). This equilateral triangle can be found by constructing the first isodynamic point of Q .

2 Parameterization

Throughout this paper, P and Q are fixed convex compact shapes in the plane, with the origin contained in the interior of P . For $\theta \in [0, 2\pi]$, let P_θ be P rotated counter-clockwise by angle θ around the origin.

We will parameterize all the scaled, rotated, and translated placements of P as points in \mathbb{R}^4 , as follows. For $\theta \in [0, 2\pi]$ and $\lambda \geq 0$, we define

$$\begin{aligned} s &:= \lambda \cos \theta, \\ t &:= \lambda \sin \theta. \end{aligned}$$

For $(s, t, x, y) \in \mathbb{R}^4$, we define the affine map $\phi = \Phi(s, t, x, y) : \mathbb{R}^2 \mapsto \mathbb{R}^2$ as follows:

$$\phi(u) = \begin{pmatrix} s & t \\ -t & s \end{pmatrix} \cdot u + \begin{pmatrix} x \\ y \end{pmatrix} = \lambda \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \cdot u + \begin{pmatrix} x \\ y \end{pmatrix}.$$

In other words, ϕ rotates by θ around the origin, scales by λ , and finally translates by (x, y) . In particular, $\phi(P) = \lambda P_\theta + (x, y)$.

We now define the set $\mathcal{F} \subset \mathbb{R}^4$ to correspond to all those ϕ such that $\phi(P) \subseteq Q$. Formally,

$$\mathcal{F} = \{(s, t, x, y) \mid \Phi(s, t, x, y)(P) \subseteq Q\}.$$

► **Lemma 1.** *The set \mathcal{F} is a compact convex set in \mathbb{R}^4 .*

Proof. For each point $p \in P$ and each closed halfspace h containing Q , define the set $H_{p,h} \subset \mathbb{R}^4$ of all the maps ϕ such that $\phi(p) \in h$. Since Φ is an affine map and h is a halfspace, $\phi(p) \in h$ is a linear constraint in the variables s, t, x, y , and so $H_{p,h}$ is a closed halfspace in \mathbb{R}^4 .

We claim that $\mathcal{F} = \bigcap H_{p,h}$. Obviously, for any $\phi \in \mathcal{F}$, $p \in P$, and $h \supset Q$ we have $\phi(p) \in h$, so it remains to prove the other direction. Consider a $\phi \notin \mathcal{F}$. This means that $\phi(P)$ contains a point $\phi(p) \notin Q$. Since Q is a compact convex set, this implies that there is a halfspace h containing Q but not containing $\phi(p)$. It follows that $\phi \notin H_{p,h}$.

As the intersection of closed convex sets, \mathcal{F} is closed and convex. Since $s^2 + t^2 = \lambda^2$ and for sufficiently large λ we can never translate λP into Q (for instance, when λ is larger than the ratio of the diameters of the two shapes), s and t are bounded. Finally, for $(x, y) \notin Q$, we always have $\phi(P) \not\subseteq Q$ because P contains the origin, and so x and y are bounded. ◀

We now define the set $\mathcal{P} \subset \mathbb{R}^2$ as the projection of \mathcal{F} onto the (s, t) -plane. By Lemma 1, \mathcal{P} is a compact convex set in the plane. Note that a pair (s, t) corresponds uniquely to a pair (λ, θ) and lies in \mathcal{P} if and only if there are $(x, y) \in \mathbb{R}^2$ such that $(s, t, x, y) \in \mathcal{F}$. In other words, whenever λP_θ can be translated into Q .

The origin lies in the interior of \mathcal{P} , since for small enough λ we can always translate P_θ into Q (for instance if λ is smaller than the ratio of the largest inscribed circle of Q and the smallest circumscribed circle of P).

Consider now a fixed $\theta \in [0, 2\pi]$. The set of (s, t) corresponding to (λ, θ) , for $\lambda \geq 0$, is the ray from the origin with orientation θ in the (s, t) -plane. Since \mathcal{P} is convex and the origin lies in its interior, this ray intersects the boundary of \mathcal{P} in a single point (s_θ, t_θ) . We set $\lambda(\theta) := \sqrt{s_\theta^2 + t_\theta^2}$. By definition of \mathcal{F} and \mathcal{P} , $\lambda(\theta)$ is the largest scaling factor such that $\lambda(\theta)P_\theta$ can be translated into Q .

► **Lemma 2.** *The function $\lambda: [0, 2\pi] \rightarrow \mathbb{R}$ is continuous.*

Proof. We show that $\theta \mapsto (s_\theta, t_\theta)$ is continuous, this implies the lemma. Assume for a contradiction, that there is a θ and an $\varepsilon > 0$ such that for every $i > 0$ there exists a θ_i with $|\theta - \theta_i| < 1/i$ such that the distance between (s_θ, t_θ) and $(s_{\theta_i}, t_{\theta_i})$ is at least ε . The points $(s_{\theta_i}, t_{\theta_i})$ all lie on the boundary of \mathcal{P} , so compactness of \mathcal{P} implies that there is a subsequence that converges to some (s, t) on the boundary of \mathcal{P} . This means we can write $(s, t) = (s_{\theta'}, t_{\theta'})$ for some $\theta' \neq \theta$. Since the origin lies in the interior of \mathcal{P} , there is a disk of radius $\rho > 0$ around the origin that lies inside \mathcal{P} . Let $\delta = |\theta - \theta'|$. For $i > 2/\delta$, we have $|\theta_i - \theta'| > \delta/2$. Since $(s_{\theta_i}, t_{\theta_i})$ and $(s_{\theta'}, t_{\theta'})$ both have distance at least ρ from the origin, their distance is lower-bounded in terms of ρ and δ , contradicting the existence of the subsequence converging to $(s_{\theta'}, t_{\theta'})$. ◀

To simplify the notation, let's define $P_\theta^* = \lambda(\theta)P_\theta$. Our goal is to find a continuous function $\tau: [0, 2\pi] \rightarrow \mathbb{R}^2$ such that $P_\theta^* + \tau(\theta) \subseteq Q$.

3 The polygonal case

In this section we will assume that P is a convex m -gon, while Q is a convex n -gon. The set \mathcal{F} can then be described as the intersection of only $n \times m$ halfspaces in \mathbb{R}^4 . Indeed, to ensure that $\phi(P) \subseteq Q$ it suffices to require that for each vertex p of P we have $\phi(p) \in Q$, which is equivalent to require that $\phi(p) \in h$ for each halfplane h supporting an edge of Q .

By the upper-bound theorem, \mathcal{F} is thus a convex polytope in \mathbb{R}^4 of complexity $O(n^2m^2)$. In fact, Agarwal et al. [1] have shown that \mathcal{F} has complexity $O(mn^2)$, and that it can be computed in time $O(mn^2 \log n)$.

The projection \mathcal{P} is therefore a convex polygon with at most $O(mn^2)$ vertices. As we have seen in the previous section, its boundary encodes the function $\lambda(\theta)$.

It remains to define the continuous function $\theta \mapsto \tau(\theta)$. We proceed as follows: Each vertex v of \mathcal{P} is of the form (s_θ, t_θ) for some θ . This gives us the value of $\lambda(\theta)$ for this θ . The vertex v is the projection of at least one vertex $(s_\theta, t_\theta, x_\theta, y_\theta)$ of \mathcal{F} . We pick one such vertex and define $\tau(\theta) = (x_\theta, y_\theta)$. Finally, we interpolate linearly between these values along the segments connecting consecutive vertices in \mathbb{R}^4 . By convexity of \mathcal{F} , these segments lie in \mathcal{F} and are therefore feasible. Since each segment projects on an edge of \mathcal{P} , each point on a segment corresponds indeed to a translation of P_θ^* . The resulting function $\tau: [0, 2\pi] \rightarrow \mathbb{R}^2$ is continuous.

We observe that Agarwal et al. gave a simpler algorithm to compute the projection \mathcal{P} in time $O(mn^2 \log n)$ without computing \mathcal{F} first. It seems that this simpler algorithm is not sufficient for our purposes, as we need a vertex of \mathcal{F} corresponding to each vertex of \mathcal{P} .

► **Theorem 3.** *Given a convex m -gon P and a convex n -gon Q , we can in time $O(mn^2 \log n)$ compute the continuous function $\lambda(\theta)$ and a continuous function $\tau: [0, 2\pi] \rightarrow \mathbb{R}^2$ such that $P_\theta^* + \tau(\theta) = \lambda(\theta)P_\theta + \tau(\theta) \subseteq Q$.*

4 The general case

The goal of this section is to generalize Theorem 3 to compact convex shapes. We start by investigating the nature of the set of feasible translations for a given angle θ .

► **Lemma 4.** *For each $\theta \in [0, 2\pi]$, the set of feasible translations $T(\theta) := \{\tau \in \mathbb{R}^2 \mid P_\theta^* + \tau \subseteq Q\}$ is either a point or a line segment.*

Proof. Since P_θ^* and Q are convex, the set $T(\theta) = \{\tau \in \mathbb{R}^2 \mid P_\theta^* + \tau \subseteq Q\}$ is convex. If $T(\theta)$ is neither a point nor a line segment, then we can find three distinct points $p, q, r \in T(\theta)$ that are not collinear. Set $\sigma = (p + q + r)/3$, and observe that $P_\theta^* + \sigma$ lies strictly in the interior of Q . It follows that there is a $\delta > 0$ such that $(1 + \delta)P_\theta^* + \sigma \subseteq Q$, contradicting the definition of $\lambda(\theta)$. ◀

► **Lemma 5.** *For any $\theta \in [0, 2\pi]$ where $T(\theta)$ is a line segment, the boundary of Q contains two line segments e, e' that are parallel to the line segment $T(\theta)$. For any point σ in the relative interior of $T(\theta)$ there is a $\delta > 0$ such that for any θ' with $|\theta - \theta'| < \delta$ there exists an $\sigma' \in T(\theta')$ such that the line $\sigma\sigma'$ is orthogonal to $T(\theta)$.*

Proof. We assume that $T(\theta)$ is horizontal, and its endpoints are $(0, 0)$ and $\tau_1 = (x_1, 0)$. Among the points with the largest y -coordinate in P_θ^* , let p_1 be the leftmost one, and let p_2 be the rightmost one (where $p_1 = p_2$ is possible as well). Among the points with the smallest y -coordinate in P_θ^* , let q_1 and q_2 be the leftmost and the rightmost one. Since $P_\theta^* \subset Q$ and $P_\theta^* + \tau_1 \subset Q$, the convex hull \mathcal{C} of P_θ^* and $P_\theta^* + \tau_1$ also lies in Q .

Let now σ be any point in the relative interior of $T(\theta)$, that is $\sigma = (x, 0)$ with $0 < x < x_1$. The set $P_\theta^* + \sigma$ lies in \mathcal{C} , and touches the boundary of \mathcal{C} only on the horizontal segments $p_1(p_2 + \tau_1)$ and $q_1(q_2 + \tau_1)$. Thus, by definition of $\lambda(\theta)$, there must be a boundary point of Q on the relative interior of both these segments. But $\mathcal{C} \subseteq Q$ implies that this is only possible if both $p_1(p_2 + \tau_1)$ and $q_1(q_2 + \tau_1)$ are part of the boundary of Q , forming the segments e and e' .

6:6 The Reverse Kakeya Problem

Let \mathcal{H} be the strip bounded by the horizontal lines through e and e' , let \mathcal{D} be the closure of $\mathcal{H} \setminus \mathcal{C}$, and let $P_0 := P_\theta^* + \sigma$. Since $\mathcal{D} \cap P_0 = \emptyset$, \mathcal{D} is closed, and P_0 is compact, there is a $\rho > 0$ such that any point in \mathcal{D} and any point in P_0 have distance at least ρ .

Let D be the diameter of P_0 , and let W be the width of the strip \mathcal{H} . Let

$$\delta < \min \left(\frac{W}{4D}, \frac{\rho}{8D}, \frac{\rho W}{32D^2} \right),$$

and consider θ' with $|\theta - \theta'| < \delta$. We will show that there is a y such that the translation $\sigma' = (x, y) \in T(\theta')$.

We first obtain $P_1 := \lambda(\theta)P_{\theta'} + \sigma$ from $P_0 = P_\theta^* + \sigma$ by a rotation around point σ by an angle smaller than δ . Let W' be the vertical width of the set P_1 . We scale P_1 by factor W/W' around σ , obtaining $P_2 = \frac{W}{W'}\lambda(\theta)P_{\theta'} + \sigma$ of vertical width W . Finally, we translate P_2 vertically by vector $(0, y)$ such that $P_3 = P_2 + (0, y) \subset \mathcal{H}$.

We claim that $P_3 = \frac{W}{W'}\lambda(\theta)P_{\theta'} + (x, y) \subset Q$. Since the vertical width of P_3 is W , this implies $\lambda(\theta') = \frac{W}{W'}\lambda(\theta)$, and so $P_3 = P_{\theta'}^* + (x, y)$, and the claim $\sigma' = (x, y) \in T(\theta')$ will follow.

During the rotation, each point of P_θ^* travels a distance of at most δD . This implies that the Hausdorff distance of P_0 and P_1 is at most $\delta D < \rho/8$. Therefore

$$\begin{aligned} W - 2\delta D &\leq W' \leq W + 2\delta D, & \text{and so} \\ 1 + \frac{2\delta D}{W - 2\delta D} &\geq \frac{W}{W'} \geq 1 - \frac{2\delta D}{W + 2\delta D} \end{aligned}$$

Using $\delta \leq W/(4D)$ we have $2\delta D \leq W/2$, and so

$$\begin{aligned} 1 + 4\delta \frac{D}{W} &\geq \frac{W}{W'} \geq 1 - \frac{4}{3}\delta \frac{D}{W}, & \text{implying} \\ 4\delta \frac{D}{W} &\geq \frac{W}{W'} - 1 \geq -\frac{4}{3}\delta \frac{D}{W}, & \text{and so} \quad \left| \frac{W}{W'} - 1 \right| \leq 4\delta \frac{D}{W}. \end{aligned}$$

During the scaling by factor W/W' , a point travels a distance of at most $|W/W' - 1| \cdot D$, which is bounded by $4\delta D^2/W < \rho/8$.

It follows that P_2 and P_0 have Hausdorff distance at most $\rho/8 + \rho/8 = \rho/4$. This implies in particular that the translation length $\tau \leq \rho/4$, which in turn implies that the Hausdorff distance of P_3 and P_0 is at most $\rho/2$.

There is no point of \mathcal{D} within distance $\rho/2$ of P_0 , so $P_3 \cap \mathcal{D} = \emptyset$. From $P_3 \subset \mathcal{H}$ then follows $P_3 \subset \mathcal{C} \subset Q$. ◀

Let $\mathcal{S}(\theta)$ be the two-dimensional affine subspace of \mathbb{R}^4 where the first two coordinates are $\lambda(\theta) \cos \theta$ and $\lambda(\theta) \sin \theta$. In other words,

$$\mathcal{S}(\theta) := \{(\lambda(\theta) \cos \theta, \lambda(\theta) \sin \theta, x, y) \mid (x, y) \in \mathbb{R}^2\}.$$

Since the first two coordinates are constant, $\mathcal{S}(\theta)$ is parallel to the (x, y) -plane in \mathbb{R}^4 . We next set $\mathcal{T}(\theta) := \mathcal{S}(\theta) \cap \mathcal{F}$. Note that $T(\theta)$ is the projection of $\mathcal{T}(\theta)$ on the (x, y) -plane, or, put differently, $\mathcal{T}(\theta)$ is $T(\theta)$ “lifted” to the two-dimensional affine subspace $\mathcal{S}(\theta)$.

We define the function $\ell: [0, 2\pi] \rightarrow \mathbb{R}$ such that $\ell(\theta)$ is the length of the segment $T(\theta)$ (and zero when $T(\theta)$ is a point). Since $\mathcal{S}(\theta)$ is parallel to the (x, y) -plane, $\ell(\theta)$ is also the length of the segment $\mathcal{T}(\theta)$. Finally, we define the function $m: [0, 2\pi] \rightarrow \mathbb{R}^4$ such that $m(\theta)$ is the midpoint of the segment $\mathcal{T}(\theta)$.

► **Lemma 6.** *The function $[0, 2\pi] \mapsto m(\theta)$ is continuous.*

Proof. Let's assume for a contradiction that the claim is false. Then there exists a $\theta \in [0, 2\pi]$ and an $\varepsilon > 0$ such that for every $i \in \{1, 2, \dots\}$ there is $\theta_i \in [0, 2\pi]$ such that $|\theta - \theta_i| < 1/i$ and $|m(\theta) - m(\theta_i)| > \varepsilon$.

Let p_i and q_i be the endpoints of $\mathcal{T}(\theta_i)$. These points lie in the set $\mathcal{S} \cap \mathcal{F}$, where $\mathcal{S} = \bigcup_{\theta \in [0, 2\pi]} \mathcal{S}(\theta)$. Since \mathcal{S} is closed and \mathcal{F} is compact, the intersection $\mathcal{S} \cap \mathcal{F}$ is compact, and so (p_i) and (q_i) have converging subsequences. To avoid double indices, we replace our sequence (θ_i) by such a subsequence where both (p_i) and (q_i) converge, and set $p := \lim_{i \rightarrow \infty} p_i$ and $q := \lim_{i \rightarrow \infty} q_i$. By compactness, $p, q \in \mathcal{S} \cap \mathcal{F}$. Since $p_i, q_i \in \mathcal{S}(\theta_i)$ and $\lim_{i \rightarrow \infty} \theta_i = \theta$, we have $p, q \in \mathcal{S}(\theta)$, and so $p, q \in \mathcal{S}(\theta) \cap \mathcal{F} = \mathcal{T}(\theta)$. By continuity of vector addition, the midpoint $m(\theta_i)$ of $p_i q_i$ converges to the midpoint m of $p q$, that is, $m = \lim_{i \rightarrow \infty} m(\theta_i)$. Because $|m(\theta) - m(\theta_i)| > \varepsilon$, we have $|m(\theta) - m| \geq \varepsilon$. Since $m(\theta)$ and m are on $\mathcal{T}(\theta)$, it follows that the segment $\mathcal{T}(\theta)$ has length at least ε .

We now pick points $(s_\theta, t_\theta, x_1, y_1)$ and $(s_\theta, t_\theta, x_2, y_2)$ from the segment $\mathcal{T}(\theta)$ at distance $\varepsilon/3$ from the two endpoints, respectively, where $s_\theta = \lambda(\theta) \cos \theta$ and $t_\theta = \lambda(\theta) \sin \theta$. Let $\sigma_1 := (x_1, y_1)$ and $\sigma_2 := (x_2, y_2)$ be the projections of the two points onto $T(\theta)$. By Lemma 5, there exists a $\delta > 0$ such that for any $\theta' \in [0, 2\pi]$ with $|\theta - \theta'| < \delta$ there exist points $\sigma'_1, \sigma'_2 \in T(\theta')$ such that $\sigma'_1 \sigma_1$ and $\sigma'_2 \sigma_2$ are orthogonal to $T(\theta)$. Since the segment $\sigma'_1 \sigma'_2$ is a subset of $T(\theta')$, this implies that

$$\ell(\theta') \geq |\sigma'_1 \sigma'_2| \geq |\sigma_1 \sigma_2| = \ell(\theta) - \frac{2\varepsilon}{3}.$$

It follows that for $i > 1/\delta$, we have

$$|p_i q_i| = \ell(\theta_i) \geq \ell(\theta) - \frac{2\varepsilon}{3}.$$

which implies that the segment $p q$ has length at least $\ell(\theta) - 2\varepsilon/3$.

Since the length of $\mathcal{T}(\theta)$ is $\ell(\theta)$ and $p q$ is a subset of $\mathcal{T}(\theta)$, the distance between the midpoint m of $p q$ and the midpoint $m(\theta)$ of $\mathcal{T}(\theta)$ is at most $\varepsilon/3$. This is a contradiction to the assumption that $|m(\theta) - m(\theta_i)| > \varepsilon$ for all i . ◀

We now obtain our theorem by defining $\tau(\theta)$ to be the midpoint of $T(\theta)$, that is, the projection of $m(\theta)$ on the (x, y) -plane.

► **Theorem 7.** *For any compact convex shapes P and Q , there exists a continuous function $\tau: [0, 2\pi] \rightarrow \mathbb{R}^2$ such that $P_\theta^* + \tau(\theta) = \lambda(\theta)P_\theta + \tau(\theta) \subseteq Q$.*


Using the same proof idea as in Lemma 6, we can also show that the function $\theta \mapsto \ell(\theta)$ is continuous.

5 Lower bounds

We now construct convex polygons P and Q such that there are many combinatorially distinct inscribed maximal similar placements of P in Q . These imply identical lower bounds for the complexity of the polygon \mathcal{P} .

► **Theorem 8.** *For any n there is a convex n -gon Q such that there are $\Theta(n^2)$ maximal placements of an equilateral triangle in Q .*

For any m and n there is a convex m -gon P and a convex n -gon Q such that there are $\Theta(mn^2)$ maximal similar placements of P in Q .



■ **Figure 1** Construction of Q with many equilateral triangles inscribed.


Proof. Let rr' be a vertical segment of length one. Let C_1 be a unit radius circle with center r , let C_2 be a unit radius circle with center r' , see Figure 1. Let $\beta = \pi/36$ and assume $n \equiv 2 \pmod{4}$. We distribute $n + 1$ points p_0, p_1, \dots, p_n regularly along the arc of C_1 from polar coordinate $4\pi/3 - \beta$ to polar coordinate $4\pi/3 + \beta$, and $n + 1$ points q_0, q_1, \dots, q_n regularly along the arc of C_1 from polar coordinate $5\pi/3 - \beta$ to polar coordinate $5\pi/3 + \beta$ (Figure 1 shows only one intermediate point). Finally, we place $n + 2$ points r_0, \dots, r_{n+1} on an arc of length ε of C_2 around r (The points are exaggerated in the figure, as they would otherwise be indistinguishable from r .)

We now construct a polygon Q as the intersection of three groups of halfplanes as follows:

- the lower halfplanes with bounding lines $r_0r_1, r_1r_2, r_2r_3, \dots, r_{n-1}r_n$;
- the upper halfplanes with bounding lines $p_0p_2, p_5p_6, p_9p_{10}, p_{13}p_{14}, \dots, p_{n-1}p_n$;
- the upper halfplanes with bounding lines $q_0q_1, q_4q_5, q_8q_9, q_{12}q_{13}, \dots, q_{n-2}q_n$.

We observe that r, p_j, q_j form an equilateral triangle of side length one for each $j \in \{0, 1, \dots, n\}$. The points p_1 and q_{n-1} lie outside Q , the points $p_3, p_4, p_7, p_8, p_{11}, p_{12}, \dots, p_{n-2}$ and the points $q_2, q_3, q_6, q_7, q_{10}, q_{11}, \dots, q_{n-3}$ lie in the interior of Q , and all other points lie on the boundary of Q .

When ε is small enough, then the circles of radius one around each r_i intersect the boundary of Q in the same edges and in points very close to the intersection points with the circle C_1 . This implies that for every $i \in \{1, 2, \dots, n\}$ and every $j \in \{3, 7, 11, 15, \dots, n - 3\}$, there is an equilateral triangle Δ_{ij} with side length one with one corner at r_i and two other corners very close to p_j and q_j in the interior of Q . If we rotate Δ_{ij} clockwise around r_i , it will ultimately hit the edge supported by the line $p_{j-2}p_{j-1}$ in a point close to p_{j-1} , see Figure 2. If we rotate Δ_{ij} counter-clockwise around r_i , then it will ultimately hit the edge supported by the line $q_{j+1}q_{j+2}$ in a point close to q_{j+1} . Let $f(\theta)$ and $g(\theta)$ be the length of the ray from r_i to the boundary of Q along the two incident edges of Δ_{ij} , as we rotate along the interval where the triangle is contained in Q . Each function evaluates to one at one extreme and to a value strictly larger than one at the other extreme, implying that there must be an intermediate angle where $f(\theta) = g(\theta) > 1$. At this orientation we can enlarge Δ_{ij} to an equilateral triangle of side length $f(\theta) = g(\theta)$ that touches the boundary of Q at every corner. One corner is at r_i , the second touches an edge supported either by $p_{j-2}p_{j-1}$



■ **Figure 2** Detail of the construction of Q (magnified).


or $p_{j+2}p_{j+3}$, and the third touches an edge supported either by $q_{j-3}q_{j-2}$ or by $q_{j+1}q_{j+2}$. It follows that there are $\Theta(n^2)$ combinatorially distinct such equilateral triangles. Each of them is maximal for the given orientation, since the slopes of the incident edges do not allow for the triangle to be translated in any way.

We will now generalize the construction to a convex m -gon P . We first observe that the construction above will work for any sufficiently small value of β . The polygon P is obtained by taking an equilateral triangle Δwu_0v_0 of side length one, rotating it m times around w by an angle $\alpha < \pi/(6n)$, and taking the convex hull of these $m + 1$ equilateral triangles Δwu_kv_k , see the red polygon in Figure 3. The polygon Q is obtained as above, but with $\beta < \alpha/12$. For every $k \in \{0, 1, \dots, m\}$, every $i \in \{1, \dots, n\}$, and every $j \in \{3, 7, 11, 15, \dots, n - 3\}$, there is a placement of P such that w coincides with r_i , u_k is very close to p_j in the interior of Q , v_k is very close to q_j in the interior of Q , and all other vertices of P lie well inside the interior of Q . As in the previous construction, we can now slightly rotate and scale this placement to obtain a maximal placement of P . ◀

6 A geometric construction for triangles and the isodynamic point

We consider now the following question: Given convex polygons P and Q , what is the largest scaling factor λ^* such that λ^*P can be fully turned inside Q ? From the considerations of Section 2 follows that $\lambda^* = \min_{0 \leq \theta < 2\pi} \lambda(\theta)$. In particular, λ^* is a local minimum of $\theta \mapsto \lambda(\theta)$. We now give a geometric characterization of these local minima that will also allow us to construct them.

First we consider the case when P and Q are triangles. For a fixed θ , $\lambda(\theta)$ can be computed by linear programming, with three variables λ, x, y for the translation and scaling, and 3×3 constraints to keep all vertices of $\lambda P + (x, y)$ inside Q . This means that there is an optimal solution where at least three constraints are tight: either the three vertices of $\lambda P + (x, y)$ lie on the three edges of Q , or a vertex of $\lambda P + (x, y)$ coincides with a vertex of Q and another vertex of $\lambda P + (x, y)$ lies on an edge of Q .



■ **Figure 3** Generalizing the lower bound to a convex m -gon P .

Consider now an angle $\theta_0 \in [0, 2\pi]$ where λ has a local minimum, and let $\lambda_0 = \lambda(\theta_0)$. That implies there is an $\varepsilon > 0$ such that for $\theta_0 - \varepsilon < \theta < \theta_0 + \varepsilon$ we have $\lambda(\theta) \geq \lambda_0$. In other words, for $\theta \in [\theta_0 - \varepsilon, \theta_0 + \varepsilon]$, we can translate $\lambda_0 P_\theta$ into Q .

If a vertex U of $\lambda_0 P_\theta + (x, y)$ coincides with a vertex of Q and another vertex V lies on an edge e of Q , this is only possible if the segment uv is orthogonal to e . In other words, an edge of $\lambda_0 P_\theta + (x, y)$ coincides with a height of Q . There are nine possible candidate placements of this form.

Otherwise, the three vertices U, V, W of $\lambda_0 P_\theta + (x, y)$ must lie on the three edges a, b, c of Q . We have the following lemma.

► **Lemma 9.** *If the vertices U, V, W of a triangle Δ lie on the three edges a, b, c of a triangle Q , but there exists an $\varepsilon > 0$ such that Δ_θ can be translated into Q for all $-\varepsilon < \theta < \varepsilon$, then the normals to a, b, c in U, V, W meet in a point inside the circumcircle of Q . There is at most one such placement that is similar to a given triangle P .*

Proof. Assume the normals do not meet in a point. Then pick a point F in the interior of the only bounded cell in the arrangement of normals. Rotating Δ around F will cause all vertices of Δ to leave Q , either for the clockwise or counter-clockwise rotation. So, for arbitrarily small θ , there is a triangle Δ_θ that cannot be translated into Q , a contradiction.

Let now A, B, C be the vertices of Q such that $a = BC, b = AC$ and $c = AB$. Let $\alpha, \beta,$ and γ be the triangle angles incident to A, B and C , respectively. We are looking for a point $F \in \mathbb{R}^2$ such that the orthogonal projections of F onto the edges $a, b,$ and c form a triangle $\bar{A}\bar{B}\bar{C}$ that is similar to the given triangle $P = \Delta UVW$, or, in other words, so that $\bar{A}\bar{B} : \bar{B}\bar{C} = UV : VW$ and $\bar{A}\bar{B} : \bar{A}\bar{C} = UV : UW$. We first observe that if F does not lie in the circumcircle of Q , then its orthogonal projection on at least one of the supporting lines of the edges of Q does not lie on the edge, so F must indeed lie inside the circumcircle.

Since $\angle \bar{A}\bar{B}F = \angle \bar{A}\bar{C}F = 90^\circ$, the points \bar{B} and \bar{C} lie on the circle with diameter AF , and so we have $\bar{B}\bar{C} = AF \sin \alpha$. With the same reasoning we obtain $\bar{C}\bar{A} = BF \sin \beta$ and

$\bar{A}\bar{B} = CF \sin \gamma$. This means that the point F is the point where the following holds:

$$\begin{aligned} UV : VW &= \bar{A}\bar{B} : \bar{B}\bar{C} = CF \sin \gamma : AF \sin \alpha \\ UV : UW &= \bar{A}\bar{B} : \bar{A}\bar{C} = CF \sin \gamma : BF \sin \beta \end{aligned}$$

It follows that the point F satisfies

$$\begin{aligned} AF : CF &= VW \sin \gamma : UV \sin \alpha \\ BF : CF &= UW \sin \gamma : UV \sin \beta \end{aligned}$$

The points X with constant $AX : CX$ lie on an Apollonius circle around a center H on the line AC . The center lies outside the segment AC (and the circle degenerates to the bisector of A and C when $AX : CX = 1 : 1$). Similarly, the points X with constant $BX : CX$ lie on a circle around a center K on the line BC , outside the segment BC . The two circles intersect in at most two points F_1 and F_2 . The two points are inverses of each other with respect to the circumcircle of Q , so only one of them can lie inside the circumcircle. ◀

Lemma 9 allows us to construct three more candidate placements for λP_θ inside Q (for the three circular assignments of vertices of P to the edges of Q). In total, we have 12 candidate placements, and the smallest feasible placement determines the bottleneck scaling factor λ^* .

An interesting special case of Lemma 9 occurs when P is an equilateral triangle. In this case the point F of the lemma is a well-known triangle center, the *first isodynamic point*. The two isodynamic points of a triangle $\triangle ABC$ (only the equilateral triangle has only one) can be defined as the points F whose orthogonal projections on the lines supporting the triangle edges form an equilateral triangle, or equivalently as the points F such that $AF : BF : CF = 1/a : 1/b : 1/c$.


The geometric construction of Lemma 9 can also be applied to arbitrary convex polygons P and Q . Again, the linear programming argument shows that for each θ , the placement of $P_\theta^* = \lambda(\theta)P_\theta$ in Q satisfies three constraints with equality, so either three corners of P_θ^* lie on three edges of Q , or a vertex of P_θ^* coincides with a vertex of Q and another vertex of P_θ^* lies on an edge of Q . We can thus look at all possible triples of corners of P and triples of edges of Q , and geometrically construct candidate placements. The smallest feasible placement again determines λ^* .

7 Minimizing the diameter of points on the edges of a triangle

Consider now the following problem: Given a triangle $Q = \triangle ABC$, find three points u, v, w , one on each edge of Q , such that the diameter of the set $\{u, v, w\}$ is minimized. Since the diameter of three points, one on each edge of Q , is a continuous function defined on a compact domain, there exists a set of three points u, v, w that realizes the minimum diameter. Let us denote this minimal diameter as $\nabla(Q)$.

We first claim that $\nabla(Q) \leq \lambda^* = \min_{0 \leq \theta < 2\pi} \lambda(\theta)$, where λ is our well-known function for P an equilateral triangle of side length one and the given Q .

Indeed, we saw in Section 6 that λ^* is attained either by an equilateral triangle that touches each edge of Q , or by a height of Q . In the first case we immediately get a candidate solution $\{u, v, w\}$ of cost λ^* . In the second case, let's assume it's the height h_A defined by A and its projection H_A on the edge $a = BC$. Then we can choose $u = H_A$ and pick v and w close to A on b and c for a candidate solution of cost smaller than λ^* .



■ **Figure 4** Point sets of minimal diameter.

► **Theorem 10.** *If triangle $Q = \triangle ABC$ has no angle larger than 120° , then $\nabla(Q) = \lambda^*$. If its angle $\alpha > 120^\circ$, then $\nabla(Q)$ is given by the points $\{u, v, w\}$, where u is the intersection of the angular bisector of α with the edge $a = BC$, and v and w are the orthogonal projections of u onto $b = AC$ and $c = AB$.*

Proof. Let's first assume that $\alpha > 120^\circ$, and let u, v, w be as stated (see Figure 4 (a)). Since $\alpha > 120^\circ$, we have $\angle vuw = 180^\circ - \alpha < 60^\circ$, and so $vw < uv = uw$, implying that the diameter of $\{u, v, w\}$ is $uv = uw$. Assume for a contradiction there is a point set $u' \in a, v' \in b, w' \in c$ whose diameter is at most uv . If u' lies between u and C , then its distance to the line AB is larger than $uv = uw$; if u' lies between B and u , then its distance to the line AC is larger than uv . In both cases we obtain a contradiction, so $u' = u$. Since v and w are the only points on b and c at distance at most uv from u , we also have $v' = v$ and $w' = w$, proving the second part of the theorem.

Assume now that Q has no angle larger than 120° , and let u, v, w be a point set attaining the optimal diameter $\nabla(Q)$. Since (by an argument as above) $\nabla(Q)$ is less than the shortest height of Q , none of u, v, w coincides with a vertex of Q .

We claim that u, v, w form an isosceles triangle such that the angle between the two equal sides is at most 60° . Assume for a contradiction that uv is the unique longest edge. Then the diameter of the set $\{u, v, w\}$ is the length of uv , however, the length of uv can be reduced by moving u and v slightly. It follows that there is not a unique longest edge, and thus u, v, w must form an isosceles triangle such that the angle between the two equal sides is at most 60° .

Without loss of generality, we can assume that $uv = uw \geq vw$, and that $u \in a = BC, v \in b = AC, w \in c = AB$. We next show that the triangle with vertices u, v, w is equilateral. Assume for a contradiction that $vw < uv = uw$, which implies $\angle vuw < 60^\circ$. If uv is not perpendicular to b , then we can reduce the length of uv by moving v slightly, and then we can move u slightly to be closer to w . The same argument holds for uw , so $uv \perp b$ and $uw \perp c$, which implies $\alpha = 180^\circ - \angle vuw > 120^\circ$, leading to a contradiction.

It follows that $vw = uv = uw$, so the points u, v, w form an equilateral triangle. Since $\triangle uvw$ is inscribed to Q , its side length is $\lambda(\theta)$ for some θ , so $\nabla(Q) \leq \lambda^*$ implies $\nabla(Q) = \lambda^*$, see Figure 4 (b). ◀

For $\alpha < 120^\circ$, λ^* is a local minimum of the form of Lemma 9, so we can find the solution by projecting the first isodynamic point on each triangle edge. If $\alpha = 120^\circ$, then the intersection of the angular bisector of α with the edge a is the first isodynamic point, and both constructions coincide.

References

- 1 P. K. Agarwal, N. Amenta, and M. Sharir. Largest placement of one convex polygon inside another. *Discrete & Computational Geometry*, 19:95–104, 1998. doi:10.1007/PL00009337.
- 2 A. S. Besicovitch. Sur deux questions de l'intégrabilité. *Journal de la Société des Math. et de Phys.*, II, 1920.
- 3 A. S. Besicovitch. On Kakeya's problem and a similar one. *Math. Zeitschrift*, 27:312–320, 1928.
- 4 J. Bourgain. Harmonic analysis and combinatorics: How much may they contribute to each other? In V. I. Arnold, M. Atiyah, P. Lax, and B. Mazur, editors, *Mathematics: Frontiers and Perspectives*, pages 13–32. American Math. Society, 2000.
- 5 A. DePano, Yan Ke, and J. O'Rourke. Finding largest inscribed equilateral triangles and squares. In *Proc. 25th Allerton Conf. Commun. Control Comput.*, 1987.
- 6 S. Kakeya. Some problems on maxima and minima regarding ovals. *The Science Report of the Tohoku Imperial University, Series 1, Mathematics, Physics, Chemistry*, 6:71–88, 1917.
- 7 I. Laba. From harmonic analysis to arithmetic combinatorics. *Bulletin (New Series) of the AMS*, 45:77–115, 2008.
- 8 G. Pál. Ein Minimumproblem für Ovale. *Math. Ann.*, 83:311–319, 1921.
- 9 T. Tao. From rotating needles to stability of waves: Emerging connections between combinatorics, analysis and PDE. *Notices of the AMS*, 48:297–303, 2001.
- 10 T. Wolff. Recent work connected with the Kakeya problem. In H. Rossi, editor, *Prospects in Mathematics*. American Math. Society, 1999.