Lecture Notes: Minimum Enclosing Disc

Yufei Tao
Department of Computer Science and Engineering
Chinese University of Hong Kong
taoyf@cse.cuhk.edu.hk

March 8, 2021

Let P be a set of n points in $\mathbb{R}^2$. We want to find a disc D with the smallest radius to cover all the points in P. We refer to D as the *minimum enclosing disc* (MED) of P and denote it as $med(P)$. The lemma below explains why calling D *the* MED is appropriate.

Lemma 1. *There is only one disc with the smallest radius covering all the points in P.*

Proof. Assume, on the contrary, that there are two such discs D_1 and D_2; see the figure below. Then, P must be covered by the shaded area. Let A and B the intersection points of the two discs. Consider the disc D centering at the midpoint o of the segment AB and having a radius equal to the length of segment oA. D covers the shaded area (and hence, also P) but is smaller than D_1 and D_2, giving a contradiction.

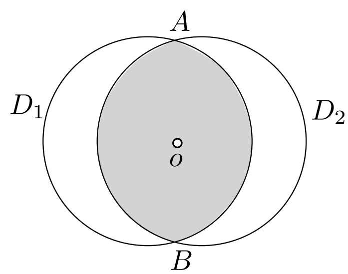

□

Today we will learn a randomized algorithm for solving the problem in $O(n)$ expected time. As we will see, this is another beautiful application of backward analysis.

1 Geometric Facts

Lemma 2. *The boundary of $med(P)$ passes at least two points of P.*

Proof. Let C be the boundary of $med(P)$. If C passes no points of P, shrink C infinitesimally to obtain a smaller disc still covering P, which contradicts the definition of C.

Suppose that C passes only one point $p \in P$. Let o be the center of C. Consider sliding a point o' from o towards p infinitesimally, and look at the circle C' centered at o' with radius equal to the length of segment $o'p$. C' is smaller than C but still contains P in the interior. This is also a contradiction. □

The following geometric fact will be useful:

Lemma 3. *Let C_1 and C_2 be two intersecting circles such that the radius of C_2 is larger than or equal to that of C_1. Let α be the area inside both circles. Let p be an arbitrary point that is in C_2 but not in C_1. Then, there exists a circle C that is smaller than C_2, larger than C_1, passes p, and contains the area α.*

The figure belows gives an illustration of the lemma, where C is the circle in dash line.

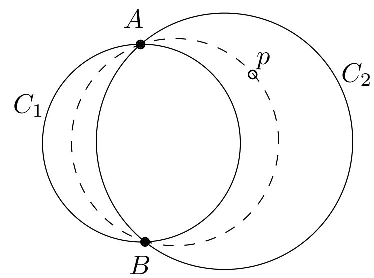

Proof. The lemma can be proved using basic geometry. We give only a sketch here (the complete proof is tedious and rudimentary).

Let us first discuss a relevant fact. Fix two distinct points A, B. Consider all the circles passing both A and B. The centers of these circles must be on the perpendicular bisector of segment AB. Every such circle C can be divided into (i) a left arc, which is the part of C on the left of segment AB, and (ii) a right arc, which is the part of C on the right. As the center o of C moves away from the midpoint m of segment AB towards right, the left arc "sweeps" towards segment AB, while the right arc "sweeps" away from the segment; furthermore, C grows continuously. The behavior is symmetric when o moves away from m towards left.

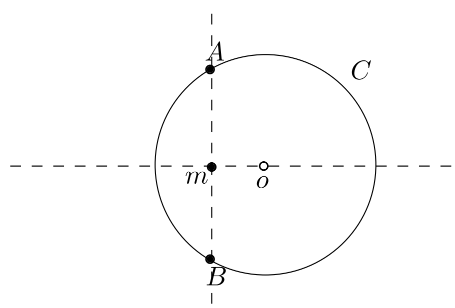

Going back to the context of the lemma, let A and B be the intersection points of C_1 and C_2. Imagine morphing a circle C from C_2 to C_1 while making sure that C passes A and B. Stop as soon as the right arc of C hits p. This is the circle we are looking for. □

2 Two Points Are Known

Let us first look at a variant of the MED problem. Let p_1, p_2 be two points in P such that there is at least one disc which has p_1, p_2 on the boundary and covers the entire P. We want to find the smallest such disc, denoted as $med(P, \{p_1, p_2\})$. Algorithm 1 presents our solution in pseudocode. Its running time is clearly $O(n)$. To prove its correctness, it suffices to show:

Lemma 4. *Define, for each $i \in [1, n]$, $P_i = \{p_1, ..., p_i\}$. For $i \geq 3$, let $D = med(P_{i-1}, \{p_1, p_2\})$. If p_i is not covered by D, then the boundary of $med(P_i, \{p_1, p_2\})$ must pass p_i.*

Algorithm 1: Two-Points-Fixed-MED(P, $\{p_1, p_2\}$)

/* suppose $P = \{p_1, p_2, ..., p_n\}$ */

$D \leftarrow$ the smallest disc covering p_1, p_2

for $i = 3$ **to** n **do**

  if p_i *not in* D **then**

    $D \leftarrow$ the disc whose boundary passes p_1, p_2, p_i

return C

Proof. Let $D' = med(P_i, \{p_1, p_2\})$. Assume on the contrary that the boundary of D' does not pass p_i. Hence, p_i falls inside D'; see the figure below. The radius of D' cannot be smaller than that of D because the latter was the MED on P_{i-1} whereas D' is just one disc covering P_{i-1}. The entire P_{i-1} must fall in the shaded area. By Lemma 3, there exists a disc smaller than D' covering P_i, giving a contradiction. □

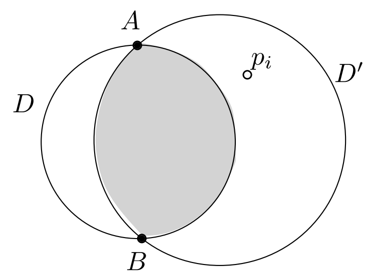

3 One Point Is Known

Next, we will generalize the two-points-fixed problem a bit. Let p_1 be a point in P such that there is at least one disc covering P whose boundary passes p_1. We want to find the smallest such circle, denoted as $med(P, \{p_1\})$.

Algorithm 2: One-Point-Fixed-MED(P, $\{p_1\}$)

/* suppose $P = \{p_1, p_2, ..., p_n\}$ */

randomly permute $p_2, ..., p_n$

$D \leftarrow$ the smallest disc covering p_1, p_2

for $i = 3$ **to** n **do**

  if p_i *not in* D **then**

    $D \leftarrow$ Two-Points-Fixed-MED($\{p_1, ..., p_i\}$, $\{p_1, p_i\}$)

return D

The algorithm's correctness is ensured by:

Lemma 5. *For $i \geq 3$, let $D = med(P_{i-1}, \{p_1\})$. If p_i is not covered by D, then the boundary of $med(P_i, \{p_1\})$ must pass p_i.*

Proof. Left as an exercise. □

Let us analyze the running time of the algorithm. Let t_i be the expected running time of the for-loop (Lines 3-5) for a specific i. Thus, the total expected running time is $O(\sum_{i=3}^{n} \mathbf{E}[t_i])$. Now, focus on t_i for a specific i. Set $D = \mathit{med}(P_i, \{p_1\})$. We know that, besides p_1, the boundary of D is determined by at most 2 other points in P — let them be π_1, π_2 (if the boundary passes more than 2 points of P other than p_1, set π_1, π_2 to 2 arbitrary points of them). Hence, if $p_i \neq \pi_1$ and $p_i \neq \pi_2$, then $t_i = O(1)$; otherwise, $t_i = O(i)$ (Lemma 4). Standard backward analysis shows that $\mathbf{E}[t_i] \leq \frac{2}{i-1} O(i) + O(1) = O(1)$. Therefore, the expected running time of Algorithm 2 is $O(n)$, which subsumes the time of random permutation at Line 1.

4 No Point Is Known

We are ready to tackle the MED problem in its most general form:

Algorithm 3: MED(P)

/* suppose $P = \{p_1, p_2, ..., p_n\}$ */

randomly permute $p_1, ..., p_n$

$D \leftarrow$ the smallest disc covering p_1, p_2

for $i = 3$ **to** n **do**

  if p_i *not in* D **then**

    $D \leftarrow$ One-Point-Fixed-MED($\{p_1, ..., p_i\}$, $\{p_i\}$)

return C

Lemma 6. *For $i \geq 3$, let $D = med(P_{i-1})$. If p_i is not covered by D, then the boundary of $med(P_i)$ passes p_i.*

Proof. Left as an exercise. □

We can once again apply backward analysis to prove that Algorithm 3 runs in $O(n)$ expected time. The details are left as an exercise.