CS 6316 Machine Learning

The Bias-Complexity Tradeoff

Yangfeng Ji

Department of Computer Science
University of Virginia

UNIVERSITY _of_ VIRGINIA | ENGINEERING

For a real-world machine learning problem, which of the following items are usually available to us?

For a real-world machine learning problem, which of the following items are usually available to us?

- ▶ Training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$
- ▶ Domain set $\mathcal{X}$
- ▶ Label set $\mathcal{Y}$

Quiz

For a real-world machine learning problem, which of the following items are usually available to us?

- ▶ Training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$
- ▶ Domain set $\mathcal{X}$
- ▶ Label set $\mathcal{Y}$
- ▶ Labeling function (the oracle) f
- ▶ Distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$
- ▶ The Bayes predictor $f_{\mathcal{D}}(x)$

Quiz

For a real-world machine learning problem, which of the following items are usually available to us?

- ▶ Training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$
- ▶ Domain set $\mathcal{X}$
- ▶ Label set $\mathcal{Y}$
- ▶ Labeling function (the oracle) f
- ▶ Distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$
- ▶ The Bayes predictor $f_{\mathcal{D}}(x)$
- ▶ The size of the hypothesis space $\mathcal{H}$
- ▶ The empirical risk of a hypothesis $h(x) \in \mathcal{H}$, $L_S(h(x))$
- ▶ The true risk of a hypothesis $h(x) \in \mathcal{H}$, $L_{\mathcal{D}}(h(x))$

Agnostic PAC Learnability

A hypothesis class $\mathcal{H}$ is agnostic PAC learnable if there exist a function $m_{\mathcal{H}} : (0, 1)^2 \to \mathbb{N}$ and a learning algorithm with the following property:

- ▶ for every distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, +1\}$ and
- ▶ for every $\epsilon, \delta \in (0, 1)$,

when running the learning algorithm on $m \geq m_{\mathcal{H}}(\epsilon, \delta)$ i.i.d. examples generated by $\mathcal{D}$, the algorithm returns a hypothesis h_S[1] such that, with probability of at least $1 - \delta$,

$$L_{\mathcal{D}}(h_S) \leq \min_{h' \in \mathcal{H}} L_{\mathcal{D}}(h') + \epsilon \tag{1}$$

---

[1]Sometimes, as $h_S(x)$ or $h(x, S)$

The Bayes Optimal Predictor

▶ The Bayes optimal predictor: given a probability distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, +1\}$, the predictor is defined as

$$f_{\mathcal{D}}(x) = \begin{cases} +1 & \text{if } \mathbb{P}[y = 1 | x] \geq \frac{1}{2} \\ -1 & \text{otherwise} \end{cases} \tag{2}$$

▶ No other predictor can do better: for any predictor h

$$L_{\mathcal{D}}(f_{\mathcal{D}}) \leq L_{\mathcal{D}}(h) \tag{3}$$

The Bayes Optimal Predictor

▶ The Bayes optimal predictor: given a probability distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, +1\}$, the predictor is defined as

$$f_{\mathcal{D}}(x) = \begin{cases} +1 & \text{if } \mathbb{P}[y = 1 | x] \geq \frac{1}{2} \\ -1 & \text{otherwise} \end{cases} \qquad (2)$$

▶ No other predictor can do better: for any predictor h

$$L_{\mathcal{D}}(f_{\mathcal{D}}) \leq L_{\mathcal{D}}(h) \qquad (3)$$

▶ Question: is $f_{\mathcal{D}} \in \text{argmin}_{h' \in \mathcal{H}} L_{\mathcal{D}}(h')$?

For illustration purpose, let us assume the gap between h_S and $f_{\mathcal{D}}$ can be visualized in the following plot

The Gap between h_S and $f_{\mathcal{D}}$

For illustration purpose, let us assume the gap between h_S and $f_{\mathcal{D}}$ can be visualized in the following plot

- $h_S = \text{argmin}_{h' \in \mathcal{H}} L_S(h')$: learned by minimizing the empirical risk
- $f_{\mathcal{D}}$: the optimal predictor if we know the data distribution $\mathcal{D}$

Q: For a given hypothesis space $\mathcal{H}$, does

$$f_{\mathcal{D}} \in \operatorname*{argmin}_{h'} L_{\mathcal{D}}(h') \tag{4}$$

hold?

Q: For a given hypothesis space $\mathcal{H}$, does

$$f_{\mathcal{D}} \in \underset{h'}{\operatorname{argmin}} \, L_{\mathcal{D}}(h') \tag{4}$$

hold?

A: it depends the selection of the hypothesis space $\mathcal{H}$, usually not.

Example: if $f_{\mathcal{D}}$ is a nonlinear classifier, while we choose to use logistic regression.

The previous example implies the error gap between h_S and $f_{\mathcal{D}}$ can be decomposed into two components

The previous example implies the error gap between h_S and $f_{\mathcal{D}}$ can be decomposed into two components

Two different perspectives of the decomposition

▶ The bias-complexity tradeoff: from the perspective of learning theory

▶ The bias-variance tradeoff: from the perspective of statistical learning/estimation

The Bias-Complexity Tradeoff

Basic Learning Procedure

The basic component of formulating a learning process

- Input/output space $\mathcal{X} \times \mathcal{Y}$
- Hypothesis space $\mathcal{H}$
- Learning via empirical risk minimization

$$h_S \in \operatorname*{argmin}_{h' \in \mathcal{H}} L_S(h') \tag{5}$$

- Goal: analyzing the true error of h_S, $L_{\mathcal{D}}(h_S)$

Example

Consider the binary classification problem with the data sampled from the following distribution

$$\mathcal{D} = \frac{1}{2}\mathcal{B}(x; 5, 1) + \frac{1}{2}\mathcal{B}(x; 1, 2) \tag{6}$$

Example (Cont.)

Given the distribution, we can compute the true risk/error of the Bayes predictor $f_{\mathcal{D}}$ as

$$
\begin{aligned}
L_{\mathcal{D}}(f_{\mathcal{D}}) &= \frac{1}{2}\mathcal{B}(x < b_{\text{Bayes}}; 5, 1) + \frac{1}{2}(1 - \mathcal{B}(x < b_{\text{Bayes}}; 1, 2)) \\
&= 0.11799 \qquad\qquad (7)
\end{aligned}
$$

The hypothesis space $\mathcal{H}$ is defined as

$$h_i(x) = \begin{cases} +1 & x > \frac{i}{N} \\ -1 & x < \frac{i}{N} \end{cases} \tag{8}$$

where $N \in \mathbb{N}$ is a predefined integer

The hypothesis space $\mathcal{H}$ is defined as

$$h_i(x) = \begin{cases} +1 & x > \frac{i}{N} \\ -1 & x < \frac{i}{N} \end{cases} \tag{8}$$

where $N \in \mathbb{N}$ is a predefined integer

- ▶ This is an unrealizable case
- ▶ The value of N is the size of the hypothesis space

The hypothesis space $\mathcal{H}$ is defined as

$$h_i(x) = \begin{cases} +1 & x > \frac{i}{N} \\ -1 & x < \frac{i}{N} \end{cases} \tag{8}$$

where $N \in \mathbb{N}$ is a predefined integer

- This is an unrealizable case
- The value of N is the size of the hypothesis space
- The best hypothesis in $\mathcal{H}$

$$h^* \in \underset{h' \in \mathcal{H}}{\operatorname{argmin}} L_{\mathcal{D}}(h') \tag{9}$$

- Very likely the best predictor in $\mathcal{H}$ is not the Bayes predictor, unless $b_{\text{Bayes}} \in \{\frac{i}{N} : i \in [N]\}$

Error Decomposition

The error gap between h_S and $f_{\mathcal{D}}$ can be decomposed as two parts

$$L_{\mathcal{D}}(h_S) - L_{\mathcal{D}}(f_{\mathcal{D}}) = \epsilon_{\text{app}} + \epsilon_{\text{est}} \qquad (10)$$

Error Decomposition

The error gap between h_S and $f_{\mathcal{D}}$ can be decomposed as two parts

$$L_{\mathcal{D}}(h_S) - L_{\mathcal{D}}(f_{\mathcal{D}}) = \epsilon_{\text{app}} + \epsilon_{\text{est}} \qquad (10)$$

▶ Approximation error ϵ_{app} caused by selecting a specific hypothesis space $\mathcal{H}$ (model bias)

▶ Estimation error ϵ_{est} caused by selecting h_S with a specific training set

Approximation Error ϵ_{app}

To reduce the approximation error ϵ_{app}, we could increase the size of the hypothesis space

The cost is that we also increase the size of training set, in order to maintain the overall error in the same level (recall the sample complexity of finite hypothesis spaces).

Approximation Error ϵ_{app}

To reduce the approximation error ϵ_{app}, we could increase the size of the hypothesis space

The cost is that we also increase the size of training set, in order to maintain the overall error in the same level (recall the sample complexity of finite hypothesis spaces).

On the other hand, if we use the same training set S, then we *may* have a larger estimation error

On the other hand, if we use the same training set S, then we *may* have a larger estimation error

The bias-complexity tradeoff: find the right balance to reduce both approximation error and estimation error.

We randomly sampled 100 examples from each class

$$\mathcal{D} = \frac{1}{2}\mathscr{B}(x; 5, 1) + \frac{1}{2}\mathscr{B}(x; 1, 2) \qquad (11)$$

Given 200 training examples, the errors with respect to different hypothesis space is the following (x axis is the size of $\mathcal{H}$)


There is a tradeoff with respect to the size of $\mathcal{H}$

Example: 2000 training examples

We randomly sampled 1000 examples from each class

$$\mathcal{D} = \frac{1}{2}\mathcal{B}(x; 5, 1) + \frac{1}{2}\mathcal{B}(x; 1, 2) \qquad (12)$$

With these 2000 training examples, the errors with respect to different hypothesis space is the following

Both errors are smaller, but the tradeoff still exists

With these 2000 training examples, the errors with respect to different hypothesis space is the following

Both errors are smaller, but the tradeoff still exists

Exercise: The bias-complexity tradeoff with a Gaussian

Summary

Three components in this decomposition

- $h_S \in \text{argmin}_{h' \in \mathcal{H}} L_S(h')$: the ERM predictor given the training set S
- $h^* \in \text{argmin}_{h' \in \mathcal{H}} L_{\mathcal{D}}(h')$: the optimal predictor from $\mathcal{H}$
- $f_{\mathcal{D}}$: the Bayes predictor given $\mathcal{D}$

Summary

Three components in this decomposition

- $h_S \in \text{argmin}_{h' \in \mathcal{H}} L_S(h')$: the ERM predictor given the training set S
- $h^* \in \text{argmin}_{h' \in \mathcal{H}} L_{\mathcal{D}}(h')$: the optimal predictor from $\mathcal{H}$
- $f_{\mathcal{D}}$: the Bayes predictor given $\mathcal{D}$

Balancing strategy:

- we can incrase the complexity of hypothesis space to reduce the bias, e.g.,
  - enlarge the hypothesis space (as in the running example)
  - replacing linear predictors with nonlinear predictors

Summary

Three components in this decomposition

- $h_S \in \mathrm{argmin}_{h' \in \mathcal{H}} L_S(h')$: the ERM predictor given the training set S
- $h^* \in \mathrm{argmin}_{h' \in \mathcal{H}} L_{\mathcal{D}}(h')$: the optimal predictor from $\mathcal{H}$
- $f_{\mathcal{D}}$: the Bayes predictor given $\mathcal{D}$

Balancing strategy:

- we can incrase the complexity of hypothesis space to reduce the bias, e.g.,
  - enlarge the hypothesis space (as in the running example)
  - replacing linear predictors with nonlinear predictors
- in the meantime, we have to increase the training size to reduce the approximation error.

The Bias-Variance Tradeoff

Let us analyze the error ϵ without the assumption of

- knowing the best predictor from $\mathcal{H}$,
  $h^* \in \mathrm{argmin}_{h' \in \mathcal{H}} L_{\mathcal{D}}(h')$
- changing the size of S

A New Perspective

Let us analyze the error ϵ without the assumption of

- knowing the best predictor from $\mathcal{H}$,
  $h^* \in \operatorname{argmin}_{h' \in \mathcal{H}} L_{\mathcal{D}}(h')$
- changing the size of S

We still need (1) the ERM predictor h_S and (2) the Bayes predictor $f_{\mathcal{D}}$

A New Way of Decomposition

. . . by considering

- ▶ the randomness in S with m training examples
- ▶ the average prediction given by $E\left[h(\boldsymbol{x}, S)\right]$ where $S \sim \mathcal{D}^m$

Data Generation Model

Consider the following *data generation model*

- ▶ $X \sim U[0, 1]$ uniform distribution
- ▶ $Y = \mathcal{N}(X + \sin(2X), \sigma^2)$ with $\sigma^2 = 0.1$

An example of S is

Hypothesis Spaces

Given S and the following hypothesis space $\mathcal{H}_1$

$$\mathcal{H}_1 = \{w_0 + w_1 x : w_0, w_1 \in \mathbb{R}\} \tag{13}$$

the regression result

# Hypothesis Spaces (Cont.)

Given S and the following hypothesis space $\mathcal{H}_3$

$$\mathcal{H}_3 = \{w_0 + w_1 x + w_2 x^2 + w_3 x^3 : w_0, w_1, w_2, w_3 \in \mathbb{R}\} \quad (14)$$

the regression result

Given S and the following hypothesis space $\mathcal{H}_{15}$

$$\mathcal{H}_{15} = \{w_0 + w_1 x + \cdots + w_{15} x^{15} : w_0, w_1, \cdots, w_{15} \in \mathbb{R}\} \quad (15)$$

Given S and the following hypothesis space $\mathcal{H}_{15}$

$$\mathcal{H}_{15} = \{w_0 + w_1 x + \cdots + w_{15} x^{15} : w_0, w_1, \cdots, w_{15} \in \mathbb{R}\} \quad (15)$$

▶ Intuitively, the degree of the polynomials indicates the potential/complexity of the hypothesis space
▶ Refer to the VC dimension section for more discussion

The difference between the best hypothesis $h(x, S)$ and the Bayes predictor $f_{\mathcal{D}}(x)$ is measured as

$$\epsilon^2 = \{h(x, S) - f_{\mathcal{D}}(x)\}^2 \qquad (16)$$

Introduce $E[h(x, S)]$ into the calculation, we have

$$\epsilon^2 = \{h(x, S) - E[h(x, S)] + E[h(x, S)] - f_{\mathcal{D}}(x)\}^2$$

Error Decomposition

The difference between the best hypothesis $h(x, S)$ and the Bayes predictor $f_{\mathcal{D}}(x)$ is measured as

$$\epsilon^2 = \{h(x, S) - f_{\mathcal{D}}(x)\}^2 \tag{16}$$

Introduce $E[h(x, S)]$ into the calculation, we have

$$
\begin{aligned}
\epsilon^2 &= \{h(x, S) - E[h(x, S)] + E[h(x, S)] - f_{\mathcal{D}}(x)\}^2 \\
&= \{h(x, S) - E[h(x, S)]\}^2 + \{E[h(x, S)] - f_{\mathcal{D}}(x)\}^2 \\
&\quad + 2\{h(x, S) - E[h(x, S)]\} \cdot \{E[h(x, S)] - f_{\mathcal{D}}(x)\}
\end{aligned}
$$

Review: Mean

Given a random variable X and its probability density function $p(x)$

- Mean: $E[X] = \int x p(x) dx$
- Approximation to the mean with samples $\{x_1, \ldots, x_m\}$

$$E[X] \approx \frac{1}{m} \sum_{i=1}^{m} x_i \qquad (17)$$

- Property: $E[\alpha X] = \alpha E[X]$ for α is determinstic
- Example: the mean of a Gaussian distribution $\mathcal{N}(x; \mu, \sigma^2)$

$$E[X] = \mu \qquad (18)$$

Review: Variance

Given a random variable X, its probability density function $p(x)$, and its mean $E[X]$

- Variance: $\text{Var}(X) = E\left[(X - E[X])^2\right]$
- Example: the variance of a Gaussian distribution $\mathcal{N}(x; \mu, \sigma^2)$

$$\text{Var}(X) = \sigma^2 \tag{19}$$

Review: Variance

Given a random variable X, its probability density function $p(x)$, and its mean $E[X]$

- Variance: $\text{Var}(X) = E\left[(X - E[X])^2\right]$
- Example: the variance of a Gaussian distribution $\mathcal{N}(x; \mu, \sigma^2)$

$$\text{Var}(X) = \sigma^2 \tag{19}$$

$$
\begin{aligned}
\text{Var}(X) &= E\left[(X - E[X])^2\right] \\
&= E\left[X^2 - 2XE[X] + E[X]^2\right] \\
&= E\left[X^2\right] - 2E[X]E[X] + E[X]^2 \\
&= E\left[X^2\right] - E[X]^2
\end{aligned}
$$

Taking the expectation of ϵ^2

$$
\begin{aligned}
E\left[\epsilon^2\right] = & \; E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right] + \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathscr{D}}(\boldsymbol{x})\}^2 \\
& + 2E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}\right] \cdot \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathscr{D}}(\boldsymbol{x})\}
\end{aligned}
$$

Taking the expectation of ϵ^2

$$
\begin{aligned}
E\left[\epsilon^2\right] &= E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right] + \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2 \\
&\quad + 2E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}\right] \cdot \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\} \\
&= E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right] + \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2 \\
&\quad + 2\{E\left[h(\boldsymbol{x}, S)\right] - E\left[h(\boldsymbol{x}, S)\right]\} \cdot \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}
\end{aligned}
$$

Taking the expectation of ϵ^2

$$
\begin{aligned}
E\left[\epsilon^2\right] &= E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right] + \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2 \\
&\quad + 2E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}\right] \cdot \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\} \\
&= E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right] + \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2 \\
&\quad + 2\{E\left[h(\boldsymbol{x}, S)\right] - E\left[h(\boldsymbol{x}, S)\right]\} \cdot \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\} \\
&= E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right] + \{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2
\end{aligned}
$$

The Bias-Variance Decomposition

The expected error is decomposed as

$$E\left[\epsilon^2\right] = \underbrace{E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right]}_{\text{variance}} + \underbrace{\{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2}_{\text{bias}^2}$$

The Bias-Variance Decomposition

The expected error is decomposed as

$$E\left[\epsilon^2\right] = \underbrace{E\left[\{h(\boldsymbol{x}, S) - E\left[h(\boldsymbol{x}, S)\right]\}^2\right]}_{\text{variance}} + \underbrace{\{E\left[h(\boldsymbol{x}, S)\right] - f_{\mathcal{D}}(\boldsymbol{x})\}^2}_{\text{bias}^2}$$

- **bias**: how far the expected prediction $E\left[h(\boldsymbol{x}, S)\right]$ diverges from the optimal predictor $f_{\mathcal{D}}(\boldsymbol{x})$

The Bias-Variance Decomposition

The expected error is decomposed as

$$E\left[\epsilon^2\right] = \underbrace{E\left[\{h(\boldsymbol{x}, S) - E[h(\boldsymbol{x}, S)]\}^2\right]}_{\text{variance}} + \underbrace{\{E[h(\boldsymbol{x}, S)] - f_{\mathcal{D}}(\boldsymbol{x})\}^2}_{\text{bias}^2}$$

- ▶ **bias**: how far the expected prediction $E[h(\boldsymbol{x}, S)]$ diverges from the optimal predictor $f_{\mathcal{D}}(\boldsymbol{x})$
- ▶ **variance**: how a hypothesis learned from a specific S diverges from the average prediction $E[h(\boldsymbol{x}, S)]$

Computing $E[h(x, S)]$

The key of computing $E[h(x, S)]$ is to eliminate the randomness introduced by S

1: **for** $k = 1, \cdots, K$ **do**
2:     Sample a traing set S_k with size m from the data generation model
3:     Find the best hypothesis via
     $h(x, S_k) \in \text{argmin}_{h'} L(h', S_k)$
4: **end for**
5: **Output**:

$$E[h(x, S)] \approx \frac{1}{K} \sum_{k=1}^{K} h(x, S_k)$$

The larger K, the better approximation

Example: Bias and Variance

With $K = 50$, $m = 100$, and $\mathcal{H}_1$, we can visualize the bias and variance of a linear regression example as following

High bias and low variance (Underfitting)

Example: Bias and Variance (Cont.)

Same training set with $\mathcal{H}_3$


Both bias and variance are fine

Same training set with $\mathscr{H}_{15}$

Low bias and high variance (Overfitting)

Same training set with $\mathscr{H}_{15}$


Low bias and high variance (Overfitting)

Exercise: The bias-variance tradeoff on linear regression with ℓ_2 regularization

The Bias-Variance Tradeoff

- **bias**: how far the expected prediction $E[h(\boldsymbol{x}, S)]$ diverges from the optimal predictor $f_{\mathscr{D}}(\boldsymbol{x})$
  - Error of this part is caused by *the selection of a hypothesis space*

The Bias-Variance Tradeoff

- **bias**: how far the expected prediction $E[h(\boldsymbol{x}, S)]$ diverges from the optimal predictor $f_{\mathscr{D}}(\boldsymbol{x})$
  - Error of this part is caused by *the selection of a hypothesis space*
- **variance**: how a hypothesis learned from a specific S diverges from the average prediction $E[h(\boldsymbol{x}, S)]$
  - Error of this part is caused by *using a particular data set S*

The VC Dimension

Infinite-size hypothesis space is learnable

Examples

- ▶ Half-space predictor
- ▶ Logistic regression predictor
- ▶ *Many others*

For a given set C and a hypothesis space $\mathcal{H}$,

▶ A dichotomy of the set is one of the possible ways of labeling the points in C using a hypothesis $h \in \mathcal{H}$

[Mohri et al., 2018, Page 36]

# Shattering

For a given set C and a hypothesis space $\mathscr{H}$,

- ▶ A dichotomy of the set is one of the possible ways of labeling the points in C using a hypothesis $h \in \mathscr{H}$
- ▶ A set C of $m \geq 1$ points is said to be shattered by a hypothesis space $\mathscr{H}$, if *all* possible dichotomies of S can be realized by $\mathscr{H}$

[Mohri et al., 2018, Page 36]

Consider the following set C and the half-space hypothesis space

$$\mathcal{H}_{\text{half}} = \{w_0 + w_1 x_1 + w_2 x_2 = 0 : w_0, w_1, w_2 \in \mathbb{R}\} \qquad (20)$$

There are $2^3 = 8$ different ways to label the points and $\mathcal{H}_{\text{half}}$ can realized all of them.

The **VC-dimension** of a hypothesis space $\mathcal{H}$, denoted VCdim($\mathcal{H}$), is the maximal size of a set $C \subset \mathcal{X}$ that can be shattered by $\mathcal{H}$.

[Shalev-Shwartz and Ben-David, 2014, Page 70]

The **VC-dimension** of a hypothesis space $\mathcal{H}$, denoted VCdim($\mathcal{H}$), is the maximal size of a set $C \subset \mathcal{X}$ that can be shattered by $\mathcal{H}$.

A: How to find the VC-dimension of a given hypothesis space?

Q: The proof consists of two parts:

- There exists a set C of size d that is shattered by $\mathcal{H}$
- Every set C of size $d + 1$ is not shattered by $\mathcal{H}$

[Shalev-Shwartz and Ben-David, 2014, Page 70]

Half Spaces

Consider a special case as following, where
VC-dim($\mathcal{H}_{\text{half}}$) = 3

$$\mathcal{H}_{\text{half}} = \{w_0 + w_1 x_1 + w_2 x_2 = 0 : w_0, w_1, w_2 \in \mathbb{R}\} \qquad (21)$$

(1) Exist a case

Half Spaces

Consider a special case as following, where
VC-dim($\mathcal{H}_{\text{half}}$) = 3

$$\mathcal{H}_{\text{half}} = \{w_0 + w_1 x_1 + w_2 x_2 = 0 : w_0, w_1, w_2 \in \mathbb{R}\} \qquad (21)$$

(1) Exist a case

(2) For any case

Axis-aligned Rectangles

Let $\mathcal{H}$ be the class of axis-aligned rectangle, formally

$$\mathcal{H} = \{h_{(a_1, a_2, b_1, b_2)} : a_1 \leq a_2 \text{ and } b_1 \leq b_2\} \qquad (22)$$

where

$$h_{(a_1, a_2, b_1, b_2)}(x_1, x_2) = \begin{cases} +1 & x_1 \in [a_1, a_2] \text{and } x_2 \in [b_1, b_2] \\ -1 & \text{otherwise} \end{cases}$$

Axis-aligned Rectangles

Let $\mathcal{H}$ be the class of axis-aligned rectangle, formally

$$\mathcal{H} = \{h_{(a_1,a_2,b_1,b_2)} : a_1 \leq a_2 \text{ and } b_1 \leq b_2\} \tag{22}$$

where

$$h_{(a_1,a_2,b_1,b_2)}(x_1, x_2) = \begin{cases} +1 & x_1 \in [a_1, a_2] \text{ and } x_2 \in [b_1, b_2] \\ -1 & \text{otherwise} \end{cases}$$

Exist a case

Let $\mathcal{H}$ be the class of axis-aligned rectangle, formally

$$\mathcal{H} = \{h_{(a_1, a_2, b_1, b_2)} : a_1 \leq a_2 \text{ and } b_1 \leq b_2\} \qquad (22)$$

where

$$h_{(a_1, a_2, b_1, b_2)}(x_1, x_2) = \begin{cases} +1 & x_1 \in [a_1, a_2] \text{ and } x_2 \in [b_1, b_2] \\ -1 & \text{otherwise} \end{cases}$$

For any case

Axis-aligned Rectangles

Let $\mathcal{H}$ be the class of axis-aligned rectangle, formally

$$\mathcal{H} = \{h_{(a_1,a_2,b_1,b_2)} : a_1 \leq a_2 \text{ and } b_1 \leq b_2\} \qquad (22)$$

where

$$h_{(a_1,a_2,b_1,b_2)}(x_1, x_2) = \begin{cases} +1 & x_1 \in [a_1, a_2] \text{ and } x_2 \in [b_1, b_2] \\ -1 & \text{otherwise} \end{cases}$$

For any case

$$\text{VC-dim}(\mathcal{H}_{\text{rect}}) = 4$$

▶ For linear predictors, the VC dimensions are equal to the numbers of parameters

$$\mathcal{H}_{\text{half}} = \{w_0 + w_1 x_1 + w_2 x_2 = 0 : w_0, w_1, w_2 \in \mathbb{R}\} \quad (23)$$


▶ However, the case is not always true. Considering the following hypothesis space

Sine Functions

The hypothesis space of sine functions is defined as

$$\mathcal{H}_{\sin} = \{\sin(\alpha \cdot x) : \alpha \in \mathbb{R}\} \tag{24}$$

- ▶ $\alpha = \frac{\pi}{4}$
- ▶ $\alpha = \frac{\pi}{2}$
- ▶ $\alpha = \pi$

Sine Functions

The hypothesis space of sine functions is defined as

$$\mathcal{H}_{\sin} = \{\sin(\alpha \cdot x) : \alpha \in \mathbb{R}\} \qquad (24)$$

- $\alpha = \frac{\pi}{4}$
- $\alpha = \frac{\pi}{2}$
- $\alpha = \pi$

Sine Functions

The hypothesis space of sine functions is defined as

$$\mathcal{H}_{\sin} = \{\sin(\alpha \cdot x) : \alpha \in \mathbb{R}\} \tag{24}$$

- $\alpha = \frac{\pi}{4}$
- $\alpha = \frac{\pi}{2}$
- $\alpha = \pi$

Sine Functions

The hypothesis space of sine functions is defined as

$$\mathscr{H}_{\sin} = \{\sin(\alpha \cdot x) : \alpha \in \mathbb{R}\} \qquad (24)$$


▶ $\alpha = \frac{\pi}{4}$

▶ $\alpha = \frac{\pi}{2}$

$$\text{VC-dim}(\mathscr{H}_{\sin}) = \infty$$

▶ $\alpha = \pi$

Reference

Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2018).
Foundations of machine learning.
MIT press.

Shalev-Shwartz, S. and Ben-David, S. (2014).
Understanding machine learning: From theory to algorithms.
Cambridge university press.