CS325 Artificial Intelligence
Ch 18a – Intro to Machine Learning

Cengiz Günay

Spring 2013

Overview

- Data-driven
- Lots of data (financial, internet, biology, etc.)?
- But also problems too difficult reflexive reasoning

Overview

- Data-driven
- Lots of data (financial, internet, biology, etc.)?
- But also problems too difficult reflexive reasoning

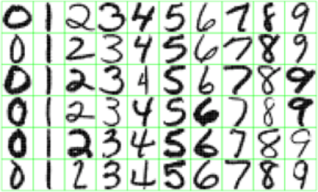

Figure 1.2: *Examples of handwritten digits from U.S. postal envelopes.*

Overview

- Data-driven
- Lots of data (financial, internet, biology, etc.)?
- But also problems too difficult reflexive reasoning

- Machine learning is inspired by the brain and neurons

The Principal Regions of the Neocortex - Side View. Figure 6.15

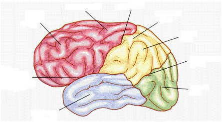

Overview

- Data-driven
- Lots of data (financial, internet, biology, etc.)?
- But also problems too difficult reflexive reasoning

- Machine learning is inspired by the brain and neurons

Jackson Beatty, Principles of Behavioral Neuroscience. Copyright © 1995 Times Mirror Higher Education Group, Inc., Dubuque, IA.

The Principal Regions of the Neocortex - Side View. Figure 6.15

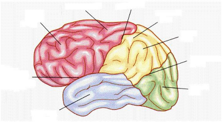

- Types of machine learning:
  - supervised (this and next class)
  - unsupervised (next week)
  - reinforcement (later)

Who uses Machine Learning (ML)?

- Compared to Bayes Nets:
    - Bayes Nets require full knowledge
    - ML is data-driven. How about sampling Bayes Nets?

Who uses Machine Learning (ML)?

- Compared to Bayes Nets:
  - Bayes Nets require full knowledge
  - ML is data-driven. How about sampling Bayes Nets?
- Companies use ML for?

Who uses Machine Learning (ML)?

- Compared to Bayes Nets:
  - Bayes Nets require full knowledge
  - ML is data-driven. How about sampling Bayes Nets?
- Companies use ML for?
  - Product recommendations: Amazon, Netflix (had an ML contest)
  - Typing: Swype keyboard, learning word suggestions
  - Pattern recognition: handwriting, OCR, audio
  - Web mining: Google page rank algorithm


amazon
Join Prime
Cengiz's Amazon.com
Today's Deals
Gift Cards
Customer Service
Shop by Department
Search
Amazon Instant Vi...
Go
Amazon Instant Video
Prime Instant Videos
Most Popular
Video Finder
Getting Started
Your Video Library
THE IT CROWD
The Complete Third Season
The IT Crowd Season 3, Ep. 1 "From Hell"
(47 customer reviews)
After Roy goes to Jen's house to retrieve £5, he sees Jen's builder who he recognises as a Builder from Hell. This convinces Jen to keep an eye on him, so he doesn't urinate in her sinks. On his way to work, Moss is harassed by a group of teenagers ...
Runtime: 24 minutes
Original air date: November 21, 2008
Network: MPI Media Group

amazon Join Prime

Cengiz's Amazon.com | Today's Deals | Gift Cards | Customer Service

Shop by **Department** | Search | Amazon Instant Vi... | Go

Amazon Instant Video | Prime Instant Videos | Most Popular | Video Finder | Getting Started | Your Video Library

The IT Crowd Season 3, Ep. 1 "From Hell"

★★★★★ (47 customer reviews)

After Roy goes to Jen's house to retrieve £5, he sees Jen's builder who he recognises as a Builder from Hell. This convinces Jen to keep an eye on him, so he doesn't urinate in her sinks. On his way to work, Moss is harassed by a group of teenagers ...

Runtime: 24 minutes

Original air date: November 21, 2008

Network: MPI Media Group

Customers Who Bought This Item Also Bought

The IT Crowd Season 2
Chris O'Dowd
★★★★★ (49)
Amazon Instant Video
$1.99 per episode
$5.99 for the season

The IT Crowd Season 1
Chris O'Dowd
★★★★☆ (121)
Amazon Instant Video
$1.99 per episode
$5.99 for the season

The Mighty Boosh Season 1
★★★★★ (32)
Amazon Instant Video
$1.99 per episode
$14.99 for the season

Black Books Season 1
Dylan Moran
★★★★★ (32)
Amazon Instant Video
$1.99 per episode
$10.99 for the season

Stanford's Stanley

ML Taxonomy

Learn what?

ML Taxonomy

Learn what? parameters, structure, hidden concepts

ML Taxonomy

Learn what? parameters, structure, hidden concepts

How?

ML Taxonomy

Learn what?	parameters, structure, hidden concepts
How?	from labels (supervised), from nature of data (unsupervised), from environmental feedback (reinforcement)

ML Taxonomy

Learn what?	parameters, structure, hidden concepts
How?	from labels (supervised), from nature of data (unsupervised), from environmental feedback (reinforcement)
Why?	

ML Taxonomy

Learn what?	parameters, structure, hidden concepts
How?	from labels (supervised), from nature of data (unsupervised), from environmental feedback (reinforcement)
Why?	future prediction (stock market), diagnosis (Bayes Nets), summarize (Google search), classify (digit recognition)

ML Taxonomy

Learn what? parameters, structure, hidden concepts

How? from labels (supervised), from nature of data (unsupervised), from environmental feedback (reinforcement)

Why? future prediction (stock market), diagnosis (Bayes Nets), summarize (Google search), classify (digit recognition)

Types: passive/active, online/offline

ML Taxonomy

Learn what?	parameters, structure, hidden concepts
How?	from labels (supervised), from nature of data (unsupervised), from environmental feedback (reinforcement)
Why?	future prediction (stock market), diagnosis (Bayes Nets), summarize (Google search), classify (digit recognition)
Types:	passive/active, online/offline

Based on output: classification vs. regression

ML Taxonomy

Learn what? parameters, structure, hidden concepts

How? from labels (supervised), from nature of data (unsupervised), from environmental feedback (reinforcement)

Why? future prediction (stock market), diagnosis (Bayes Nets), summarize (Google search), classify (digit recognition)

Types: passive/active, online/offline

Based on output: classification vs. regression

Based on behavior: generative vs. discriminative

Where did it all start?

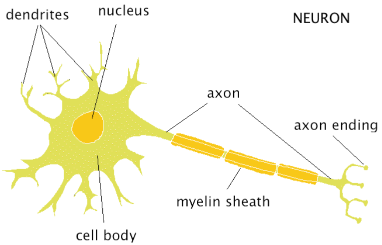

Inputs on *dendrites*, **outputs** from *axon*

Where did it all start?

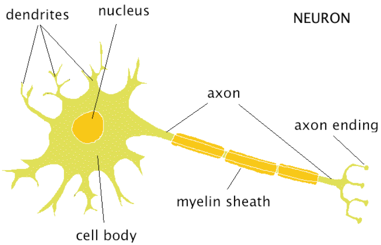

Inputs on *dendrites*, **outputs** from *axon*

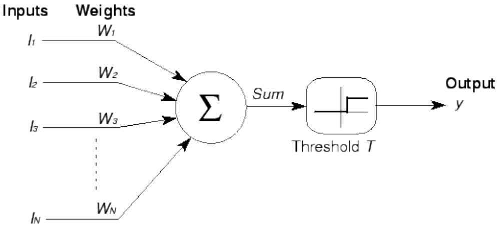

the **perceptron**

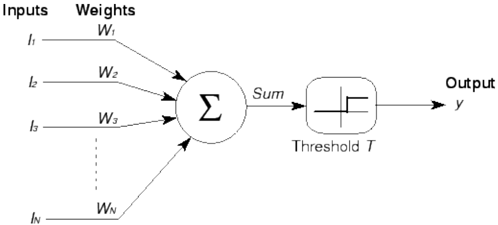

Separate Tom from Jerry based on car preference?

	Tom	Jerry
Trucks	1	0
Sedans	0	1
Hybrids	0	1
SUVs	1	0

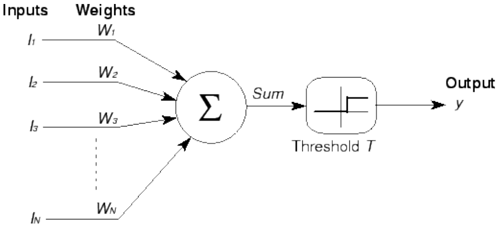

Separate Tom from Jerry based on car preference?

	Tom	Jerry
Trucks	1	0
Sedans	0	1
Hybrids	0	1
SUVs	1	0

$$\times \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix}$$

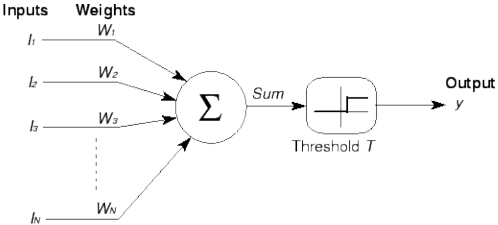

Separate Tom from Jerry based on car preference?

	Tom	Jerry
Trucks	1	0
Sedans	0	1
Hybrids	0	1
SUVs	1	0

$$\times \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix} \Rightarrow \sum_i I_i W_i = ?$$

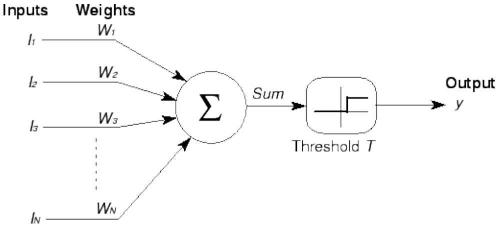

Separate Tom from Jerry based on car preference?

	Tom	Jerry
Trucks	1	0
Sedans	0	1
Hybrids	0	1
SUVs	1	0

$$\times \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix} \Rightarrow \sum_i I_i W_i = ?$$

In general...

- ML learns mapping between features and labels:

$$x_1 \dots x_n \rightarrow y_n$$

- Can be applied to different problems as long as can be vectorized (e.g., images)
- Need multiple examples (or samples)
- Question is to find function for each sample, m:

$$f(X_m) = Y_m$$

Occam's razor:

Occam's razor: prefer simplest solution

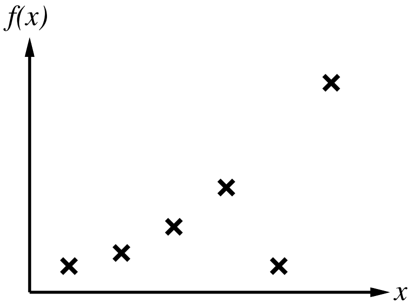

Problem with choosing type of function: complexity

Occam's razor: prefer simplest solution

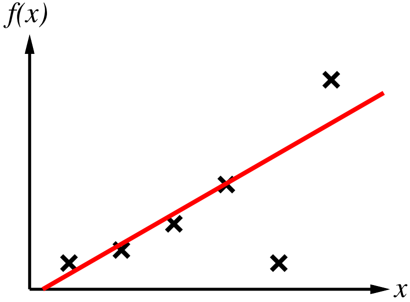

Problem with choosing type of function: complexity

Occam's razor: prefer simplest solution

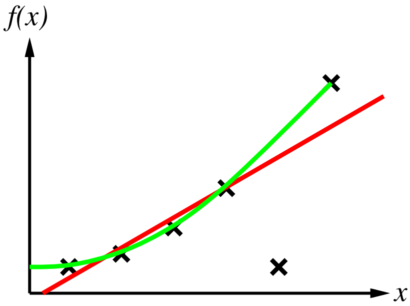

Problem with choosing type of function: complexity

Occam's razor: prefer simplest solution

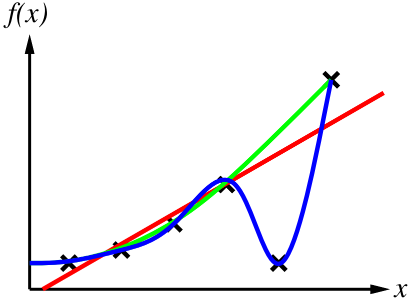

Occam's razor: prefer simplest solution