

11/13/23

Lecture 19: Cryptography + Learning

- ① Hardness of LPN \rightarrow Hardness of agnostically learning halfspaces
 - ② Daniely-Vardi lifting: Crypto hardness for learning MLP's.
-

① Recall agnostic learning: for function class \mathcal{C} ,

Given: $(x_1, y_1), \dots, (x_n, y_n) \sim \mathcal{D}$ over $\mathbb{R}^d \times \{\pm 1\}$

Goal: output $f \in \mathcal{C}$ st.

$$\mathbb{E}_{x, y} [f(x) \neq y] \leq \min_{f \in \mathcal{C}} \mathbb{E}_{x, y} [f(x) \neq y]$$

When $\mathcal{D}_x = \text{Unif}(\{0, 1\}^d)$ and \mathcal{C} is {halfspaces},
how hard is this task?

[Kalai-Klivans-Mansour-Nisan '06]:

- A) Halfspaces approx'd by deg $\text{poly}(1/\epsilon)$ polynomials,
so can learn in time $d^{\text{poly}(1/\epsilon)}$.
- B) This is qualitatively tight, assuming hardness of LPN

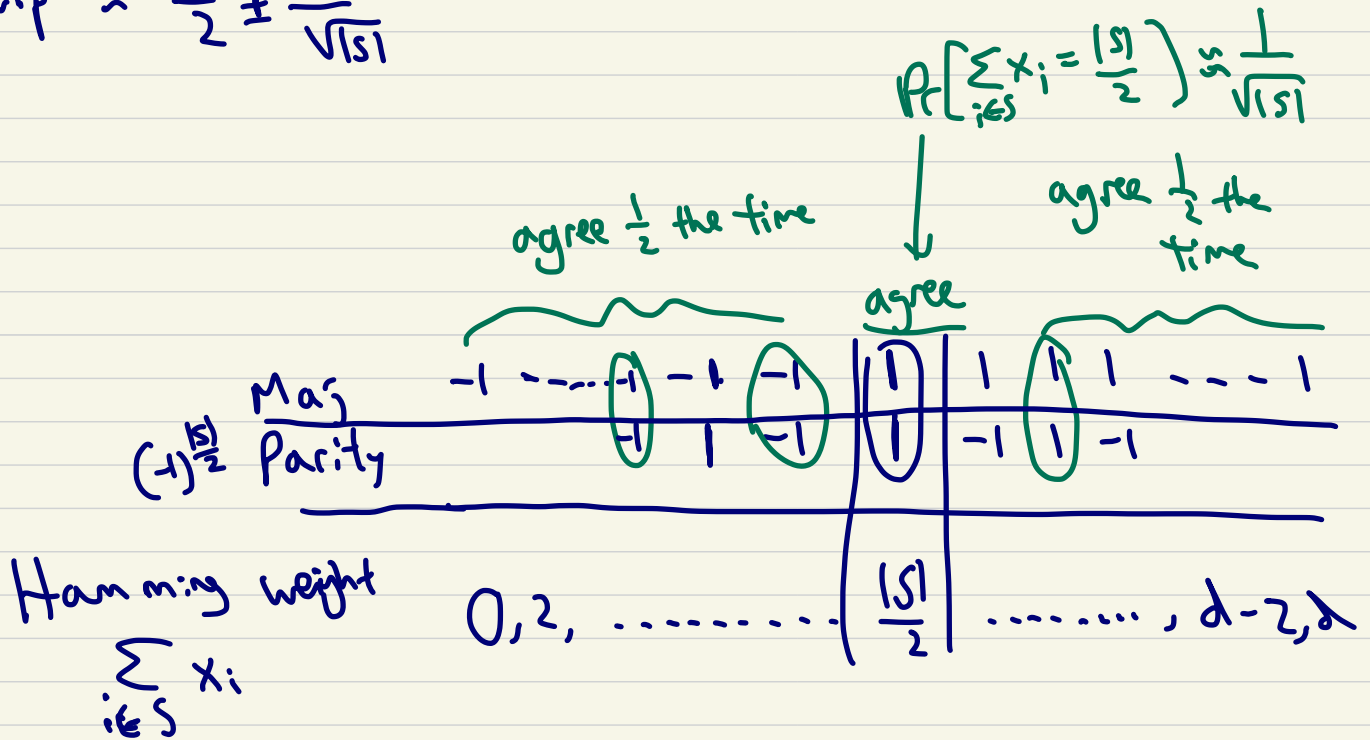
Pf of B):

Given $S \subseteq [d]$ of even size, note that

$$\text{Maj}_S(x) = \begin{cases} 1 & \text{if } \sum_{i \in S} x_i \geq \frac{|S|}{2} \\ -1 & \text{o.w.} \end{cases}$$

agrees with $\text{parity}_S(x) = \begin{cases} 1 & \text{if } \sum_{i \in S} x_i \text{ even} \\ -1 & \text{o.w.} \end{cases}$

$$\text{wip} \approx \frac{1}{2} \pm \frac{1}{\sqrt{|S|}}$$



So majority and parity have correlation $\frac{1}{\sqrt{|S|}}$.

noise in parity reduces this correlation to $\frac{1-2\gamma}{\sqrt{|S|}} \approx \frac{1}{\sqrt{n}}$

So if we could agnostically learn majorities

(special case of halfspaces) to error $\epsilon = \frac{1}{\sqrt{n}}$, would get alg. for noisy parity.

e.g. if alg for the former ran in time $n^{O(1/\epsilon^{2-\beta})}$, would imply $2^{O(n^\beta)}$ alg for noisy parity.

② Parity - Vardi lifting:

For $n \ll m$, a function $F: \{0,1\}^n \rightarrow \{0,1\}^m$

is a pseudorandom generator if no poly-time adversary can distinguish a sample from

$\text{Unif}(\{0,1\}^m)$ from a sample from

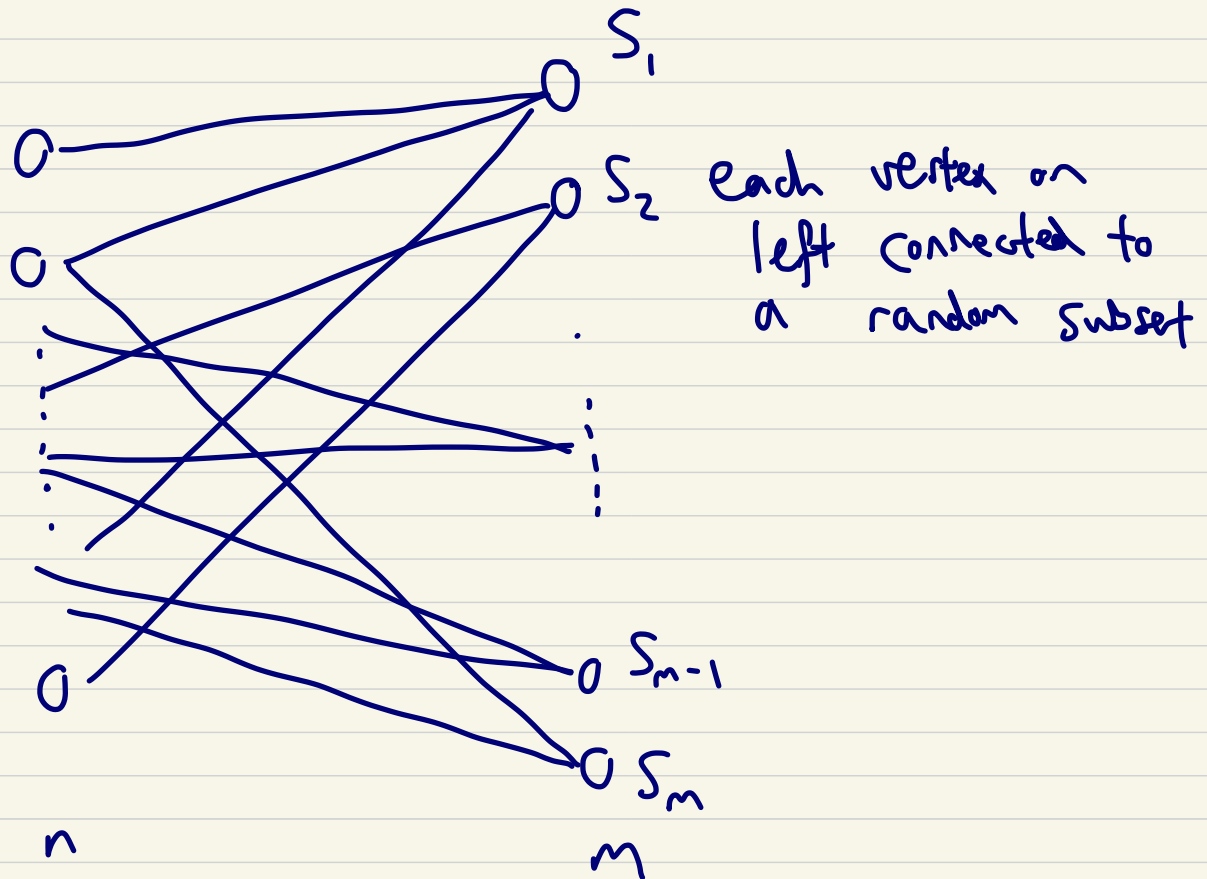
$F(\text{Unif}(\{0,1\}^n))$ with non-negl. success prob.

Goldreich's PRG:

Let $P: \{0,1\}^k \rightarrow \{0,1\}$ be a predicate, e.g.

$$P_k = \text{XOR-MAJ}_{a,b}(z) = (z_1 \oplus \dots \oplus z_a) \oplus \text{Maj}(z_{a+1}, \dots, z_{a+b}).$$

For some constant k , sample m random subsets S_1, \dots, S_m of $[n]$ of size k



Define

$$F(z) = \left(\underbrace{P(z|S_1)}_{\substack{\text{restriction} \\ \text{of } x \in \{0,1\}^n \\ \text{to bits in } S_1}}, \dots, P(z|S_m) \right)$$

"Goldreich's PRG" / "Local PRG"

Crypto assumption: For every constant $s > 1$, there is a constant k and predicate $P: \{0,1\}^k \rightarrow \{0,1\}$ s.t. that Goldreich's construction is a valid pseudorandom generator.

We will use this to prove hardness of learning MLP's.
Strategy:

- ① Hardness over $\{0,1\}^n$
 - ② Naive lifting
 - ③ Daniely-Vardi gadget
-

① Let's first show that under the above assumption, MLPs are hard to learn over $\{0,1\}^n$.

If we have a sample $F(z)$ from PRG, can regard it as a dataset of pairs

$$(S_i, P(z | S_i))$$

each S_i is a random subset of size k ,
and we want to encode this into a
sample from $\{0,1\}^n$

Given $S = \{i_1, \dots, i_k\}$, define $x^S \in \{0,1\}^{kn}$ via:

$$x^S : \underbrace{\text{||||| 0 |||||}}_{i_1} \underbrace{\text{||||||| 0 ||}}_{i_2} \dots \underbrace{\text{|| 0 |||||}}_{i_k}$$

i.e. j th block of S is e_{i_j}

Claim: \exists MLP $N: \{0,1\}^{kn} \rightarrow \{0,1\}$ s.t.

$$N(x^S) = P(z|S).$$

Pf: (tedious, included for completeness):

the function $x^S \mapsto P(z|S)$ can be implemented

as a DNF:

$$\bigvee_{\substack{b \in \{0,1\}^k \\ \text{s.t. } P(b)=1}} \bigwedge_{\substack{j \in [k] \\ \ell: z_\ell \neq b_j}} x_{j,\ell}$$

Can implement as a relu: if
there are M literals in this
conjunction, then take

$$\text{ReLU}\left(\sum_{j \in (k)} \sum_{l: z_l \neq b_j} x_{j,l} - (M-1)\right)$$

Note: at most one conjunction satisfied, so
can implement by simply summing the neurons:

$$\sum_{\substack{b \in \{0,1\}^k \\ P(b)=1}} \text{ReLU}(\dots)$$

□

Implies learning one-hidden-layer MLPs over
the distribution over strings

$$x^S : \underbrace{\text{||||| 0 |||||}}_{i_1} \underbrace{\text{||||| 0 ||}}_{i_2} \dots \underbrace{\text{|| 0 |||||}}_{i_k} \in \{0,1\}^{kn}$$

is hard (if learner achieves nontrivial test error
we know that we are in the pseudorandom scenario
and can distinguish).

II) Naive lifting

Now want to show hardness over Gaussian inputs.

Initial idea:

$$\underbrace{\text{dist. over}}_{\text{||||| 0 |||||}} \underbrace{\text{|||||}}_n \approx \text{Ber}\left(\frac{n-1}{n}\right)^n$$

So given sample x^S , can "Gaussianize" as follows:

if in block j , i th entry of x^S is

0, then draw $g_{j,i}^S \sim N(0,1) \mid \geq t$

1, then draw $g_{j,i}^S \sim N(0,1) \mid < t$

for t s.t. $\Pr_{g \sim N(0,1)} [g \geq t] = \frac{1}{n}$, so that

if we apply $\text{thres}(g) \triangleq \mathbb{1}[g \geq t]$ to $g \sim N(0, \text{Id}_{kn})$, we get a sample from $\text{Ber}(\frac{n-1}{n})^{\otimes kn}$.

Naive attempt: define N' by

$$N'(g) = N(-\text{thres}(g))$$

(applied entrywise)

FAILS for several reasons, one of which is that g need not encode a subset

Given $(S_1, P(z|S_1)), \dots, (S_m, P(z|S_m))$, define

Gaussian dataset: for every $i \in [m]$,

- sample $g \sim N(0, \text{Id}_{kn})$

- if $\text{thres}(g)$ is valid encoding of a subset S

(i.e. has exactly one 0 in each block);

- permute its entries so that $\text{thres}(g)$ encodes S_i
 - add example $(g, P(z|S_i))$ to the dataset
- otherwise, add $(g, 0)$ to the dataset

Claim: \exists MLP N_{naive} that labels this dataset perfectly (caveat, requires Sign activations)

PF: First, note \exists MLP N_{encode} s.t.

$$N_{\text{encode}}(g) = \begin{cases} 0 & \text{if } \text{thres}(g) \text{ is not valid encoding,} \\ \text{large o.w.} \end{cases}$$

Take

$$N_{\text{encode}}(g) \stackrel{\Delta}{=} \prod_{j \in [k]} \sum_{i \in \text{block } j} (\text{thres}(g)_{j,i} - (n-2))$$

some big factor

(this solves problem of g not necessarily encoding a subset)

So $N_{\text{naive}}(g) \stackrel{\Delta}{=} \text{ReLU}(N(\text{thres}(g)) - N_{\text{encode}}(g))$

Correctly labels the Gaussian dataset. \square

One more (major) issue: **thres is a discontinuous function!**

would need infinite weights to implement w/ ReLUs, or super-poly-sized weights to approximate sufficiently well...

III Daniely-Vardi lifting:

Instead of



Consider

ramp:



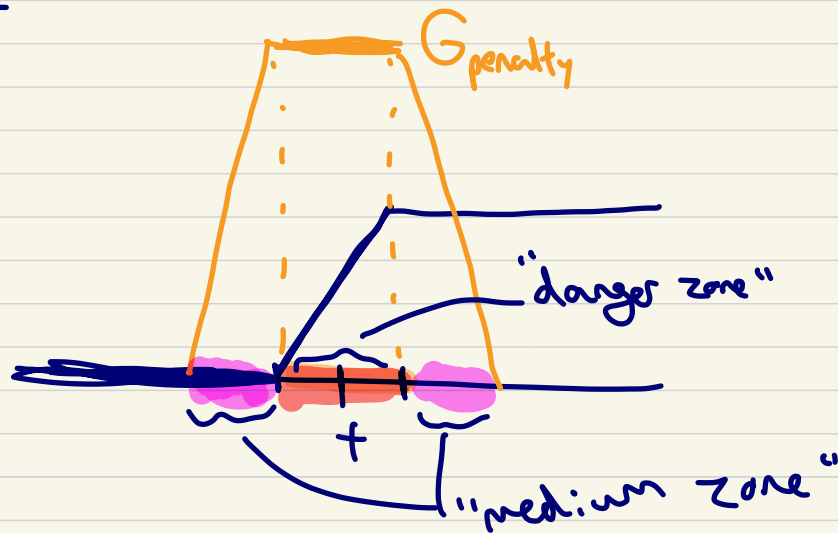
Let N'_{encode} be N_{encode} with $\text{thres} \rightarrow \text{ramp}$.

because we can only afford to use

p_{dy} -sized weights, the intermediate interval is $\geq \frac{1}{p_{dy}(n)}$ wide and we will see

some g 's w/ coordinates landing in the interval, and replacing thres w/ ramp in naive lifting would fail.

Key idea: consider



$G_{\text{penalty}}(x)$ is large whenever g has a coordinate landing in "danger zone".

Consider

$$N_{\text{final}} = \text{ReLU}(N(\text{ramp}(g)) - N'_{\text{encode}}(g) - \sum_{j \in [k], i \in [n]} G_{\text{penalty}}(g_{j,i}))$$

What happens if g falls in "medium zone"?

Idea: in our Gaussian dataset, modify the labels as follows:

- if all coordinates of g outside of danger and medium zones, keep label as before, i.e. $P(z|s_i)$
- if \exists coordinate in danger zone, set label to 0
- if no coordinate in danger zone but some in medium zone, set label to

$$\text{ReLU}(P(z|s_i) - \sum_{j,i} G_{\text{penalty}}(g_{j,i}))$$