

Shape Matching under Rigid Motion*

Siu-Wing Cheng[†]

Chi-Kit Lam[†]

Abstract

We present improved algorithms to match two polygonal shapes P and Q to approximate their maximum overlap. Let n be their total number of vertices. Our first algorithm finds a translation that approximately maximizes the overlap area of P and Q under translation in $\tilde{O}(n^2\varepsilon^{-3})$ time. The error is additive and it is at most $\varepsilon \cdot \min\{\text{area}(P), \text{area}(Q)\}$ with probability $1 - n^{-O(1)}$. We also obtain an algorithm that approximately maximizes the overlap of P and Q under rigid motion in $\tilde{O}(n^3\varepsilon^{-4})$ time. The same error bound holds with probability $1 - n^{-O(1)}$. We also show how to improve the running time to $\tilde{O}(n + \varepsilon^{-3})$ for the translation case when one of the polygons is convex.

*Research supported by the Research Grant Council, Hong Kong, China (project no. 611711).

[†]Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong. Emails: scheng@cse.ust.hk, george.lam@family.ust.hk

1 Introduction

1.1 Background

A common task in object recognition is to find a translation or rigid motion that minimizes the dissimilarity measure between two objects. For two-dimensional shapes, a robust similarity measure is the area of the symmetric difference [3]. Minimizing the area of symmetric difference is equivalent to maximizing the overlap—the area of the intersection of the two shapes.

Given two polygonal shapes P and Q with a total of n vertices, Mount et al. [14] gave an algorithm to compute their maximum overlap under translation in $O(n^4)$ time. If P and Q are convex, de Berg et al. [8] presented an algorithm to find their maximum overlap under translation in $O(n \log n)$ time. For convex polygons, Ahn et al. [2] presented algorithms to find a $(1 - \varepsilon)$ -approximate maximum overlaps under translation and under rigid motion, assuming that the polygon vertices are stored in arrays in clockwise order around the polygon boundaries. The running times are $O(\varepsilon^{-1} \log n + \varepsilon^{-1} \log(1/\varepsilon))$ for the translation case and $O(\varepsilon^{-1} \log n + \varepsilon^{-2} \log(1/\varepsilon))$ for the rigid motion case. By representing the overlap function as a sum of algebraic functions, Vigneron [16] devised an algorithm to compute a $(1 - \varepsilon)$ -approximate maximum overlap in $\tilde{O}(n^6 \varepsilon^{-3})$ time.¹

If the maximum overlap under rigid motion is not $\Omega(\max\{\text{area}(P), \text{area}(Q)\})$, then P and Q are hardly similar and knowing this is often sufficient for shape matching. This motivates the approximation of the maximum overlap to within an additive error of $\varepsilon \cdot \max\{\text{area}(P), \text{area}(Q)\}$ or $\varepsilon \cdot \min\{\text{area}(P), \text{area}(Q)\}$. Cheong et al. [6] proposed algorithms to approximate the maximum overlap such that the additive error is $\varepsilon \cdot \min\{\text{area}(P), \text{area}(Q)\}$ with probability $1 - n^{-O(1)}$. The running times are $O(n^2 \varepsilon^{-4} \log^2 n)$ time for the translation case and $O(n^3 \varepsilon^{-8} \log^5 n)$ time for the rigid motion case², assuming that $n \geq \varepsilon^{-1}$. Recently, Alt et al. [4] also obtained some probabilistic results with additive error $\varepsilon \cdot \min\{\text{area}(P), \text{area}(Q)\}$, but their running times depend on some geometric parameters, including the areas and perimeters of P and Q .

De Berg et al. [7] presented algorithms to align a set of disjoint unit disks with another set of disjoint unit disks to obtain a $(1 - \varepsilon)$ -approximate maximum overlap. Let n be the total number of disks. The running times are $O(n^2 \varepsilon^{-2} \log(n/\varepsilon))$ for the translation case and $O(n^4 \varepsilon^{-3} \log n)$ for the rigid motion case. When the overlap is $\Omega(n)$, they also presented a probabilistic algorithm that runs in $O(n^2 \varepsilon^{-4} \log(n/\varepsilon) \log^2 n)$ time.

In \mathbb{R}^d for $d \geq 3$, Ahn et al. [1] showed that the maximum overlap of two convex polytopes under translation can be computed in $O(n^{\lfloor d/2 \rfloor + 1} \log^d n)$ time, where n is the total number of bounding hyperplanes, and the running time can be improved to $O(n \log^{3.5} n)$ in \mathbb{R}^3 . Vigneron's method [16] computes a $(1 - \varepsilon)$ -approximate maximum overlap of two possibly non-convex polytopes under rigid motion in $O\left(\left(\frac{n^2}{\varepsilon} \log \frac{n}{\varepsilon}\right)^{d^2/2 + d/2 + 1}\right)$ time, assuming that the two polytopes are represented by a union of n interior-disjoint d -simplices.

1.2 Our contributions and overview

We build upon Cheong et al.'s framework [6] to approximate the maximum overlap of two polygonal shapes P and Q , which may have multiple connected components and holes. We do not assume that the supporting lines of the edges of P and Q are in general position.

Let n be the total number of vertices in P and Q . The running times of our algorithms are $O(n^2 \varepsilon^{-3} \log^{1.5} n \log(n/\varepsilon))$ for the translation case, and $O(n^3 \varepsilon^{-4} \log^{5/3} n \log^{5/3}(n/\varepsilon))$ for the rigid motion case. The error is additive and it is at most $\varepsilon \cdot \text{area}(P)$ with probability $1 - n^{-O(1)}$.

¹We use $\tilde{O}(\cdot)$ to hide multiplicative factors that are polynomial in the logarithm of n and $1/\varepsilon$.

²There is a typo in [6] in the running time bound in the case of rigid motion as noted in [16].

If Q is convex, the running time for the translation case can be improved to $O(n \log n + \varepsilon^{-3} \log^{2.5} n \log \frac{\log n}{\varepsilon})$. When both P and Q are general polygonal shapes, we can switch the roles of P and Q , so the error bound $\varepsilon \cdot \text{area}(P)$ is equivalent to $\varepsilon \cdot \min\{\text{area}(P), \text{area}(Q)\}$.

In comparison with the results of Cheong et al. [6], our running times have almost the same dependence on n (differing by some polylog factors) but lower polynomial dependence on $1/\varepsilon$. In the case of rigid motion, we obtain an improvement from ε^{-8} to ε^{-4} . Our bounds are free of geometric parameters as opposed to the result of Alt et al. [4]. For the rigid motion case, the running time of Vigneron’s algorithm [16] has a lower dependence on ε (ε^{-3} versus ε^{-4}) and a significantly higher dependence on n (n^6 versus n^3), but it returns a $(1 - \varepsilon)$ -approximate maximum overlap even if the maximum overlap is tiny compared to $\min\{\text{area}(P), \text{area}(Q)\}$.

When P or Q has multiple connected components and/or holes, we can preprocess it in $O(n \log n)$ time as follows so that the shape boundary is a single non-self-crossing polygonal curve. To deal with holes, we can compute a non-self-crossing spanning tree of the holes and the outer boundary (e.g. minimum spanning tree), and then split the tree edges into narrow channels of negligible area. These new channels have negligible effect on the maximum overlap. The splitting of the tree edges into narrow channels can also be simulated symbolically. Similarly, if a shape has multiple connected components, we can connect them by a non-self-crossing spanning tree to form a shape with a single connected boundary. Hence, we can always assume that the boundary of input shape is a single non-self-crossing polygonal curve.

We briefly sketch the framework of Cheong et al. [6] upon which we build our results. The set of all possible translations in the plane is just \mathbb{R}^2 because every translation is specified by a vector t . For every $t \in \mathbb{R}^2$, define $\mu(t) = \text{area}((P + t) \cap Q) / \text{area}(P)$. Cheong et al. proposed to sample a set S of points from P uniformly at random, and count the number of sample points contained in a translated copy of Q . For every $t \in \mathbb{R}^2$, define $\mu_S(t) = |(S + t) \cap Q| / |S|$. Their idea is to find $\tilde{t} = \text{argmax}_{t \in \mathbb{R}^2} \mu_S(t)$ and show that $\mu(\tilde{t}) \geq \mu(\hat{t}) - \varepsilon$, where $\hat{t} = \text{argmax}_{t \in \mathbb{R}^2} \mu(t)$. The dependence of $\mu_S(t)$ on t is best interpreted by forming an arrangement of translated copies of Q . For every $s \in S$ and every $t \in \mathbb{R}^2$, $s + t$ belongs to Q if and only if $t \in Q - s$. Notice that $Q - s$ is obtained by applying the translation vector $-s$ to Q . Therefore, if we form the arrangement of $\mathcal{T} = \{Q - s : s \in S\}$, then for every cell f of the arrangement and every pair of translations $t, t' \in f$, $(S + t) \cap Q = (S + t') \cap Q$. (A face of Q can be the interior of Q , the interior of an edge of Q , or a vertex of Q . A cell of the arrangement of \mathcal{T} is the common intersection of some faces of some copies of Q . So a cell of positive dimension is a relative open set.) It follows that $\text{argmax}_{t \in \mathbb{R}^2} \mu_S(t)$ can be computed by examining all vertices of the arrangement of \mathcal{T} , which is the algorithm of Cheong et al. [6]. Rigid motion is treated similarly by adding one more dimension to model the angle of rotation. The arrangement becomes three-dimensional.

We achieve our bounds by proving some results on the depth variation in the arrangements. The strategy is to show that there are plenty vertices of similar depth near the deepest arrangement vertex, so we can sample the arrangement vertices to “approximate” the deepest vertex without constructing the whole arrangement. The depth variation result in the translation case applies to any collection of compact shapes in \mathbb{R}^2 that satisfy some mild conditions. Compact shapes other than polygons are allowed, and the collection is more general than the one for the shape matching problem. We establish some geometric properties of the 3D arrangement for the rigid motion case that allows us to build the depth variation result upon the one for the translation case. Still, this yields only an $\tilde{O}(n^3 \varepsilon^{-5})$ running time. We employ a dynamic planar point location data structure to reduce the running time to $\tilde{O}(n^3 \varepsilon^{-4})$.

2 Matching under translation

Since the basic framework considers how different translated copies of Q contain different subsets of the random sample S , it is related to the notion of *range space* and *VC-dimension* [15]. A set X and a set \mathcal{R} of subsets of X form a range space (X, \mathcal{R}) . A finite subset $Y \subseteq X$ is *shattered* by \mathcal{R} if every subset of Y equals to $X \cap R$ for some $R \in \mathcal{R}$. The VC-dimension of (X, \mathcal{R}) is the cardinality of the largest finite subset $Y \subseteq X$ that can be shattered by \mathcal{R} .

Lemma 2.1. *Let K be a union of k interior-disjoint d -simplices in \mathbb{R}^d . Let X be a subset of \mathbb{R}^d . Let $\mathcal{R} = \{X \cap (K - t) : t \in \mathbb{R}^d\}$. The range space (X, \mathcal{R}) has VC-dimension $O(d \log k)$.*

Proof. Let Z be a finite subset of points in X that is shattered by \mathcal{R} . It suffices to prove an $O(d \log k)$ bound on $|Z|$. Let \mathcal{A} be the arrangement of $\{K - z : z \in Z\}$. \mathcal{A} has cells of dimensions from 0 to d , and the total number of cells is $O(k^d |Z|^d)$ [10]. Notice that for every $z \in Z$, a translation t belongs to $K - z$ if and only if z belongs to $K - t$. For every pair of translations $t, t' \in \mathbb{R}^d$ that belong to the same cell of \mathcal{A} , $t \in K - z \iff t' \in K - z$, which implies that $z \in K - t \iff z \in K - t'$. Hence, $Z \cap (K - t) = Z \cap (K - t')$. The cardinality of $\{Z \cap (K - t) : t \in \mathbb{R}^d\}$ is thus at most the total number of cells in \mathcal{A} . On the other hand, the cardinality of $\{Z \cap (K - t) : t \in \mathbb{R}^d\}$ is $2^{|Z|}$ because Z is shattered by \mathcal{R} by assumption. Hence, $2^{|Z|} = O(k^d |Z|^d)$, which implies that $|Z| = O(d \log k)$. \square

Lemma 2.1 implies the following result that for every $t \in \mathbb{R}^2$, $\mu_S(t)$ is a good approximation of $\mu(t)$.

Lemma 2.2. *For every constant $r > 0$, there exists a constant $c > 0$ such that if S is a set of points sampled uniformly at random from P and $|S| \geq c\varepsilon^{-2} \log_2 n$, then with probability at least $1 - n^{-r}$, $|\mu_S(t) - \mu(t)| \leq \varepsilon$ for every $t \in \mathbb{R}^2$.*

Proof. Let (X, \mathcal{R}) be a range space with finite positive VC-dimension ν . For every $R \in \mathcal{R}$, let ρ_R denote the probability that a point drawn from X uniformly at random belongs to R . The ε -approximation result [12, 13] says that there exists a constant $c_0 > 0$ such that for every $\varepsilon, q \in (0, 1)$, if we draw a subset S uniformly at random from X with $|S| \geq c_0 \varepsilon^{-2} (\nu + \ln(1/q))$, then it holds with probability at least $1 - q$ that for every $R \in \mathcal{R}$, $||S \cap R|/|S| - \rho_R| \leq \varepsilon$. Take $X = P$, $\mathcal{R} = \{P \cap (Q - t) : t \in \mathbb{R}^2\}$ and $q = n^{-r}$. Notice that $\rho_{P \cap (Q - t)} = \mu(t)$ and $|S \cap (P \cap (Q - t))|/|S| = |S \cap (Q - t)|/|S| = |(S + t) \cap Q|/|S| = \mu_S(t)$. By Lemma 2.1, $\nu = O(\log n)$, and thus the desired bound follows. \square

2.1 Depth variation in an arrangement

Define a *shape* to be a compact subset of \mathbb{R}^2 . Given a shape F , ∂F denotes its boundary. Let \mathcal{F} be a family of shapes. Among the vertices of the arrangement of \mathcal{F} , we use $V(\mathcal{F})$ to denote the subset that consists of intersections between boundaries of distinct shapes in \mathcal{F} . The *depth* of a point t in the arrangement of \mathcal{F} is $\text{depth}(t, \mathcal{F}) = |\{F \in \mathcal{F} : t \in F\}|$. We call the family \mathcal{F} *simple* if the following conditions are satisfied.

- \mathcal{F} contains a finite number of shapes.
- For every shape $F \in \mathcal{F}$, ∂F is a non-self-crossing closed curve. (Recall that polygons with multiple connected components and holes can be accommodated by constructing appropriate spanning trees as discussed in the overview.)
- For every pair of distinct shapes $F_1, F_2 \in \mathcal{F}$, if $F_1 \cap F_2 \neq \emptyset$, then $\partial F_1 \cap \partial F_2$ is a finite point set and $|\partial F_1 \cap \partial F_2| \geq 2$. For every triple of distinct shapes in \mathcal{F} , the common intersection of their boundaries is empty.




Figure 1: Copies of Q in \mathcal{T} are shown. The white dot is an intersection in $\partial F_1 \cap \partial F_2$. The black dots on ∂F_1 are the t_F^+ 's in V^+ . The grey dots on ∂F_1 are the t_F^- 's in V^- . Let $k = 2$. Since the depth of the white dot is $4 = k + 2$, U^+ consists of the first $k = 2$ black dots in anticlockwise order around ∂F_1 from the white dot, and U^- consists of the first $k = 2$ grey dots in clockwise order around ∂F_1 from the white dot.

Given a deep enough vertex t in the arrangement of \mathcal{F} , we show that there are plenty vertices with depth similar to that of t . Lemma 2.4 shows a weaker version of this result which will be used to prove a stronger version—Lemma 2.5. We need the following technical result.

Lemma 2.3. *Let \mathcal{F} be a simple family of shapes in \mathbb{R}^2 . Let F and F' be two distinct shapes in \mathcal{F} . If one can draw a curve from $t \in F$ to $t' \in F'$ that crosses the boundaries of exactly j shapes in $\mathcal{F} \setminus \{F, F'\}$, then $|\text{depth}(t, \mathcal{F}) - \text{depth}(t', \mathcal{F})| \leq j + 1$.*

Proof. By assumption, we can draw a curve that crosses the boundaries of exactly j shapes in $\mathcal{F} \setminus \{F, F'\}$. Therefore, $|\text{depth}(t, \mathcal{F} \setminus \{F, F'\}) - \text{depth}(t', \mathcal{F} \setminus \{F, F'\})| \leq j$. Since $t \in F$ and $t' \in F'$, we have the identities $\text{depth}(t, \mathcal{F}) = \text{depth}(t, \mathcal{F} \setminus \{F, F'\}) + \text{depth}(t, \{F'\}) + 1$ and $\text{depth}(t', \mathcal{F}) = \text{depth}(t', \mathcal{F} \setminus \{F, F'\}) + \text{depth}(t', \{F\}) + 1$. Thus, $|\text{depth}(t, \mathcal{F}) - \text{depth}(t', \mathcal{F})| \leq |\text{depth}(t, \mathcal{F} \setminus \{F, F'\}) - \text{depth}(t', \mathcal{F} \setminus \{F, F'\})| + |\text{depth}(t, \{F'\}) - \text{depth}(t', \{F\})| \leq j + 1$. \square

Lemma 2.4. *Let t_0 be a vertex in $V(\mathcal{F})$, where \mathcal{F} is a simple family of shapes in \mathbb{R}^2 . Let F_1 and F_2 be the shapes in \mathcal{F} such that $t_0 \in \partial F_1 \cap \partial F_2$. If $\text{depth}(t_0, \mathcal{F}) \geq k + 2$ for some integer $k \geq 0$, then there exists a subset $U \subseteq V(\mathcal{F})$ in $\partial F_1 \setminus \partial F_2$ such that $|U| = 2k$ and for every $t \in U$, $|\text{depth}(t, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq k$.*

Proof. For each shape $F \in \mathcal{F} \setminus \{F_1, F_2\}$ that intersects F_1 , there are at least two points in $\partial F \cap \partial F_1$ because \mathcal{F} is a simple family, and let t_F^+ and t_F^- be the first and last intersection points in $\partial F \cap \partial F_1$ with respect to the anticlockwise traversal of ∂F_1 starting from t_0 . Let $V^+ = \{t_F^+ : F \in \mathcal{F} \setminus \{F_1, F_2\} \wedge F \cap F_1 \neq \emptyset\}$. Let $V^- = \{t_F^- : F \in \mathcal{F} \setminus \{F_1, F_2\} \wedge F \cap F_1 \neq \emptyset\}$. Figure 1 shows an example.

Since $\text{depth}(t_0, \mathcal{F}) \geq k + 2$, at least $k + 2$ shapes in \mathcal{F} contain t_0 (including F_1 and F_2). Thus, at least k shapes in $\mathcal{F} \setminus \{F_1, F_2\}$ intersect F_1 , including those containing t_0 and possibly more. By the simplicity of \mathcal{F} again, for every $F \in \mathcal{F}$ that intersects F_1 , neither t_F^+ nor t_F^- lies in the boundary of any shape in $\mathcal{F} \setminus \{F, F_1\}$. Therefore, neither t_F^+ nor t_F^- can be $t_{F'}^+$ or $t_{F'}^-$ for some $F' \neq F$. It follows that $|V^+| = |V^-| \geq k$, and $V^+ \cap V^- = \emptyset$.

Order the points in V^+ according to the anticlockwise traversal of ∂F_1 starting from t_0 . Collect the first k points in V^+ in this order and put them in the ordered list U^+ . For the i th

point $t_F^+ \in U^+$, the anticlockwise traversal of ∂F_1 from t_0 to t_F^+ crosses the boundaries of $i - 1$ shapes in $\mathcal{F} \setminus \{F, F_2\}$. Therefore, Lemma 2.3 implies that $|\text{depth}(t_F^+, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq i \leq k$. Symmetrically, we order the points in V^- according to the clockwise traversal of ∂F_1 starting from t_0 , collect the first k points in V^- in this order, and put them in the ordered set U^- . Then, $|\text{depth}(t_F^-, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq k$ for every $t_F^- \in U^-$. The sets U^+ and U^- are disjoint as $V^+ \cap V^- = \emptyset$. Thus, the set $U = U^+ \cup U^-$ satisfies the lemma. \square

Lemma 2.4 shows that for every vertex t , there are at least $2k$ vertices with depth similar to that of t . The next lemma improves this bound to $(k + 1)(k + 2)/2$.

Lemma 2.5. *Let t_0 be a vertex in $V(\mathcal{F})$, where \mathcal{F} is a simple family of shapes in \mathbb{R}^2 . If $\text{depth}(t_0, \mathcal{F}) \geq k + 2$ for some integer $k \geq 0$, then there exists a subset $U \subseteq V(\mathcal{F})$ such that $|U| = (k + 1)(k + 2)/2$ and for every $t \in U$, $|\text{depth}(t, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq k$.*

Proof. Let F_0 be a shape in \mathcal{F} whose boundary contains t_0 . Order $V(\mathcal{F}) \cap \partial F_0$ in anticlockwise order along ∂F_0 starting from t_0 . Since $\text{depth}(t_0, \mathcal{F}) \geq k + 2$, at least $k + 2$ shapes in \mathcal{F} contain t_0 . Therefore, at least $k + 1$ shapes in $\mathcal{F} \setminus \{F_0\}$ intersect F_0 , including those containing t_0 and possibly more. For every shape $F \in \mathcal{F} \setminus \{F_0\}$ that intersects F_0 , let t_F denote the first intersection point between ∂F and ∂F_0 in anticlockwise order along ∂F_0 starting from t_0 . Let $t_{F_1}, \dots, t_{F_{k+1}}$ be the first $k + 1$ such intersection points (including t_0) in anticlockwise order around ∂F_0 starting from t_0 . Thus, $t_{F_1} = t_0$.

We claim that for $i \in [1, k + 1]$, $|\text{depth}(t_{F_i}, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq i - 1$. The claim is trivially true for $i = 1$ as $t_{F_1} = t_0$. For $i \in [2, k + 1]$, the anticlockwise traversal of ∂F_0 from $t_0 = t_{F_1}$ to t_{F_i} crosses the boundaries of $i - 2$ shapes in $\mathcal{F} \setminus \{F_1, F_i\}$. By Lemma 2.3, $|\text{depth}(t_{F_i}, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq i - 1$.

Our claim implies that $\text{depth}(t_{F_i}, \mathcal{F}) \geq \text{depth}(t_0, \mathcal{F}) - (i - 1) \geq k - i + 3$. Applying Lemma 2.4 to t_{F_i} gives a subset $U_i \subseteq V(\mathcal{F}) \cap (\partial F_i \setminus \partial F_0)$ such that $|U_i| = 2(k - i + 1)$ and for every $t \in U_i$, $|\text{depth}(t, \mathcal{F}) - \text{depth}(t_{F_i}, \mathcal{F})| \leq k - i + 1$. Then, for every $t \in U_i$, $|\text{depth}(t, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq |\text{depth}(t, \mathcal{F}) - \text{depth}(t_{F_i}, \mathcal{F})| + |\text{depth}(t_{F_i}, \mathcal{F}) - \text{depth}(t_0, \mathcal{F})| \leq (k - i + 1) + (i - 1) = k$.

Let $U = \bigcup_{i=1}^{k+1} U_i \cup \{t_{F_i} : i \in [1, k + 1]\}$. The sets $\bigcup_{i=1}^{k+1} U_i$ and $\{t_{F_i} : i \in [1, k + 1]\}$ are disjoint because every U_i is disjoint from ∂F_0 , but every t_{F_i} belongs to ∂F_0 . As \mathcal{F} is a simple family, each point in $\bigcup_{i=1}^{k+1} U_i$ belongs to the intersection of exactly two shapes' boundaries. It follows that each point in $\bigcup_{i=1}^{k+1} U_i$ is contained in at most two distinct U_i 's. Hence, $|U| \geq (k + 1) + \frac{1}{2} \sum_{i=1}^{k+1} 2(k - i + 1) = (k + 1)(k + 2)/2$. \square

2.2 Algorithm for translation

Recall that $\mathcal{T} = \{Q - s : s \in S\}$. Because S is a random sample, it is clear that \mathcal{T} is a simple family with probability 1. We introduce a procedure DEPTHSAMPLE in Algorithm 1 to sample a set $W \subseteq V(\mathcal{T})$. Lemma 2.6 proves its correctness and probability bound.

Lemma 2.6. *Let $\varepsilon_0 \in (0, 1)$. Let $\mathring{t} = \arg \max_{t \in \mathbb{R}^2} \mu(t)$. Let $\tilde{t} = \arg \max_{t \in W} \mu_S(t)$, where W is the point set returned by DEPTHSAMPLE. If $W = \emptyset$, let $\max_{t \in W} \mu_S(t) = 0$ and let $\tilde{t} = (0, 0)$. There exists a constant $c \geq 3$ such that if $|S| \geq c\varepsilon_0^{-2} \ln n$, then:*

- (i) $\Pr(\mu_S(\tilde{t}) \geq \mu_S(\mathring{t}) - \varepsilon_0) \geq 1 - n^{-r}$, and
- (ii) $\Pr(\mu(\tilde{t}) \geq \mu(\mathring{t}) - 3\varepsilon_0) \geq 1 - 2n^{-r}$.

Proof. Let $k = \lfloor \varepsilon_0 |S| \rfloor - 2$. Since $c \geq 3$, $\varepsilon_0 |S| \geq 3$ and so $k \geq 1$. If $k > \text{depth}(\mathring{t}, \mathcal{T}) - 2$, then $\mu_S(\mathring{t}) = \text{depth}(\mathring{t}, \mathcal{T})/|S| < (k + 2)/|S| \leq \varepsilon_0$. So $\max_{t \in W} \mu_S(t) \geq 0 > \mu_S(\mathring{t}) - \varepsilon_0$, satisfying (i).

Algorithm 1 Construct a random sample W of the vertex set $V(\mathcal{T})$

```
1: function DEPTHSAMPLE( $n, S, Q, \varepsilon_0$ )
2:    $w \leftarrow \lceil (9rn^2\varepsilon_0^{-2} \ln n)/2 \rceil$ , where  $\varepsilon_0 \in (0, 1)$  and  $r$  is some positive value fixed a priori.
3:   initialize  $W$  to be an empty list.
4:   for  $i = 1$  to  $w$  do
5:     pick distinct pair  $s_i, s'_i$  uniformly at random from  $S$ .
6:     pick an edge  $e_i$  uniformly at random from the edge set of  $Q - s_i$ .
7:     pick an edge  $e'_i$  uniformly at random from the edge set of  $Q - s'_i$ .
8:     if  $e_i$  and  $e'_i$  intersect then
9:       compute the intersection point  $t_i$  of  $e_i$  and  $e'_i$ .
10:      append  $t_i$  to  $W$ .
11:     end if
12:   end for
13:   return  $W$ 
14: end function
```

Suppose that $k \leq \text{depth}(\hat{t}, \mathcal{T}) - 2$. Let \hat{t} be a vertex of the cell containing \hat{t} in the arrangement of \mathcal{T} . Then, $\text{depth}(\hat{t}, \mathcal{T}) \geq \text{depth}(\hat{t}, \mathcal{T}) \geq k + 2$. Since $\text{depth}(\hat{t}, \mathcal{T}) > 1$, we can choose \hat{t} to be a vertex from $V(\mathcal{T})$. By Lemma 2.5, there exists a subset $U \subseteq V(\mathcal{T})$ such that $|U| = (k + 1)(k + 2)/2$ and for every $t \in U$, $|\text{depth}(t, \mathcal{T}) - \text{depth}(\hat{t}, \mathcal{T})| \leq k$. Let E_i be the event that the edges e and e' sampled in the i th iteration in DEPTHSAMPLE intersect, and $t_i = e \cap e' \in U$. Observe that if $t_i \in U$ for some i between 1 and $\lceil (9rn^2\varepsilon_0^{-2} \ln n)/2 \rceil$, then the lemma is satisfied because $\mu_S(t_i) \geq \mu_S(\hat{t}) - |\mu_S(t_i) - \mu_S(\hat{t})| = (\text{depth}(\hat{t}, \mathcal{T}) - |\text{depth}(t_i, \mathcal{T}) - \text{depth}(\hat{t}, \mathcal{T})|)/|S| \geq (\text{depth}(\hat{t}, \mathcal{T}) - k)/|S| \geq \mu_S(\hat{t}) - \varepsilon_0$.

It remains to bound the probability that E_i does not happen for all i . There are no more than $n^2|S|^2/2$ pairs of edges from the shapes in \mathcal{T} . Among these pairs, $|U|$ of them yield intersection points in U . Thus, $\Pr(E_i) \geq 2|U|/(n^2|S|^2) = (k + 1)(k + 2)/(n^2|S|^2) \geq (\varepsilon_0|S| - 2)(\varepsilon_0|S| - 1)/(n^2|S|^2) \geq 2\varepsilon_0^2/(9n^2)$. The last inequality follows from the fact that $|S| \geq 3/\varepsilon_0$. Hence, the probability that E_i does not happen for all i is at most $(1 - 2\varepsilon_0^2/(9n^2))^{(9rn^2\varepsilon_0^{-2} \ln n)/2} \leq e^{-r \ln n} = n^{-r}$. This completes the proof of (i).

By Lemma 2.2, with probability at least $1 - n^{-r}$, $|\mu_S(\hat{t}) - \mu(\hat{t})| \leq \varepsilon_0$ and $|\mu_S(\tilde{t}) - \mu(\tilde{t})| \leq \varepsilon_0$. Then, (i) implies that $\mu(\tilde{t}) \geq \mu(\hat{t}) - 3\varepsilon_0$ with probability at least $1 - 2n^{-r}$. This proves (ii). \square

The shape matching algorithm under translation works as follows. Call DEPTHSAMPLE with $\varepsilon_0 = \varepsilon/3$ to obtain the set W . By Lemma 2.6(ii), it suffices to count $|(S + t) \cap Q|$ for every $t \in W$ and report $\tilde{t} = \text{argmax}_{t \in W} \mu_S(t)$. Partition $\mathcal{T} = \{Q - s : s \in S\}$ into m subfamilies \mathcal{T}_i each consisting of $|S|/m$ copies of Q , where $m = \sqrt{\log n}/\varepsilon$. Compute the arrangement of \mathcal{T}_i and the depths of its cells by a plane sweep. Also, compute a point location data structure of the arrangement of \mathcal{T}_i . For every $t \in W$ and every $i \in [1, m]$, we locate t in the arrangement of \mathcal{T}_i to obtain the depth $d_i(t)$ of t . As a result, $\mu_S(t) = \sum_{i=1}^m d_i(t)/|S|$ for every $t \in W$, and then we select $\text{argmax}_{t \in W} \mu_S(t)$.

Theorem 2.7. *Let P and Q be two polygonal shapes with a total of n vertices. Let opt be the maximum overlap of P and Q under translation. For every $\varepsilon \in (0, 1)$, one can compute a translation \tilde{t} such that $\text{area}((P + \tilde{t}) \cap Q) \geq \text{opt} - \varepsilon \cdot \text{area}(P)$ with probability $1 - n^{-O(1)}$ in $O(n^2\varepsilon^{-3} \log^{1.5} n \log(n/\varepsilon))$ time.*

Proof. The correctness and probability bound follow from Lemma 2.6(ii) and the setting of $\varepsilon_0 = \varepsilon/3$. To construct S , we triangulate P and then sample points from

the triangles with probabilities proportional to their areas. This takes $O(n \log n + |S|)$ time. Calling DEPTHSAMPLE takes $O(n^2 \varepsilon^{-2} \log n)$ time. Processing the arrangement \mathcal{A}_i of \mathcal{T}_i takes $O(|\mathcal{A}_i| \log |\mathcal{A}_i|) = O(n^2 |S|^2 m^{-2} \log(n|S|/m))$ time. Then, the computation of $d_i(t)$ for all $t \in W$ takes $O(|W| \log(n|S|/m))$ time. The total running time is thus $O(n^2 \varepsilon^{-4} m^{-1} \log^2 n \log(n \log n / (\varepsilon^2 m)) + mn^2 \varepsilon^{-2} \log n \log(n \log n / (\varepsilon^2 m)))$. Setting $m = \sqrt{\log n} / \varepsilon$ gives an $O(n^2 \varepsilon^{-3} \log^{1.5} n \log(n/\varepsilon))$ bound. \square

If Q is convex, we can reduce the running time further because an arrangement of translates of Q has a lower complexity.

Theorem 2.8. *Let P be a polygonal shape and let Q be a convex polygon with a total of n vertices. Let opt be the maximum overlap of P and Q under translation. For every $\varepsilon \in (0, 1)$, one can compute in $O(n \log n + \varepsilon^{-3} \log^{2.5} n \log \frac{\log n}{\varepsilon})$ time a translation \tilde{t} such that $\text{area}((P + \tilde{t}) \cap Q) \geq \text{opt} - \varepsilon \cdot \text{area}(P)$ with probability $1 - n^{-O(1)}$.*

Proof. Since Q is convex, for every distinct pair $s, s' \in S$, either $(Q - s) \cap (Q - s') = \emptyset$ or the boundaries of $Q - s$ and $Q - s'$ intersect at exactly two points. It means that the arrangement of $\{Q - s : s \in S\}$ has at most $|S|^2$ vertices that are intersections of boundaries of shapes in $\{Q - s : s \in S\}$. This is a significant reduction from the $n^2 |S|^2 / 2$ bound when Q is a general polygonal shape. We introduce changes in DEPTHSAMPLE and the handling of \mathcal{T}_i to exploit this property.

We make two changes in DEPTHSAMPLE. First, after sampling distinct $s, s' \in S$, we do not sample edges from $Q - s$ and $Q - s'$. Instead, we directly compute the intersections between the boundaries of $Q - s$ and $Q - s'$ in $O(\log n)$ time [9], and if the intersection is non-empty, randomly return one of the two intersections. Second, DEPTHSAMPLE iterates $w = \lceil 9r\varepsilon^{-2} \ln n \rceil$ times instead of $\lceil (9rn^2\varepsilon^{-2} \ln n) / 2 \rceil$ times. The number w of iterations is only needed in bounding $\prod_{i=1}^w (1 - \Pr(E_i))$ from above towards the end of the proof of Lemma 2.6. Since we now sample from at most $|S|^2$ arrangement vertices, one can adapt the derivation for $\Pr(E_i)$ straightforwardly and obtain $\Pr(E_i) \geq \varepsilon^2 / 9$. Then, we can show that $\prod_{i=1}^w (1 - \Pr(E_i)) \leq n^{-r}$ as before. We conclude that the running time of DEPTHSAMPLE is reduced to $O(\varepsilon^{-2} \log^2 n)$.

We also change the handling of \mathcal{T}_i . Instead of building a point location data structure for \mathcal{T}_i , we locate all points in W in a batch by a plane sweep over \mathcal{T}_i . We sweep a line from left to right over \mathcal{T}_i . Since Q is convex, its boundary can be partitioned into upper and lower convex chains by splitting at the leftmost and rightmost vertices. The sweep events include the endpoints of the convex chains, the boundary intersections, and the points in W . The sweep status structure stores the chains that currently intersect the sweep line. Given two chains whose intersections with the sweep line are adjacent, we can compute the intersections between these two chains in $O(\log n)$ time [9]. Therefore, the time needed becomes $O((|S|^2 m^{-2} + |W|) \log(|S|^2 m^{-2} + |W|) \log n)$.

As a result, the total running time is now reduced to

$$\begin{aligned} & O(n \log n + \varepsilon^{-2} \log^2 n) + \\ & O(|S|^2 m^{-1} \log(|S|^2 m^{-2} + |W|) \log n + m|W| \log(|S|^2 m^{-2} + |W|) \log n) \\ = & O(n \log n + \varepsilon^{-4} m^{-1} \log^3 n \log(|S|^2 m^{-2} + |W|) + m\varepsilon^{-2} \log^2 n \log(|S|^2 m^{-2} + |W|)). \end{aligned}$$

Setting $m = \sqrt{\log n} / \varepsilon$ gives an $O(n \log n + \varepsilon^{-3} \log^{2.5} n \log \frac{\log n}{\varepsilon})$ running time. \square

3 Matching under rigid motion

We use \mathbb{M} to denote the configuration space $\mathbb{R}^2 \times [0, 2\pi)$ of rigid motion. For every subset $X \subseteq \mathbb{R}^2$ and every $\theta \in [0, 2\pi)$, X_θ denotes the rotated copy of X around the origin by angle




Figure 2: Let Q be a quadrilateral as shown in the figure. Two translates of Q are rotated and moved upward simultaneously, forming two twisted pillars. The dashed-dotted lines at the cross-section at $\theta = \phi$ form the set \mathcal{L}^ϕ . The line ℓ at $\theta = 0$ sweeps out the shaded surface ℓ_* as it rotates and moves upward to ℓ_ϕ . As θ varies, the two supporting lines ℓ and ℓ' from the same copy of Q stays concurrent, and $\ell_\theta \cap \ell'_\theta$ is equal to $(\ell \cap \ell')_\theta$, sweeping out the bold curve shown.

θ in the anticlockwise direction. For every $(t, \theta) \in \mathbb{M}$, the corresponding rigid motion rotates Q around the origin by θ in the anticlockwise direction, and then translates P by t . Define $\mu(t, \theta) = \text{area}((P + t) \cap Q_\theta) / \text{area}(P)$. We will work with a point set S sampled uniformly at random from P , and we define $\mu_S(t, \theta) = |(S + t) \cap Q_\theta| / |S|$.

We use L to denote the set of distinct supporting lines of the edges of Q . We do not make any general position assumption, therefore, L may contain parallel lines and three or more lines in L may have a common intersection. However, although two edges of Q may have the same supporting line, L does not store any duplicate. Using the random sample $S \subset P$, we define $\mathcal{L} = \{\ell - s : \ell \in L \wedge s \in S\}$, which is a refinement of $\mathcal{T} = \{Q - s : s \in S\}$.

If we rotate L and Q around the origin by angle θ in the anticlockwise direction, we get $\mathcal{L}^\theta = \{\ell_\theta - s : \ell \in L \wedge s \in S\}$ and $\mathcal{T}^\theta = \{Q_\theta - s : s \in S\}$. Notice that \mathcal{L}^θ is not obtained by rotating \mathcal{L} around the origin by angle θ . Similarly, \mathcal{T}^θ is not obtained by rotating \mathcal{T} . For every $t \in \mathbb{R}^2$, define the depth of t in the arrangements of \mathcal{T}^θ and \mathcal{L}^θ to be $\text{depth}(t, \mathcal{T}^\theta) = \text{depth}(t, \mathcal{L}^\theta) = |\{Q_\theta - s \in \mathcal{T}^\theta : t \in Q_\theta - s\}|$.

Treat the θ -axis of \mathbb{M} as the vertical axis. For every subset $X \subset \mathbb{M}$ and every $s \in \mathbb{R}^2$, we use $X(s)$ as the shorthand for $X - (s, 0)$, i.e. translate X by the vector $(-s, 0)$. Define $Q_* = \{(x, y, \theta) \in \mathbb{M} : \theta \in [0, 2\pi) \wedge (x, y) \in Q_\theta\}$, which looks like a twisted pillar (Figure 2). Define \mathcal{T}^* to be the set of twisted pillars $\{Q_*(s) : s \in S\}$ obtained by sliding Q_* horizontally by different translations in S . For every $\ell \in L$, define $\ell_* = \{(x, y, \theta) \in \mathbb{M} : \theta \in [0, 2\pi) \wedge (x, y) \in \ell_\theta\}$, which is the curved surface swept by ℓ as we rotate it and slide it vertically. Sliding the surfaces ℓ_* by different translations in S gives a collection of surfaces $\mathcal{L}^* = \{\ell_*(s) : \ell \in L \wedge s \in S\}$. Notice that \mathcal{L}^* is a refinement of \mathcal{T}^* . For every $(t, \theta) \in \mathbb{M}$, define the depth of (t, θ) in the arrangements of \mathcal{T}^* and \mathcal{L}^* to be $\text{depth}((t, \theta), \mathcal{T}^*) = \text{depth}((t, \theta), \mathcal{L}^*) = |\{Q_*(s) \in \mathcal{T}^* : (t, \theta) \in Q_*(s)\}|$.

$V(\mathcal{T}^\theta)$ is the subset of vertices in the arrangement of \mathcal{T}^θ that are intersections of boundaries of distinct shapes in \mathcal{T}^θ . Notice that \mathcal{T}^θ and \mathcal{L}^θ are horizontal cross-sections of \mathcal{T}^* and \mathcal{L}^* , respectively, at the value θ . Hence, $\text{depth}((t, \theta), \mathcal{T}^*) = \text{depth}((t, \theta), \mathcal{L}^*) = \text{depth}(t, \mathcal{T}^\theta) = \text{depth}(t, \mathcal{L}^\theta)$.

For every rigid motion $(t, \theta) \in \mathbb{M}$, $s + t \in Q_\theta \iff t \in Q_\theta - s \iff (t, \theta) \in Q_*(s)$. Thus, the deepest vertex in the arrangement of \mathcal{L}^* maximizes $\mu_S(t, \theta)$. The next result bounds the VC-dimension of an appropriate range space in the rigid motion case as in Lemma 2.1 in the translation case.

Lemma 3.1. *Let X be a subset of \mathbb{R}^2 . Let $\mathcal{R} = \{X \cap (Q_\theta - t) : t \in \mathbb{R}^2 \wedge \theta \in [0, 2\pi)\}$. The range space (X, \mathcal{R}) has VC-dimension $O(\log n)$.*

Proof. Let Z be a finite subset of points in X that is shattered by \mathcal{R} . It suffices to prove an $O(\log n)$ bound on $|Z|$. Let \mathcal{A} be the arrangement of $\{\ell_*(z) : \ell \in L \wedge z \in Z\}$, which is a refinement of the arrangement of $\{Q_*(z) : z \in Z\}$. Take a line $\ell \in L$. Let its equation be $(a, b) \cdot t = c$, where $t = (t_x, t_y) \in \mathbb{R}^2$ and a, b and c are constants. Then, the equation of ℓ_* is

$$(a \ b) \cdot \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \cdot \begin{pmatrix} t_x \\ t_y \end{pmatrix} = c.$$

If we replace $\cos \theta$ by a variable α and $\sin \theta$ by a variable β , the system becomes $a\alpha t_x + a\beta t_y - b\beta t_x + b\alpha t_y = c$, which is a multivariate polynomial of degree two. As a result, \mathcal{A} is the arrangement of the zero-sets of degree-two polynomials in four variables restricted to the subset that satisfies $\alpha^2 + \beta^2 = 1$. By the result in [5], \mathcal{A} has at most $O(n^3|Z|^3)$ cells. By the same argument in the proof of Lemma 2.1, $2^{|Z|} = O(n^3|Z|^3)$ as Z is shattered by \mathcal{R} . It follows that $|Z| = O(\log n)$. \square

Next, we prove that $\mu_S(t, \theta)$ is a good approximation of $\mu(t, \theta)$, which is analogous to Lemma 2.2 in the translation case.

Lemma 3.2. *For every constant $r > 0$, there exists a constant $c > 0$ such that if S is a set of points sampled uniformly at random from P and $|S| \geq c\varepsilon^{-2} \log_2 n$, then with probability at least $1 - n^{-r}$, $|\mu_S(t, \theta) - \mu(t, \theta)| \leq \varepsilon$ for every $(t, \theta) \in \mathbb{M}$.*

Proof. Take $X = P$, $\mathcal{R} = \{P \cap (Q_\theta - t) : (t, \theta) \in \mathbb{M}\}$ and $q = n^{-r}$. Let ρ_R be the probability that a point drawn from P uniformly at random belongs to R . If $R = P \cap (Q_\theta - t)$, then $\rho_R = \mu(t, \theta)$ because $\mu(t, \theta) = \text{area}((P + t) \cap Q_\theta) / \text{area}(P) = \text{area}(P \cap (Q_\theta - t)) / \text{area}(P)$. Also, $|S \cap R| / |S| = |S \cap (P \cap (Q_\theta - t))| / |S| = |(S + t) \cap Q_\theta| / |S| = \mu_S(t, \theta)$. The VC-dimension of $O(\log n)$ in Lemma 3.1 and the ε -approximation result [12, 13] says that there exists a constant $c_0 > 0$ such that for every $\varepsilon, q \in (0, 1)$, if we draw a subset S uniformly at random from X with $|S| \geq c_0\varepsilon^{-2}(O(\log n) + \ln(1/q))$, then it holds with probability at least $1 - q = 1 - n^{-r}$ that for every $R \in \mathcal{R}$, $||S \cap R| / |S| - \rho_R| \leq \varepsilon$, equivalently, $|\mu_S(t, \theta) - \mu(t, \theta)| \leq \varepsilon$. \square

3.1 Depth variation in the configuration space

We will show that it is possible to sample vertices from the arrangement of \mathcal{L}^* to find one that is approximately deepest. We first prove two technical results in Lemmas 3.3 and 3.4 about the surfaces in \mathcal{L}^* .

Lemma 3.3. *Let ℓ_1 and ℓ_2 be two possibly non-distinct lines in \mathbb{R}^2 . Let s_1 and s_2 be two possibly non-distinct points in \mathbb{R}^2 .*

- (i) *If ℓ_1 and ℓ_2 are not parallel, then $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2)$ is a strictly θ -monotone curve, i.e. any plane orthogonal to the θ -axis intersects the curve at at most one point.*
- (ii) *If ℓ_1 and ℓ_2 are parallel and $s_1 \neq s_2$, then $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2)$ is non-empty for at most two values of θ . At each such value of θ , $\ell_{1,\theta} - s_1 = \ell_{2,\theta} - s_2$. Furthermore, if $\ell_1 = \ell_2$ and $s_1 \neq s_2$, then $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2)$ is non-empty for exactly two values of θ .*

Proof. Consider (i). Since ℓ_1 and ℓ_2 are not parallel, $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ are not parallel for all θ . Thus, they intersect at exactly one point for each $\theta \in [0, 2\pi)$.

Consider (ii). Since ℓ_1 and ℓ_2 are parallel, $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ are parallel for all θ . Thus, $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ intersect if and only if they coincide. Let the equation of ℓ_i be $(a, b) \cdot t = c_i$, where $t = (t_x, t_y) \in \mathbb{R}^2$, and a, b and c_i are constants. The line $\ell_{i,\theta} - s_i$ is always tangent to the circle $C_i : \|t - s_i\| = |c_i|(a^2 + b^2)^{-1/2}$.

If c_1 and c_2 have the same sign, the parallel lines $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ intersect if and only if they coincide at a common outer tangent of C_1 and C_2 . No such common outer tangent exists if C_1 is inside C_2 or vice versa. If $\ell_1 = \ell_2$, then c_1 is equal to c_2 and the circles C_1 and C_2 have the same radius, so they are not nested and have exactly two common outer tangents.

If c_1 and c_2 have different signs, then $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ intersect if and only if they coincide at a common inner tangent of C_1 and C_2 . No such common inner tangent exists if C_1 and C_2 intersect at two points, or if one is nested inside the other. \square

Lemma 3.4. *Let ℓ_1, ℓ_2 and ℓ_3 be three possibly non-distinct lines in \mathbb{R}^2 . Let s_1, s_2 and s_3 be three possibly non-distinct points in \mathbb{R}^2 . For $i \in [1, 3]$, let $s_i = (s_{i,x}, s_{i,y})$, and let the equation of ℓ_i be $(a_i, b_i) \cdot (t_x, t_y) = c_i$, where $(t_x, t_y) \in \mathbb{R}^2$, and a_i, b_i and c_i are constants. If a point $(t_x, t_y, \theta) \in \mathbb{M}$ belongs to $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2) \cap \ell_{3,*}(s_3)$, then θ satisfies the following equation.*

$$\begin{vmatrix} a_1 & b_1 & a_1 s_{1,x} + b_1 s_{1,y} \\ a_2 & b_2 & a_2 s_{2,x} + b_2 s_{2,y} \\ a_3 & b_3 & a_3 s_{3,x} + b_3 s_{3,y} \end{vmatrix} \cos \theta + \begin{vmatrix} a_1 & b_1 & a_1 s_{1,y} - b_1 s_{1,x} \\ a_2 & b_2 & a_2 s_{2,y} - b_2 s_{2,x} \\ a_3 & b_3 & a_3 s_{3,y} - b_3 s_{3,x} \end{vmatrix} \sin \theta = \begin{vmatrix} a_1 & b_1 & c_1 \\ a_2 & b_2 & c_2 \\ a_3 & b_3 & c_3 \end{vmatrix} \quad (1)$$

Furthermore, if ℓ_1 and ℓ_2 are not parallel and s_3 is picked uniformly at random independent from s_1 and s_2 from some subset of \mathbb{R}^2 with positive area, then with probability 1, equation (1) has at most two solutions for $\theta \in [0, 2\pi)$.

Proof. For $i \in [1, 3]$, since $(t, \theta) \in \ell_{i,*}(s_i)$, the point t belongs to $\ell_{i,\theta} - s_i$. Equivalently, $t_{-\theta} \in \ell_i - (s_i)_{-\theta}$. The equation of $\ell_i - (s_i)_{-\theta}$ is given by

$$\begin{aligned} & a_i(t_x + s_{i,x} \cos \theta + s_{i,y} \sin \theta) + b_i(t_y - s_{i,x} \sin \theta + s_{i,y} \cos \theta) = c_i \\ \iff & a_i t_x + b_i t_y - c_i + (a_i s_{i,x} + b_i s_{i,y}) \cos \theta + (a_i s_{i,y} - b_i s_{i,x}) \sin \theta = 0 \end{aligned}$$

Let $t_{-\theta} = (x', y')$. (Recall that $t_{-\theta}$ denotes the point t rotated around the origin by angle $-\theta$ in the anticlockwise direction.) We write these equations in matrix form:

$$\begin{pmatrix} a_1 & b_1 & -c_1 + (a_1 s_{1,x} + b_1 s_{1,y}) \cos \theta + (a_1 s_{1,y} - b_1 s_{1,x}) \sin \theta \\ a_2 & b_2 & -c_2 + (a_2 s_{2,x} + b_2 s_{2,y}) \cos \theta + (a_2 s_{2,y} - b_2 s_{2,x}) \sin \theta \\ a_3 & b_3 & -c_3 + (a_3 s_{3,x} + b_3 s_{3,y}) \cos \theta + (a_3 s_{3,y} - b_3 s_{3,x}) \sin \theta \end{pmatrix} \begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

The square matrix on the left hand side of the equation must not be invertible; otherwise, multiplying its inverse to both sides give $(x', y', 1) = (0, 0, 0)$, a contradiction. Therefore,

$$\begin{vmatrix} a_1 & b_1 & -c_1 + (a_1 s_{1,x} + b_1 s_{1,y}) \cos \theta + (a_1 s_{1,y} - b_1 s_{1,x}) \sin \theta \\ a_2 & b_2 & -c_2 + (a_2 s_{2,x} + b_2 s_{2,y}) \cos \theta + (a_2 s_{2,y} - b_2 s_{2,x}) \sin \theta \\ a_3 & b_3 & -c_3 + (a_3 s_{3,x} + b_3 s_{3,y}) \cos \theta + (a_3 s_{3,y} - b_3 s_{3,x}) \sin \theta \end{vmatrix} = 0$$

The linearity of determinant gives equation (1).

Now, suppose that ℓ_1 and ℓ_2 are not parallel and s_3 is picked uniformly at random independent of s_1 and s_2 from a subset of \mathbb{R}^2 with positive area. To show that equation (1) has at most two solutions for $\theta \in [0, 2\pi)$, it suffices to show that the coefficient of $\cos \theta$ is non-zero. Expanding the coefficient of $\cos \theta$ by its last column, we get

$$\begin{aligned} \begin{vmatrix} a_1 & b_1 & a_1 s_{1,x} + b_1 s_{1,y} \\ a_2 & b_2 & a_2 s_{2,x} + b_2 s_{2,y} \\ a_3 & b_3 & a_3 s_{3,x} + b_3 s_{3,y} \end{vmatrix} &= (a_3 s_{3,x} + b_3 s_{3,y}) \begin{vmatrix} a_1 & b_1 \\ a_2 & b_2 \end{vmatrix} + K \\ &= \begin{vmatrix} a_1 & b_1 \\ a_2 & b_2 \end{vmatrix} (a_3, b_3) \cdot s_3 + K \end{aligned} \quad (2)$$

where K denotes the remaining terms that are independent of s_3 . Since ℓ_1 and ℓ_2 are not parallel, $\begin{vmatrix} a_1 & b_1 \\ a_2 & b_2 \end{vmatrix} \neq 0$. Since (a_3, b_3) is a non-zero vector orthogonal to ℓ_3 , $(a_3, b_3) \neq (0, 0)$. Hence, (2) is non-zero with probability 1 as s_3 is picked uniformly at random from a subset of \mathbb{R}^2 with positive area. \square

We use Lemma 3.3 and 3.4 above to prove some combinatorial properties of the arrangement of \mathcal{L}^* . The first property is that three surfaces in \mathcal{L}^* intersect in at most two points.

Lemma 3.5. *Let $\ell_{i,*}(s_i)$ for $i \in [1, 3]$ be three distinct surfaces in \mathcal{L}^* such that s_1, s_2 and s_3 are not all the same. The common intersection of these surfaces contains at most two points with probability 1.*

Proof. Suppose that ℓ_1, ℓ_2 and ℓ_3 are parallel and s_1, s_2 and s_3 are distinct. By Lemma 3.3(ii), $\ell_{1,*}(s_1)$ and $\ell_{2,*}(s_2)$ intersect at at most two values of θ , and for each such value of θ , $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ coincide. But then, since $\ell_{3,\theta}$ is also parallel to $\ell_{1,\theta}$, $\ell_{3,\theta} - s_3$ does not intersect $\ell_{1,\theta} - s_1$ with probability 1 because s_3 is chosen randomly from P independent of s_1 and s_2 . Thus, the surfaces $\ell_{1,*}(s_1), \ell_{2,*}(s_2)$ and $\ell_{3,*}(s_3)$ have no common intersection in this case with probability 1.

Suppose that ℓ_1, ℓ_2 and ℓ_3 are parallel but s_1, s_2 and s_3 are not all distinct. Assume that $s_1 = s_2$. Since $\ell_{1,*}(s_1)$ and $\ell_{2,*}(s_2)$ are distinct surfaces by assumption, the lines $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ are distinct and parallel for all θ . So the surfaces $\ell_{1,*}(s_1), \ell_{2,*}(s_2)$ and $\ell_{3,*}(s_3)$ have no common intersection in this case.

Suppose that exactly two of the ℓ_i 's are parallel, say ℓ_1 and ℓ_2 . If $s_1 = s_2$, then we can argue as in the previous paragraph that $\ell_{1,*}(s_1), \ell_{2,*}(s_2)$ and $\ell_{3,*}(s_3)$ have no common intersection. If $s_1 \neq s_2$, then by Lemma 3.3 (ii), $\ell_{1,*}(s_1)$ and $\ell_{2,*}(s_2)$ intersect at at most two values of θ . For each such value of θ , $\ell_{3,\theta}$ and $\ell_{1,\theta}$ are non-parallel as ℓ_1 and ℓ_3 are non-parallel, so $\ell_{3,\theta} - s_3$ intersects $\ell_{1,\theta} - s_1$ at exactly one point.

Suppose that no two ℓ_i 's are parallel. By Lemma 3.4, the three surfaces intersect at at most two values of θ with probability 1. At each such value of θ , the intersection is a single point because $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2)$ is a strictly θ -monotone curve by Lemma 3.3(i). \square

Let $V(\mathcal{L}^*)$ denote the subset of the vertices of the arrangement of \mathcal{L}^* such that each vertex in $V(\mathcal{L}^*)$ lies in the common intersection of three distinct surfaces $\ell_{i,*}(s_i)$, where $i \in [1, 3]$, for some $\ell_1, \ell_2, \ell_3 \in L$ and $s_1, s_2, s_3 \in S$ such that s_1, s_2 and s_3 are distinct. Although a vertex in $V(\mathcal{L}^*)$ may belong to more than three surfaces in \mathcal{L}^* by definition, we show that this happens with probability zero.

Lemma 3.6. *For every $(t, \theta) \in \mathbb{M}$, if $(t, \theta) \in V(\mathcal{L}^*)$, then with probability 1, (t, θ) belongs to exactly three surfaces in \mathcal{L}^* .*

Proof. Take a vertex $(t_0, \theta_0) \in V(\mathcal{L}^*)$. By definition, (t_0, θ_0) belongs to three distinct surfaces $\ell_{i,*}(s_i)$, where $i \in [1, 3]$, for some distinct $s_1, s_2, s_3 \in S$. Consider any surface $\ell_{4,*}(s_4) \in \mathcal{L}^* \setminus \{\ell_{1,*}(s_1), \ell_{2,*}(s_2), \ell_{3,*}(s_3)\}$. Since s_1, s_2 and s_3 are distinct, s_4 must be different from at least two of them. Assume that $s_4 \notin \{s_2, s_3\}$. By Lemma 3.5, $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2) \cap \ell_{4,*}(s_4)$ contains at most two intersection points. For each such intersection point (t, θ) , the line $\ell_{3,\theta} - s_3$ does not contain t with probability 1 because s_3 is randomly chosen from P independent of s_1, s_2 and s_4 . Hence, $(t_0, \theta_0) \in \bigcap_{i=1}^4 \ell_{i,*}(s_i)$ with probability zero. \square

Given a deep point in the arrangement of \mathcal{L}^* , we show that there are plenty vertices in $V(\mathcal{L}^*)$ nearby with similar depths.

Lemma 3.7. *Let θ_0 be a value in $[0, 2\pi)$ such that \mathcal{T}^{θ_0} is a simple family of shapes. Let t_0 be a vertex in $V(\mathcal{T}^{\theta_0})$. If $\text{depth}(t_0, \mathcal{T}^{\theta_0}) \geq k + 2$ for some integer $k \geq 0$, then with probability 1, there exists a subset $U \subseteq V(\mathcal{L}^*)$ such that $|U| \geq k(k+1)(k+2)/3$ and for every $(t, \theta) \in U$, $|\text{depth}((t, \theta), \mathcal{L}^*) - \text{depth}(t_0, \mathcal{T}^{\theta_0})| \leq 2k + 2$.*

Proof. By Lemma 2.5, there is a set U_0 of $(k+1)(k+2)/2$ vertices in $V(\mathcal{T}^{\theta_0})$ such that for every $t \in U_0$, $|\text{depth}(t, \mathcal{T}^{\theta_0}) - \text{depth}(t_0, \mathcal{T}^{\theta_0})| \leq k$.

Take a vertex $t \in U_0$. Assume that $t \in \partial(Q_{\theta_0} - s_1) \cap \partial(Q_{\theta_0} - s_2)$ for two distinct sample points $s_1, s_2 \in S$. Let $\ell_{1, \theta_0} - s_1$ and $\ell_{2, \theta_0} - s_2$ be the supporting lines of the two edges of $Q_{\theta_0} - s_1$ and $Q_{\theta_0} - s_2$ that intersect at t , respectively. Since ℓ_{1, θ_0} and ℓ_{2, θ_0} are not parallel, neither are ℓ_1 and ℓ_2 . By Lemma 3.3(i), $\ell_{1,*}(s_1) \cap \ell_{2,*}(s_2)$ is a strictly θ -monotone curve ξ_t in \mathbb{M} .

For each $s \in S \setminus \{s_1, s_2\}$, the line $\ell_{1, \theta} - s$ is parallel to $\ell_{1, \theta_0} - s_1$ for all θ . Furthermore, Lemma 3.3(ii) implies that $\ell_{1, \theta} - s$ and $\ell_{1, \theta_0} - s_1$ must coincide at exactly two values of θ . Equivalently, the surface $\ell_{1,*}(s)$ must intersect the curve ξ_t at exactly two points. Since s_1, s_2 and s are distinct, these two intersection points belong to $V(\mathcal{L}^*)$. Similarly, the surface $\ell_{2,*}(s)$ also intersects the curve ξ_t at exactly two intersection points in $V(\mathcal{L}^*)$. By Lemma 3.6, with probability 1, no intersection point in $\ell_{1,*}(s) \cap \xi_t$ or $\ell_{2,*}(s) \cap \xi_t$ is contained in a fourth surface. As a result, ξ_t contains at least $4|S| - 8$ vertices in $V(\mathcal{L}^*)$ that are the intersections between ξ_t and the supporting surfaces of $Q_*(s)$, where $s \in S \setminus \{s_1, s_2\}$. Observe that these vertices on ξ_t contain as a subset the intersections between ξ_t and the boundaries of copies of Q_* in $\mathcal{T}^* \setminus \{Q_*(s_1), Q_*(s_2)\}$.

Walk upward from (t, θ_0) along ξ_t . Collect a vertex in $V(\mathcal{L}^*)$ whenever we meet one. When passing such a vertex, our depth in $\mathcal{T}^* \setminus \{Q_*(s_1), Q_*(s_2)\}$ may increment or decrement. Thus, if we collect k such vertices, our depth in \mathcal{T}^* at any of the vertex collected differs from $\text{depth}(t, \mathcal{T}^{\theta_0})$ by at most $k + 2$, and hence from $\text{depth}(t_0, \mathcal{T}^{\theta_0})$ by at most $2k + 2$. An exception may occur if θ increases to 2π before k vertices are collected; in this case, the upward traversal along ξ_t wraps around to $\theta = 0$ and continue from there. Symmetrically, we walk downward from (t, θ_0) along ξ_t to collect another k vertices in $V(\mathcal{L}^*)$. As $|S| \geq \text{depth}(t_0, \mathcal{T}^{\theta_0}) \geq k + 2$, we get $4|S| - 8 \geq 4k$. Therefore, the bidirectional traversals of ξ_t collect exactly $2k$ vertices.

Let U be the set of all vertices collected this way along ξ_t over all $t \in U_0$. By Lemma 3.6, each vertex in $V(\mathcal{L}^*)$ is at the intersections of exactly three surfaces with probability 1, so it is at the intersection of three curves, each formed by a pair from the triple of surfaces. It follows that when we collect vertices along ξ_t over all $t \in U_0$, a vertex can be collected up to three times. Hence, the total number of distinct vertices identified is at least $2k|U_0|/3 = k(k+1)(k+2)/3$. \square

3.2 Algorithm for rigid motion

The algorithm for rigid motion resembles the one that allows translation only. We introduce a procedure RIGIDDEPTHSAMPLE in Algorithm 2 to sample a set W of vertices from $V(\mathcal{L}^*)$. Following the analysis of Lemma 2.6, we obtain Lemma 3.8 below.

Lemma 3.8. *Let ε_0 be a value in $(0, 1)$. Let $(\mathring{t}, \mathring{\theta}) = \arg \max_{(t, \theta) \in \mathbb{M}} \mu(t, \theta)$. Let $(\tilde{t}, \tilde{\theta}) = \arg \max_{(t, \theta) \in W} \mu_S(t, \theta)$, where W is the point set returned by RIGIDDEPTHSAMPLE. If $W = \emptyset$, let $\max_{(t, \theta) \in W} \mu_S(t, \theta) = 0$ and let $(\tilde{t}, \tilde{\theta}) = (0, 0, 0)$. There exists a constant $c \geq 6$ such that if $|S| \geq c\varepsilon_0^{-2} \ln n$, then:*

- (i) $\Pr(\mu_S(\tilde{t}, \tilde{\theta}) \geq \mu_S(\mathring{t}, \mathring{\theta}) - \varepsilon_0) \geq 1 - n^{-r}$, and
- (ii) $\Pr(\mu(\tilde{t}, \tilde{\theta}) \geq \mu(\mathring{t}, \mathring{\theta}) - 3\varepsilon_0) \geq 1 - 2n^{-r}$.

Algorithm 2 Construct a random sample W of the vertex set $V(\mathcal{L}^*)$

```

1: function RIGIDDEPTHSAMPLE( $n, S, Q, \varepsilon_0$ )
2:    $w \leftarrow \lceil 36rn^3\varepsilon_0^{-3} \ln n \rceil$ , where  $\varepsilon_0 \in (0, 1)$  and  $r$  is any positive value fixed a priori.
3:   for  $i = 1$  to  $w$  do
4:     pick distinct triple  $s_1, s_2, s_3$  at random from  $S$ .
5:     pick three supporting lines  $\ell_1, \ell_2, \ell_3$  of  $Q$  uniformly at random.
6:     if the surfaces  $\ell_{1,*}(s_1), \ell_{2,*}(s_2)$  and  $\ell_{3,*}(s_3)$  intersect at some vertices then
7:       randomly pick a vertex  $(t, \theta)$  from the intersection.
8:       add  $(t, \theta)$  to  $W$ .
9:     end if
10:  end for
11:  return  $W$ 
12: end function

```

Proof. Let $k = \lfloor \varepsilon_0 |S| / 2 \rfloor - 1$. Since $c \geq 6$, $\varepsilon_0 |S| \geq 6$ and so $k > 1$. If $k > \text{depth}(\mathring{t}, \mathring{\theta}, \mathcal{L}^*) - 2$, then $\mu_S(\mathring{t}, \mathring{\theta}) = \text{depth}(\mathring{t}, \mathring{\theta}, \mathcal{L}^*) / |S| < (k + 2) / |S| \leq \varepsilon_0$. So $\max_{(t, \theta) \in W} \mu_S(t, \theta) \geq 0 > \mu_S(\mathring{t}, \mathring{\theta}) - \varepsilon_0$, satisfying (i).

Suppose that $k \leq \text{depth}(\mathring{t}, \mathring{\theta}, \mathcal{L}^*) - 2$. Since S is chosen uniformly at random, $\mathcal{T}^{\mathring{\theta}}$ is a simple family of shapes with probability 1. Let \hat{t} be a vertex of the cell containing \mathring{t} in the arrangement of $\mathcal{T}^{\mathring{\theta}}$. Since $\text{depth}(\hat{t}, \mathcal{T}^{\mathring{\theta}}) = \text{depth}(\mathring{t}, \mathring{\theta}, \mathcal{L}^*) \geq k + 2 > 1$, we can choose \hat{t} to be a vertex from $V(\mathcal{T}^{\mathring{\theta}})$. Notice that $\text{depth}(\hat{t}, \mathcal{T}^{\mathring{\theta}}) \geq \text{depth}(\mathring{t}, \mathring{\theta}, \mathcal{L}^*) \geq k + 2$. By Lemma 3.7, there exists a subset $U \subseteq V(\mathcal{L}^*)$ such that $|U| = k(k + 1)(k + 2) / 3$ and for every $(t, \theta) \in U$, $|\text{depth}((t, \theta), \mathcal{L}^*) - \text{depth}(\hat{t}, \mathcal{T}^{\mathring{\theta}})| \leq 2k + 2$. Let E_i the event that the three surfaces sampled in the i th iteration in RIGIDDEPTHSAMPLE intersect, and the intersection point (t_i, θ_i) selected belongs to U . Observe that if $(t_i, \theta_i) \in U$ for some i , then the lemma is satisfied because $\mu_S(t_i, \theta_i) \geq \mu_S(\hat{t}, \mathring{\theta}) - |\mu_S(t_i, \theta_i) - \mu_S(\hat{t}, \mathring{\theta})| = (\text{depth}(\hat{t}, \mathcal{T}^{\mathring{\theta}}) - |\text{depth}((t_i, \theta_i), \mathcal{L}^*) - \text{depth}(\hat{t}, \mathcal{T}^{\mathring{\theta}})|) / |S| \geq (\text{depth}(\hat{t}, \mathcal{T}^{\mathring{\theta}}) - 2k - 2) / |S| \geq \mu_S(\mathring{t}, \mathring{\theta}) - \varepsilon_0$.

It remains to bound the probability that E_i does not happen for all i . There are no more than $n^3 |S|^3 / 6$ triples of surfaces from \mathcal{L}^* generated by distinct triple of points from S . By Lemma 3.5, each such triple of surfaces intersect in at most two points. Thus, $\Pr(E_i) \geq 3|U| / (n^3 |S|^3) = k(k + 1)(k + 2) / (n^3 |S|^3) \geq (\varepsilon_0 / 2 - 2 / |S|)(\varepsilon_0 / 2 - 1 / |S|)(\varepsilon_0 / 2) / n^3 \geq \varepsilon_0^3 / (36n^3)$. The last inequality follows from the fact that $|S| \geq 6 / \varepsilon_0$. Hence, the probability that E_i does not happen for all i is at most $(1 - \varepsilon_0^3 / (36n^3))^{36rn^3\varepsilon_0^{-3} \ln n} \leq e^{-r \ln n} = n^{-r}$. This completes the proof of (i).

By Lemma 3.2, with probability at least $1 - n^{-r}$, $|\mu_S(\mathring{t}, \mathring{\theta}) - \mu(\mathring{t}, \mathring{\theta})| \leq \varepsilon_0$ and $|\mu_S(\tilde{t}, \tilde{\theta}) - \mu(\tilde{t}, \tilde{\theta})| \leq \varepsilon_0$. Then, it follows from (i) that $\mu(\tilde{t}, \tilde{\theta}) \geq \mu(\mathring{t}, \mathring{\theta}) - 3\varepsilon_0$ holds with probability at least $1 - 2n^{-r}$. This proves (ii). \square

The shape matching algorithm under rigid motion works as follows. We call RIGIDDEPTH-SAMPLE with $\varepsilon_0 = \varepsilon / 3$ to obtain a sample set W . For each $(t, \theta) \in W$, one can count $|(S + t) \cap Q_\theta| = |(S + t)_{-\theta} \cap Q|$ by answering a point location query for $(s + t)_{-\theta}$ for every $s \in S$ as follows. First, construct the arrangement of the supporting planes of Q and mark the cells in the arrangement that lies in Q . This can be done in $O(n^3)$ time using an incremental algorithm [10]. Second, build a point location data structure for this arrangement [11], which uses $O(n^3 \log n)$ space and preprocessing time, and answers a point location query in $O(\log^2 n)$ time. Finally, for each $s \in S$ and each $(t, \theta) \in W$, we issue a point location query to decide whether $(s + t)_{-\theta} \in Q$. This allows us to compute $|(S + t) \cap Q_\theta|$. The total running time is




Figure 3: The figures show the cross-section arrangements for different values of θ . The sweep plane hits one or more vertices, when the shaded region collapses to a single point. As the sweep plane moves away from these vertices, new shaded region appears.

$\tilde{O}(n^3\varepsilon^{-5})$, which already compares favorably with the $\tilde{O}(n^3\varepsilon^{-8})$ running time in [6]. We can do better as follows.

Partition S into m subfamilies S_1, \dots, S_m , each consisting of $|S|/m$ points. We will perform a space sweep over each arrangement \mathcal{A}_j of $\{\ell_*(s) : \ell \in L \wedge s \in S_j\}$. Observe that for each $\theta \in [0, 2\pi)$, the *cross-section arrangement* $\{\ell_\theta - s : \ell \in L \wedge s \in S_j\}$ of \mathcal{A}_j is an arrangement of lines. As we sweep a plane from $\theta = 0$ to $\theta = 2\pi$, the geometry of the cross-section arrangement changes continuously; however, there is no topological change before the sweep plane hits the next vertex of \mathcal{A}_j .

We maintain a data structure of Goodrich and Tamassia that answers point location queries in a dynamic monotone planar subdivision [11]. The cross-section arrangement can be viewed as a directed planar graph with every edge directed from its upper endpoint to its lower endpoint. By the result in [11], if the geometry of the subdivision is changed such that the topology of the directed planar graph does not change, the point location data structure needs not be updated at all. Let N be the complexity of the subdivision. The data structure answers a point location query in $O(\log^2 N)$ time and an update can be done in $O(\log N)$ time.

When sweeping over \mathcal{A}_j , there are three types of events when the topology of the directed planar graph changes.

- Type 1: Refer to Figure 3. At some $\phi \in [0, 2\pi)$, there exists $k \geq 3$ lines $\ell_{i,\phi} - s_i$ for $i \in [1, k]$ such that: (i) $\bigcap_{i=1}^k \ell_{i,\phi} - s_i \neq \emptyset$, (ii) ℓ_1, \dots, ℓ_k are mutually non-parallel, (iii) $s_1 \notin \{s_2, \dots, s_k\}$ and $\bigcap_{i=2}^k \ell_{i,\theta} - s_i \neq \emptyset$ for all $\theta \in [0, 2\pi)$.
- Type 2: Refer to Figure 3. At some $\phi \in [0, 2\pi)$, there exists two lines $\ell_{i,\phi} - s_i$ for $i \in [1, 2]$

such that: $\ell_{1,\phi} - s_1 = \ell_{2,\phi} - s_2$ and for every $\theta \neq \phi$ in an arbitrarily small neighborhood of ϕ , the following conditions hold: (i) $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ are distinct parallel lines, (ii) no two lines in the cross-section arrangement intersect in the interior of the strip bounded by $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$, and (iii) for every edge in the cross-section arrangement that lies in the strip bounded by $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$, at least one of edge endpoints has vertex degree four.

- Type 3: At some $\phi \in [0, 2\pi)$, some line ℓ_ϕ becomes horizontal. Hence, for every $s \in S_j$, $\ell_\phi - s$ becomes horizontal.

Events of types 1 and 2 happen when the sweep plane hits one or more vertices of \mathcal{A}_j . If L satisfies the general position assumption, there are exactly three lines involved in a type 1 event, and for every edge that lies in the strip in a type 2 event, both endpoints of that edge have vertex degree four. But we do not make the general position assumption. An event of type 3 happens when the directions of some edges in the directed planar graph switch.

Before proving the completeness of the above list of events, we first establish two technical lemmas.

Lemma 3.9. *Let ℓ_1, ℓ_2, ℓ_3 and ℓ_4 be four possibly non-distinct lines in \mathbb{R}^2 such that ℓ_1 and ℓ_3 are not parallel and ℓ_2 and ℓ_4 are not parallel. Let s_1 and s_2 be two possibly non-distinct points in \mathbb{R}^2 . Let t_1 and t_2 be two points in \mathbb{R}^2 such that $t_1 \in \ell_1 \cap \ell_3$ and $t_2 \in \ell_2 \cap \ell_4$. If $\ell_{1,\theta} - s_1, \ell_{2,\theta} - s_2, \ell_{3,\theta} - s_1$, and $\ell_{4,\theta} - s_2$ have a non-empty common intersection for some θ , then $\|s_1 - s_2\| = \|t_1 - t_2\|$.*

Proof. Since $t_1 \in \ell_1 \cap \ell_3$ and $t_2 \in \ell_2 \cap \ell_4$, we get $t_{1,\theta} - s_1 \in (\ell_{1,\theta} - s_1) \cap (\ell_{3,\theta} - s_1)$ and $t_{2,\theta} - s_2 \in (\ell_{2,\theta} - s_2) \cap (\ell_{4,\theta} - s_2)$. If $\ell_{1,\theta} - s_1, \ell_{2,\theta} - s_2, \ell_{3,\theta} - s_1$, and $\ell_{4,\theta} - s_2$ have a non-empty common intersection for some θ , we can set $t_{1,\theta} - s_1 = t_{2,\theta} - s_2$ for that particular θ . It follows that $(t_1 - t_2)_\theta = s_1 - s_2$ and hence $\|s_1 - s_2\| = \|t_1 - t_2\|$. \square

Lemma 3.10. *Let $\ell_{1,*}(s_1), \ell_{2,*}(s_2), \ell_{3,*}(s_3)$, and $\ell_{4,*}(s_4)$ be four distinct surfaces in \mathcal{L}^* . Let E denote the event that at least three points in $\{s_1, s_2, s_3, s_4\}$ are identical. Let E' denote the event that $\bigcap_{i=1}^4 \ell_{i,*}(s_i) = \emptyset$. Then, $\Pr(E \vee E') = 1$.*

Proof. If E does not happen, then there are only two other possibilities:

- E_1 : Some triple of points in $\{s_1, s_2, s_3, s_4\}$ are distinct.
- E_2 : There are exactly two pairs of identical points in $\{s_1, s_2, s_3, s_4\}$, but the four points are not all equal.

We show below that $\Pr(\overline{E'}|E_1) = 0$ if $\Pr(E_1) \neq 0$ and $\Pr(\overline{E'}|E_2) = 0$ if $\Pr(E_2) \neq 0$. This implies that $\Pr(E_1 \wedge \overline{E'}) = \Pr(E_2 \wedge \overline{E'}) = 0$, and therefore, $\Pr(E \vee E') = 1 - \Pr(\overline{E} \wedge \overline{E'}) = 1 - \Pr(E_1 \wedge \overline{E'}) - \Pr(E_2 \wedge \overline{E'}) = 1$ as stated in the lemma.

Suppose s_1, s_2 , and s_3 are distinct, i.e. E_1 happens. By Lemma 3.5, there exist at most two (t, ϕ) 's in $\mathbb{R}^2 \times [0, 2\pi)$ such that (t, ϕ) belongs to the common intersection of $\ell_{1,*}(s_1), \ell_{2,*}(s_2)$, and $\ell_{3,*}(s_3)$. It means that $(t, \phi) \in V(\mathcal{L}^*)$. Then, Lemma 3.6 says that $(t, \phi) \notin \ell_{4,*}(s_4)$ with probability 1. Therefore, $\Pr(E'|E_1) = 1$, or equivalently, $\Pr(\overline{E'}|E_1) = 0$.

Suppose that $s_1 = s_3$ and $s_2 = s_4$ but $s_1 \neq s_2$, i.e. E_2 happens. If ℓ_1 and ℓ_3 are parallel, then $\ell_{1,\theta} - s_1$ and $\ell_{3,\theta} - s_1$ are parallel and distinct for all θ because $\ell_{1,*}$ and $\ell_{3,*}$ are distinct surfaces by assumption. Similarly, if ℓ_2 and ℓ_4 are parallel, then $\ell_{2,\theta} - s_2$ and $\ell_{4,\theta} - s_2$ are parallel and distinct for all θ . The remaining possibility is that ℓ_1 and ℓ_3 are non-parallel, and ℓ_2 and ℓ_4 are non-parallel. By Lemma 3.9, $\ell_{1,\theta} - s_1, \ell_{2,\theta} - s_2, \ell_{3,\theta} - s_1$, and $\ell_{4,\theta} - s_2$ have a

non-empty common intersection for some θ only if s_1 and s_2 are at some fixed distance apart. This happens with probability zero because s_1 and s_2 are random independent samples from P . In summary, $\Pr(E'|E_2) = 1$, or equivalently, $\Pr(\overline{E'}|E_2) = 0$. \square

Type 3 events are only needed for because the data structure of Goodrich and Tammasia represents the cross-section arrangement as a directed planar graph in which every edge is directed from its upper endpoint to its lower endpoint. We prove below that Type 1 and Type 2 events capture all combinatorial changes in the cross-section arrangement.

Lemma 3.11. *With probability 1, every combinatorial change to the cross-section arrangement during the sweep is of either type 1 or type 2.*

Proof. A combinatorial change happens only when certain region in the cross-section arrangement collapses at some angle ϕ . Consider the bounding lines of the collapsed region.

Case 1: The bounding lines are pairwise non-parallel. Let $\ell_{1,\phi} - s_1, \dots, \ell_{k,\phi} - s_k$, where $k \geq 3$, be the bounding lines of the collapsed region. We claim that s_1, \dots, s_k cannot be all equal. Otherwise, since $\bigcap_{i=1}^k (\ell_{i,\phi} - s_i)$ is a single point, $\bigcap_{i=1}^k (\ell_{i,\theta} - s_i)$ is a single point for all θ , meaning that they cannot bound any region to be collapsed just before θ becomes ϕ , a contradiction. This proves our claim.

If $k = 3$, we may assume that $s_1 \notin \{s_2, s_3\}$. Since $\ell_{2,\phi} - s_2$ and $\ell_{3,\phi} - s_3$ are non-parallel by assumption, $\ell_{2,\theta} - s_2$ and $\ell_{3,\theta} - s_3$ are non-parallel for all θ . We conclude that this event is of type 1. Consider the case that $k \geq 4$. By our claim, we can assume that $s_1 \neq s_i$ for some $i \in [2, k]$, say s_2 . Applying Lemma 3.10 to $\ell_{i,\phi} - s_i$ for $i \in [1, 4]$ shows that $s_1 = s_3 = s_4$ or $s_2 = s_3 = s_4$. Without loss of generality, assume that $s_2 = s_3 = s_4$. Applying Lemma 3.10 again to $\ell_{1,\phi} - s_1, \ell_{3,\phi} - s_3, \ell_{4,\phi} - s_4$, and $\ell_{5,\phi} - s_5$ shows that $s_3 = s_4 = s_5$. By repeating this analysis, we conclude that s_2, \dots, s_k are all equal, but s_1 is not equal to them. Since $\bigcap_{i=2}^k (\ell_{i,\phi} - s_i)$ is a single point, $\bigcap_{i=2}^k (\ell_{i,\theta} - s_i)$ is a single point for all θ . This shows that the event is of type 1.

Case 2: A pair of distinct bounding lines are parallel. Let the parallel bounding lines be $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$. Since $\ell_{1,\phi} - s_1 = \ell_{2,\phi} - s_2$ but $\ell_{1,\theta} - s_1 \neq \ell_{2,\theta} - s_2$ just before θ becomes ϕ , the points s_1 and s_2 must be different. There are two things to be proved in order to show that the event is of type 2. First, we need to show that for every θ in an arbitrarily small neighborhood of ϕ , no two distinct lines cross inside the strip bounded by $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$. Second, every edge in the strip has at least one endpoint with vertex degree four.

Assume to the contrary that there exist $\ell_3, \ell_4 \in L$ and $s_3, s_4 \in S$ such that $\ell_{3,\theta} - s_3$ and $\ell_{4,\theta} - s_4$ cross in the interior of the strip. By continuity, $\bigcap_{i=1}^4 (\ell_{i,\phi} - s_i) \neq \emptyset$. By Lemma 3.10, at least three points in $\{s_1, s_2, s_3, s_4\}$ are identical, say $s_2 = s_3 = s_4$. Since $\bigcap_{i=2}^4 (\ell_{i,\phi} - s_i) \neq \emptyset$, $\bigcap_{i=2}^4 (\ell_{i,\theta} - s_i) \neq \emptyset$, for all θ . But then $\ell_{3,\theta} - s_3$ cannot cross $\ell_{4,\theta} - s_4$ in the interior of the strip just before θ becomes ϕ because such a crossing would have to lie on $\ell_{2,\theta} - s_2$, a contradiction. This proves the first condition.

Assume to the contrary that there exists an edge e in the strip such that every endpoint of e has vertex degree greater than 4. Let $\ell_{3,\theta} - s_3$ be the line containing e . In this case, there exist two distinct lines $\ell_{4,\theta} - s_4$ and $\ell_{5,\theta} - s_5$ that pass through the intersection points $(\ell_{1,\theta} - s_1) \cap (\ell_{3,\theta} - s_3)$ and $(\ell_{2,\theta} - s_2) \cap (\ell_{3,\theta} - s_3)$, respectively. Because θ is an arbitrary value in an arbitrarily small neighborhood of ϕ , it means that $\ell_{1,\theta} - s_1, \ell_{3,\theta} - s_3$, and $\ell_{4,\theta} - s_4$ are concurrent for some range of θ . This happens only if $s_1 = s_3 = s_4$. We can similarly conclude that $s_2 = s_3 = s_5$. Hence, $s_1 = s_2$. But then since $\ell_{1,\phi} - s_1 = \ell_{2,\phi} - s_2$, the lines $\ell_{1,\theta} - s_1$ and $\ell_{2,\theta} - s_2$ must be equal for all θ , contradicting the fact that they bound a strip to be collapsed just before θ becomes ϕ . This proves the second condition. \square

In summary, the improved shape matching algorithm under rigid motion works by sweeping

the three-dimensional curved arrangement \mathcal{A}_j for $j \in [1, m]$. Throughout the sweep, the cross-section arrangement is represented and maintained using the dynamic point location structure of Goodrich and Tamassia [11]. This data structure is queried for each point $(t, \theta) \in W$ encountered during the sweep, so that we obtain the depth $d_j(t, \theta)$ of each $(t, \theta) \in W$ in \mathcal{A}_j . Finally, $\mu_S(t, \theta) = \sum_{j=1}^m d_j(t, \theta)/|S|$ and we return $\arg \max_{(t, \theta) \in W} \mu_S(t, \theta)$.

We are ready to prove the performance of the shape matching algorithm under rigid motion.

Theorem 3.12. *Let P and Q be two polygonal shapes with a total of n vertices. Let opt be the maximum overlap of P and Q under rigid motion. For every $\varepsilon \in (0, 1)$, one can compute a rigid motion $(\tilde{t}, \tilde{\theta})$ such that $\text{area}((P + \tilde{t}) \cap Q_{\tilde{\theta}}) \geq \text{opt} - \varepsilon \cdot \text{area}(P)$ with probability $1 - n^{-O(1)}$ in $O(n^3 \varepsilon^{-4} \log^{5/3} n \log^{5/3}(n/\varepsilon))$ time.*

Proof. The correctness and probability bound follow from Lemma 3.8(ii) and the setting of $\varepsilon_0 = \varepsilon/3$. Locating all points in W during the sweep over \mathcal{A}_j takes $O(|W| \log^2(n|S|/m))$ time.

A type 1 event is simulated by $\Theta(k)$ vertex deletions and insertions, and edge deletions and insertions. We charge these updates to the $\Omega(k)$ triples of surfaces $\{\ell_{1,*} - s_1, \ell_{i,*} - s_i, \ell_{i+1,*} - s_{i+1}\}$ for $i \in [2, k-1]$. By Lemma 3.5, every such triple induces at most two vertices, meaning that every such triple is charged $O(1)$ times. Handling all type 1 events thus takes $O(n^3 |S|^3 m^{-3} \log(n|S|/m))$ time. Similarly, a type 2 event induced by two lines $\ell_{1,\phi} - s_1$ and $\ell_{2,\phi} - s_2$ can be simulated by $O(n|S|/m)$ updates. (Since every edge inside the strip has at least one edge of vertex degree four by the definition of a type 2 event, the only topological change that happens is the collapse of the strip.) We charge these updates to the pair $\{\ell_{1,*}(s_1), \ell_{2,*}(s_2)\}$. By Lemma 3.3(ii), the pair $\{\ell_{1,*}(s_1), \ell_{2,*}(s_2)\}$ is charged by at most two events of type 2. Handling all type 2 events thus takes $O(n^3 |S|^3 m^{-3} \log(n|S|/m))$ time. When a line ℓ_ϕ becomes horizontal in a type 3 event, $\ell_\phi - s$ becomes horizontal for every $s \in S_j$. There are $O(n|S|/m)$ edges on every such line $\ell_\phi - s$, so the batch of type 3 events for ℓ_ϕ can be simulated by $O(n|S|^2 m^{-2})$ edge deletions and insertions. Hence, it takes $O(n^2 |S|^2 m^{-2} \log(n|S|/m))$ time to handle all type 3 events during the space sweep.

In the event (with probability zero) that the space sweep encounters a violation of Lemma 3.6 or a combinatorial change not of type 1 or 2, we can just halt and return $(\tilde{t}, \tilde{\theta}) = (0, 0, 0)$.

Summing over all \mathcal{A}_j 's gives an $O(n^3 |S|^3 m^{-2} \log(n|S|/m) + m|W| \log^2(n|S|/m))$ running time, which becomes $O(n^3 \varepsilon^{-4} \log^{5/3} n \log^{5/3}(n/\varepsilon))$ by setting $m = \varepsilon^{-1} \log^{2/3} n \log^{-1/3}(n/\varepsilon)$. \square

4 Conclusion

We presented improved algorithms to find the maximum overlap of two polygonal shapes under translation and rigid motion, respectively. They improve the previous best running times by Cheong et al. [6] from $\tilde{O}(n^2 \varepsilon^{-4})$ to $\tilde{O}(n^2 \varepsilon^{-3})$ in the translation case, and from $\tilde{O}(n^3 \varepsilon^{-8})$ to $\tilde{O}(n^3 \varepsilon^{-4})$ in the rigid motion case. Moreover, degeneracy and disconnected shapes are allowed, which should make the results more applicable in practice. It is open whether the dependence on n can be reduced. It would also be interesting to develop fast shape matching algorithms in three dimensions.

References

- [1] H.-K. Ahn, S.-W. Cheng, and I. Reinbacher. Maximum overlap of convex polytopes under translation. *Proceedings of the Annual International Symposium on Algorithms and Computations*, 97–108, 2010.

- [2] H.-K. Ahn, O. Cheong, C.-D. Park, C.-S. Shin, and A. Vigneron. Maximizing the overlap of two planar convex sets under rigid motions. *Computational Geometry: Theory and Applications*, 37 (2007), 3–15.
- [3] H. Alt, J. Blömer, and H. Wagener. Approximation of convex polygons. *Proceedings of the International Colloquium on Automata, Languages and Programming*, 703–716, 1990.
- [4] H. Alt, L. Scharf, and D. Schymura. Probabilistic matching of planar regions. *Computational Geometry: Theory and Applications*, 43(2):99–114, 2010.
- [5] S. Basu, R. Pollak, and M.-F. Roy. On the number of cells defined by a family of polynomials on a variety. *Mathematika*, 43(01):120–126, 1996.
- [6] O. Cheong, A. Efrat, and S. Har-Peled. Finding a guard that sees most and a shop that sells most. *Discrete & Computational Geometry*, 37:545–563, 2007.
- [7] M. de Berg, S. Cabello, P. Giannopoulos, C. Knauer, R. van Oostrum, and R. Veltkamp. Maximizing the area of overlap of two unions of disks under rigid motion. *Proceedings of the Scandinavian Workshop on Algorithm Theory*, 138–149, 2004.
- [8] M. de Berg, O. Devillers, M. van Kreveld, O. Schwarzkopf, and M. Teillaud. Computing the maximum overlap of two convex polygons under translations. *Theory of Computing Systems*, 31 (1998), 613–628.
- [9] D.P. Dobkin and D.G. Kirkpatrick. Determining the separation of preprocessed polyhedra—a unified approach. *Proceedings of the 17th International Colloquium on Automata, Languages and Programming*, 1990, 400–413.
- [10] H. Edelsbrunner. *Algorithms in combinatorial geometry*, Springer-Verlag, Berlin, 1987.
- [11] M. T. Goodrich and R. Tamassia. Dynamic trees and dynamic point location. *SIAM Journal on Computing*, 28(2):612–636, 1998.
- [12] S. Har-Peled and M. Sharir. Relative (p, ε) -approximations in geometry. *Discrete & Computational Geometry*, 45:462–496, 2011.
- [13] Y. Li, P. M. Long, and A. Srinivasan. Improved bounds on the sample complexity of learning. *Journal of Computer and System Sciences*, 62(3):516–527, 2001.
- [14] D. M. Mount, R. Silverman, and A. Y. Wu. On the area of overlap of translated polygons. *Computer Vision and Image Understanding*, 64(1):53–61, 1996.
- [15] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [16] A. Vigneron. Geometric optimization and sums of algebraic functions. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, 906–917, 2010.