http://www.cs.washington.edu/education/courses/cse546/14au/

What's learning? Point Estimation

Machine Learning – CSE546

Carlos Guestrin
University of Washington

September 25, 2014

What is Machine Learning ?

Machine Learning

Study of algorithms that

- improve their performance
- at some task
- with experience

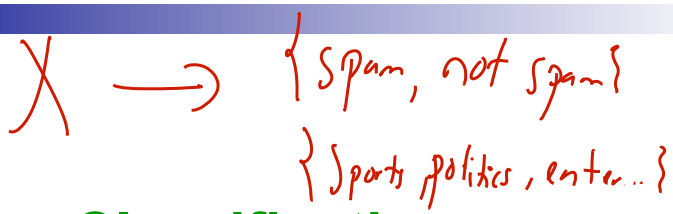

Classification

from data to discrete classes

Spam filtering

data → f → prediction

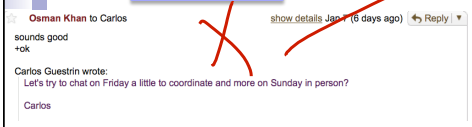

f(x) → not spam

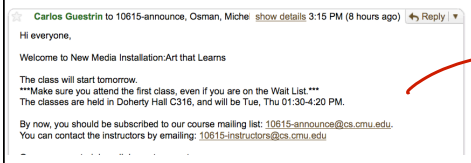

f(x) → not spam

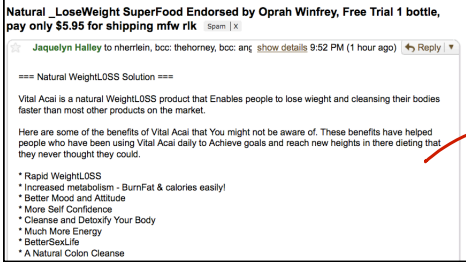

f(x) → Spam

Text classification

→ Company home page

vs

Personal home page

vs

Univeristy home page

vs

…

Object detection

(Prof. H. Schneiderman)

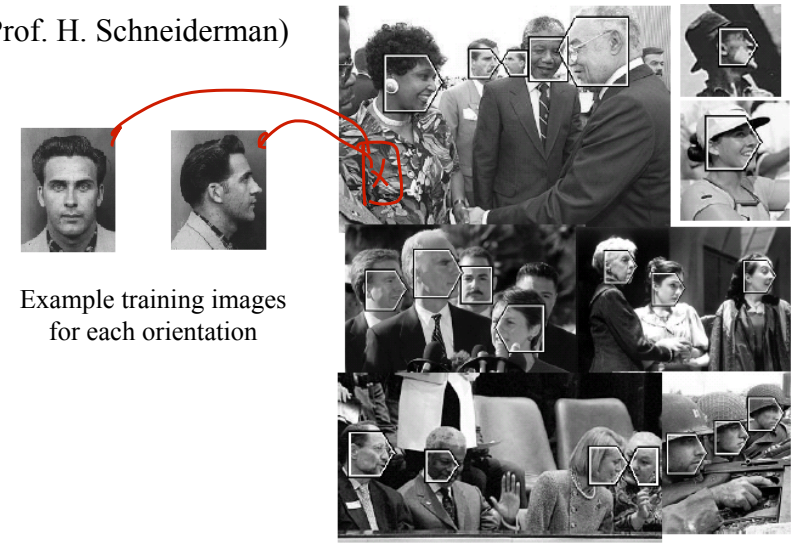

Example training images
for each orientation

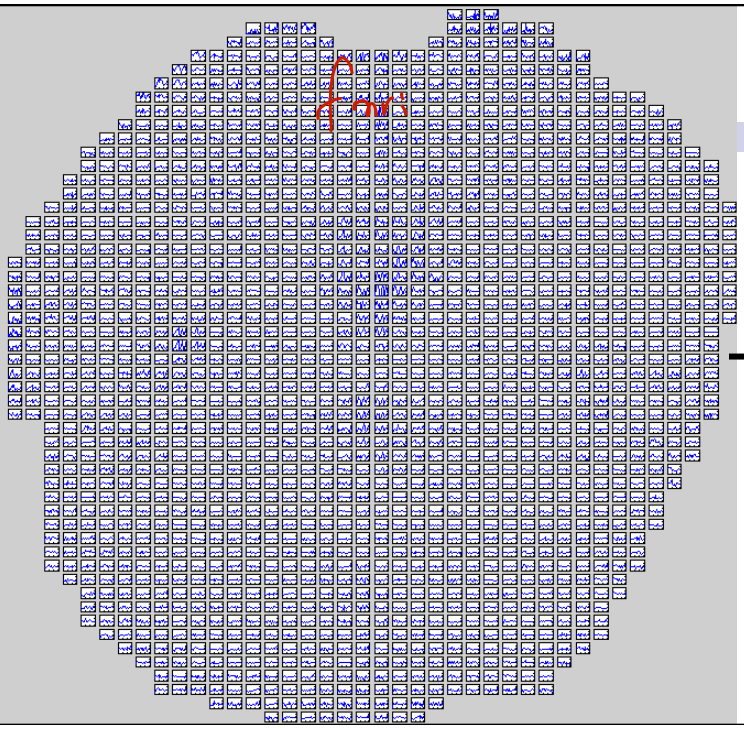

Reading
a noun
(vs verb)

[Rustandi et al.,
2005]

Weather prediction

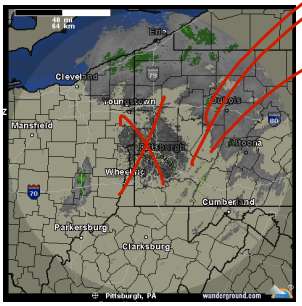

The classification pipeline

Training

training data

$\{(x_1, y_1), (x_2, y_2) \ldots\}$

email x_1 label y_1

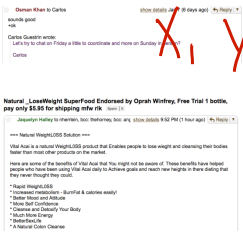

learn $f : X \rightarrow Y$

Testing

/prediction/using/...

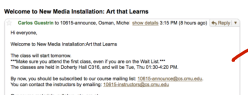

new email x_{new}

prediction $f(x_{new}) = \hat{y}$

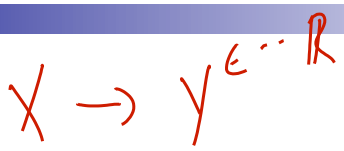

Regression

predicting a numeric value

Stock market

Weather prediction revisted

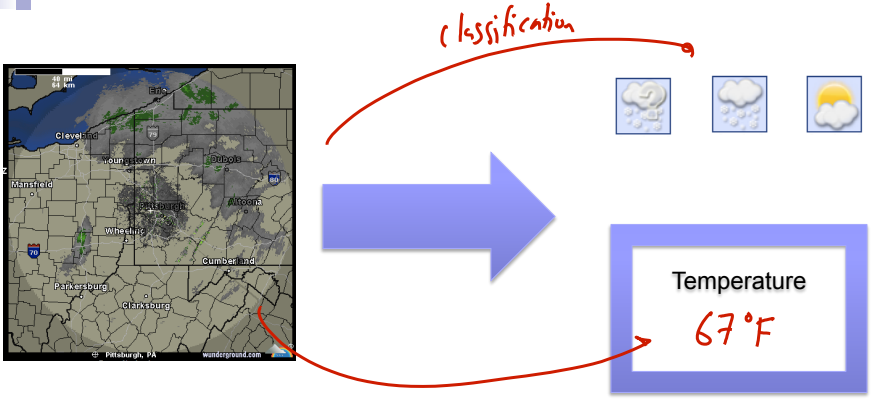

Modeling sensor data

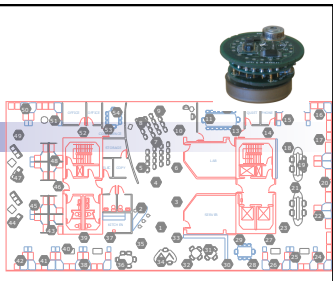

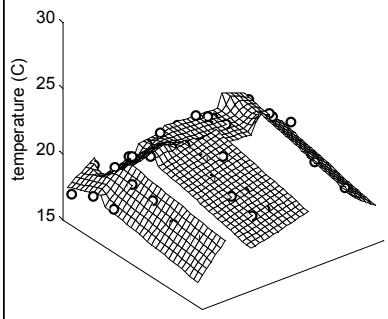

- Measure temperatures at some locations
- Predict temperatures throughout the environment

[Guestrin et al. '04]

Unsupervised

Similarity

finding data

Given image, find similar images

same input

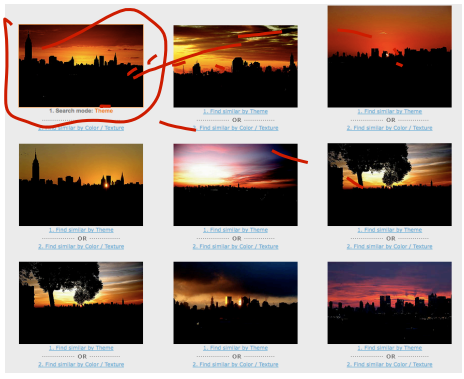

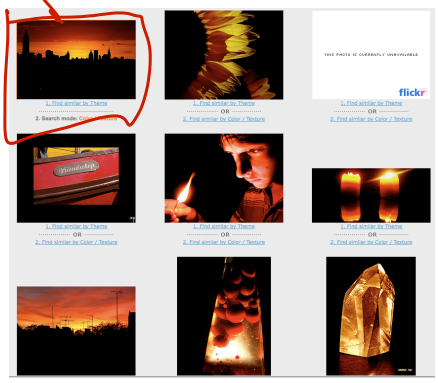

Sunsets

Color

http://www.tiltomo.com/

Similar products

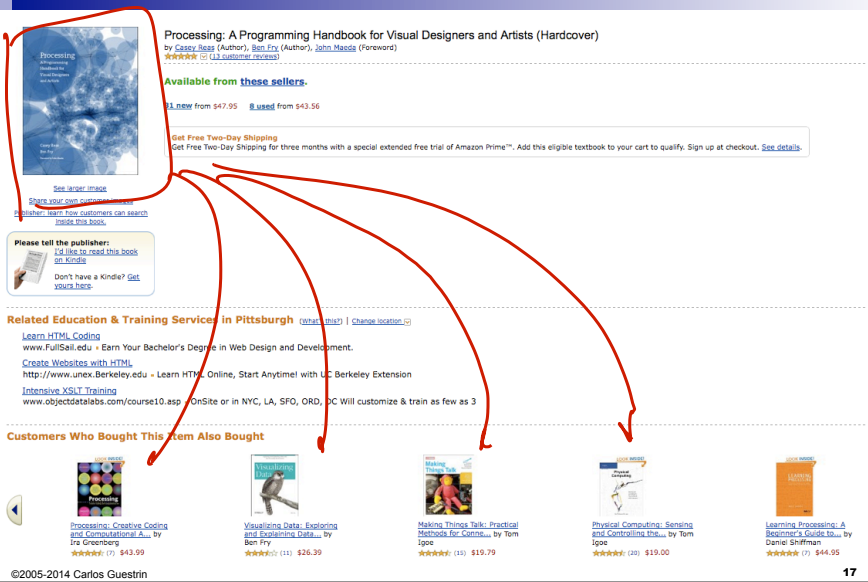

Clustering

discovering structure in data

Clustering Data: Group similar things

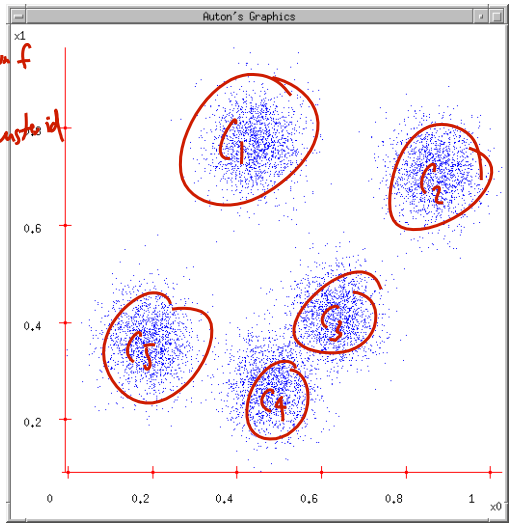

Learning a function f

f : location → clusterid

diff

Supervised learning

training data had labeled classes

Clustering images

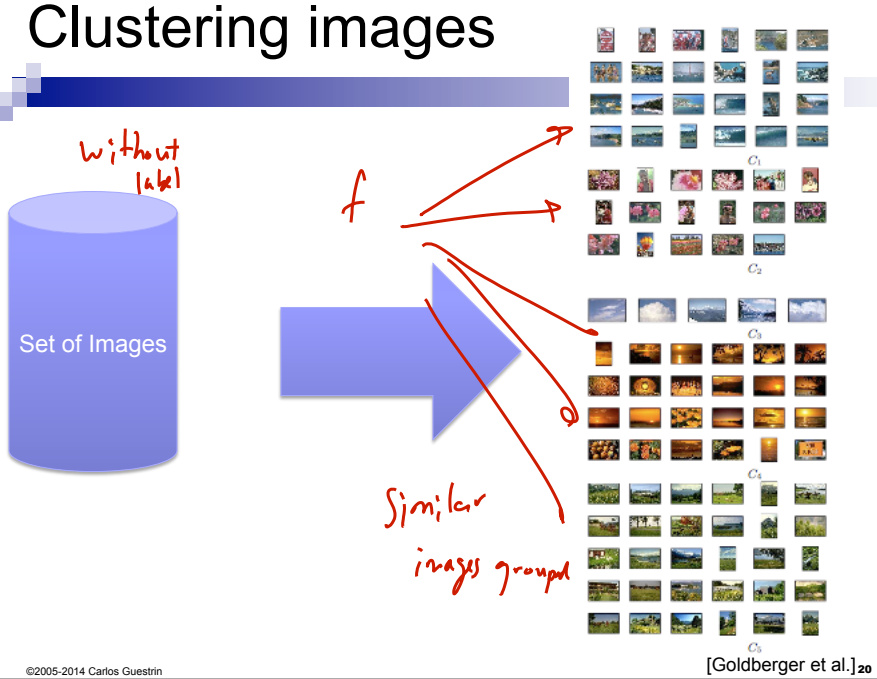

[Goldberger et al.]

Clustering web search results

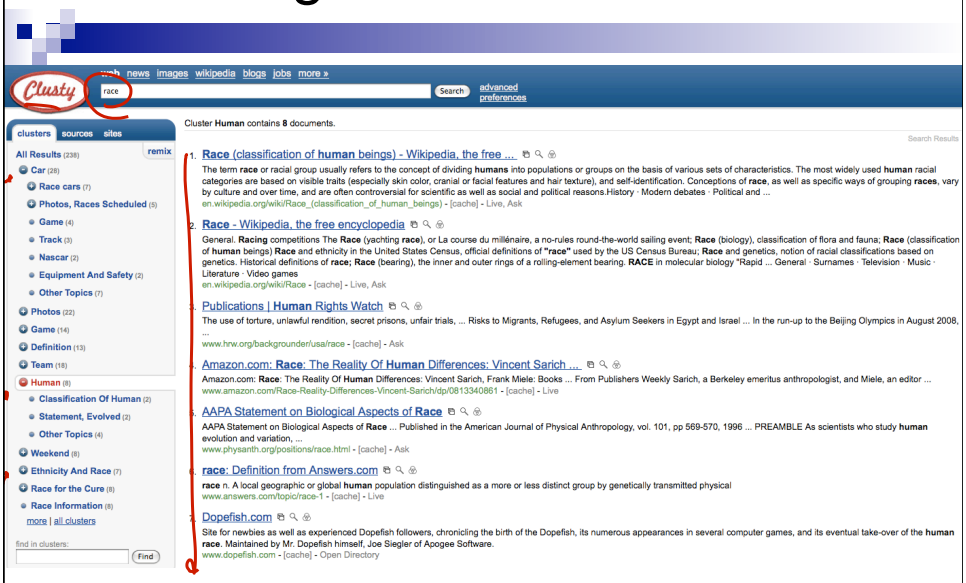

Embedding

visualizing data

Embedding images

1 000 000 dims → 2 dims

Images have thousands or millions of pixels.

Can we give each image a coordinate, such that similar images are near each other?

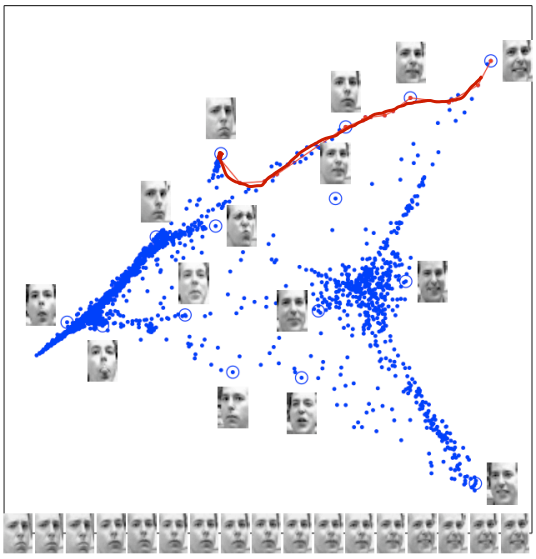

[Saul & Roweis '03]

Embedding words

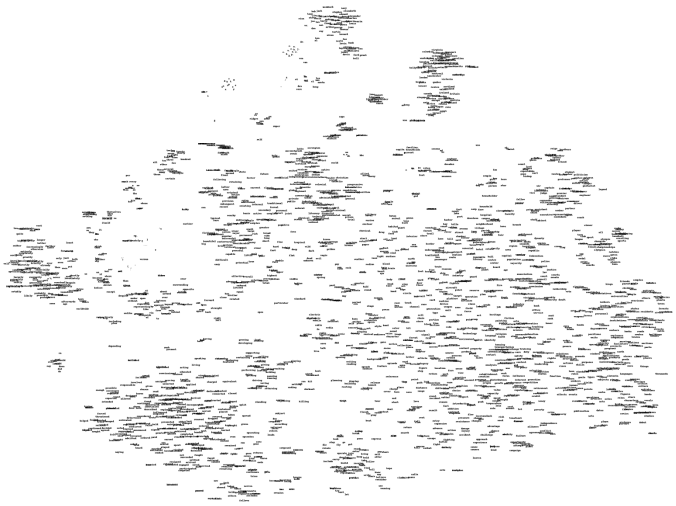

[Joseph Turian]

Embedding words (zoom in)

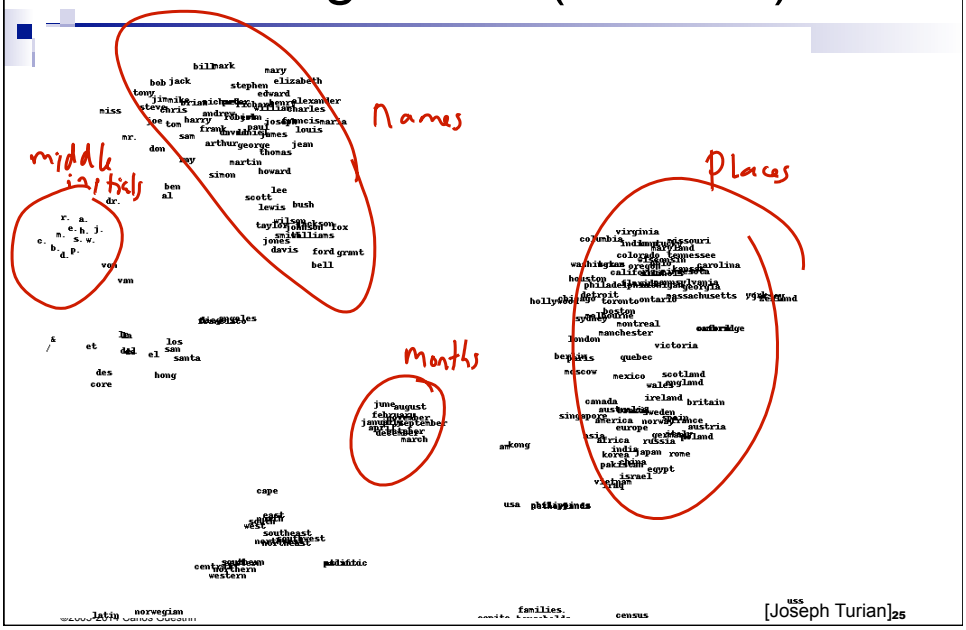

[Joseph Turian]25

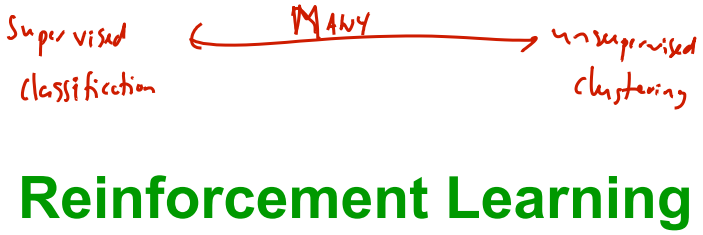

Reinforcement Learning

training by feedback

Learning to act

- Reinforcement learning
- An agent
  - Makes sensor observations
  - Must select action
  - Receives rewards
    - positive for “good” states
    - negative for “bad” states

[Ng et al. ’05]

Bringing it all together…

Combining video, text and audio

HURLEY: Uh ... the Chinese people have water.
(Sayid and Kate go to check it out.)

[EXT. BEACH - CRASH SITE]

(Sayid holds the empty bottle in his hand and questions Sun.)

SAYID: (quietly) Where did you get this?
(He looks at her.)

[EXT. JUNGLE]

(Sawyer is walking through the jungle. He reaches a spot. He kneels down and looks back to check that no one's followed him.

SAYID

SUN

HOLDING

Taskar et al.

Automatically Discovered and Labeled Actions

shout

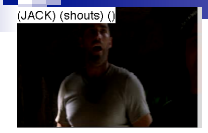

sit down

smile

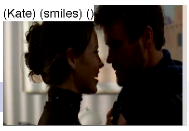

wake

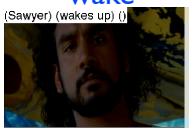

follow

swim

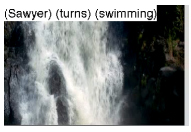

grab

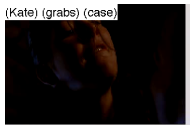

kiss

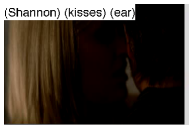

open door

point

Growth of Machine Learning

One of the most sought for specialties in industry today!!!!

- Machine learning is preferred approach to
  - Speech recognition, Natural language processing
  - Computer vision
  - Medical outcomes analysis
  - Robot control
  - Computational biology
  - Sensor networks
  - …
- This trend is accelerating, especially with **Big Data**
  - Improved machine learning algorithms
  - Improved data capture, networking, faster computers
  - Software too complex to write by hand
  - New sensors / IO devices
  - Demand for self-customization to user, environment

Syllabus

- Covers a wide range of Machine Learning techniques – from basic to state-of-the-art
- You will learn about the methods you heard about:
  - Point estimation, regression, naïve Bayes, logistic regression, nearest-neighbor, decision trees, boosting, perceptron, overfitting, regularization, dimensionality reduction, PCA, error bounds, VC dimension, SVMs, kernels, margin bounds, K-means, EM, mixture models, semi-supervised learning, HMMs, graphical models, active learning, reinforcement learning…
- Covers algorithms, theory and applications
- **It's going to be fun and hard work ☺**

Prerequisites

- Formally:
  - STAT 341, STAT 391, or equivalent
- Probabilities
  - Distributions, densities, marginalization…
- Basic statistics
  - Moments, typical distributions, regression…
- Algorithms
  - Dynamic programming, basic data structures, complexity…
- Programming
  - Python will be very useful, but we'll help you get started
- We provide some background, but the class will be fast paced
- Ability to deal with "abstract mathematical concepts"

Recitations & Python

- We'll run an ***optional*** recitations:
  - Time/Location ~~TBD~~ EEB 045 Wednesdays @ 5-6:20pm
- We are recommending Python for homeworks!
  - There are many resources to get started with Python online
  - We'll run an ***optional*** tutorial:
    - First recitation: next week

Staff

- Three Great TAs: Great resource for learning, interact with them!
  - **April Shen**
    Office hours:
  - **Harley Montgomery**
    Office hours:
  - **Tianqi Chen**
    Office hours:
  - Prof: **Carlos Guestrin**
    Office hours: Thursdays 12:30-1:30pm (starting next week)

Communication Channels

- Only channel for announcements, questions, etc. – Catalyst Group:
  - Subscribe!
  - All non-personal questions should go here
  - Answering your question will help others
  - Feel free to chime in
- For e-mailing instructors about personal issues, use:
  - cse546-instructors@cs.washington.edu

Text Books

- Required Textbook:
  - Machine Learning: a Probabilistic Perspective; Kevin Murphy
- Optional Books:
  - Pattern Recognition and Machine Learning; Chris Bishop
  - The Elements of Statistical Learning: Data Mining, Inference, and Prediction; Trevor Hastie, Robert Tibshirani, Jerome Friedman
  - Machine Learning; Tom Mitchell
  - Information Theory, Inference, and Learning Algorithms; David MacKay

Grading

- 4 homeworks (35%)
  - First one goes out next week
    - Start early, Start early
- Final project (30%)
  - Full details out around 10/9
  - Projects done individually, or groups of two students
- Midterm (15%)
  - Thu., 10/30 in class
- Final (20%)
  - By registrar: Monday, December 08 1030-1220, EEB 045

Homeworks

- Homeworks are hard, start early ☺
- Due in the beginning of class
- 33% subtracted per late day
- You have 3 LATE DAYS to use for homeworks only throughout the quarter
  - Please plan accordingly and after that don't be about deadlines, travel,… ☺
- All homeworks **must be handed in**, even for zero credit
- Use Catalyst to submit homeworks

- Collaboration
  - You may **discuss** the questions
  - Each student writes their own answers
  - Write on your homework anyone with whom you collaborate
  - Each student must write their own code for the programming part
  - **Please don't search for answers on the web, Google, previous years' homeworks, etc.**
    - please ask us if you are not sure if you can use a particular reference

Projects

- An opportunity to exercise what you learned and to learn new things
- Individually or groups of two
- Must involve real data
  - Must be data that you have available to you by the time of the project proposals
- Must involve machine learning
- It's encouraged to be related to your research, but must be something new you did this quarter
  - Not a project you worked on during the summer, last year, etc.

- Full details in a couple of weeks

- Tue., October 21 at 10:30am: **Project Proposals**
- Tue., November 11 at 10:30am: **Project Milestone**
- Thu., December 4, 2:30-4:30pm: **Poster Session**
- Fri., December 12 at 10:30am: **Project Report**

Enjoy!

- ML is becoming ubiquitous in science, engineering and beyond
- It's one of the hottest topics in industry today
- This class should give you the basic foundation for applying ML and developing new methods
- The fun begins…

Your first consulting job

- A billionaire from the suburbs of Seattle asks you a question:
  - He says: I have thumbtack, if I flip it, what's the probability it will fall with the nail up?
  - You say: Please flip it a few times:

  - You say: The probability is:
  - **He says: Why???**
  - You say: Because…

Thumbtack – Binomial Distribution

- P(Heads) = θ, P(Tails) = 1-θ

$P(D|\theta) = P(HHTTH|\theta) = \theta\theta(1-\theta)(1-\theta)\theta$

$= \theta^3(1-\theta)^2$

independent identically distributed

IID

- Flips are i.i.d.:
  - Independent events
  - Identically distributed according to Binomial distribution

HHTTH

HHHTT

- Sequence *D* of α_H Heads and α_T Tails

$$P(\mathcal{D} \mid \theta) = \theta^{\alpha_H}(1-\theta)^{\alpha_T}$$

Maximum Likelihood Estimation

ln is a monotonic fn. $\arg\max f(\theta) = \arg\max \ln f(\theta)$

- **Data:** Observed set *D* of α_H Heads and α_T Tails
- **Hypothesis:** Binomial distribution
- Learning θ is an optimization problem
  - What's the objective function?

$\max_\theta P(D|\theta) = \max_\theta \theta^{\alpha_H}(1-\theta)^{\alpha_T}$

- MLE: Choose θ that maximizes the probability of observed data:

$$\hat{\theta}_{MLE} = \arg\max_\theta P(\mathcal{D} \mid \theta)$$

$$= \arg\max_\theta \ln P(\mathcal{D} \mid \theta)$$

$\arg\max f(x) = \arg\min -f(x)$

- loss function

Your first learning algorithm

$$\hat{\theta} = \arg\max_{\theta} \ln P(\mathcal{D} \mid \theta)$$
$$= \arg\max_{\theta} \ln \theta^{\alpha_H}(1-\theta)^{\alpha_T}$$

- Set derivative to zero: $\frac{d}{d\theta} \ln P(\mathcal{D} \mid \theta) = 0$

$\frac{d}{d\theta} \ln \theta^{\alpha_H}(1-\theta)^{\alpha_T} = \frac{d}{d\theta}\left(\alpha_H \ln\theta + \alpha_T \ln(1-\theta)\right) = \alpha_H \frac{d \ln\theta}{d\theta} + \alpha_T \frac{d \ln(1-\theta)}{d\theta}$

$= \frac{\alpha_H}{\theta} - \frac{\alpha_T}{1-\theta} = 0 \Rightarrow \theta = \frac{\alpha_H}{\alpha_H + \alpha_T} \Rightarrow \frac{3}{5}$!!

$\ln ab = b \ln a$

$\ln ab = \ln a + \ln b$

$\frac{d}{d\theta} \ln\theta = \frac{1}{\theta}$

How many flips do I need?

$$\hat{\theta}_{MLE} = \frac{\alpha_H}{\alpha_H + \alpha_T}$$

- Billionaire says: I flipped 3 heads and 2 tails.
- You say: $\hat{\theta}$ = 3/5, I can prove it!
- He says: What if I flipped 30 heads and 20 tails?
- You say: Same answer, I can prove it!
- **He says: What's better?**
- You say: Humm… The more the merrier???
- He says: Is this why I am paying you the big bucks???

Simple bound (based on Hoeffding's inequality)

- For $N = \alpha_H + \alpha_T$, and $\hat{\theta}_{MLE} = \frac{\alpha_H}{\alpha_H + \alpha_T}$

$|\theta^* - \hat{\theta}_{MLE}|$?

- Let θ^* be the true parameter, for any $\varepsilon > 0$:

diff: error — too big

$$P(|\hat{\theta} - \theta^*| \geq \epsilon) \leq 2e^{-2N\epsilon^2}$$

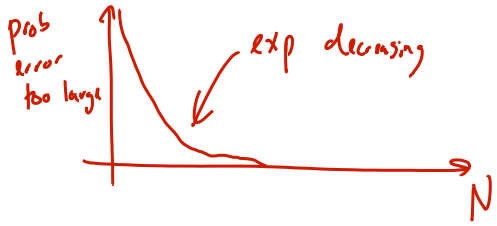

PAC Learning

- PAC: Probably Approximate Correct
- Billionaire says: I want to know the thumbtack parameter θ, within $\varepsilon = 0.1$, with probability at least $1-\delta = 0.95$. How many flips?

prob lose my job

$$P(|\hat{\theta} - \theta^*| \geq \epsilon) \leq 2e^{-2N\epsilon^2} \leq \delta$$

$$\ln \delta \geq \ln 2 - 2N\varepsilon^2$$

$$\Rightarrow N \geq \frac{\ln 2/\delta}{2\varepsilon^2}$$

if $\delta = 0.05$, $\varepsilon = 0.1$

$N \geq 184.4$ flips

loose bound

What about continuous variables?

- Billionaire says: If I am measuring a continuous variable, what can you do for me?
- **You say: Let me tell you about Gaussians…**

$$P(x \mid \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}}$$

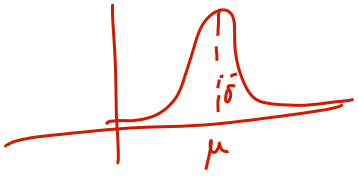

Some properties of Gaussians

- affine transformation (multiplying by scalar and adding a constant)
  - X ~ $N(\mu,\sigma^2)$
  - Y = aX + b ➔ Y ~ $N(a\mu+b,a^2\sigma^2)$
- Sum of Gaussians
  - X ~ $N(\mu_X,\sigma^2_X)$
  - Y ~ $N(\mu_Y,\sigma^2_Y)$
  - Z = X+Y ➔ Z ~ $N(\mu_X+\mu_Y, \sigma^2_X+\sigma^2_Y)$

Learning a Gaussian

- Collect a bunch of data
  - Hopefully, i.i.d. samples
  - e.g., exam scores
- Learn parameters
  - Mean
  - Variance

$$P(x \mid \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}}$$

MLE for Gaussian

- Prob. of i.i.d. samples $D=\{x_1,\ldots,x_N\}$:

$$P(\mathcal{D} \mid \mu, \sigma) \overset{\text{iid}}{=} \left(\frac{1}{\sigma\sqrt{2\pi}}\right)^N \prod_{i=1}^{N} e^{\frac{-(x_i-\mu)^2}{2\sigma^2}}$$

$$\hat{\mu}_{MLE}, \hat{\sigma}^2_{MLE} = \arg\max_{\mu,\sigma} P(D \mid \mu, \sigma) = \arg\max_{\mu,\sigma} \ln P(D \mid \mu, \sigma)$$

- Log-likelihood of data:

$$\ln P(\mathcal{D} \mid \mu, \sigma) = \ln\left[\left(\frac{1}{\sigma\sqrt{2\pi}}\right)^N \prod_{i=1}^{N} e^{\frac{-(x_i-\mu)^2}{2\sigma^2}}\right]$$

$$\max_{\mu,\sigma} = -N \ln \sigma\sqrt{2\pi} - \sum_{i=1}^{N} \frac{(x_i-\mu)^2}{2\sigma^2}$$