MACHINE LEARNING

1

Why study machine learning

- AI is the enterprise of design and analysis of intelligent agents.
- Intelligent behavior requires knowledge (e.g., model of the environment)
- Explicitly specifying the knowledge needed for specific tasks is hard, and often infeasible
- How to acquire knowledge?
  - Learning

2

Why study machine learning

- Learning modifies the agent's decision mechanisms to improve performance
- Environment changes over time – adapt to changes
- Learning is essential for unknown environments,
  - i.e., when designer lacks omniscience

3

Why study machine learning

- Applications
  - Medical diagnosis/image analysis (e.g. pneumonia, pap smears)
  - Scientific discovery
  - Spam Filtering, fraud detection (e.g. credit cards, phone calls)
  - Search and recommendation (e.g. google, amazon)
  - Automatic speech recognition & speaker verification
  - Locating/tracking/identifying objects in images & video (e.g. faces)

4

Example

Handwritten Digit Recognition

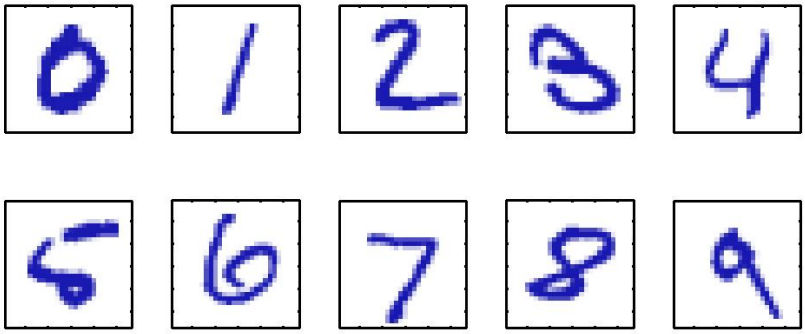

5

Learning problem

- Input vector x
- A training example is a pair (x, t), t: *target value* (category of the digit)
- *Training set* {(x,t)}
- *Learning/training* phase: find a model h(x)
- Use learned h(x) to categorize new examples

6

Learning problem

- Learn a function from examples
- *f* is the unknown target function
- An example is a pair (*x*, *f(x)*)
- Problem: find a hypothesis *h*
  - such that $h \approx f$
  - given a training set of examples
- Must choose hypothesis space
  - The set of polynomials, Decision trees
  - Bayesian networks, Neural networks

7

Canonical Learning Problems

- *Supervised Learning*: given examples of inputs and corresponding desired outputs, predict outputs on future inputs.
  - *Classification*: target has finite domain - categories
  - *Regression*: target has continuous domain
- *Unsupervised Learning*: given only inputs, automatically discover representations, features, structure, etc.
  - *Clustering*
- *Reinforcement Learning*: occasional rewards

8

Learning as Bayesian Inference

- Formulate the learning task as a process of probabilistic inference
- By updating beliefs about hypotheses based on data, we can learn about the world.
- Bayesian framework provides a sound probabilistic basis for understanding many learning algorithms and designing new algorithms
- Bayesian Decision Theory
  - A fundamental statistical approach to the problem of pattern classification

9