## 10-301/10-601 Introduction to Machine Learning

Machine Learning Department
School of Computer Science
Carnegie Mellon University

## PAC Learning

Slides from Matt Gormley
Lecture 14
Mar. 11, 2024

LEARNING THEORY

## Questions for today (and next lecture)

1. Given a classifier with zero training error, what can we say about true error (aka. generalization error)? (Sample Complexity, Realizable Case)
2. Given a classifier with low training error, what can we say about true error (aka. generalization error)?
(Sample Complexity, Agnostic Case)
3. Is there a theoretical justification for regularization to avoid overfitting? (Structural Risk Minimization)

PAC/SLT Model for Supervised ML $\underbrace{x^{(1)} \sim p^{*}(\cdot)}_{x^{(3)}}$


## PAC/SLT Model for Supervised ML

- Problem Setting
- Set of possible inputs, $\mathbf{x} \in \mathcal{X}$ (all possible patients)
- Set of possible outputs, $y \in \mathcal{Y}$ (all possible diagnoses)
- Distribution over instances, p*(•)
- Exists an unknown target/function, c* $: \mathcal{X} \rightarrow \mathcal{Y}$ (the doctor's brain) labeliy
- Set, $\mathcal{H}_{\mathcal{H}}$ of candidate hypothesis functions, $\mathrm{h}: \mathcal{X} \rightarrow \mathcal{Y}$ (all possible decision trees)
- Learner is given $N$ training examples
$\mathrm{D}=\left\{\left(\mathbf{x}^{(1)}, \mathrm{y}^{(1)}\right),\left(\mathbf{x}^{(2)}, \mathrm{y}^{(2)}\right), \ldots,\left(\mathbf{x}^{(\mathrm{N})}, \mathrm{y}^{(\mathrm{N})}\right)\right\}$
where $x^{(i)} \sim p^{*}(\cdot)$ and $y^{(i)}=c^{*}\left(x^{(i)}\right)$
(history of patients and their diagnoses)
- Learner produces a hypothesis function, $\hat{y}=h(x)$, that best approximates unknown target function $y=c^{*}(x)$ on the training data


## IMPORTANT NOTE

## In our discussion of PAC Learning, we are only concerned with the problem of binary classification

There are other theoretical frameworks (including PAC) that handle other learning settings, but this provides us with a representative one.

## PAC/SLT Model for Supervised ML



## Two Types of Error

## 1. True Error (aka. expected risk)

$$
R(h)=\underbrace{}_{\mathbf{x} \sim p^{*}(\mathbf{x})}\left(c^{*}(\mathbf{x}) \neq h(\mathbf{x})\right)
$$

2. Train Error (aka. empirical risk)


$$
\begin{aligned}
\hat{R}(h) & =P_{\underbrace{\mathbf{x} \sim \mathcal{S}}} \\
& =\frac{1}{N} \sum_{i=1}^{N} \underbrace{\left.c^{c^{*}\left(\mathbf{x}^{(i)}\right)} \neq h\left(\mathbf{x}^{(i)}\right)\right)}_{y^{*}(\mathbf{x}) \neq h(\mathbf{x})} \\
& =\frac{1}{N} \sum_{i=1}^{N} \mathbb{1}\left(y^{(i)} \neq h\left(\mathbf{x}^{(i)}\right)\right)
\end{aligned}
$$

where $\left.\mathcal{S} \equiv\left\{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}\right)\right\}_{i=1}^{N}$ is the training data set, and $\mathbf{x} \sim$ $\mathcal{S}$ denotes that x is sampled from the empirical distribution.

## PAC / SLT Model

1. Generate instances from unknown distribution $p^{*}$

$$
\begin{equation*}
\mathbf{x}^{(i)} \sim p^{*}(\mathbf{x}), \forall i \tag{1}
\end{equation*}
$$

2. Oracle labels each instance with unknown function $c^{*}$

$$
\begin{equation*}
y^{(i)}=c^{*}\left(\mathbf{x}^{(i)}\right), \forall i \tag{2}
\end{equation*}
$$

3. Learning algorithm chooses hypothesis $h \in \mathcal{H}$ with low(est) training error, $\hat{R}(h)$

$$
\begin{equation*}
\hat{h}=\underset{\sim}{\operatorname{argmin}} \hat{R}(h) \tag{3}
\end{equation*}
$$

4. Goal: Choose an $h$ with low generalization error $R(h)$

## Three Hypotheses of Interest

The true function $\widehat{c}_{*}^{*}$ is the one we are trying to learn and that labeled the training data:

$$
\begin{equation*}
y^{(i)}=c^{*}\left(\mathbf{x}^{(i)}\right), \forall i \tag{1}
\end{equation*}
$$

The expected risk minimizer has lowest true error:

$$
\begin{equation*}
\text { best-in-class } \quad h^{*}=\underset{h \in(t)}{\operatorname{argmin}} \underbrace{R}_{1}(h) \tag{2}
\end{equation*}
$$

The empirical risk minimizer has lowest training error:

$$
\begin{equation*}
\hat{h}=\underset{h \in \mathcal{H}}{\operatorname{argmin}} \hat{R}(h) \tag{3}
\end{equation*}
$$

## Three Hypotheses of Interest

$$
y^{(i)}=c^{*}\left(\mathbf{x}^{(i)}\right), \forall i
$$

$$
h^{*}=\underset{h \in \mathcal{H}}{\operatorname{argmin}} R(h)
$$

## Question: True or False: h* and c* are always equal.

## Answer:

## Three Hypotheses of Interest

$$
y^{(i)}=c^{*}\left(\mathbf{x}^{(i)}\right), \forall i
$$

$$
h^{*}=\underset{h \in \mathcal{H}}{\operatorname{argmin}} R(h)
$$

Question: True or False: h* and c* are always equal.
Answer:



PAC LEARNING

## PAC Learning

- Q : Can we bound $\mathrm{R}(\mathrm{h})$ in terms of $\hat{R}(\mathrm{~h})$ ?
- A:Yes!
- PAC stands for

> Probably $\longrightarrow$ canlidence Approximately $\longrightarrow$ estimation

## Correct

A PAC Learner yields a hypothesis $h \in \mathcal{H}$ which is... approximately correct
$R(h) \approx \neq R\left(h^{*}\right) \leftarrow$ with high probability $\operatorname{Pr}\left(R(h) \approx \emptyset^{\prime}\right) \approx 1$
$R\left(h^{*}\right)$

Probably Approximately Correct (PAC) Learning

PAC Criterion
$\forall h \in \$ \notin \operatorname{Pr}(|R(h)-\hat{R}(h)|<\varepsilon) \geqslant 1-\delta$

$R(h)$
$\hat{R}(h)$ is defined u.r.t. D and $D$ is randan sample from $\rho^{*}$

Sample Complexity is the min number of train examples $N(\varepsilon, \delta)$ sit. the PAC criterion is satisfied for $\varepsilon_{1} \delta^{\circ}$

Consistent Learner $C^{\star} \in \notin$ A hypothesis $h \in \mathcal{F} \in \dot{b}$ consistent with training data $D$ if $\underset{\underset{D}{\vec{R}}(h)}{\underset{\sim}{n}(h)}=0$

SAMPLE COMPLEXITY RESULTS

## Sample Complexity Results

Definition 0.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1 ). best-in-class $h^{*}$

Four Cases we care about...

> We'll start with the finite case...

## Probably Approximately Correct (PAC) Learning

Theorem 1: Realizable Case, Finite $|\mathrm{H}|$

## Sample Complexity Results

Definition 0.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1).

Four Cases we care about...

|  | Realizable Agnostic |  |
| :---: | :---: | :---: |
| Finite $\|\mathcal{H}\|$ | Thm. $1 \sqrt{N \geq \frac{1}{\epsilon}\left[\log (\|\mathcal{H}\|)+\log \left(\frac{1}{\delta}\right)\right]} \operatorname{la}$ beled examples are sufficient so that with probability $(1-\delta)$ all $h \in \mathcal{H}$ with $\hat{R}(h)=0$ have $R(h) \leq \epsilon$. |  |
| Infinite $\mid \mathcal{H}$ |  |  |

## Example: Conjunctions



Thm. $1 \quad N \geq \frac{1}{\epsilon}\left[\log (\mathcal{H})+\log \left(\frac{1}{\delta}\right)\right]$ labeled examples are sufficient so that with probability $(1-\delta)$ all $h \in \mathcal{H}$ with $\hat{R}(h)=0$ have $R(h) \leq \epsilon$.

## Answer:

A. $10 *(2 * \ln (10)+\ln (100)) \approx 92$
B. $10^{*}\left(3^{*} \ln (10)+\ln (100)\right) \approx 116$
C. $10 *(10 * \ln (2)+\ln (100)) \approx 116$
D. $10 *(10 * \ln (3)+\ln (100)) \approx 156 \leftarrow$
E. $100 *(2 * \ln (10)+\ln (10)) \approx 691$
F. $100 *\left(3^{*} \ln (10)+\ln (10)\right) \approx 922$
G. $100 *(10 * \ln (2)+\ln (10)) \approx 924$
H. $100 *(10 * \ln (3)+\ln (10)) \approx 1329 \longrightarrow$ Toxic

$$
x_{1}, x_{2}, \cdots, x_{10}
$$

$$
\stackrel{x_{1}}{\sim x_{1}} \stackrel{\operatorname{lf}\left(=3^{10} \quad \ln \left(3^{10}\right)=10 \ln 3\right.}{ }
$$

## Sample Complexity Results

Definition 0.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1).

Four Cases we care about...


## Background: Contrapositive

- Definition: The contrapositive of the statement

$$
(A) \Rightarrow(B)
$$

is the statement

$$
\neg B \Rightarrow(A
$$

and the two are logically equivalent (i.e. they share all the same truth values in a truth table!)

- Proof by contrapositive: If you want to prove $A \Rightarrow B$, instead prove $\neg B \Rightarrow \neg A$ and then conclude that $A \Rightarrow B$
- Caution: sometimes negating a statement is easier said than done, just be careful!

Proof of Theorem 1

- Assume we have $k$ bad hypotheses in H where a bad model $h_{i}$ is consirtant $\left(\hat{R}\left(h_{i}\right)=0\right)$ but $R\left(h_{i}\right)>\varepsilon$.
- Pick bad hypothesis $h_{i}$. The prob that $h_{i}$ is consistent with $\left(x^{(1)}, y^{(1)}\right)$ is $(1-\varepsilon)$

$$
\begin{aligned}
& \text {-u- } \quad\left\{\left(x^{(1)}, y^{(1)}\right),\left(x^{(2)}, y^{(2)}\right)\right\} \leqslant(1-\varepsilon)^{2} \\
& \vdots-\quad D \quad \leqslant(1-\varepsilon)^{N}
\end{aligned}
$$

- Prob that at least one bad $h$ is consistent with $D \leqslant k(1-\varepsilon)^{N} \leqslant \mid \mathcal{L} /(1-\varepsilon)^{N}$

$$
\text { Union bound: } \Phi(A \cup B) \leqslant P(A)+P(B)
$$

$$
L=P(A)+P(B)-P(A A B)
$$

Proof of Theorem 1

* Prob of a bad hypothesis looking gad empirically $\leq|\$|(1-\varepsilon)^{N}$

Known fact: $\forall x:(1-x) \leqslant \exp (-x) \longleftarrow \leqslant 1$ 拉 $\mid \operatorname{exN)}$

$$
\begin{aligned}
& \mid f(\exp (-\varepsilon N) \leqslant \delta \\
\Leftrightarrow & \left.\frac{|f|}{\delta} \right\rvert\, \leqslant \exp (\lg N) \\
\Leftrightarrow & \log (|\mathcal{L}| \mid)+\log \left(\frac{1}{\delta}\right) \leqslant \varepsilon N \\
\Leftrightarrow & \left(\frac{1}{\varepsilon}\right)\left[\log (|f+|)+\log \left(\frac{1}{8}\right)\right] \leqslant \underbrace{N}
\end{aligned}
$$

## Proof of Theorem 1

## Sample Complexity Results

Definition 0.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1).

## Four Cases we care about...

|  | Realizable | Agnostic |
| :---: | :---: | :---: |
| Finite $\|\mathcal{H}\|$ | Thm. $1 \quad N \geq\left(\frac{1}{\frac{1}{2}} \log (\right.$ (TR) $\left.)+\log \left(\frac{1}{2}\right)\right]$ la beled examples are sufficient so that witt probability $(1-\delta)$ all $h \in \mathcal{H}$ with $R(h)=0$ have $R(h) \leq \epsilon$. |  labeled examples are suffticient so that with probability $(1-\delta)$ for all $h \in \mathcal{H}$ we have that $\|R(h)-\hat{R}(h)\| \leq \epsilon$ |
| Infinite $\mid \mathcal{H}$ |  |  |



## Finite vs. Infinite $|\mathrm{H}|$

## Finite $|\mathrm{H}|$

- Example: $\mathrm{H}=$ the set of all decision trees of depth D over binary feature vectors of length $M$

- Example: $\mathrm{H}=$ the set of all conjunctions over binary feature vectors of length $M$


## Infinite |H|

- Example: $\mathrm{H}=$ the set of all linear decision boundaries in $M$ dimensions

- Example: $\mathrm{H}=$ the set of all neural networks with 1-hidden layer with length M inputs


## Sample Complexity Results

Definition 0.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1).

Four Cases we care about...


## Sample Complexity Results

Definition 0.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1).

## Four Cases we care about...

## Realizable

Thm. $1 \quad N \geq \frac{1}{\epsilon}\left[\log (|\mathcal{H}|)+\log \left(\frac{1}{\delta}\right)\right]$ labeled examples are sufficient so that with probability $(1-\delta)$ all $h \in \mathcal{H}$ with $\hat{R}(h)=0$ have $R(h) \leq \epsilon$.

Thm. $3 N=O\left(\frac{1}{\epsilon}\left[\mathrm{VC}(\mathcal{H}) \log \left(\frac{1}{\epsilon}\right)+\log \left(\frac{1}{\delta}\right)\right]\right)$ labeled examples are sufficient so that with probability $(1-\delta)$ all $h \in \mathcal{H}$ with $\hat{R}(h)=0$ have $R(h) \leq \epsilon$.

Agnostic
Thm. $2 \quad N \geq \frac{1}{2 \epsilon^{2}}\left[\log (|\mathcal{H}|)+\log \left(\frac{2}{\delta}\right)\right]$ labeled examples are sufficient so that with probability $(1-\delta)$ for all $h \in \mathcal{H}$ we have that $|R(h)-\hat{R}(h)| \leq \epsilon$.

Thm. $\left.4 \quad N=O\left(\frac{1}{\epsilon^{2}} \quad \mathrm{VC}(\mathcal{H})+\log \left(\frac{1}{\delta}\right)\right]\right)$ labeled examples are sufficient so that with probability $(1-\delta)$ for all $h \in \mathcal{H}$ we have that $|R(h)-\hat{R}(h)| \leq \epsilon$.

