

CS 2710 Foundations of AI
Lecture 20-b

Learning

Milos Hauskrecht
milos@cs.pitt.edu
5329 Sennott Square

CS 2710 Foundations of AI

Machine Learning

- The field of **machine learning** studies the design of computer programs (agents) capable of learning from past experience or adapting to changes in the environment
- The need for building agents capable of learning is everywhere
 - Predictions in medicine, text classification, speech recognition, image/text retrieval, commercial software
- Machine learning is not only the deduction but induction of rules from examples that facilitate prediction and decision making

CS 2710 Foundations of AI

Learning

Learning process:

Learner (a computer program) takes data D representing past experiences and tries to either:

- to develop an appropriate response to future data, or
- describe in some meaningful way the data seen

Example:

Learner sees a set of past patient cases (patient records) with corresponding diagnoses. It can either try:

- to predict the presence of a disease for future patients
- describe the dependencies between diseases, symptoms (e.g. builds a Bayesian network for them)

Types of learning

- **Supervised learning**
 - Learning mapping between inputs x and desired outputs y
 - Teacher gives me y 's for the learning purposes
- **Unsupervised learning**
 - Learning relations between data components
 - No specific outputs given by a teacher
- **Reinforcement learning**
 - Learning mapping between inputs x and desired outputs y
 - Critic does not give me y 's but instead a signal (reinforcement) of how good my answer was
- **Other types of learning:**
 - **Concept learning, explanation-based learning, etc.**

Supervised learning

Data: $D = \{d_1, d_2, \dots, d_n\}$ a set of n examples

$$d_i = \langle \mathbf{x}_i, y_i \rangle$$

\mathbf{x}_i is input vector, and y is desired output (given by a teacher)

Objective: learn the mapping $f : X \rightarrow Y$


$$\text{s.t. } y_i \approx f(x_i) \quad \text{for all } i = 1, \dots, n$$

Two types of problems:


- **Regression:** X discrete or continuous \rightarrow
 Y is **continuous**
- **Classification:** X discrete or continuous \rightarrow
 Y is **discrete**

Supervised learning examples

- **Regression:** Y is **continuous**



- **Classification:** Y is **discrete**



Handwritten digit (array of 0,1s)