CSC 580 Principles of Machine Learning

02 Limits of Learning

Chicheng Zhang

Department of Computer Science

THE UNIVERSITY OF ARIZONA

*slides credit: built upon CSC 580 Fall 2021 lecture slides by Kwang-Sung Jun

Motivation

- Machine learning is a general & useful framework

- Understand when machine learning will and will not work

Optimal classification with known D

- Suppose:
    - Binary classification, 0-1 loss $\ell(y, \hat{y}) = I(y \neq \hat{y})$
    - D is *known:* for every (x, y), $P_D(x, y)$ can be computed

- What is the f that minimizes $L_D(f) = P_{(x,y)\sim D}(y \neq f(x))$?

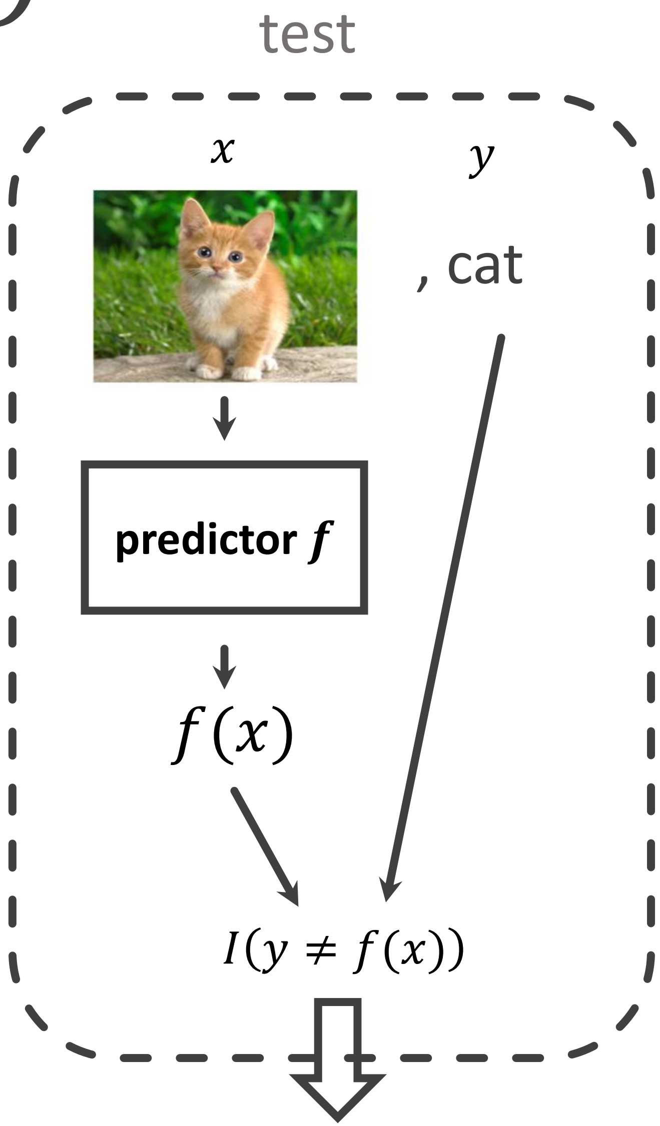

Generalization error: $L_D(f) = \mathrm{E}_{(x,y)\sim D} I(y \neq f(x)) = P_{(x,y)\sim D}\,(y \neq f(x))$

Simple case: discrete domain $\mathcal{X}$

$P_D(x, y)$	$x = 1$	$x = 2$	$x = 3$
$y = -1$	0.2	0.2	0.15
$y = +1$	0.1	0.3	0.05

Which classifier is better?

- $f_1(1) = -1, f_1(2) = -1, f_1(3) = -1 \quad \Rightarrow \quad L_D(f_1) = 0.1 + 0.3 + 0.05$
- $f_2(1) = -1, f_2(2) = +1, f_2(3) = -1 \quad \Rightarrow \quad L_D(f_2) = 0.1 + 0.2 + 0.05$

- What is the best classifier?
- For any x, should choose y that has higher value of $P_D(x, y)$
- $f^*(1) = -1, f^*(2) = +1, f^*(3) = -1$

Bayes optimal classifier

- $f_{BO}(x) = \arg\max_{y \in \mathcal{Y}} P_D(X = x, Y = y) = \arg\max_{y \in \mathcal{Y}} P_D(Y = y \,|X = x)\,, \forall x \in \mathcal{X}$
- Theorem: f_{BO} achieves the smallest error rate among all functions.
- Bayes error rate: $L_D(f_{BO})$

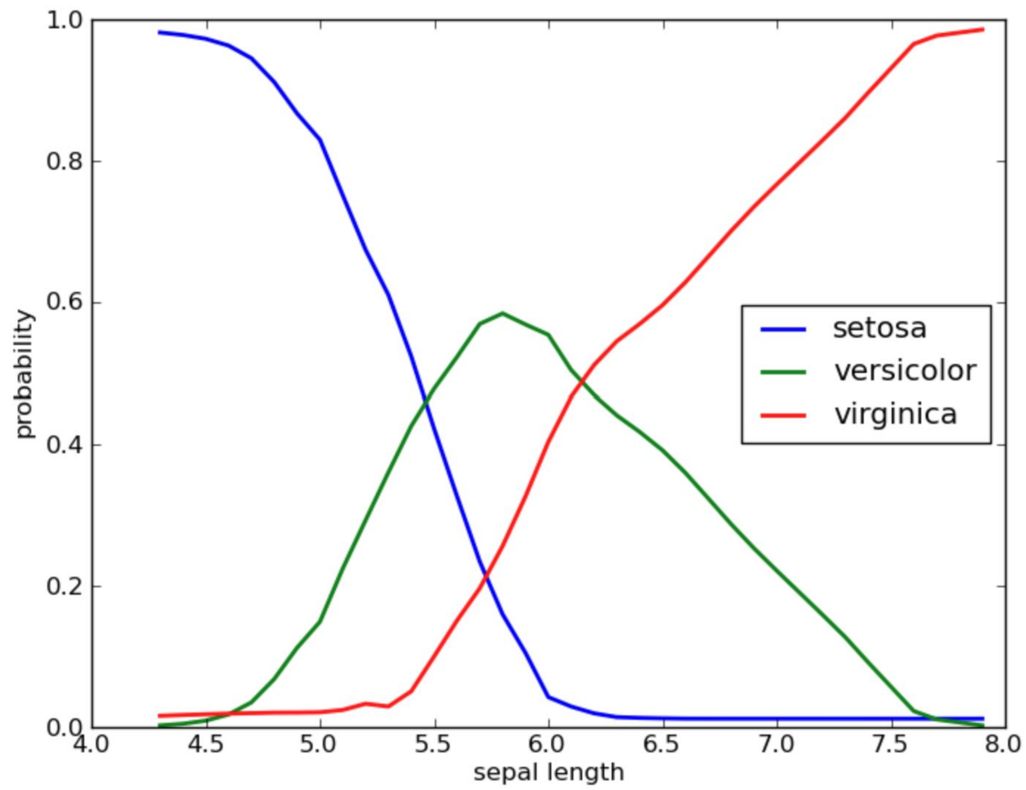

Proof of theorem

- Step 1: consider accuracy:
  - $A_D(f) = 1 - L_D(f) = P_D\big(Y = f(X)\big) = \sum_x P_D\big(X = x, Y = f(x)\big)$
  - Suffices to show f_{BO} has the highest accuracy

- Step 2: comparison:

$$A_D(f_{BO}) - A_D(f) = \sum_x P_D\big(X = x, Y = f_{BO}(x)\big) - P_D\big(X = x, Y = f(x)\big) \geq 0$$

$$\boxed{f_{BO}(x) = \arg\max_{y \in \mathcal{Y}} P_D(X = x, Y = y)}$$

- Remark: similar reasoning can be used to prove the theorem with continuous domain $\mathcal{X}$ (sum -> integral)

Bayes error rate: alternative form

$$
\begin{aligned}
L_D(f_{BD}) &= P_D\big(Y \neq f_{BD}(X)\big) \\
&= \textstyle\sum_x P_D(Y \neq f_{BD}(x) \mid X = x)\, P_D(X = x) \\
&= \textstyle\sum_x (1 - P_D(Y = f_{BD}(x) \mid X = x))\, P_D(X = x) \\
&= \textstyle\sum_x \left(1 - \max_y P_D(Y = y \mid X = x)\right) P_D(X = x) \\
&= \mathrm{E}\left[1 - \max_y P_D(Y = y \mid X)\right]
\end{aligned}
$$

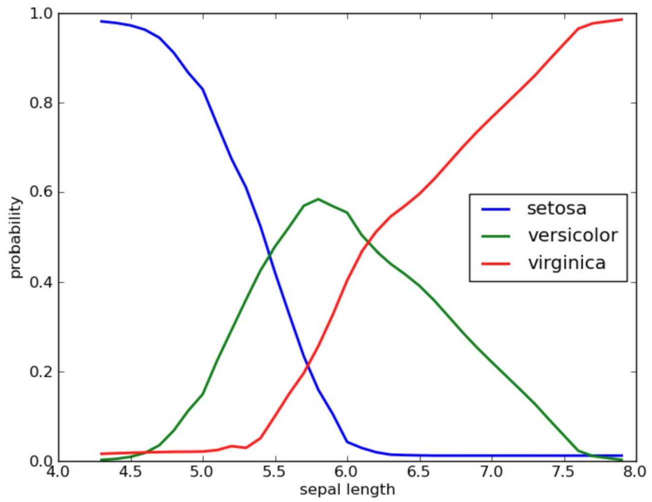

- Special case: binary classification
  - $L_D(f_{BD}) = \sum_x P_D(Y \neq f_{BD}(x), X = x)$
    $= \sum_x \min(\, P_D(Y = +1, X = x)\,, P_D(Y = -1, X = x))$

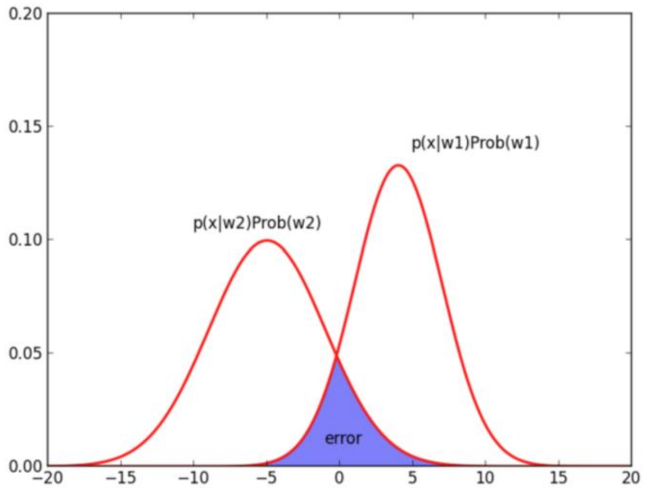

When is the Bayes error rate nonzero?

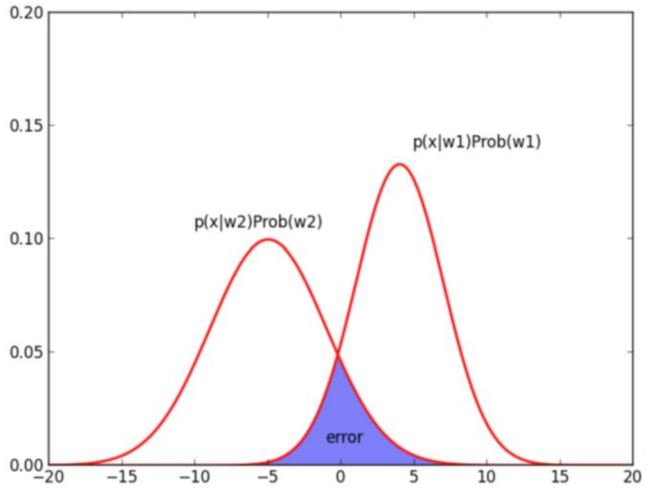

- $L_D(f_{BO}) = \sum_x \min(\, P_D(Y = +1, X = x)\,, P_D(Y = -1, X = x))$

- Limited feature representation

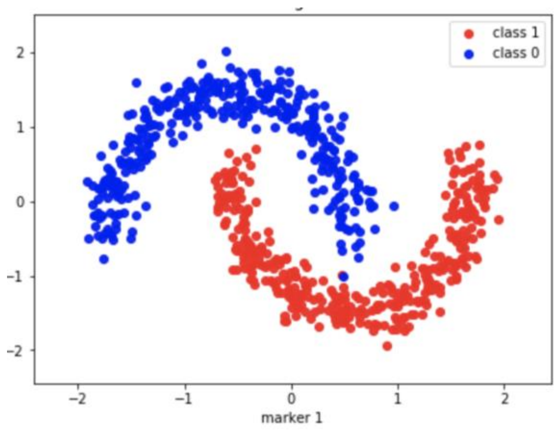

- feature noise
  - Sensor failure
  - Typo in reviews for sentiment classification

- label noise
  - Crowdsourcing settings

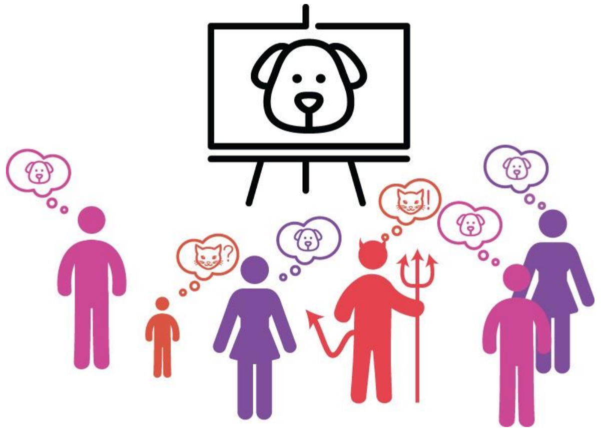

{Over,Under}-fitting

- Q: should I train a shallow or deep decision tree?

- Shallow tree:

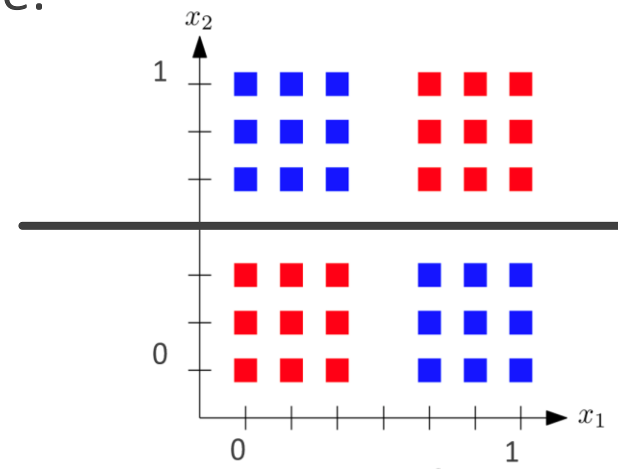

Deep tree:

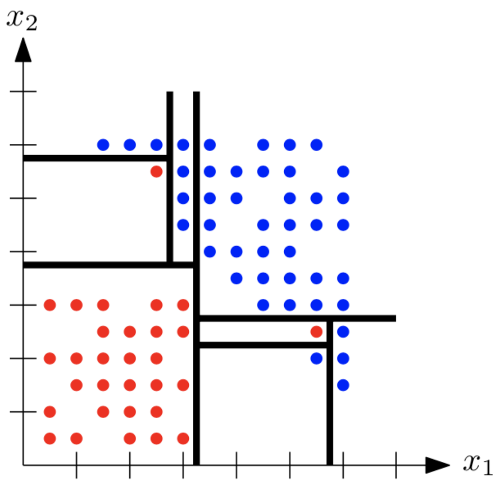

- Underfitting: have the opportunity to learn something but didn't
- Overfitting: pay too much attention to idiosyncrasies to training data, and do not generalize well
- A model that neither overfits nor underfits is expected to do best

Unbiased model evaluation using test data

- Your boss says: I will allow your recommendation system to run on our website only if the error is <= 10%!
- How to prove it?
- Idea: reserve some data as test data for evaluating predictors

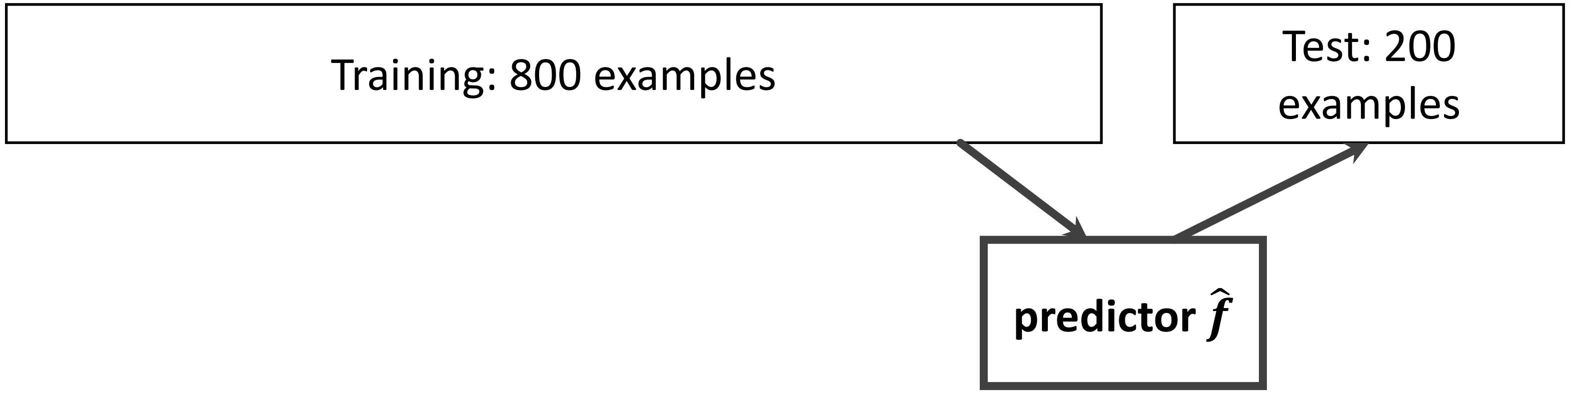

- $L_{\text{test}}(\hat{f}) = \frac{1}{|S_{\text{test}}|} \sum_{(x,y)\in S_{\text{test}}} I(y \neq \hat{f}(x))$
- Law of large numbers $\Rightarrow L_{\text{test}}(\hat{f}) \to L_D(\hat{f})$

Law of large numbers (LLN)

- Suppose $v_1, \ldots, v_n$ are independent random variables that are identically distributed, the sample average $\bar{v} = \frac{1}{n}\sum_{i=1}^{n} v_i$ converges to $\mathrm{E}[v_1]$ as $n \to \infty$

- Useful in e.g. election poll
- Foundations of statistics

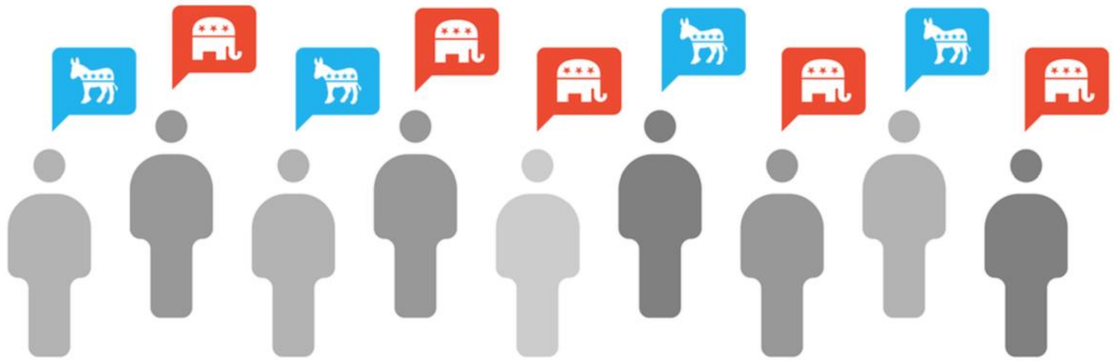

Training: 800 examples

Test: 200 examples

predictor $\hat{f}$

- Can we apply LLN to conclude that
- $L_{\text{train}}(\hat{f}) = \frac{1}{|S_{\text{train}}|}\sum_{(x,y)\in S_{\text{train}}} I(y \neq \hat{f}(x)) \to L_D(\hat{f})$ as $|S_{\text{train}}| \to \infty$?

Never touch your test data

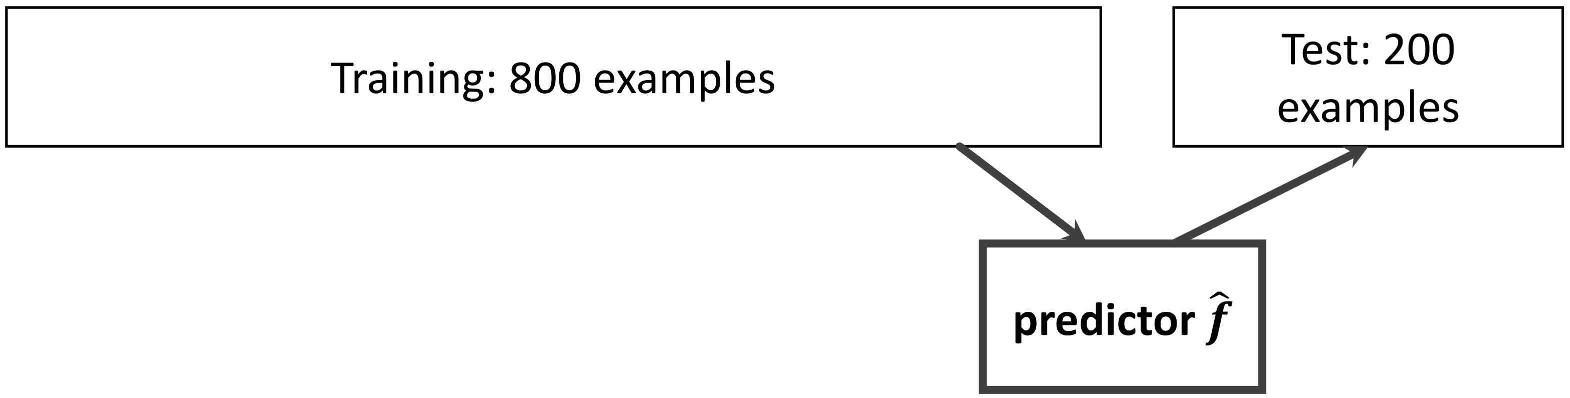

- If $\hat{f}$ depends on test examples, $L_{\text{test}}(\hat{f})$ may no longer estimate $L_D(\hat{f})$ accurately
- E.g. indirect dependence:
  - adaptive data analysis – choose a new learning algorithm based on seeing that the previous algorithm produces an high-test-error model

Class Participation

- Asking review questions on Piazza (3pts)
  - Every week, I will ask two of you to post questions (related to the past week's material) on Piazza
  - 3 questions per student
- Other in-class / Piazza discussions (e.g. asking/answering in-class questions; Piazza Q&As)
- Extra credit: Catching errors in the CIML book
  - Post on Piazza; we'll discuss and confirm together, and hopefully send these back to the author
  - 1pt for every error found

Supervised learning setup

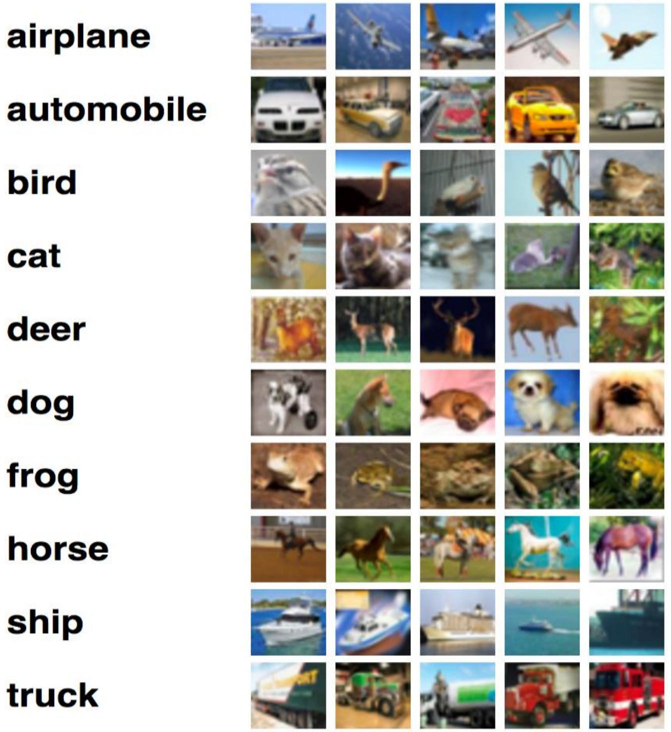

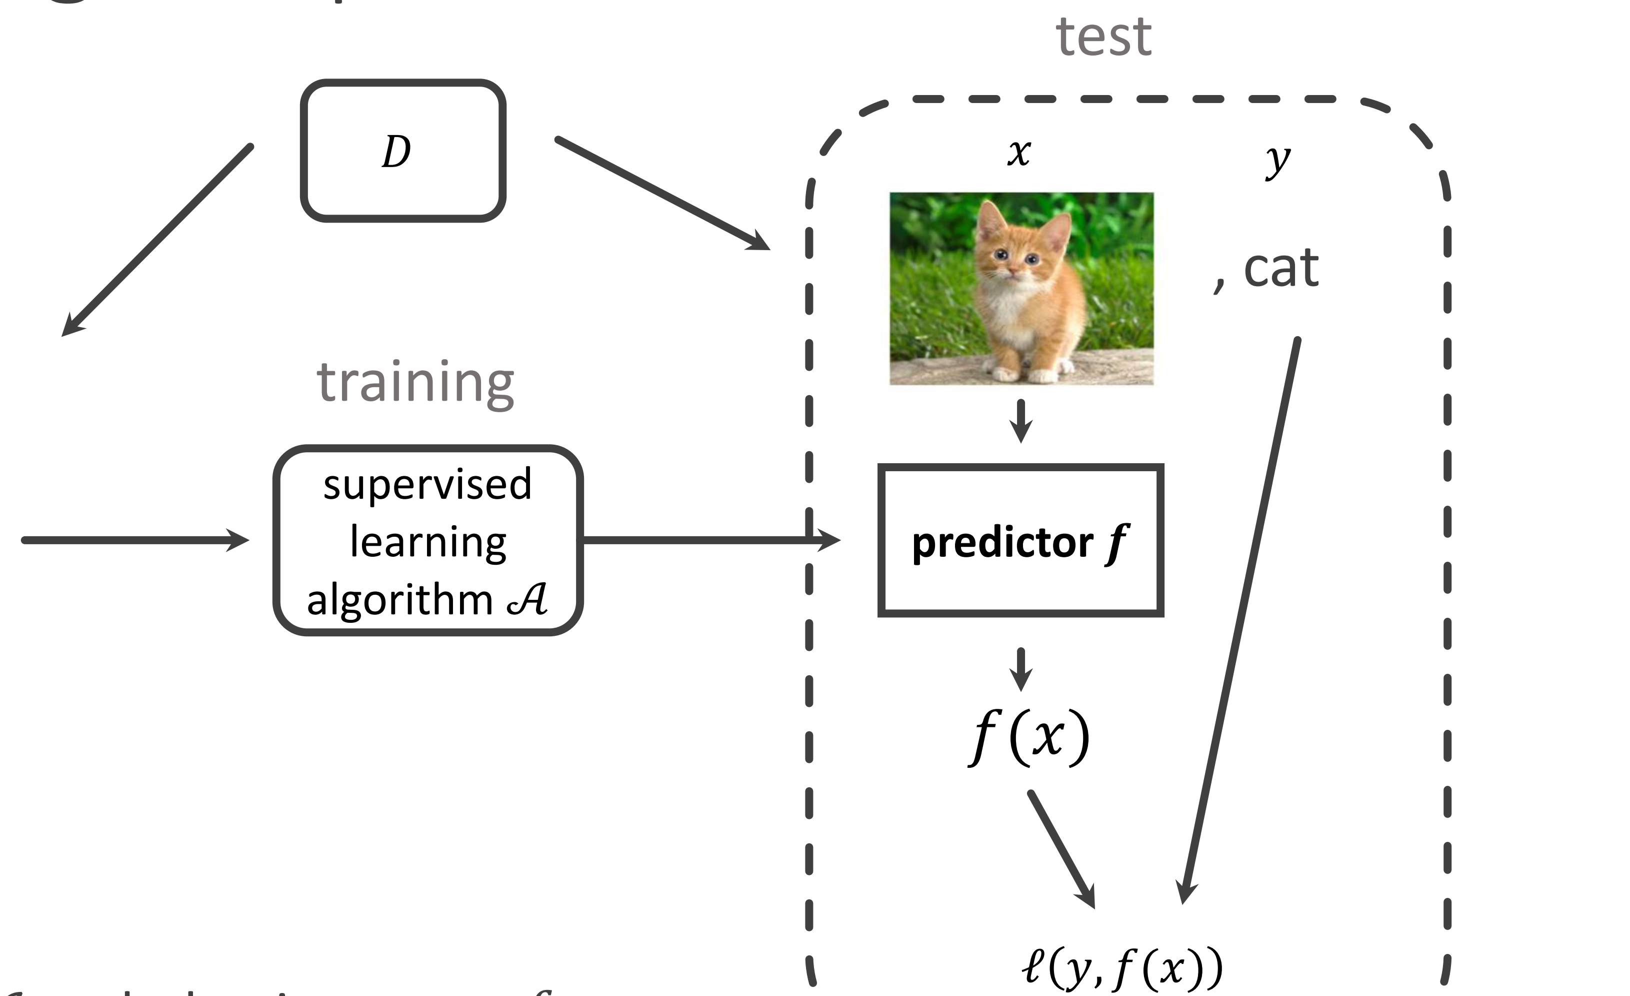

- Goal: design learning algorithm $\mathcal{A}$ such that its output f on iid training data S has low generalization error

Generalization error: $L_D(f) = \mathrm{E}_{(x,y)\sim D}\, \ell(y, f(x))$

Terminologies

- Model: the predictor $\hat{f}$
  - Often from a model class $\mathcal{F}$,
  - e.g. $\mathcal{F} = \{\text{decision trees}\}, \{\text{linear classifiers}\}$

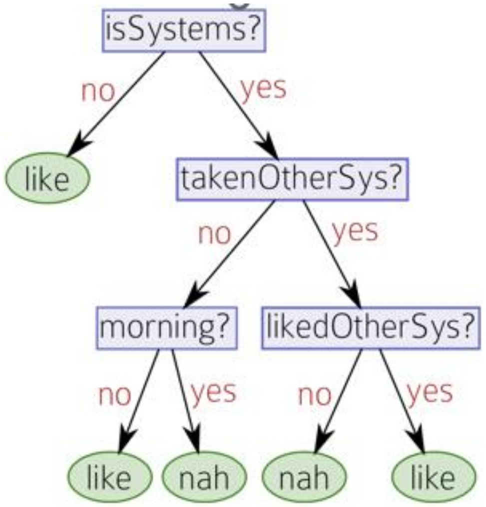

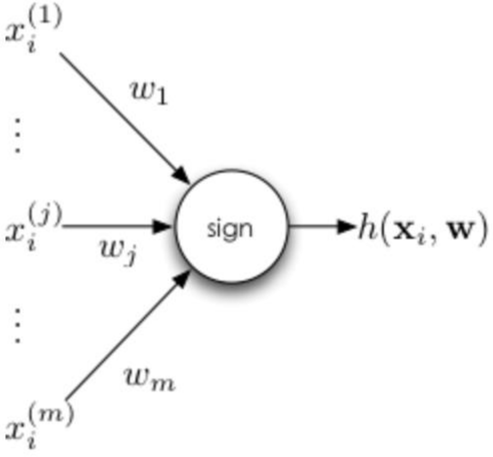

- Parameter: specifics of $\hat{f}$
  - E.g. for decision tree $\hat{f}$: tree structure, questions in nodes, labels in leaves
  - For linear classifier: linear coefficients

- Hyperparameter: specifics of learning algorithm $\mathcal{A}$
  - E.g. in DecisionTreeTrain, constrain to output tree of depth $\leq h$
  - Tuning hyperparameters often results in {over, under}-fitting

Hyperparameter tuning using validation set

- E.g. in decision tree training, how to choose tree depth $h \in \{1, \ldots, H\}$?

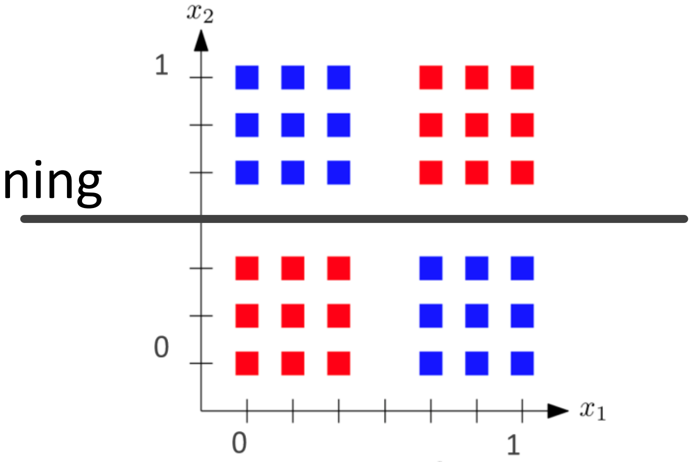

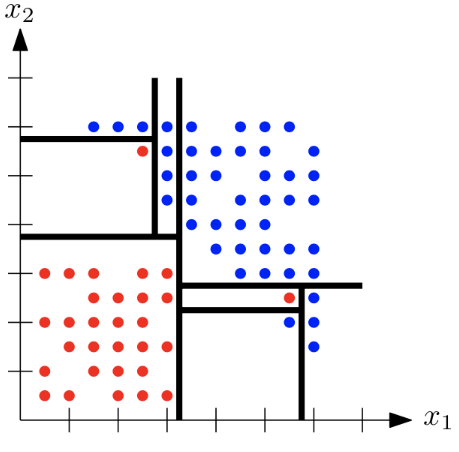

- For each hyperparameter $h \in \{1, \ldots, H\}$:
  - Train Tree_h using DecisionTreeTrain by constraining the tree depth to be h
- Choose one from $\mathrm{Tree}_1, \ldots, \mathrm{Tree}_H$

- Idea 1: choose Tree_h that minimizes training error
- Idea 2: choose Tree_h that minimizes test error
- Idea 3: further split training set to training set and validation set (development/hold-out set), (1) train Tree_h's using the (new) training set; (2) choose Tree_h that minimizes validation error

Training: 700 examples	Dev:100 examples	Test: 200 examples

Hyperparameter tuning using validation set

- E.g. in decision tree training, how to choose tree depth $h \in \{1, \ldots, H\}$?

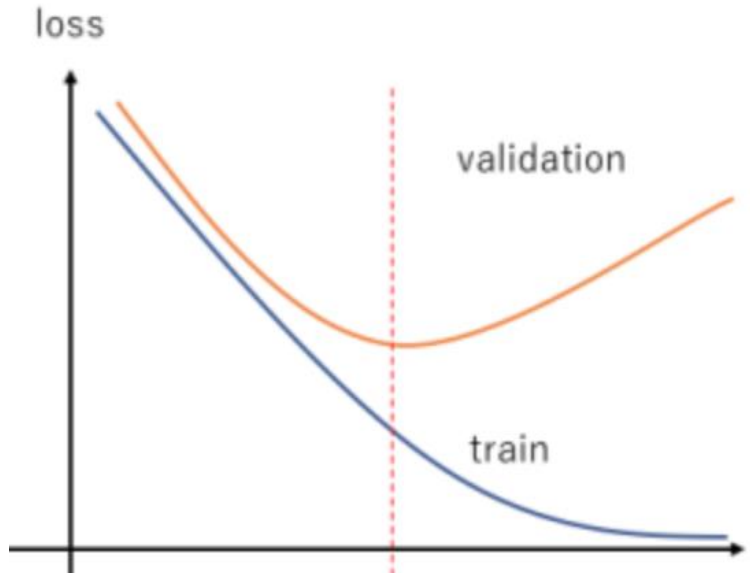

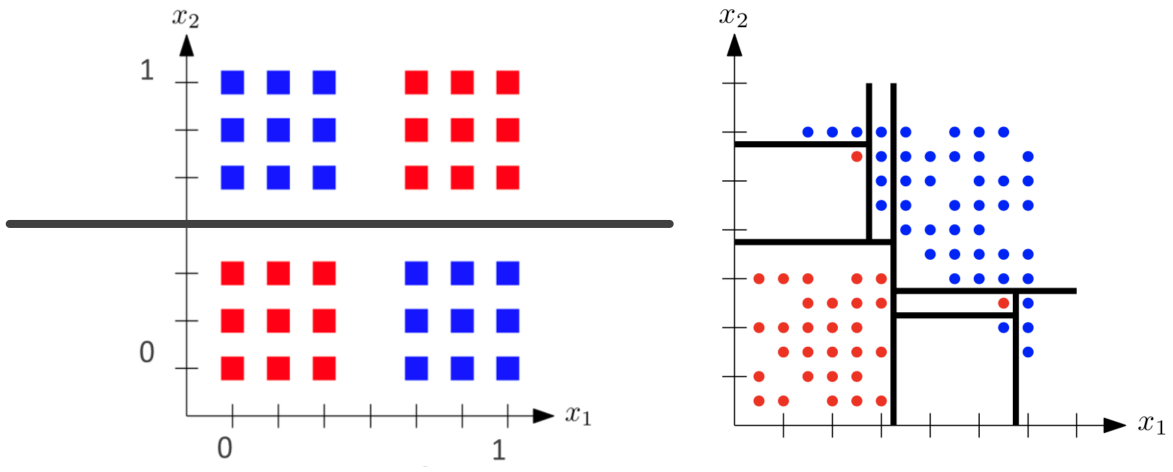

hyperparameter

- Law of large numbers => Validation error closely approximates test error & generalization error

Hyperparameter tuning: cross-validation

- Main idea: split the training / validation data in multiple ways

- For hyperparameter $h \in \{1, \dots, H\}$
  - For $k \in \{1, \dots, K\}$
    - train $\hat{f}_k^h$ with $S \setminus \mathrm{fold}_k$
    - measure error rate $e_{h,k}$ of $\hat{f}_k^h$ on fold_k
  - Compute the average error of the above: $\widehat{\mathrm{err}}^h = \frac{1}{K} \sum_{k=1}^{K} e_{h,k}$
- Choose $\hat{h} = \arg\min_h \widehat{\mathrm{err}}^h$
- Train $\hat{f}$ using S (all the training points) with hyperparameter $\hat{h}$

- $k = |S|$: leave one out cross validation (LOOCV)

Training set S

$\mathrm{fold}_1, \quad \dots, \quad \mathrm{fold}_5$

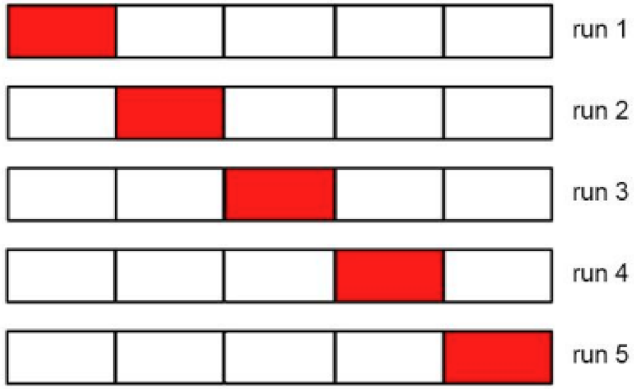

Inductive bias

- What classification problem is class A vs. class B?
    - Birds vs. Non-birds
    - Flying animals vs. non-flying animals

- Definition of **<u>inductive bias</u>**: in the absence of data that narrow down the target concept, what type of solutions are we likely to prefer?

- What is the inductive bias of learning shallow decision trees?

An example real-world machine learning pipeline

- Any step can go wrong
  - E.g. data collection, data representation
- Debugging pipeline: run oracle experiments
  - Assuming the downstream tasks are perfectly done, is this step achieving what we want?
- General suggestions:
  - Build the stupidest thing that could possibly work
  - Decide whether / where to fix it

1	real world goal	increase revenue
2	real world mechanism	better ad display
3	learning problem	classify click-through
4	data collection	interaction w/ current system
5	collected data	query, ad, click
6	data representation	bow^2, $\pm$ click
7	select model family	decision trees, depth 20
8	select training data	subset from april'16
9	train model & hyperparams	final decision tree
10	predict on test data	subset from may'16
11	evaluate error	zero/one loss for $\pm$ click
12	deploy!	(hope we achieve our goal)

Next lecture (8/31)

- Geometric view of machine learning; nearest neighbor methods

- Assigned reading: CIML Chap. 3 (Geometry and Nearest Neighbors)

- HW1 will be assigned