Machine Learning
CS489/698
Lecture 1: Jan 4th, 2017

Pascal Poupart

Associate Professor

David R. Cheriton School of Computer Science

University of Waterloo

Machine Learning

- Arthur Samuel (1959): Machine learning is the field of study that gives computers the ability to learn without being explicitly programmed.

- Tom Mitchell (1998): A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

Three categories

Supervised learning

Reinforcement learning

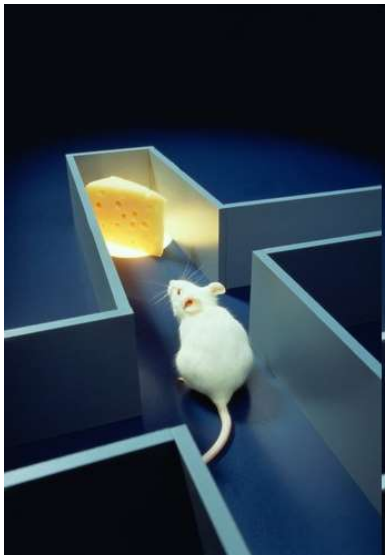

Unsupervised learning

Supervised Learning

- Example: digit recognition (postal code)

- Simplest approach: memorization

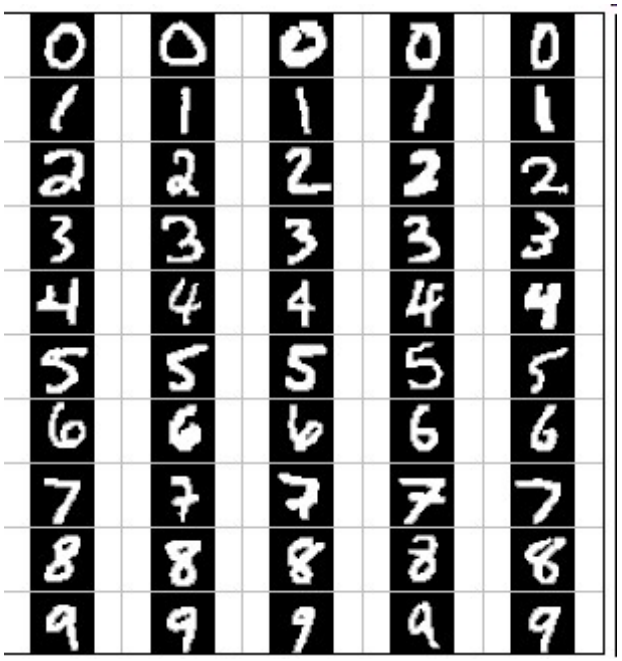

Supervised Learning

- Nearest neighbour:

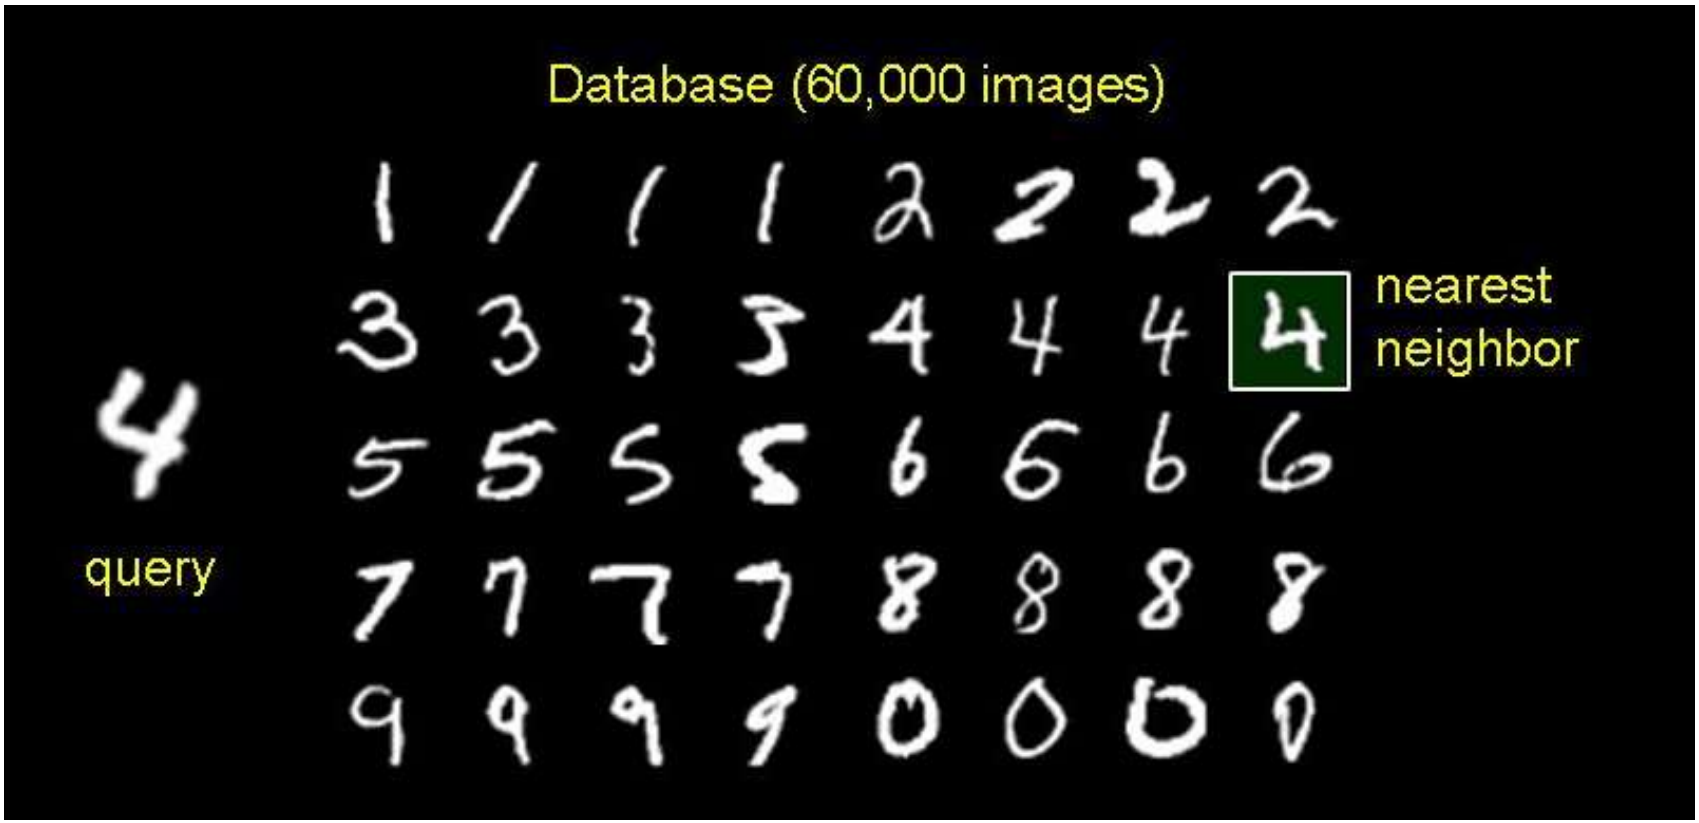

More Formally

- Inductive learning:
  - Given a training set of examples of the form (x,f(x))
    - x is the input, f(x) is the output
  - Return a function h that approximates f
    - h is called the hypothesis

Prediction

- Find function **h** that fits f at instances x

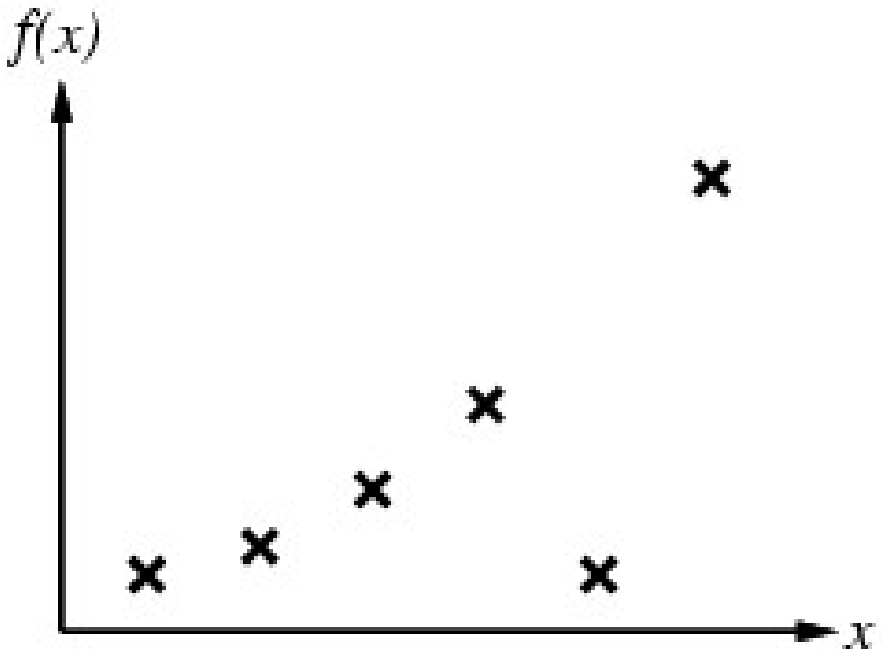

Prediction

- Find function **h** that fits f at instances x

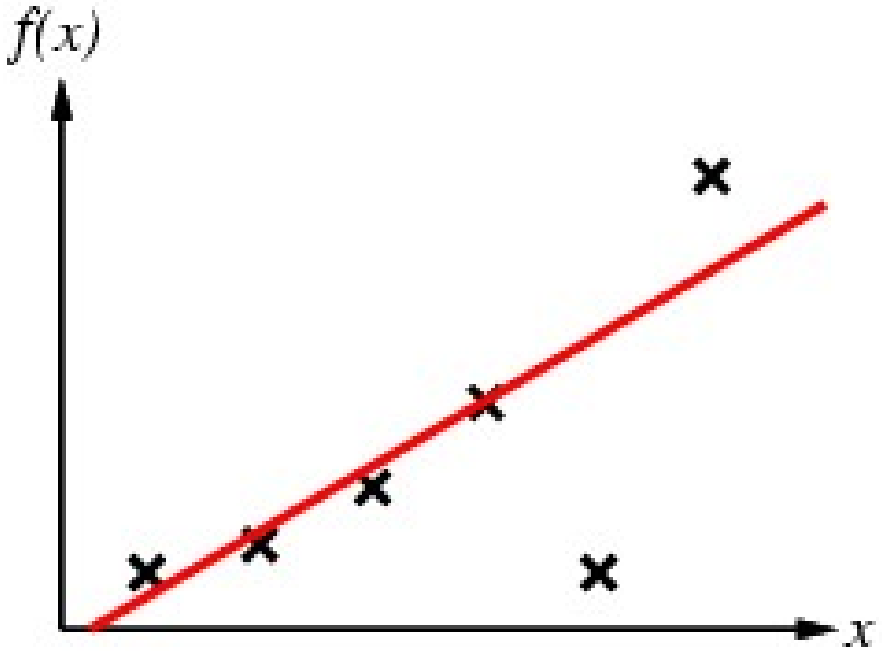

Prediction

- Find function **h** that fits f at instances x

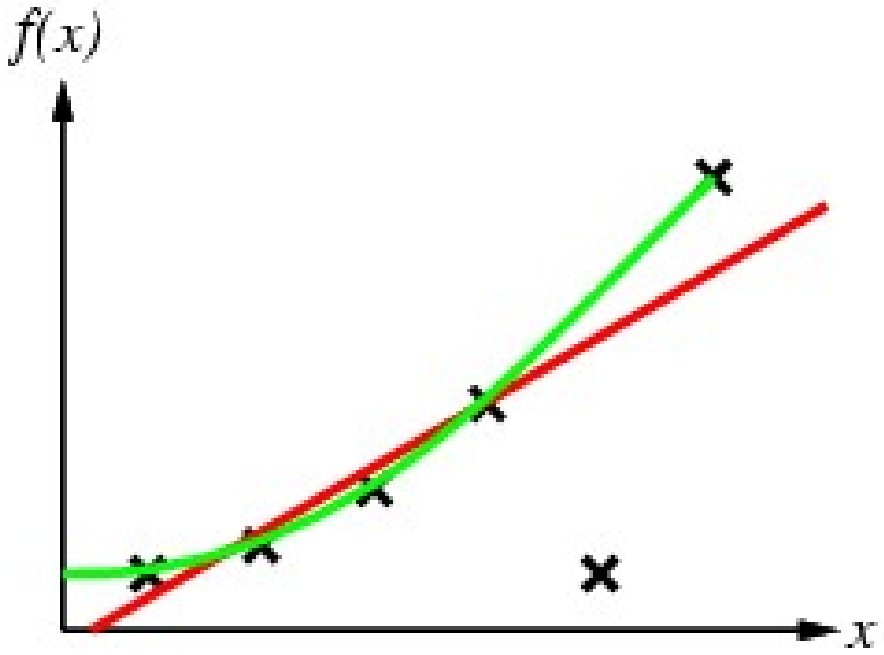

Prediction

- Find function **h** that fits f at instances x

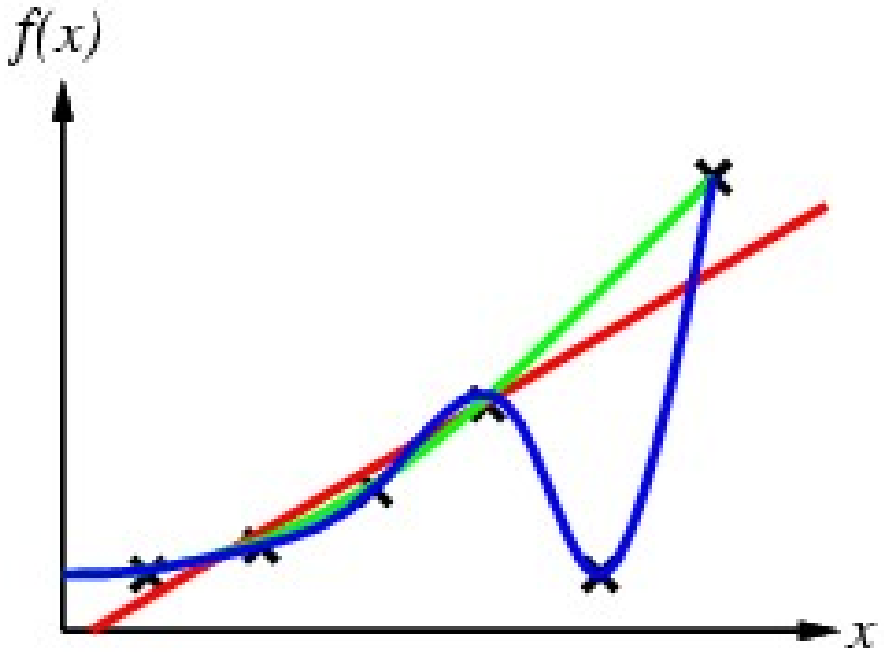

Prediction

- Find function **h** that fits f at instances x

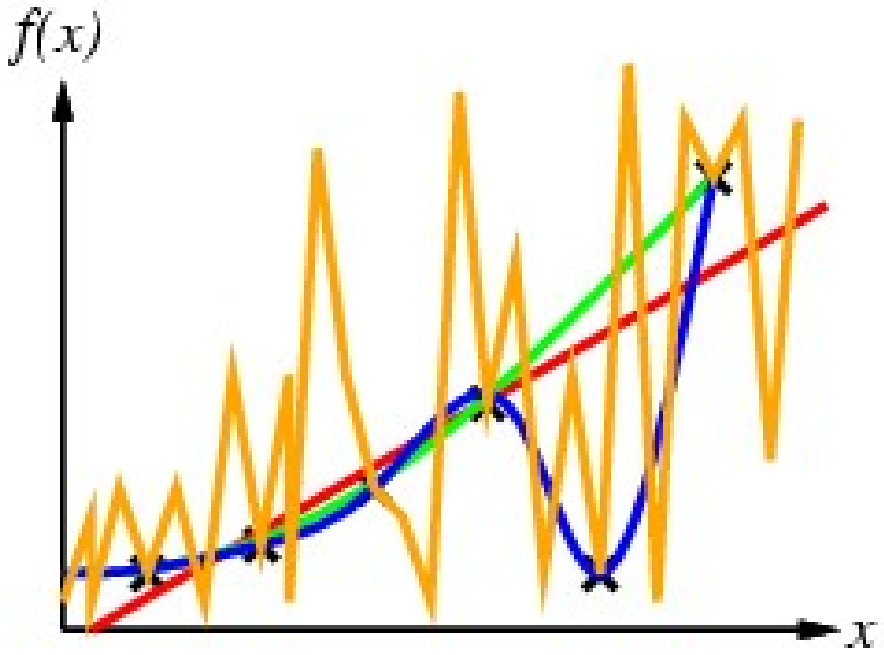

Generalization

- Key: a good hypothesis will generalize well (i.e. predict unseen examples correctly)

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

ImageNet Classification

- 1000 classes
- 1 million images

- Deep neural networks (supervised learning)

Autoencoders

- Unsupervised learning
- Compress and then reconstruct input

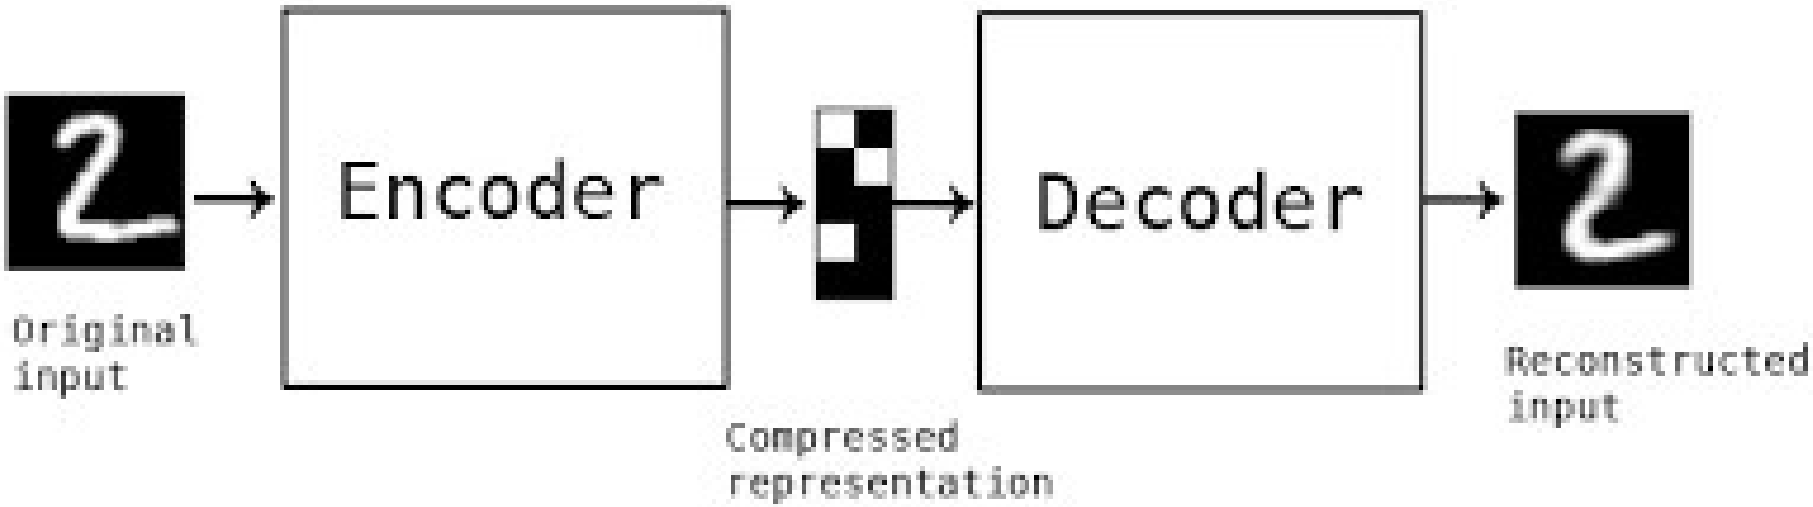

Unsupervised Feature Generation

- Encoder trained on large number of images

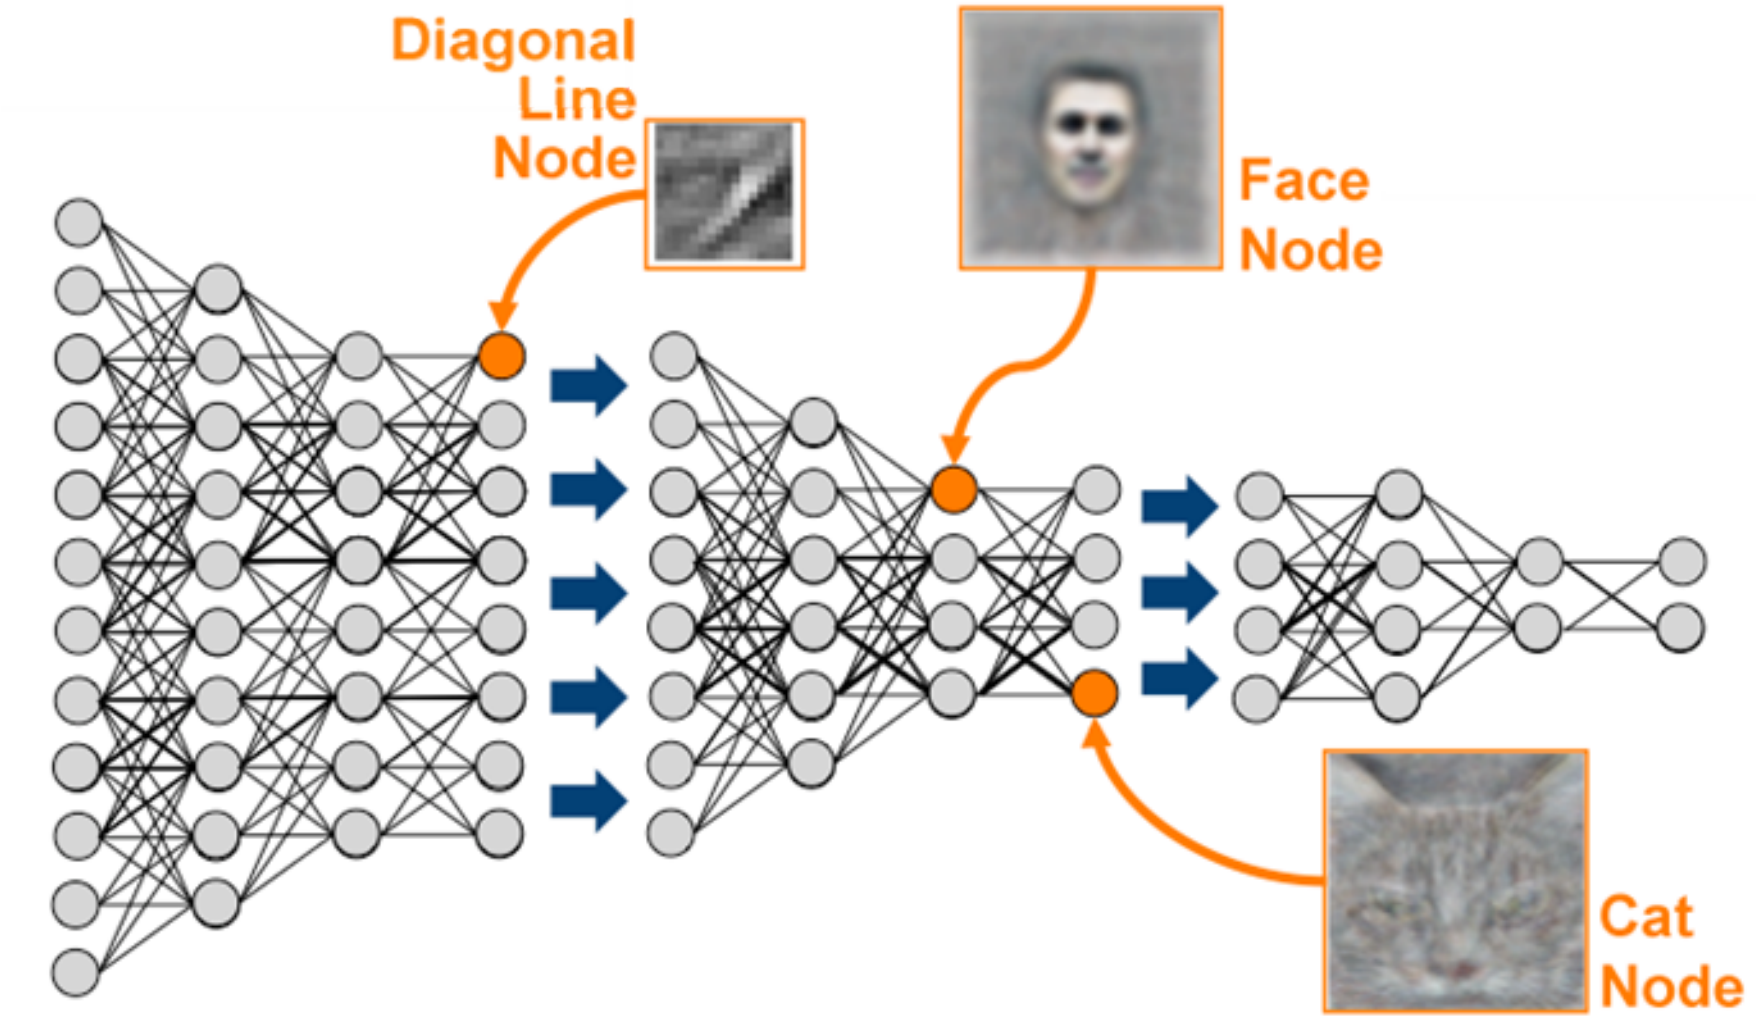

Reinforcement Learning

- Differs from supervised learning

Animal Psychology

- Negative reinforcements:
  - Pain and hunger
- Positive reinforcements:
  - Pleasure and food
- Reinforcements used to train animals

- Let's do the same with computers!

Helicopter Control

- Difficult to control:
  - Highly unstable

- Andrew Ng (Stanford, 2006):
  - Autonomous control by reinforcement learning
  - Step 1: learn neural net simulator based on flight data with human pilot
  - Step 2: optimize controller based on reinforcements for following a predefined trajectory

Smart Walker

- UW Researchers: Farheen Omar, Richard Hu, Adam Hartfiel, Mathieu Sinn, James Tung, Pascal Poupart

+

Force sensors
Load sensors
Video cameras
Microphone
Speech synthesizer
Servo-brakes
etc.

assist

caregivers

users

walker

devices

Smart walker

Research Goals

- Long-term goals:
  - Identify context and triggers of falls
  - Improved policies for wheelchair prescription & assisted living
  - Assess balance control and stability
  - Diagnose movement disorders
- Research performed:
  - Automated activity recognition (context)
  - 3D pose modeling (balance assessment, movement disorders)

Activity Recognition

- State of the art: kinesiologists hand label sensor data by looking at video feeds
  - Time consuming and error prone!

Backward view

Forward view

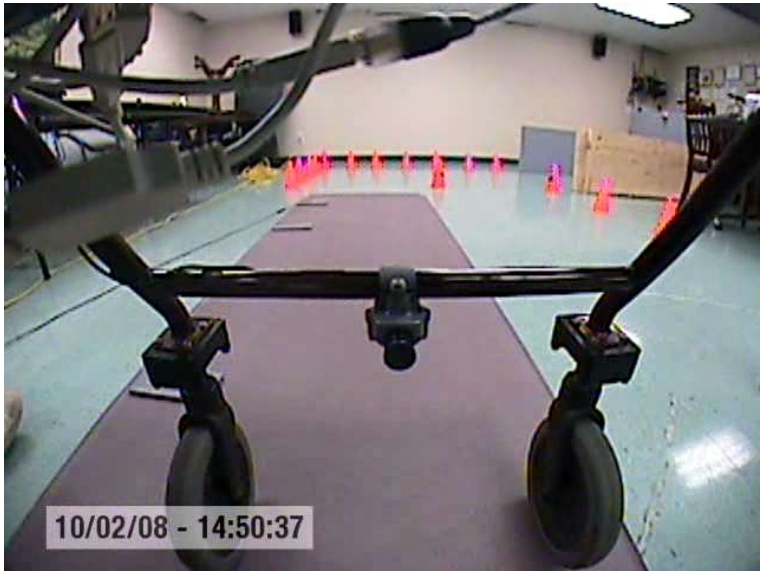

Raw Sensor Data

- 8 channels:
  - Forward acceleration
  - Lateral acceleration
  - Vertical acceleration
  - Load on left rear wheel
  - Load on right rear wheel
  - Load on left front wheel
  - Load on right front wheel
  - Wheel rotation counts (speed)

- Data recorded at 50 Hz and digitized (16 bits)

Experiment

- 8 walker users at Winston Park (84-97 years old)
- 12 older adults (80-89 years old) in the Kitchener-Waterloo area who do not use walkers

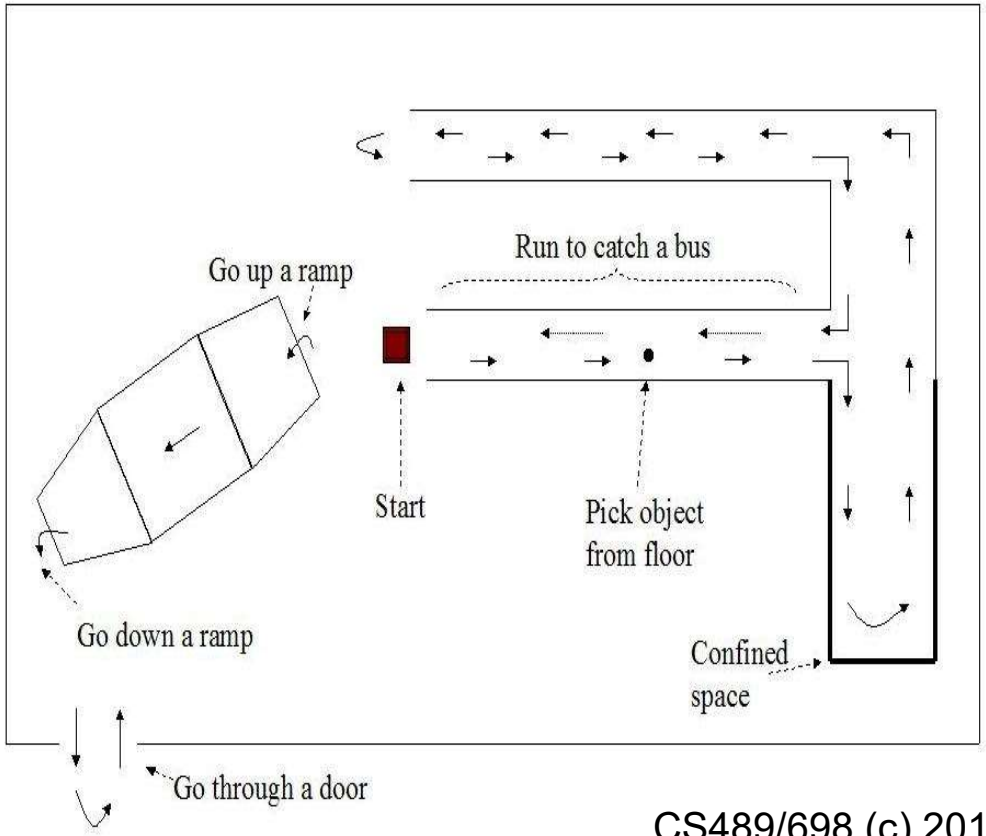

Activities

- Not Touching Walker (NTW)
- Standing (ST)
- Walking Forward (WF)
- Turning Left (TL)
- Turning Right (TR)
- Walking Backwards (WB)
- Sitting on the Walker (SW)
- Reaching Tasks (RT)
- Up Ramp/Curb (UR/UC)
- Down Ramp/Curb (DR/DC)

Probabilistic Models

- Hidden Markov Model (HMM)
  - Supervised
    - Maximum likelihood (ML)
  - Unsupervised
    - Expectation maximization (EM)
    - Bayesian Learning
- Conditional Random Field (CRF)
  - Supervised
    - Maximum conditional likelihood
    - Automated feature extraction

Demo


Visualization
File Playback
Video Data Annotation FullTimeseries read
Video
WP18_foot_course.AVI
File
WPSubject18_Time15-55-35DataAnno
Video
Previous Second
Next Second
Synch 1
Edit Text
Edit Text
SensorFlip
Sensor Viewing
Figure 1: Predictions
Predicted Behaviour
Correct Behaviour
13: going down curb
12: going up curb
11: going down ramp
10: going up ramp
9: sitting on walker
8: reaching tasks
7: transfer
6: walking backward
5: turn right
4: turn left
3: walking forward
2: stop/standing
1: not touching walker
Figure 2: Sensor Measurements
ch 0 - Accelerometer Y-axis
ch 1 - Accelerometer X-axis
ch 2 - Rear Right Load Cell
ch 3 - Front Right Load Cell
ch 4 - Front Left Load Cell
ch 5 - Rear Left Load Cell
ch 6 - Accelerometer Z-axis
ch 7 - Wheel encoder

Applications of Machine Learning

- Speech recognition
  - Siri, Cortana
- Natural Language Processing
  - Conversational agents
- Computer vision
  - Image and video analysis
- Robotic Control
  - Autonomous vehicles
- Intelligent assistants
  - Activity recognition, recommender systems

Vision

- Meta-programming: program computers to learn by themselves
- Lifelong machine learning: machines that continuously learn
- Transfer learning: machines that generalize their experience to new situations

- Challenges:
  - Knowledge representation
  - Computational complexity
  - Sample complexity