## Recap: Basic Text Analysis

- Represent text in terms of "features" (e.g., how often each word/phrase appears, whether it's a named entity, etc)
- Can repeat this for different documents: represent each document as a "feature vector"
"Sentence":



In general (not just text): first represent data as feature vectors

## Example: Representing an Image



## Example: Representing an Image



## Example: Representing an Image



## Example: Representing an Image



Image source: starwars.com

## Back to Text

Unigram bag of words model is already quite powerful:

- Enough to learn topics (each text doc: raw word counts without stopwords)
- Enough to learn a simple detector for email spam

These are HW2 problems

## Finding Possibly Related Entities

## f) $\because$ in $\triangle$



The solar batteries have reportedly been spotted in San Juan's airport.
By John Patrick Pullen Octoher 16, 2017
暗

Exactly one week after Tesla $¢ E O$ Elon Musk sugested his company could help with Puerto Rico's electricity crisis in the aftermath of Hurricane Maria, more of the company's Powerwall battery packs have arrived on the island, according to a photo snapped at San Juan airport Friday, Oct. 13.

Source: http://fortune.com/2017/10/16/elon-musks-tesla-powerwalls-have-landed-in-puerto-rico/

## Co-Occurrences

For example: count \# news articles that have different named entities co-occur


What does it mean for a named entity to co-occur with itself?
Example: could count \# articles in which word appears $\geq 2$ times

## Different Ways to Count

- Just saw: for all doc's, count \# of doc's in which two named entities co-occur
- This approach ignores \# of co-occurrences within a specific document (e.g., if 1 doc has "Elon Musk" and "Tesla" appear 10 times, we count this as 1)
- Could instead add \# co-occurrences, not just whether it happened in a doc
- Instead of looking at \# doc's, look at co-occurrences within a sentence, or a paragraph, etc


## Bottom Line

- There are many ways to count co-occurrences
- You should think about what makes the most sense/is reasonable for the problem you're looking at


# We aim to find interesting relationships by looking at co-occurrences 

## Black and white frequently co-occur, but is this relationship interesting?



How l'm counting: For each pixel, look at neighboring 4 pixels and compare their values (1 of "green green", "green white", "green black", "white white", "white black", "black black")


|  | Green | White | Black |
| :---: | :---: | :---: | :---: |
| Green | 1000 | 200 | 200 |
| White |  | 2000 | 350 |
| Black |  |  | 2000 |

Probability of drawing "White, Black"? 350/5750

Probability of drawing a card that has "White" on it?

$$
(200+2000+350) / 5750
$$

1000 of these cards:
Green, Green

200 of these cards:
Green, White

200 of these cards:
Green, Black
$P($ Green, White $)=\frac{200}{5750}$

$$
P(\text { Green, Black })=\frac{200}{5750}
$$

$P($ White, Black $)=\frac{350}{5750}$

$$
P(\text { Green })=\frac{1400}{5750}
$$

$$
P(\text { White })=\frac{2550}{5750}
$$

$$
P(\text { Black })=\frac{2550}{5750}
$$

## Measuring Association: Pointwise Mutual Information (PMI)

$$
\mathrm{PMI}(\mathrm{~A}, \mathrm{~B})=\log _{2} \frac{\mathrm{P}(\mathrm{~A}, \mathrm{~B})}{\mathrm{P}(\mathrm{~A}) \mathrm{P}(\mathrm{~B})}
$$

Base of log doesn't really matter (we'll use base 2)

PMI can be positive or negative

Higher PMI $\rightarrow$ more "interesting"

$$
\begin{aligned}
& \text { PMI(Green, White })=\log _{2} \frac{200 / 5750}{(1400 / 5750)(2550 / 5750)}=1.63 \ldots \text { bits } \\
& \mathrm{PMI}(\text { Green, Black })=\log _{2} \frac{200 / 5750}{(1400 / 5750)(2550 / 5750)}=1 .-1.63 \ldots \text { bits } \\
& \mathrm{PMI}(\text { White, Black })=\log _{2} \frac{350 / 5750}{(2550 / 5750)(2550 / 5750)}=-1.69 \ldots \text { bits }
\end{aligned}
$$

$P($ Green, White $)=\frac{200}{5750}$

$$
P(\text { Green, Black })=\frac{200}{5750}
$$

$$
\mathrm{P}(\text { White, Black })=\frac{350}{5750}
$$

$$
P(\text { White })=\frac{2550}{5750}
$$

$$
P(\text { Black })=\frac{2550}{5750}
$$

## What is PMI Measuring?

## Probability of $A$ and $B$ co-occurring

Probability of just A occurring
Probability of just B occurring
If A and B were "independent"
$\rightarrow$ probability of $A$ and $B$ co-occurring would be $P(A) P(B)$

## What is PMI Measuring?

## Probability of $A$ and $B$ co-occurring


if equal to 1
$\rightarrow A$, B are indep.

Probability of $A$ and $B$ co-occurring if they were independent
PMI measures (the log of) a ratio that says how far $A$ and $B$ are from being independent

There are lots of connections of information
theory to prediction
Rough intuition:
Something surprising $\leftrightarrow$ less predictable $\leftrightarrow$ more bits to store

## Looking at All Pairs of Outcomes

- $P M I$ measures how $P(A, B)$ differs from $P(A) P(B)$ using a log ratio
- Log ratio isn't the only way to compare!
- Another way to compare:

$$
\left.\begin{array}{l}
\frac{[P(A, B)-P(A) P(B)]^{2}}{P(A) P(B)} \\
\text { Phi-square }=\sum_{A, B} \frac{[P(A, B)-P(A) P(B)]^{2}}{P(A) P(B)} \\
\text { Chi-square }=N \times \text { Phi-square }
\end{array}\right] \begin{gathered}
\text { between } 0 \text { and } 1 \\
0 \rightarrow \text { pairs are all } \\
\text { indep. }
\end{gathered}
$$

Phi-square is
$\mathrm{N}=$ sum of all co-occurrence counts (in upper right of triangle earlier)

## Example: Phi-Square Calculation



Chi-square $=\mathrm{N} \times$ Phi-square
$\mathrm{N}=$ sum of all co-occurrence counts (in upper right of triangle earlier)

## Example: Phi-Square Calculation



Phi-square $=\sum \frac{[P(A, B)-P(A) P(B)]^{2}}{P(A) P(B)}$

Add these up to get: Phi-square $=0.6470 .$.

A, B Interpretation: neighboring pixels not close to being indep.

## Back to Earlier Example



Often we know what kind of named entities are found
Example: Elon Musk and Tim Cook are people, Tesla and Apple are companies
$\rightarrow$ can ask what people are related to what companies

## Back to Earlier Example

|  | Tesla | Apple |
| :---: | :---: | :---: |
| Elon Musk | 300 | 1 |
| Tim Cook | 1 | 195 |

PMI, phi-square, chi-square calculations done same way as before
Main things to calculate first:

P(Elon Musk, Tesla)
P(Elon Musk, Apple)
P(Tim Cook, Tesla)
P(Tim Cook, Apple)

P(Elon Musk)
P(Tim Cook)
P(Tesla)
P(Apple)

The math here is actually a bit easier to think about because the rows and columns aren't indexing the same items

## Back to Earlier Example



## Recap: Co-Occurrences

- Joint probability $\mathrm{P}(\mathrm{A}, \mathrm{B})$ can be poor indicator of whether A and B co-occurring is "interesting"
- Find interesting relationships between pairs of items by looking at PMI
- Intuition: "Interesting" co-occurring events should occur more frequently than if they were to co-occur independently
- In practice: some times it is helpful to generalize PMI and look instead at

$$
\begin{gathered}
\mathrm{PMI}_{\rho}(\mathrm{A}, \mathrm{~B})=\log _{2} \frac{\mathrm{P}(\mathrm{~A}, \mathrm{~B})^{\rho}}{\mathrm{P}(\mathrm{~A}) \mathrm{P}(\mathrm{~B})} \quad \begin{array}{c}
\text { Tune parameter } \\
\rho>0
\end{array} \\
\text { (we'll talk about parameter tuning later in the course) }
\end{gathered}
$$

## Example Application of PMI: Image Segmentation



Phillip Isola, Daniel Zoran, Dilip Krishnan, and Edward H. Adelson. Crisp boundary detection using pointwise mutual information. ECCV 2014.

## Example Application of PMI: Word Embeddings



Image source: https://deeplearning4j.org/img/countries_capitals.png

## Continuous Measurements

- So far, looked at relationships between discrete outcomes
- For pair of continuous outcomes, use a scatter plot
Computing Improvements: Transistors Per Circuit


Image source: https://plot.Iy/~MattSundquist/5405.png

## The Importance of Staring at Data




## Correlation



Beware: Just because two variables appear correlated doesn't mean that one can predict the other

## Correlation $\neq$ Causation



# Important: At this point in the course, we are finding possible relationships between two entities 

We are not yet making statements about prediction (we'll see prediction later in the course)

We are not making statements about causality (beyond the scope of this course)

## Causality



Studies in 1960's: Coffee drinkers have higher rates of lung cancer
Can we claim that coffee is a cause of lung cancer?
Back then: coffee drinkers also tended to smoke more than non-coffee drinkers (smoking is a confounding variable)
To establish causality, groups getting different treatments need to appear similar so that the only difference is the treatment

## Establishing Causality

If you control data collection


Compare outcomes of two groups
Randomized controlled trial (RCT) also called A/B testing

Example: figure out webpage layout to maximize revenue (Amazon)
Example: figure out how to present educational material to improve learning (Khan Academy)

If you do not control data collection
In general: not obvious establishing what caused what

## Course Outline

Part 1: Identify structure present in "unstructured" data
Exploratory data analysis

- Frequency and co-occurrences Basic probability \& statistics
- Clustering

Unsupervised learning

- Topic modeling (special kind of clustering)

Part 2: Make predictions using structure found in part 1
Predictive data analysis
Supervised learning

- Introduction to classification
- Adaptive nearest neighbor methods
- Deep learning models for classification

