Today

4-16-26

- Review Interior Point
- Algorithms for interior point
- Intro to PAC Learning
- Interior point for learning thresholds

Tuesday: using interior point to compute any function

X = totally ordered domain

ex: $X = \mathbb{R}$, $X \subseteq [0,1]$, $X = \{1, 2, \dots, |X|\}$

Def An algorithm $A : X^n \to X$ solves the interior point problem on X with error β if $\forall X \in X^n$,

$$\Pr\left[\min_{x_i \in X} x_i \le A(X) \le \max_{x_i \in X} x_i\right] \ge 1 - \beta.$$

Thm [BNSV15] $\varepsilon, \beta \in (0, 1/4)$ and $\delta(n) \le \frac{1}{20n}$. For every $n \in \mathbb{N}$, solving the interior point problem on X with error β requires

$$n = \Omega(\log^* |X|)$$

$\log^*(z)$ = iterated logarithm

= # times repeat $\log(\log(\cdots \log(z))$ to get ≤ 1

Slow-growing: $\log^*(z^{65536}) = 5$

Q: Lower bound applies for worst-case datasets. Usually we care about iid data. Does it apply?

A: Yes. Suppose $A(X)$ worked whp for data iid from any distribution. But for any dataset $X \sim D_n$ ← hard dist over datasets

$$P_X = \frac{1}{n}\sum_{i=1}^{n} \delta_{x_i}$$ is a distribution.

Can avoid lbd with assumptions on distribution (eg Gaussian, or weaker: Aliakbarpour et al., 2023 https://arxiv.org/abs/2305.13440)

Algorithms for Interior Point

https://arxiv.org/abs/2211.06387

Theorem [Cohen, Lyu, Nelson, Sarlós, Stemmer 2022]
There is an (ε, δ)-DP algorithm for interior point on X with error δ using $n = O\left(\log^* |X| \cdot \frac{(\log 1/\delta)^2}{\varepsilon}\right)$ samples.

Complicated!

We've already seen one algorithm for the task: exponential mechanism for median.

utility $u(y; X) = \left|\frac{n}{2} - \#\{i : x_i \le y\}\right|$

Sample $p(y) \propto \exp\left(-\frac{\varepsilon}{2}\left|\frac{n}{2} - \#\{i : x_i \le y\}\right|\right)$

ε-DP b/c exponential mechanism.

Utility analysis: at least one y^* has

$$\exp\left[-\frac{\varepsilon}{2}\left|\frac{n}{2}-\frac{n}{2}\right|\right) = e^0$$

$$\Pr[A(X) \text{ not IP}] \le \frac{\Pr[A(X) \text{ not IP}]}{\Pr[A(X)=y^*]}$$

$$\le \frac{\#\{\text{points outside data}\} \cdot \exp\left(-\frac{\varepsilon}{2}\frac{n}{2}\right)}{1}$$

$$\le |X| \cdot \exp\left(-\frac{\varepsilon n}{4}\right)$$

less than β when $n \ge \frac{4\log(|X|/\beta)}{\varepsilon}$.

Optimal for pure-DP.

PAC Learning

PAC = "probably approximately correct"

Simple, well-studied learning setting.

binary classification.

Data domain X

Loss $L_D(f) = \Pr_{(x,y)\sim D}[f(x) \neq y]$

Hypothesis class $\mathcal{H} \subseteq \{f : X \to \{0,1\}\}$

Def Distribution D over $X \times \{0,1\}$ is realizable by $\mathcal{H}$ if $\exists f^* \in \mathcal{H}$ s.t.

$$\Pr_{(x,y)\sim D}[f^*(x) = y] = 1$$

Def An algorithm $A : X^n \to \mathcal{H}$ is an (α, β) PAC learner if for all realizable D, with probability at least $1-\beta$, for $X \sim D^{\otimes n}$ $A(X) = f$ satisfies

$$L_D(f) \leq \alpha.$$

Example Threshold functions.

X totally ordered.

$$\mathcal{H} = \{ f_t^{\text{thresh}} : t \in X \}$$

where $$f_t^{\text{thresh}}(x) = \begin{cases} 1 & \text{if } x \geq t \\ 0 & \text{if } x < t \end{cases}$$

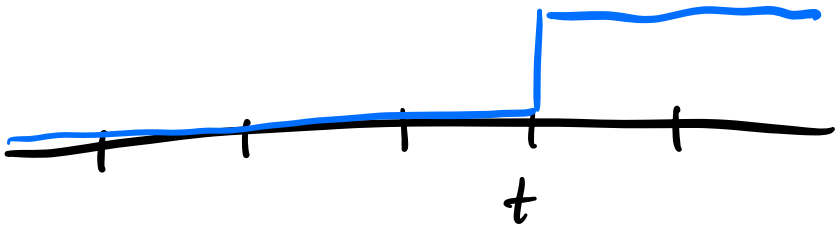

Without privacy, class is (α, β)-PAC learnable with $n = \Theta\left(\frac{\log(1/\beta)}{\alpha}\right)$ samples.

Idea: w.p. $1-\beta$, must output function that correctly labels $1-\alpha$ fraction of all probability mass

Learning Thresholds with DP

Interior point $\Leftrightarrow$ threshold learning are equivalent (almost)

$\Rightarrow$ Suppose A is (ε, δ)-DP and an (α, β)-PAC learner for thresholds on samples of size $\approx \frac{n}{\alpha}$

Given n examples, want an interior point.

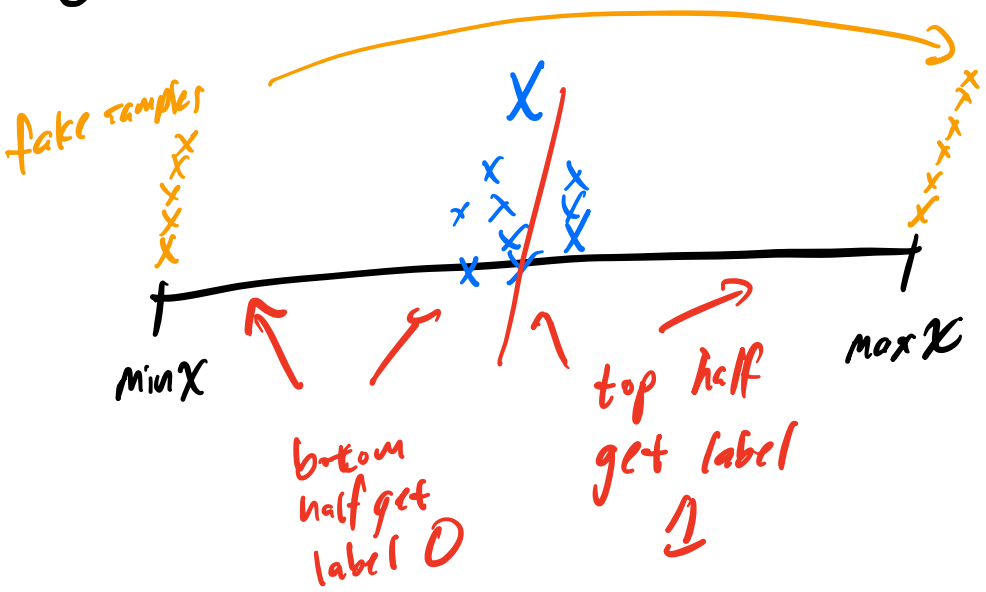

n real samples

$\frac{n}{\alpha} - n = n\left(\frac{1}{\alpha} - 1\right) = \frac{n}{\alpha}(1-\alpha)$ fake samples.

Observe: if A doesn't give an IP for real samples, will miss $\approx \frac{1}{\alpha}$ fraction of overall mass

$\Rightarrow$ not allowed for (α, β)-PAC learning.

(tough)

Exercise | Show how to use an algorithm for interior point on n points can be used for threshold learning on $O(n/\alpha)$ points.

References:

Links above, plus BNSV15 & Bun's dissertation (linked last lecture)