CMSC 754: Fall 2021 Dave Mount

Homework 4: Sampling, Motion Planning, and More

Handed out Tue, Nov 30. Due, Tue, Dec 14, 9:30am. (There is no class that day, but I'll record pretend lecture where I will discuss the solutions.) Late homeworks are not accepted (unless an extension has been prearranged) so please turn in whatever you have completed by the due date. Unless otherwise specified, you may assume that all inputs are in *general position*. Whenever asked to give an algorithm running in $O(f(n))$ time, you may give a *randomized algorithm* whose expected running time is $O(f(n))$.

Problem 1. The objective of this problem is to investigate the *VC-dimension* of some range spaces. Recall that a *range space* Σ is a pair $(X, \mathcal{R})$, where X is a (finite or infinite) set, called *points*, and $\mathcal{R}$ is a (finite or infinite) family of subsets of X, called *ranges*.

For each of the following range spaces, derive its VC-dimension and prove your result. (Note that in order to show that the VC-dimension is k, you need to give an example of a k-element subset that is shattered and prove that no set of size $k+1$ can be shattered.) Throughout, you may assume that points are in general position.

(a) $\Sigma = (\mathbb{R}^2, W)$, where W consists of all horizontal wedges. A *horizontal wedge* is the region bounded between two rays, one horizontal and the other of arbitrary (finite) slope (see Fig. 1(a)). (It is my intention that the wedge angle is acute, but if you want to consider obtuse wedges instead, that is fine with me.)

(b) $\Sigma = (\mathbb{R}^2, W')$, where W' consists of all horizontal double wedges. A *horizontal double wedge* is the region bounded between two lines, one horizontal and the other of arbitrary (finite) slope. It consists of the points that lie above one line and below the other (see Fig. 1(b)).

Hint: Getting a tight upper bound on the VC-dimension seems to be very difficult. For this part, it suffices to come up with any constant upper bound, by whatever method you like. (See, e.g., the material on *canonical shapes* from the Lecture 20 Notes.) Deriving a better bound is left as the Challenge Problem.

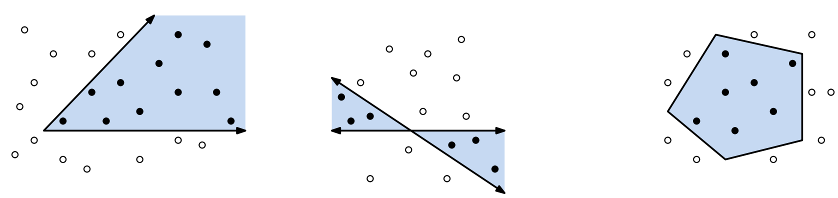

(a) (b) (c)

Figure 1: VC-Dimension of some range spaces.

Also, answer the following.

(c) Prove that the range space $\Sigma = (\mathbb{R}^2, C)$ consisting of convex polygons in the plane has *unbounded* VC-dimension (see Fig. 1(c)). That is, show that for any $n > 0$, there exists an n-element point set that is shattered by the range space of convex polygons.

Here is a sample solution for parts (a) and (b), to give you some idea of what I am looking for. (It would be even nicer if you include some figures.)

Example: Consider the range space $\Sigma = (\mathbb{R}^2, H)$ where H consists of all closed horizontal halfspaces, that is, halfplanes of the form $y \geq y_0$ or $y \leq y_0$. We claim that $\mathrm{VC}(\Sigma) = 2$.

VC(Σ) ≥ 2 : Consider the points $a = (0, -1)$ and $b = (0, 1)$. The ranges $y \geq 2$, $y \geq 0$, $y \leq 0$ and $y \leq 2$ generate the subsets $\{\emptyset, \{a\}, \{b\}, \{a, b\}\}$, respectively. Therefore, there is a set of size two that is shattered.

VC(Σ) < 3 : Consider any three element set $\{a, b, c\}$ in the plane. Let us assume that these points have been given in increasing order of their y-coordinates. Observe that any horizontal halfplane that contains b, must either contain a or c. Therefore, no 3-element point set can be shattered.

Problem 2. In this problem we will explore an idea for constructing a *weak ε-net* for a set of P points in the plane. By a "weak" ε-net, we mean a set of points that satisfies the standard definition of an ε-net, but it can be formed from *any* set of points, not just the points of P. We'll give the broad outline, and you will fill in the details.

You are given n points P in $\mathbb{R}^2$. Let us make the general-position assumption that there are no duplicate x- or y-coordinates. We construct a set $S \subset \mathbb{R}^2$ as follows. We first compute an integer $k \geq 1$, and let $m = \lfloor n/k \rfloor$. Next, we sort the points of S in increasing order according to their x-coordinates, and let $\langle x_1, \ldots, x_n \rangle$ denote the resulting sorted sequence of coordinates. We take every kth element from this sorted sequence:

$$X = \{x_k, x_{2k}, x_{3k}, \ldots, x_{mk}\}.$$

We repeat the same process for the y-coordinates, by first sorting them in increasing order as $\langle y_1, \ldots, y_n \rangle$, and setting

$$Y = \{y_k, y_{2k}, y_{3k}, \ldots, y_{mk}\}.$$

(Note that we use the same value of k and m in defining both X and Y.) Finally, we set $S = X \times Y$, that is, for $1 \leq i, j \leq m$, we include the point (x_i, y_j) into S. Clearly, S has m^2 elements. Observe that while they are based on the coordinates of the points of P, the points of S need not belong to P.

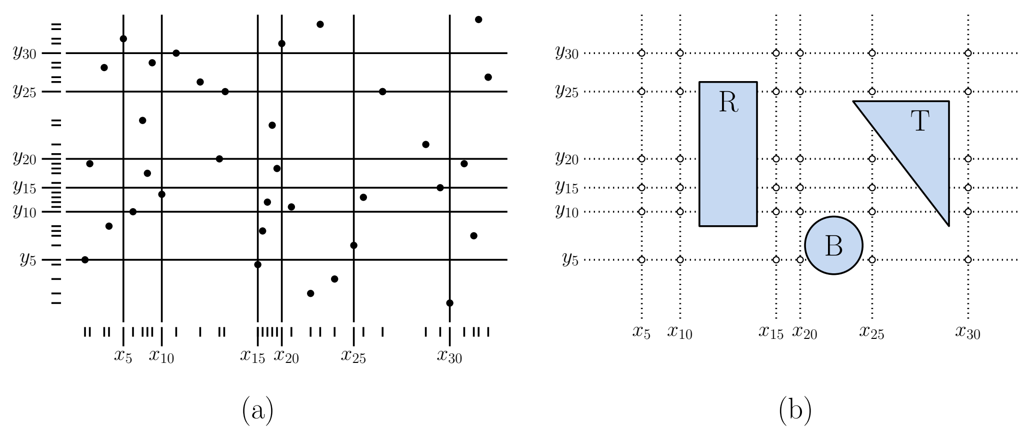

Figure 2: Weak ε-net construction (for $k = 5$).

(a) For each of the following range spaces, answer the following question. If a range from the set contains no elements of S, what is the maximum number of elements of P that it might contain? (That is, if $Q \cap S = \emptyset$ then how large can $|Q \cap P|$ be?) In each case justify your answer. (Express your answer precisely, not as an asymptotic function.)

(R) Axis-aligned rectangles (of any width and height).

(T) Axis-aligned right triangles. This is defined to be any right triangle such that the legs (i.e., the sides incident to the 90° angle) are parallel to the x- and y-axes. The hypotenuse can have any slope.

(B) Euclidean balls (of any radius).

(b) Suppose you are given a parameter $\varepsilon > 0$. Based on your answers to (R), (T), and (B) above, what value should we set k to (as a function of n and ε) in the above construction so that the resulting set S is an ε-net (in the weak sense) for P. We want k to be as large as possible, so that the resulting ε-net is as small as possible.

(c) Suppose that further, we would like a weak ε-sample. For each of the range spaces, (R), (T), and (B) above, is there any value of k such that the resulting set S is an ε-sample of P, where the size of S is a function of ε (but not n)? If so, give this value and justify its correctness. If not, explain why the resulting set S's size must depend on n.

Problem 3. This problem has many parts, but each answer is very short, and sentence, formula, or short derivation suffices for most.

Iterative reweighting (seen in Lecture 19) can be used to solve a number of geometric optimization involving set systems of constant VC-dimension. In this problem, we will derive an algorithm for computing the circular disk enclosing a set of n points in $\mathbb{R}^2$. (We will derive an $O(n \log n)$ time solution, but there is actually an $O(n)$ time more efficient solution based on a randomized incremental approach.)

Given a set of points $P \subset \mathbb{R}^2$, define the *minimum enclosing ball*, $\mathrm{MEB}(P)$, to be the circular disk B of smallest radius that contains P. Generally, $\mathrm{MEB}(P)$ is determined by a constant number of points P^* on the boundary of the ball, called the *basis* (see Fig. 3(a)). (There at most three in $\mathbb{R}^2$ and generally at most $d+1$ points in $\mathbb{R}^d$). The following two facts are easy to prove:

- For any set S such that $P^* \subseteq S \subseteq P$, $\mathrm{MEB}(S) = \mathrm{MEB}(P^*)$ (see Fig. 3(b)).
- For any set S such that $S \subseteq P$, radius$(\mathrm{MEB}(S)) \leq$ radius$(\mathrm{MEB}(P))$ (see Fig. 3(c)). (We cannot infer that $\mathrm{MEB}(S) \subseteq \mathrm{MEB}(P)$, since one disk might "bulge out" from the other.)

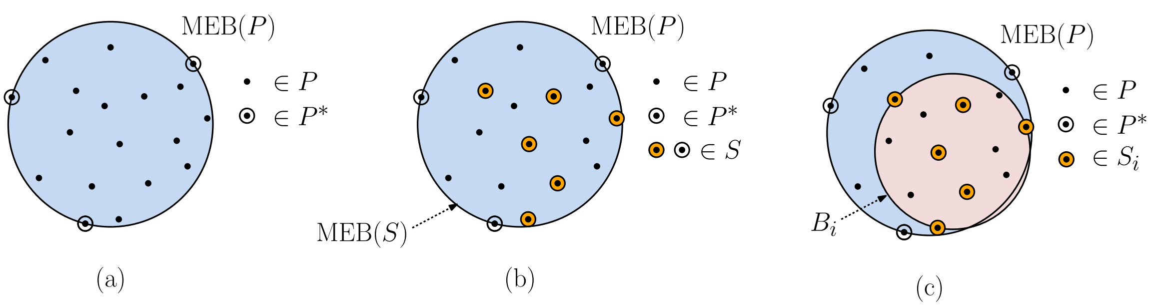

Figure 3: Reweighting algorithm for MEB.

Suppose that we have access to a black-box algorithm for computing the MEB, but it is very slow and can only be applied to sets of constant size. Our algorithm operates by computing an ε-net S of constant size for P with respect to the set system where the ranges are the points of P that lie *outside* of a circular disk. (We will see that this set system has constant VC-dimension.) We then invoke our slow MEB algorithm to compute the MEB for S. If the resulting ball contains P, we are done. If not, we double the weights of all the points of P outside this ball, and repeat. Here is the algorithm:

(1) Set $\varepsilon \leftarrow$ "fill this in later". For each $p \in P$, set $w_0(p) = 1$.

(2) Repeat for $i \leftarrow 1, 2, \ldots$

  (a) Let S_i be a weighted ε-net for P of size $O((1/\varepsilon)\log(1/\varepsilon))$ with respect to the weights w_{i-1} for ranges consisting of points lying *outside* a circular disk.

  (b) Let $B_i \leftarrow \mathrm{MEB}(S_i)$

(c) If $P \subseteq B_i$, then return B_i as the MEB (success)

(d) Otherwise, for each $p \in P \setminus B_i$, set $w_i(p) \leftarrow 2w_{i-1}(p)$.

We will prove that this algorithm is correct (subject to a suitable choice of ε).

(a) Consider any set system $(X, \mathcal{R})$. The *complimentary* set system is $(X, \overline{\mathcal{R}})$, where $\overline{\mathcal{R}}$ is defined as follows. For each $R \in \mathcal{R}$, $\overline{\mathcal{R}}$ contains the compliment set $\overline{R} = X \setminus R$.

Prove that if $(X, \mathcal{R})$ has VC-dimension d, then its compliment $(X, \overline{\mathcal{R}})$ also has VC-dimension d. (Given the fact that the set system of Euclidean ball ranges has constant VC-dimension, this implies that the set system of points outside a Euclidean ball also has the same VC-dimension.)

(b) To establish (partial) correctness, prove that if in step 2(c), if the ball $B_i = \mathrm{MEB}(S_i)$ contains all the points of P, then $B_i = \mathrm{MEB}(P)$.

(c) Prove that after i (unsuccessful) iterations, $W_i(P) \leq n(1+\varepsilon)^i$. (Using the fact that $1 + x \leq e^x$, it further follows that $W_i(P) \leq ne^{\varepsilon i}$.)

(d) Suppose that the ith iteration is not successful. Prove that there exists at least one point of the basis P^* that does not lie in B_i.

(e) Using (d), prove that after i (unsuccessful) iterations, $W_i(P^*) \geq 2^{i/3}$. (**Hint:** You may use the fact given in Lecture 19 that the growth rate is slowest if the points of P^* are doubled equally often.)

(f) Based on (c) and (e), prove that $i \leq \lg n/((1/3) - \varepsilon \lg e)$.

(g) Based on (f), specify a (constant) value of ε such that the algorithm is guaranteed to terminate within $O(\log n)$ iterations.

(h) Assuming that an ε-net can be computed in $O(n)$ time, show that this algorithm runs in time $O(n \log n)$.

Hint: The analysis is structurally similar to the one of Lecture 19 on the iterative reweighting hitting set algorithm. I would suggest consulting the latex lecture notes, which has a bit more details than the hand-written notes.

Problem 4. In this problem, we will be planning the motion of a line-segment robot $\mathcal{R}$ in the plane amidst a collection of obstacles consisting of n disjoint obstacles $\mathcal{P} = \{P_1, \ldots, P_n\}$. Each obstacle is an axis-parallel rectangle. In particular $P_i = [x_i^-, x_i^+] \times [y_i^-, y_i^+]$. The robot is 2-units long, and its reference position oriented vertically with its midpoint at the origin (see Fig. 4(a)). The robot is restricted to two types of motion:

Translation: It can translate through a vector $t = (t_x, t_y)$, moving its reference point from its current position p to $p + t$.

Rotation: The robot can rotate about its reference point by either $+90°$ (that is, counterclockwise) or $-90°$ (that is, clockwise). When a rotation is performed, the entire circular arc swept out by the any point on the segment must be free of any obstacles. (Think of the segment as spinning in the plane—not picking it up, rotating it, and putting it down.)

Therefore, the robot's *configuration* consists of a point (x, y) where its center point is located and its orientation $o \in \{V, H\}$, where "V" indicates that the robot is parallel with the y-axis and "H" indicates that it is parallel with the x-axis.

A motion plan consists of a sequence of translations and rotations. Note that the robot either translates or rotates. It cannot translate while rotating.

(a) For each of the possible orientations of the robot ("H" or "V"), describe the shape of the corresponding collision obstacle $\mathcal{C}_{\mathcal{R}}(P_i)$ (in terms of the parameters $x_i^-, x_i^+, y_i^-, y_i^+$).

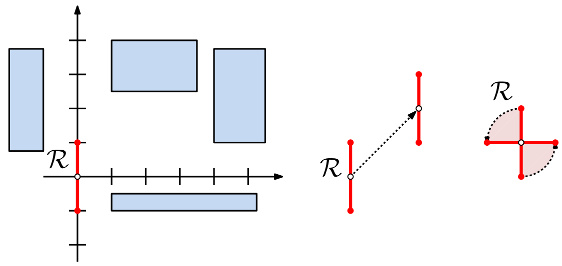

(a) (b)

Figure 4: Motion planning for a rotating/translating segment.

(b) For each of the possible orientations of the robot ("H" or "V"), describe the shape of the corresponding collision obstacle $\mathcal{C}_{\mathcal{R}}(P_i)$ in the cases of $+90°$ (counterclockwise) and $-90°$ (clockwise) rotation (in terms of the parameters $x_i^-, x_i^+, y_i^-, y_i^+$). There are four shapes in all and their boundaries will involve circular arcs.

(c) Given your answers to (a) and (b), present an algorithm to determine whether there exists a motion plan from an arbitrary starting placement configuration $s = (x_s, y_s, o_s)$ to a given target $t = (x_t, y_t, o_t)$. You may assume that both s and t are collision-free.

Hint: I'm really looking for a high level description of how to combine reachability among the various collision obstacles. Efficiency is not a huge consideration, but your solution should run in polynomial time in n.

Challenge Problem. Provide a better upper bound on the VC-dimension for horizontal double wedges. (This is intentionally open-ended, and you don't need to have the upper bound exact to get full credit for this problem.)