Lecture 18

Fourier-based learning algorithms

- the low degree algorithm
- Fourier Concentration
- Noise sensitivity

Recall Favier Trunstormi

$$
\begin{aligned}
& x_{s}(x)=\prod_{i \in s} x_{i} \\
& \left\langle f_{i} g\right\rangle=\frac{1}{2^{n}} \sum_{x} f(x) g(x) \\
& \hat{f}(s)=\left\langle f, x_{s}\right\rangle=1-2 \cdot \operatorname{Pr}\left[f(x) \neq X_{s}(x)\right] \\
& \forall f, f(x)=\sum \hat{f}(s) X_{s}(x) \\
& \text { Plancherel }\langle f, g\rangle=\sum_{s} \hat{f}(s) \hat{g}(s)
\end{aligned}
$$

Learning via Farrier Representation
will look at learning algorithms that are based on estimating Fourier representation of fath $f$ (similar to polynomial interpolation)

Approximating one Fourier coefficient:
lemma for any $S \leq[n]$, can approx 7 no $\hat{f}(s)$ to within additive $\gamma$

$$
\text { (ie. } \mid \text { output }-\hat{f}(s) \mid \leq \gamma \text { ) }
$$

with prob $\geq 1-\delta$ in $O\left(\frac{1}{\gamma^{2}} \log \frac{1}{\delta}\right)$
samples.
(Proved last time)

The low degree algorithm
definition of fetus for which low degree
Fourier coeffs pretty much suffice to describe fats:
def $f:\{ \pm 1\}^{n} \rightarrow \mathbb{R}$ has $\alpha(\varepsilon, n)$-Fourier concentration

$$
\text { if } \sum_{\substack{s \leq[n]}} \hat{f}(s)^{2} \leq \varepsilon \quad \forall 0<\varepsilon<1
$$

st.

$$
|s|>\alpha(\varepsilon, n)
$$

for Boolean $f$ this implies

$$
\sum_{\substack{s \leq[n] \\ \text { st. } \\|s| \leq \alpha(\varepsilon, n)}} \hat{f}(s)^{2} \geq 1-\varepsilon
$$

examples

1) fate $f$ which depends on $\leq k$ vars $\}$ if $f$ doesit depend

$$
\text { has } \sum_{\substack{s \\|s|>|. \\|s|>k}} \hat{f}(s)^{2}=0
$$ on $X_{i}$ then all $\hat{f}(s)$ for which its satisfy $\hat{f}(s)=0$

Low degree algorithm
approximates fetus with $d \equiv \alpha(\varepsilon, h)$ Fourier concentrations
$\begin{array}{rll}\text { Given: } & d & \text { degree } \\ & \tau & \text { accuracy } \\ & \delta & \text { confidence }\end{array}$

Algorithm:

- Take $m=O\left(\frac{n^{d}}{\tau} \ln \frac{n^{d}}{\delta}\right)$ samples

- For each $S$ st. $|S| \leq d$ :

$$
C_{s} \leftarrow \text { estimate of } \hat{f}(S)
$$

- let $h(x) \equiv \sum_{|s| \leq d} C_{s} \cdot X_{s}(x)$
- output sign (h) as hypothesis

Why does this work?
Two stages:

1) Show that $f$ has low F.C.

$$
\Rightarrow E_{x}\left[(f(x)-h(x))^{2}\right] \text { small }
$$

2) Show that $\operatorname{Pr}[f(x) \neq \operatorname{sign}(h(x))] \leq E_{x}\left[(f(x)-h(x))^{2}\right]$.


First "stage":
 together: $f$ has low FCc. $\Rightarrow \operatorname{sigh}(h(x)$ is good approximation off
$h$ satisfies $E_{x}\left[(f(x)-h(x))^{2} \leq \varepsilon+\tau\right.$ with prob $\geq 1-\delta$ (Proved last time)
$2^{\text {nd }}$ "Stage":
The $2 f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$

$$
h:\{ \pm 1\}^{n} \rightarrow \mathbb{R}
$$

then $\operatorname{Pr}[f(x) \neq \operatorname{sign}(h(x))] \leq E\left[(f(x)-h(x))^{2}\right]$

Proof.

$$
\begin{aligned}
& E\left[(f(x)-h(x))^{2}\right]=\frac{1}{2^{n}} \sum_{x}(f(x)-h(x))^{2} \quad \text { deft. } \\
& \operatorname{Pr}[f(x) \neq \operatorname{sign}(h(x))]=\frac{1}{2^{n}} \sum_{x} 1_{\left\{f(x) \neq \operatorname{sign}^{2}(h(x))\right\}} \begin{array}{l}
\text { compare these the to to the } \\
\text { ter. }
\end{array}
\end{aligned}
$$

Consider " $(f(x)-h(x))^{2}$ "vs. "1 $\left\{_{\{f(x) \neq \operatorname{sign}(h(x))\} \text { " }}\right.$
Case 1 if $f(x)=\operatorname{sign}(h(x))$ :

$$
\begin{aligned}
& 1_{f(x) \neq \sin (h(x))}=0 \\
& (f(x)-h(x))^{2} \geq 0
\end{aligned}
$$

Case 2 if $f(x) \neq \operatorname{sign}(h(x))$ : why? e.g.

$$
\begin{aligned}
& 1_{f(x) \neq \operatorname{sign}(h(x))=1} \\
& (f(x)-h(x))^{2} \geq 1
\end{aligned}
$$

if $f(x)=+1$ them in this axe $h(x)<0$ :


So, $\forall x$

$$
(f(x)-h(x))^{2} \geq 1_{f(x) \neq \operatorname{sign}(h(x))}
$$ analogocos)

Correctness of learning algorithm

The if $C$ has fourier concentration $d=\alpha(\varepsilon, n)$ then there is a $q=O\left(\frac{n^{d}}{\varepsilon} \log \frac{n^{d}}{\delta}\right)$ sample uniform distribution learning algorithm for $C$ ie. algorithm gets $q$ samples + with prob $\geq 1-\delta$ outputs $h^{\prime}$ sit. $\operatorname{Pr}\left[f \neq h^{\prime}\right] \leqslant 2 \cdot \varepsilon$

Pf.
run low degree alg with $\tau=\varepsilon$ $t_{\text {th }} \Rightarrow$ get $\quad h$ s.t. $E\left[(f-h)^{2}\right] \leq \varepsilon+\varepsilon=2 \varepsilon$ output $h^{\prime}=\operatorname{sign}_{\uparrow}(h)$
thu $2 \Rightarrow h^{\prime}$ has error $\leq 2 . \varepsilon$

Applications

1) Bounded depth decision trees

$$
\begin{aligned}
& f(x)=\sum_{l \in \text { laves }} \underbrace{f_{l}(x)}_{\substack{\text { of } T}} \cdot \underbrace{\operatorname{val}(l)}_{\substack{\text { cost } \\
\text { fhhinh } \\
\text { dependson } \\
\text { step on many } \\
\text { vars }}} \\
& \hat{f}(s)=\sum_{\text {val }} \text { vel } \underbrace{\hat{f}_{l}(s)}_{\substack{0 \text { for } \\
|S|>\text { depth }}} \quad \text { linearity } \\
& \Rightarrow \quad \forall S \text { st. }|s|>\operatorname{depth}, \quad \hat{f}(s)=0
\end{aligned}
$$

so $O\left(\frac{n}{}_{\text {depth }^{\varepsilon}} \log \frac{n^{\text {depth }}}{\delta}\right)$ suffices
2) Constant depth cats
def. "Boolean CRt C" is DAG
gates: $\underbrace{1_{1}, v_{1}}_{\text {operations }}, \underbrace{1,0,}_{\text {cons ts }} \underbrace{x_{1} \cdots x_{n}}_{\text {vars }}$
how many inputs? const, poly, unbounded?

can we compute parity of $n$ bits in const depth?

Yes! can compute any fetn on $n$ bits in const depth "Karnaugh maps"
parity in const depth, poly size?
no! [Furst Save Sipser] lemon Switching lemma
lemons $\Rightarrow$ kmonade:
Th m [Hastad, Linial Mansour Nisan]
If computable via sizes depth $d$ ckts

$$
\sum_{|s|>t} \hat{f}^{2}(s) \leq \alpha \quad \text { for } \quad t=O\left(\log \frac{s}{\alpha}\right)^{d-1}
$$

take $s=\operatorname{pol}(n)\}$

$$
\left.\begin{array}{l}
s=\text { poly }(n) \\
d=\text { cons } \\
\alpha=O(\varepsilon)
\end{array}\right\} \Rightarrow t=f\left(\log ^{d}\left(\frac{n}{\varepsilon}\right)\right)
$$

yields $n^{O\left(\log ^{d}\left(\frac{n}{\varepsilon}\right)\right)}$ sample algorithm
(can improve to $n^{\text {O(loglogn } n}[$ Jackson $]$ )
(recall parity of $s$ will hare 1 legree $|s|$ lane, tourer coff of
3) Learning halfspaces
def. $h(x)=\operatorname{sign}(w \cdot x-\theta)$ is "halfspace function"

The Let $h$ be halfspace over $\{ \pm 1\}^{n}$ then $h$ has f.c. $\alpha(\varepsilon)=\frac{c}{\varepsilon^{2}}$
(ie. $\sum_{|s| \geq \frac{c}{\varepsilon^{2}}} \hat{h}(s)^{2} \leq \varepsilon$ )
(will prove soon)
Corr low degree alg learns half spaces under unif dist with $n^{0\left(1 / \varepsilon^{2}\right)}$ unit. samples.

Cactually $O\left(n^{5}\right)$ sample algorithms exist, but this approach will have "big win" soon)

Key idea:
Noise Sensitivity use to bound Fourier concentration
def. "Noise operator" $0<\varepsilon<1 / 2$
$N_{\varepsilon}(x)=$ randomly flip each bit of $x$ with prob \&
def "Noise sensitivity"

$$
N S_{\varepsilon}(f)=\operatorname{Pr}_{\substack{x \in\left\{ \pm 1 j^{n} \\\right. \text { noise }}}\left[f(x) \neq f \mathbb{N}_{\varepsilon}(x)\right]
$$

Examples

1. $f(x)=x_{1} \quad n s_{\varepsilon}(f)=\varepsilon$
2. 

$$
\begin{aligned}
f(x)=x_{1} x_{2} \ldots x_{k} \quad n s_{\varepsilon}(f) & =\operatorname{Pr}\left[f(x)=F+f\left(N_{\varepsilon}(x)\right)=T\right] \\
& +\operatorname{Pr}\left[f(x)=T \alpha f\left(N_{\varepsilon}(x)\right)=F\right] \\
\text { by symmetry } \rightarrow & =2 \cdot \operatorname{Pr}\left[f(x)=T \alpha f\left(N_{\varepsilon}(x)\right)=F\right] \\
\text { if } \varepsilon \ll \frac{1}{k} \approx \frac{1}{2^{k-k}(\varepsilon k)} \rightarrow & \frac{2}{2^{k}}\left(1-(1-\varepsilon)^{k}\right)
\end{aligned}
$$

3. 

$$
\begin{aligned}
& f(x)=\operatorname{Maj}\left(x_{1} \cdots x_{n}\right) \\
& n s_{\varepsilon}(f)=O(\sqrt{\varepsilon})
\end{aligned}
$$

Sketch:
$\operatorname{Maj}(x) \sim$ random walk on line starting at 0

egg. $x=\left(\begin{array}{cccc}11 & -1 & -1 & -1\end{array}\right) \quad$ well known fact;

$N_{\varepsilon}(x) \sim \operatorname{random}$ walk on $\underbrace{\varepsilon n \text { bits }}_{\text {each flip }}$ displaces by

$$
E[\text { displacement }]=2 \sqrt{\varepsilon n}
$$

$$
\begin{aligned}
& (-1 \rightarrow+1 \text { or } \\
& +1 \rightarrow-1)
\end{aligned}
$$

Process: take walk specified by $x+$ continue walk according to

$$
N_{\varepsilon}(x) \cdot 2
$$

heuristic argument:
pretend first walk leaves us at $\sqrt{n}$
$\operatorname{Pr}\left[2^{\text {nd }}\right.$ walk takes us across 0$]$

$$
\begin{aligned}
& =\frac{1}{2} \operatorname{Pr}[2^{\text {nd }} \text { displacement }>\underbrace{\sqrt{\varepsilon n}}_{\frac{1}{2 \sqrt{n}} \cdot 2 \sqrt{\varepsilon n}} \\
& \leq 2 \sqrt{\varepsilon} \quad \text { by Markov's } \neq
\end{aligned}
$$

4. any LTF (1/2 space)

$$
T h_{m}(\text { Pres }) \quad N S_{\varepsilon}(L T F)<8.8 \sqrt{\varepsilon}
$$

best possible since $N S_{\varepsilon}\left(M_{a_{j}}\right)=\theta(\sqrt{\varepsilon})$
5. Parity fetus $X_{s}(x)$ for $|s|=k$

$$
\begin{aligned}
n s_{\varepsilon}(f) & \left.=\operatorname{Pr} \text { [odd \# bits in } S \text { flipped by } N_{\varepsilon}\right] \\
& =\frac{1-(1-2 \varepsilon)^{k}}{2} \text { for }||s|=1: \varepsilon
\end{aligned}
$$

