

Lecture 21 (Part I)

average vs. worst case complexity

Last time:

Weak vs. Strong Learning

Def. Algorithm A "weakly PAC learns" concept class \mathcal{C} if $\exists \gamma > 0$

st. $\forall c \in \mathcal{C} + \forall$ dists \mathcal{D}

$\forall \delta > 0$ \leftarrow ($\delta = \frac{1}{4}$ or $\frac{1}{n^2}$ doesn't affect)

with prob $\geq 1 - \delta$

given examples of c (labelled)

in "strong" learning this is $1 - \epsilon$

A outputs h s.t. $\Pr_{\mathcal{D}} [h(x) = c(x)] \geq \frac{1}{2} + \gamma$

not good compared to $1 - \epsilon$ or 99%

↑ advantage over guessing

Assume $c \in \mathcal{C}$ maps $X \rightarrow \{0, 1\}$
← domain

\mathcal{C} is concept class where s bounds size of WL-description of concepts
← parametrized by "size" n of inputs
← fctn of n

Thm \mathcal{C} weakly learnable $\Rightarrow \exists$ efficient algorithm

Strong learner

- using $\frac{\text{poly}(n, s, \log \frac{1}{\epsilon}, \log \frac{1}{\delta})}{\epsilon}$ samples & time
- Outputs hypotheses of size $\text{poly}(n, s, \log \frac{1}{\epsilon})$ (useful later)

(\star can evaluate in poly time)

Why? given \mathcal{A} learning \mathcal{C}
use \mathcal{A} with $\epsilon_0 = 1/4$
boost \mathcal{A} to arbitrary ϵ

Boosting \Rightarrow

average case complexity insights:

Corr \mathcal{C} learnable \Rightarrow all concepts $c \in \mathcal{C}$
have poly sized ckt

Pf idea

$\forall c \in \mathcal{C}$, use $\epsilon < \frac{1}{2^n}$ (e.g. no error)
in boosting

will output consistent hypothesis of

poly size in $n (= \log \frac{1}{\epsilon}) + |c|$

\leftarrow # bits used by WL
to describe c

that is poly time
evaluable.

\Rightarrow poly sized ckt.



Thm. Suppose f cannot be computed by poly sized ckt's. Then there is a sequence of distributions $\{\mathcal{D}_n^*\}_{n=1}^\infty$ st. f is "average-

case hard" on $\{\mathcal{D}_n^*\}_{n=1}^\infty$

no poly sized ckt gets f right more than $\frac{1}{2} + \frac{1}{\text{poly}(n)}$ of time

need hard dist for each input size to make it well defined

why nontrivial?
 output for any single input can be hardwired into ckt. \Rightarrow no "hard" inputs. how to find a set of "collectively hard" inputs?

Pf idea

if not, f can be weakly-learned by poly sized ckt's

- (boosting) \Rightarrow can strongly learn in size $\text{poly}(\log \frac{1}{\epsilon}, \dots)$
- (previous corollary) \Rightarrow can learn f with 0 error
- \Rightarrow f computable by poly-sized ckt's