

CS 188 SECTION 11

These slides are on Piazza! Search for “Daylen’s slides”

MIDTERM 2 IS OVER!

UPCOMING DEADLINES

- HW 5 due **Friday** @ 11:59
- Project 5 due **Monday 8/1** @ 5pm
- HW 6 due next week Wednesday

WHERE WE ARE

- Part 1: Making Decisions
 - Using search and planning to take actions in environments
 - A* search, Minimax, MDPs, Reinforcement Learning
- Part 2: Probabilistic Reasoning
 - Trying to learn about the world from partial information
 - Bayes Nets, Decision Networks, HMMs
- Part 3: Machine Learning
 - Building classifiers to make predictions based on data
 - Perceptrons, Deep Learning, Neural Networks

GET EXCITED ABOUT MACHINE LEARNING!

- Machine learning does lots of mundane but useful things: email spam classification, envelope address reading, credit card fraud detection...
- And a lot of cool things! Visual question answering, self-driving cars, image captioning, conversational agents


WHAT IS MACHINE LEARNING?

- Mostly it's about building **classifiers**
- You are given boatloads of **data** and their **labels**
- Each data **sample** has many **features**
- You learn a **function** f that takes in data and spits out the label
- Hopefully your function **generalizes** to new data that you haven't seen before

CHOICES FOR THE FUNCTION F


- Decision Trees
- Neural Networks
- Support Vector Machines
- Gaussian Classifiers
- Perceptrons
- Logistic Regression
- Nearest Neighbors
- Naive Bayes

NAIVE BAYES



$$P(Y | F_1 \dots F_N) \propto P(Y, F_1 \dots F_N) = P(Y) \prod_i P(F_i | Y)$$

NAIVE BAYES



$$P(Y | F_1 \dots F_N) \propto P(Y, F_1 \dots F_N) = P(Y) \prod_i P(F_i | Y)$$

$P(Y)$

ham	: 0.66
spam	: 0.33

$P(W|\text{spam})$

the	:	0.0156
to	:	0.0153
and	:	0.0115
of	:	0.0095
you	:	0.0093
a	:	0.0086
with:		0.0080
from:		0.0075
...		

$P(W|\text{ham})$

the	:	0.0210
to	:	0.0133
of	:	0.0119
2002:		0.0110
with:		0.0108
from:		0.0107
and	:	0.0105
a	:	0.0100
...		

DATA HYGIENE AND OVERFITTING

- Usually, we split our data into 3 sets:
 - Training data: Use this to learn f
 - Validation data: Use this to score f
 - Test data: What you actually want to do well on
- We want our function f to **generalize** well to the test data
- If we aren't careful, our function f can **overfit** to the training data (i.e. our training accuracy will be high, but our validation and test accuracy will be low)

OVERFITTING EXAMPLE

- During training, just because we saw lots of spam and no ham that used the word “money” doesn’t mean that’s true at test time
- Solution (for Naive Bayes): Smoothing

$$P_{LAP,k}(x) = \frac{c(x) + k}{N + k|X|}$$

WORKSHEET