Artificial Intelligence and Machine Learning

Introduction — AI & Learning

Logistics

- Course home page: http://csci.viu.ca/~liuh/479
- General Information
- Late Policy: Tell me before the submit deadline and let's discuss a plan
- Exam Policy: One piece of letter-sized, double-sided paper with notes

Outline

- What is AI?
- What is learning?
- Types of learning
- What is machine learning?
- How does machine learning work?
- What can go wrong with machine learning
- Course overview

What is Intelligence

- the ability to understand and learn things;
- the ability to think and understand instead of doing things by instinct or automatically.

AI Definition

- A good general definition of AI could be:

  AI is the part of computer science concerned with designing intelligent computer systems, that is, computer systems that exhibit the characteristics we associate with intelligence in human behaviour - understanding language, learning, reasoning and solving problems.

Learning

- Learning is essential for unknown environments,
  - i.e., when designer lacks omniscience
- Learning is useful as a system construction method,
  - i.e., expose the agent to reality rather than trying to write it down
- Learning modifies the agent's decision mechanisms to improve performance

Learning Methods

Prediction Methods

- Use some variables to predict unknown or future values of other variables.
  - Classification [Predictive]
  - Regression [Predictive]
  - Outlier Detection [Predictive]

Description Methods

- Find human-interpretable patterns that describe the data.
  - Clustering [Descriptive]
  - Association Rule Discovery [Descriptive]

Types of Learning

Directly import from higher beings

Inductive Learning

- Supervised learning: correct answer for each example
- Unsupervised learning: correct answers not given
- Reinforced learning: occasional rewards

Supervised Machine Learning

Supervised Machine Learning techniques automatically learn a model from a set of historical examples.

Typical steps:

- Build a model (learning)
- Calibrate the model (validation)
- Use the model (applying)

What does a model look like?

A mathematical function

A table

A set of rules/logical expression

A tree

A diagram

A program

…

Mathematical Representation of Inductive learning

- Simplest form: learn a function from examples

f is the target function

An example is a pair (x, f(x))

Problem: find a hypothesis function h
such that h ≈ f
given a training set of examples

- (This is a highly simplified model of real learning:
  - Ignores prior knowledge
  - Assumes examples are given)

Inductive learning method

- Construct/adjust model h by searching through a set of possible models for the model that best captures the relationship between the input and the output

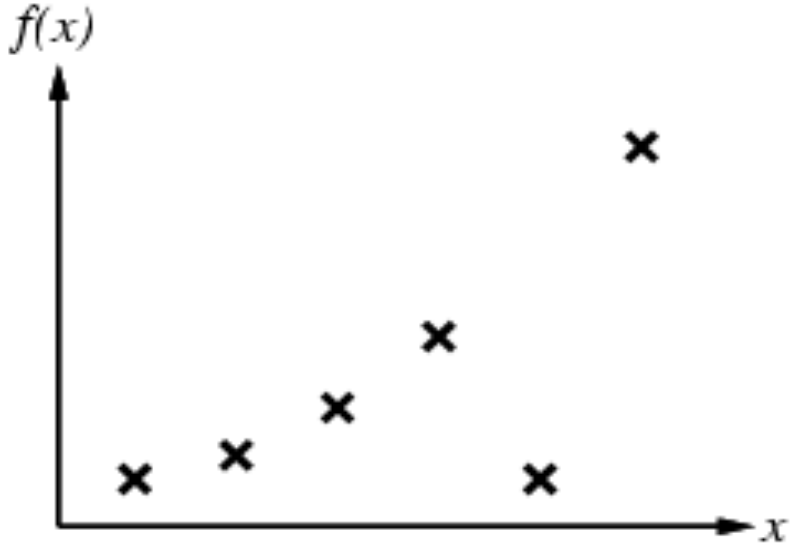

Inductive learning method

- h is consistent if it agrees with f on all examples

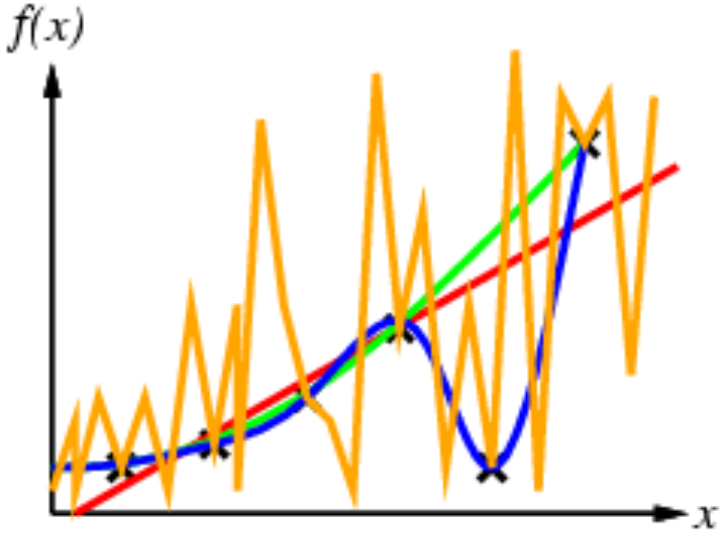

Inductive learning method

- Consistent model while ignoring noise in the dataset

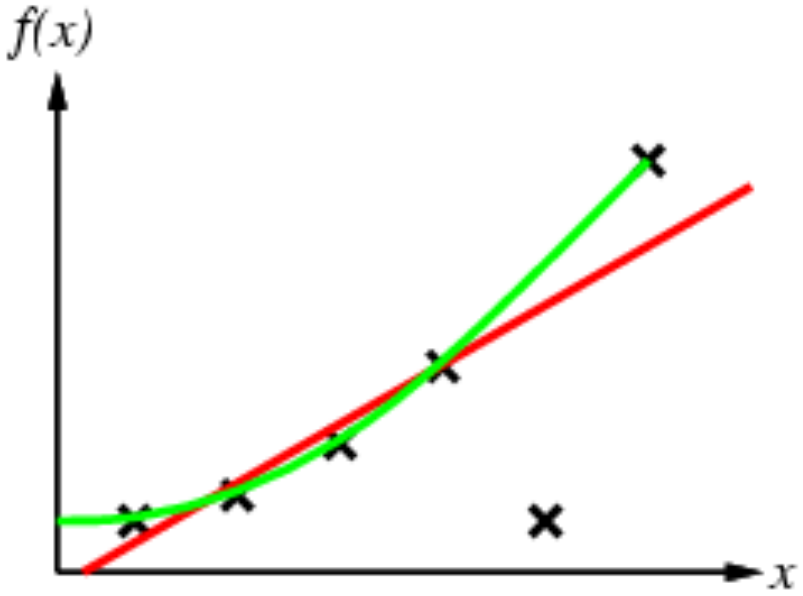

Inductive learning method

- Consistent model assuming there is no noise in the given dataset

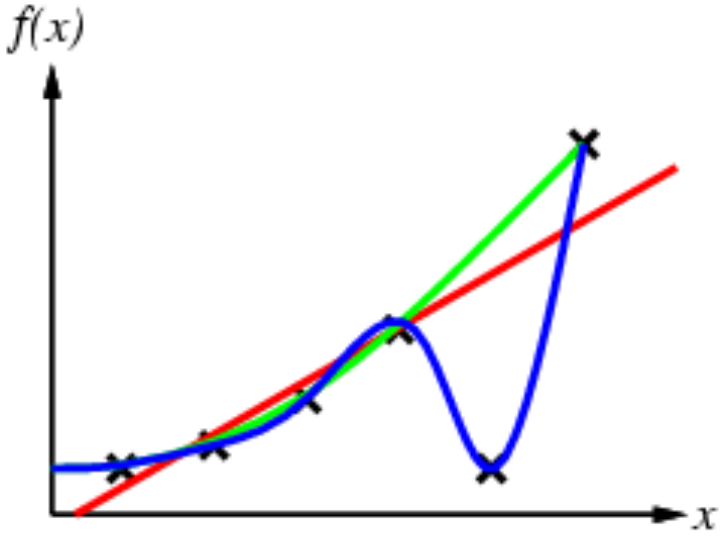

What can go wrong?

- Under-fitting

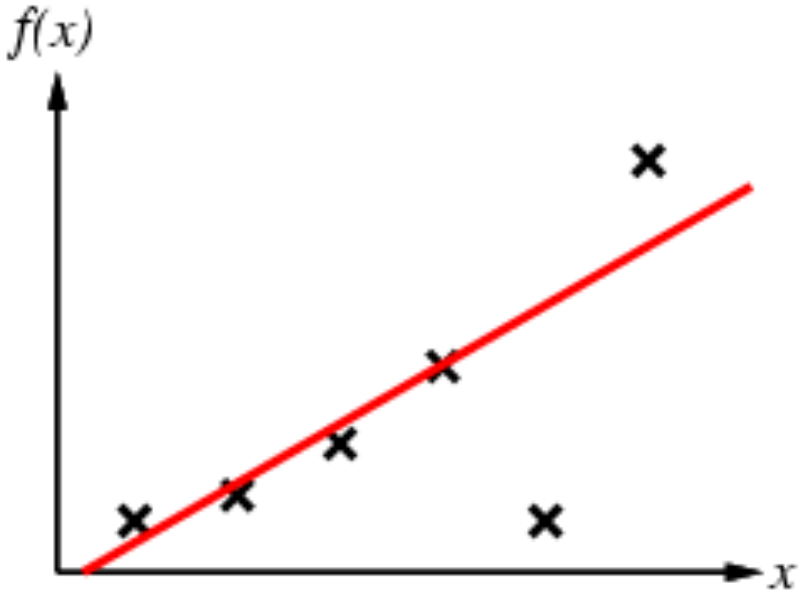

What can go wrong?

- Overfitting

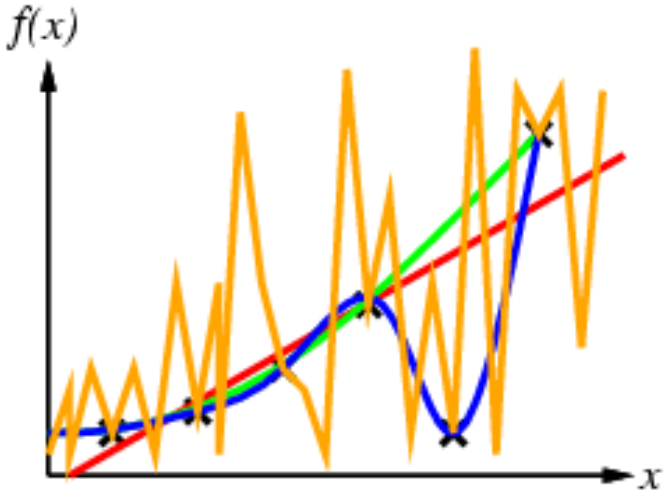

Ultimate Goal of Machine Learning

Construct a model that generalizes beyond the given dataset and that isn't influenced by the noise in the dataset.

How to achieve this goal? (Topics of our course)

Inductive Bias

The set of assumptions that define the model selection criteria of an machine learning algorithm

There are two types of bias:

- Restriction bias
- Preference bias
  - Ockham's razor: prefer the simplest hypothesis consistent with data

Summary

- Learning performance (of prediction methods) = prediction accuracy measured on test set
- Learning performance (for description methods) = description accuracy + interestingness
- Machine learning algorithms (for predictive models) work by searching through sets of potential models.
- There are two sources of information that guide this search:
  - The training data
  - The inductive bias of the algorithm
- Striking the right balance between model complexity and simplicity (between under-fitting and overfitting) is the hardest part of machine learning.