Introduction to Machine Learning

Alice Gao

Lecture 6

Readings: RN 18.1, 18.2. PM 7.1, 7.2.

Outline

Learning Goals

By the end of the lecture, you should be able to

- Identify reasons for building an agent that can learn.
- Describe different types of learning.
- Define supervised learning, classification, and regression.
- Define bias, variance, and describe the trade-off between the two.
- Describe how prevent over-fitting by performing cross validation.

Applications

- Medical diagnosis
- Spam filtering
- Facial recognition
- Speech understanding
- Handwriting recognition

Agents that learn

Learning is the ability of an agent to improve its performance on future tasks based on experience.

We want an agent to

- Do more
- Do things better
- Do things faster

Why would we want an agent to learn?

The Learning Architecture

- Problem/Task
- Experiences/Data
- Background knowledge/Bias
- Measure of improvement

Types of learning problems

- Supervised learning:
  Given input features, target features, and training examples, predict the value of the target features for new examples given their values on the input features.

- Unsupervised learning:
  Learning classifications when the examples do not have targets defined.

- Reinforcement Learning:
  Learning what to do based on rewards and punishments.

CQ: Supervised or Unsupervised Learning

CQ: We are given information on a user's credit card transactions. We would like to detect whether some of the transactions are fraudulent by finding some transactions that are different from the other transactions. We have no information on whether any particular transaction is fraudulent or not.

Is this a supervised or unsupervised learning problem?

(A) Supervised learning

(B) Unsupervised learning

Two types of supervised learning problems

- Classification: target features are discrete.
- Regression: target features are continuous.

CQ: Classification or regression

CQ: Is the following problem classification or regression?

You are given historical data on the weather condition (sunny, cloudy, rain, or snow) on a particular day of the year. You want to predict the weather condition of this day next year.

(A) Classification

(B) Regression

(C) This is not supervised learning.

CQ: Classification or regression

CQ: Is the following problem classification or regression?

You are given historical data on the price of a house at several points in time. You want to predict the price of this house next month.

(A) Classification

(B) Regression

(C) This is not supervised learning.

Supervised Learning

- Given training examples of the form $(x, f(x))$

- Return a function h (a.k.a a hypothesis)
  that approximates the true function f.

Learning as a search problem

- Given a hypothesis space, learning is a search problem.
- Search space is prohibitively large for systematic search.
- ML techniques are often some forms of local search.

Example: A prediction task

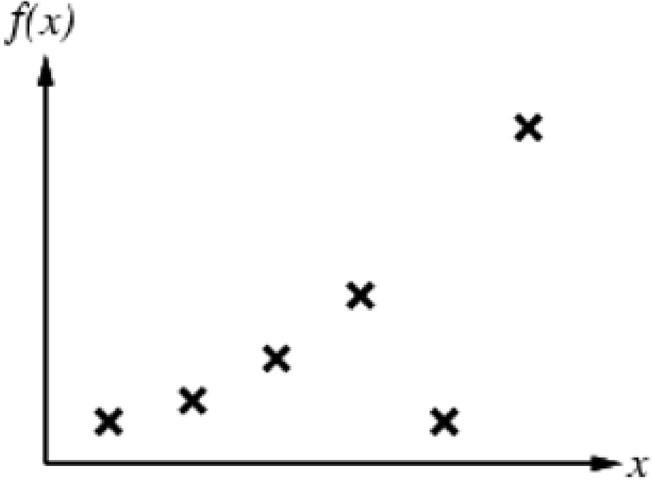

Example: A prediction task

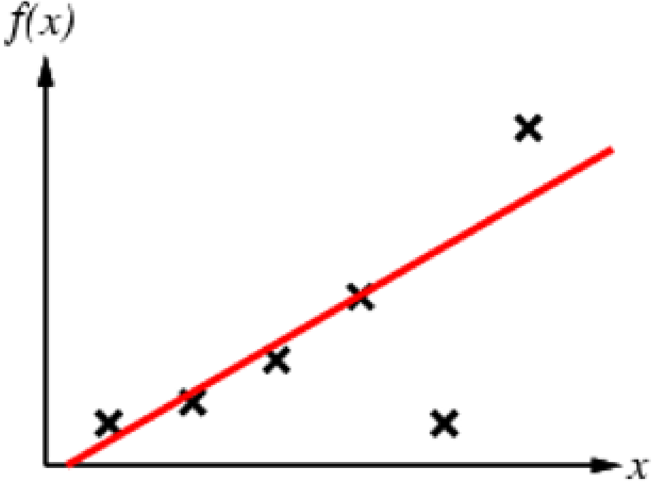

Example: A prediction task

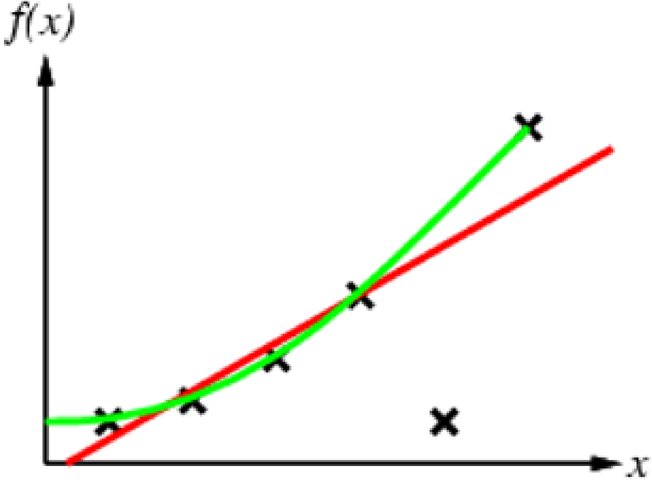

Example: A prediction task

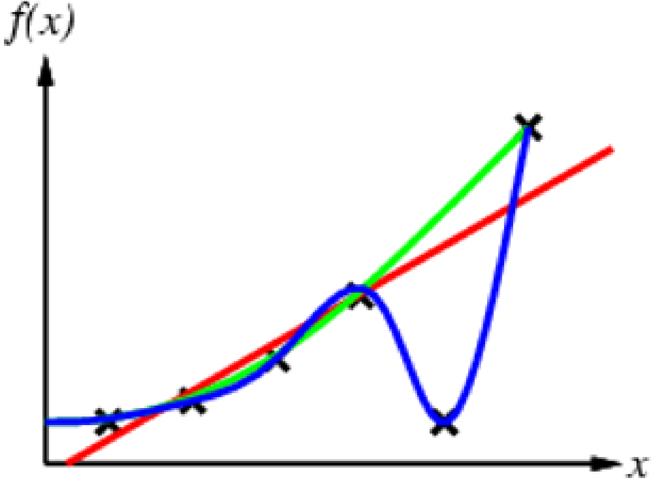

Example: A prediction task

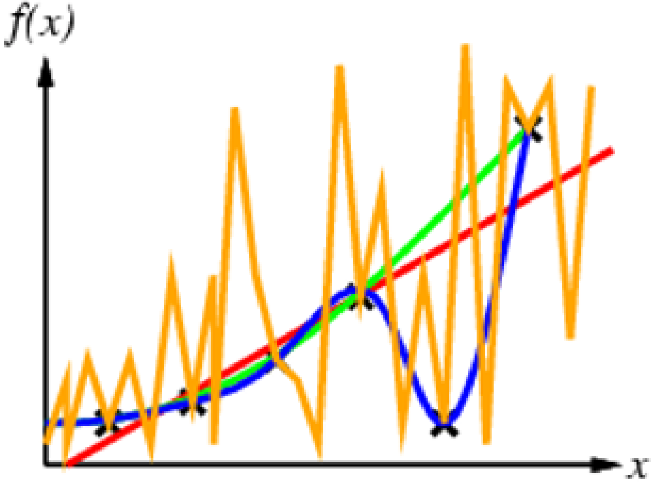

Generalization

- Goal of ML is to find a hypothesis
  that can predict unseen examples correctly.
- How do we choose a hypothesis that generalizes well?
  - Ockham's razor
  - Cross validation
- A trade-off between
  - complex hypotheses that fit the training data well
  - simpler hypotheses that may generalize better

Bias-Variance Trade-off

How well does the hypothesis fit the data
as the hypothesis becomes more complex?

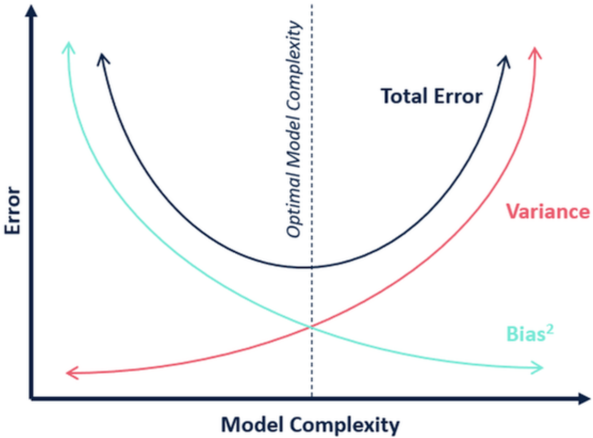

Bias-Variance Trade-off

- Bias: If I have infinite data, how well can I fit the data with my learned hypothesis?

  A hypothesis with high bias: makes strong assumptions, too simplistic, has few degrees of freedom, does not fit the training data well.

- Variance: How much does the learned hypothesis vary given different training data?

  A hypothesis with high variance: has a lot of degrees of freedom. is very flexible. whenever the training data changes, the hypothesis changes a lot. fits the training data very well.

Bias-Variance Trade-off

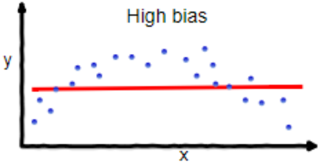

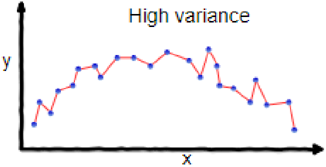

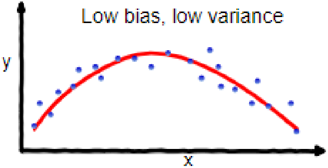

Cross Validation

How do we find a hypothesis that has low bias and low variance?
Use cross validation.

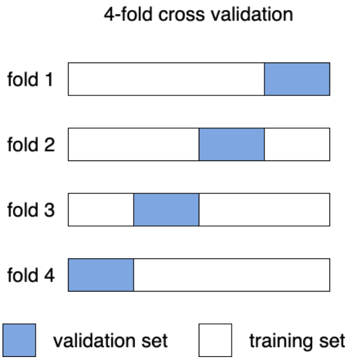

Cross Validation

- Break training data into K equally sized partitions.
- Train a learning algorithm on $K - 1$ partitions (training set).
- Test on the remaining 1 partition (validation set).
- Do this K times, each time testing on a different partition.
- Calculate the average error on the K validation sets.

After cross validation

After running cross-validation, you can

- Select one of the K trained hypotheses as your final hypothesis.
- Train a new hypothesis on all of the data, using parameters selected by cross-validation.

Over-fitting

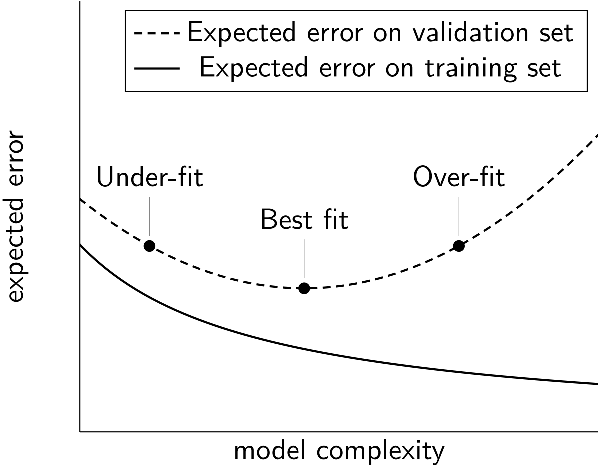

Revisiting the Learning Goals

By the end of the lecture, you should be able to

- Identify reasons for building an agent that can learn.
- Describe different types of learning.
- Define supervised learning, classification, and regression.
- Define bias, variance, and describe the trade-off between the two.
- Describe how prevent over-fitting by performing cross validation.