

## Admin

## Assignment 6a

$\square$ How'd it go?
$\square$ Which option/extension are you picking?

## Quiz \#3 next Monday

## No hours today

## Machine Learning is...

## Machine Learning is...

Machine learning is programming computers to optimize a performance criterion using example data or past experience.

> -- Ethem Alpaydin

The goal of machine learning is to develop methods that can automatically detect patterns in data, and then to use the uncovered patterns to predict future data or other outcomes of interest.
-- Kevin P. Murphy

The field of pattern recognition is concerned with the automatic discovery of regularities in data through the use of computer algorithms and with the use of these regularities to take actions.
-- Christopher M. Bishop

## Machine Learning is...

Machine learning is about predicting the future based on the past. -- Hal Daume III


## Machine Learning is...

Machine learning is about predicting the future based on the past. -- Hal Daume III







| Regression Example |
| :--- |
| Price of a used car |
| x : car attributes <br> (e.g. mileage) <br> $y$ : price |


| Regression applications |
| :--- |
| How many clicks will a particular website, ad, etc. get? |
| Predict the readability level of a document |
| Predict pause between spoken sentences? |
| Economics/Finance: predict the value of a stock |
| Car/plane navigation: angle of the steering wheel, acceleration, ... |
| $\ldots$ |

Supervised learning: ranking

## NLP Ranking Applications

reranking N -best output lists (e.g. parsing, machine translation, ...)

Rank possible simplification options
flight search (search in general)
...


## Unsupervised learning applications

learn clusters/groups without any label

- cluster documents
- cluster words (synonyms, parts of speech, ...)


## compression

bioinformatics: learn motifs
...

| Reinforcement learning |  |
| :--- | :---: |
| left, right, straight, left, left, left, straight <br> left, straight, straight, left, right, straight, straight | GOoD |
| left, right, straight, left, left, left, straight <br> left, straight, straight, left, right, straight, straight | -38.5 |
| Given a sequence of examples/states and a reward after <br> completing that sequence, learn to predict the action to take in <br> for an individual example/state |  |



| Reinforcement learning example |
| :---: |
|  |
| hiltps：／／www．youtube．com／watch？$=$＝x／MM99xPQC8 |
|  |
|  |


| Text classification |  |  |
| :---: | :---: | :---: |
|  | label |  |
| 三 | spam | For this class，l＇m mostly going to focus on classification |
| 三 | not spam | I＇ll use text classification as a running example |
| 三 | not spam |  |

## Other learning variations

What data is available：
－Supervised，unsupervised，reinforcement learning
■ semi－supervised，active learning，．．．

How are we getting the data： －online vs．offline learning

Type of model：
－generative vs．discriminative －parametric vs．non－parametric
Representing examples


Feature examples
$\square$ ロRaw

| Raw data | Features |
| :---: | :---: |
| 三 | Clinton said banana repeatedly last week on tv， ＂banana，banana，banana＂ |
|  |  |
| 三 | Occurrence of words（unigrams） |

Features
Clinton said banana
repeatedly last week on tv，
＂banana，banana，banana＂
$(1,1,1,0,0,1,0,0, \ldots)$


Occurrence of words（unigrams）


## Lots of other features

POS: occurrence, counts, sequence

Constituents

Whether ' $V$ lagra' occurred 15 times

Whether 'banana' occurred more times than 'apple'

If the document has a number in it
...

Features are very important, but we're going to focus on the model


| Classification | revisited |
| :---: | :---: |
| Training data <br> examples <br> red, round, leaf, 3oz, ... <br> green, round, no leaf, 4oz, ... <br> yellow, curved, no leaf, 4oz, .. <br> green, curved, no leaf, 5oz, ... | Test set <br> label <br> apple <br> apple red, round, no leaf, $40 z, \ldots$ ? <br> banana <br> banana |


| Classification revisited |  |  |
| :---: | :---: | :---: |
| Training data Test set <br> examples  |  |  |
| red, round, leaf, 3oz, ... apple |  |  |
| green, round, no leaf, 40 , $\ldots$ | apple | red, round, no leaf, 4oz, ...? |
| yellow, curved, no leaf, 4oz, ... banana <br> Learning is about generalizing from the training data <br> green, curved, no leaf, $502, \ldots$ banana <br> What does this assume about the training and test set? |  |  |
|  |  |  |



More technically...

We are going to use the probabilistic model of learning

There is some probability distribution over example/label pairs called the data generating distribution

Both the training data and the test set are generated based on this distribution



## Probabilistic models

Probabilistic models define a probability distribution over features and labels:


## Probabilistic models: classification

Probabilistic models define a probability distribution over features and labels:


Given an unlabeled example: yellow, curved, no leaf, boz predict the label
How do we use a probabilistic model for classification/prediction?

## Probabilistic models

Probabilistic models define a probability distribution over features and labels:
yellow, curved, no leaf, 6oz, banana $\longrightarrow$ probabilistic
yellow, curved, no leaf, 6oz, apple

For each label, ask for the probability under the model Pick the label with the highest probability


## Probabilistic models

Probabilities are nice to work with
$\square$ range between 0 and 1
$\square$ can combine them in a well understood way

- lots of mathematical background/theory

Provide a strong, well-founded groundwork

- Allow us to make clear decisions about things like smoothing
- Tend to be much less "heuristic"
- Models have very clear meanings


## Probabilistic models: big questions

1. Which model do we use, i.e. how do we calculate p(feature, label)?
2. How do train the model, i.e. how to we we estimate the probabilities for the model?
3. How do we deal with overfitting (i.e. smoothing)?


## Step 1: picking a model

What we're really trying to do is model the data generating distribution, that is how likely the feature/label combinations are


| Some math |
| :---: |
| p(features,label) |
|  |
|  |
|  |
| $=p\left(x_{1}, x_{2}, \ldots, x_{m}, y\right)$ |
| What rule? |
|  |


| Some math |
| :--- |
| $p($ features, label $)=p\left(x_{1}, x_{2}, \ldots, x_{m}, y\right)$ |
|  |
| $=p(y) p\left(x_{1}, x_{2}, \ldots, x_{m} \mid y\right)$ |
|  |
| $=p(y) p\left(x_{1} \mid y\right) p\left(x_{2}, \ldots, x_{m} \mid y, x_{1}\right)$ |
|  |
| $=p(y) p\left(x_{1} \mid y\right) p\left(x_{2} \mid y, x_{1}\right) p\left(x_{3}, \ldots, x_{m} \mid y, x_{1}, x_{2}\right)$ |
|  |
| $=p(y) \prod_{j=1}^{m} p\left(x_{i} \mid y, x_{1}, \ldots, x_{i-1}\right)$ |

## Step 1: pick a model

$$
p(\text { features,label })=p(y) \prod_{j=1}^{m} p\left(x_{i} \mid y, x_{1}, \ldots, x_{i-1}\right)
$$

So, far we have made NO assumptions about the data

$$
p\left(x_{m} \mid y, x_{1}, x_{2}, \ldots, x_{m-1}\right)
$$

How many entries would the probability distribution table have if we tried to represent all possible values and we had 7000 binary features?

Full distribution tables

| $x_{1}$ | $x_{2}$ | $x_{3}$ | $\ldots$ | $y$ | $p()$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | $\ldots$ | 0 | $*$ |
| 0 | 0 | 0 | $\ldots$ | 1 | $*$ |
| 1 | 0 | 0 | $\ldots$ | 0 | $*$ |
| 1 | 0 | 0 | $\ldots$ | 1 | $*$ |
| 0 | 1 | 0 | $\ldots$ | 0 | $*$ |
| 0 | 1 | 0 | $\ldots$ | 1 | $*$ |

All possible combination of features!

Table size: $2^{7000}=$ ?


| Full distribution tables |
| :---: |
|  |
| Storing a table of that size is impossible! <br> How are we supposed to learn/estimate each entry in the table? |

## Step 1: pick a model

$$
p(\text { features,label })=p(y) \prod_{j=1}^{m} p\left(x_{i} \mid y, x_{1}, \ldots, x_{i-1}\right)
$$

Naïve Bayes assumption
$p($ features, label $)=p(y) \prod_{j=1}^{m} p\left(x_{i} \mid y, x_{1}, \ldots, x_{i-1}\right)$

So, far we have made NO assumptions about the data
Model selection involves making assumptions about the data
We've done this before, n -gram language model, parsing, etc.
These assumptions allow us to represent the data more compactly and to estimate the parameters of the model
Naive Bayes assumption
$\quad p($ features, label $)=p(y) \prod_{j=1}^{m} p\left(x_{i} \mid y, x_{1}, \ldots, x_{i-1}\right)$

$$
p\left(x_{i} \mid y, x_{1}, x_{2}, \ldots, x_{i-1}\right)=p\left(x_{i} \mid y\right)
$$

Assumes feature i is independent of the the other
features given the label
Is this true for text, say, with unigram features?

