

 uitana




 regular numbertest eq aquation
 npy lexical GO mandidx
 osservition
bumblition

## LSA 311

 Computational Lexical SemanticsDan Jurafsky<br>Stanford University

Lecture 2: Word Sense Disambiguation

## Word Sense Disambiguation (WSD)

- Given
- A word in context
- A fixed inventory of potential word senses
- Decide which sense of the word this is
- Why? Machine translation, QA, speech synthesis
- What set of senses?
- English-to-Spanish MT: set of Spanish translations
- Speech Synthesis: homographslike bass and bow
- In general: the senses in a thesaurus like WordNet


## Two variants of WSD task

- Lexical Sample task
- Small pre-selected set of target words (line, plant)
- And inventory of senses for each word
- Supervised machine learning: train a classifier for each word
- All-words task
- Every word in an entire text
- A lexicon with senses for each word
- Data sparseness: can't train word-specific classifiers


## WSD Methods

- Supervised Machine Learning
- Thesaurus/Dictionary Methods
- Semi-Supervised Learning
 Word mintaisistic sense discurlse case parsing discourse lane 童 를 mandulue nell mintan

 uitare





 corpus anild smintis nguar numbertert equation


$\underset{\substack{\text { observetion } \\ \text { bimplition }}}{ }$


## Word Sense Disambiguation

## Supervised Machine Learning

## Supervised Machine Learning Approaches

- Supervised machine learning approach:
- a training corpus of words tagged in context with their sense
- used to train a classifier that can tag words in new text
- Summary of what we need:
- the tag set ("sense inventory")
- the training corpus
- A set of features extracted from the training corpus
- A classifier


## Supervised WSD 1: WSD Tags

- What's a tag?


## A dictionary sense?

- For example, for WordNet an instance of "bass" in a text has 8 possible tags or labels (bass1 through bass8).


## 8 senses of "bass" in WordNet

1. bass-(the lowest part of the musical range)
2. bass, bass part - (the lowest part in polyphonic music)
3. bass, basso-(an adult male singer with the lowest voice)
4. sea bass, bass - (flesh of lean-fleshed saltwater fish of the family Serranidae)
5. freshwater bass, bass - (any of various North American lean-fleshed freshwater fishes especially of the genus Micropterus)
6. bass, bass voice, basso-(the lowest adult male singing voice)
7. bass - (the member with the lowest range of a family of musical instruments)
8. bass-(nontechnical name for any of numerous edible marine and freshwater spiny-finned fishes)

## Inventory of sense tags for bass

| WordNet <br> Sense | Spanish <br> Translation | Roget <br> Category | Target Word in Context |
| :--- | :--- | :--- | :--- |
| bass $^{4}$ | lubina | FISH/INSECT | ...fish as Pacific salmon and striped bass and... |
| bass $^{4}$ | lubina | FISH/INSECT | ... produce filets of smoked bass or sturgeon... |
| bass $^{7}$ | bajo | MUSIC | ...exciting jazz bass player since Ray Brown... |
| bass $^{7}$ | bajo | MUSIC | $\ldots$ play bass because he doesn't have to solo... |

## Supervised WSD 2: Get a corpus

- Lexical sample task:
- Line-hard-serve corpus - 4000 examples of each
- Interest corpus - 2369 sense-tagged examples
- All words:
- Semantic concordance: a corpus in which each open-class word is labeled with a sense from a specific dictionary/thesaurus.
- SemCor: 234,000 words from Brown Corpus, manually tagged with WordNet senses
- SENSEVAL-3 competition corpora - 2081 tagged word tokens


## SemCor

<wf pos=PRP>He</wf>
<wf pos=VB lemma=recognize wnsn=4 lexsn=2:31:00::>recognized</wf> <wf pos=DT>the</wf>
<wf pos=NN lemma=gesture wnsn=1 lexsn=1:04:00::>gesture</wf> <punc>.</punc>

## Supervised WSD 3: Extract feature vectors Intuition from Warren Weaver (1955):

"If one examines the words in a book, one at a time as through an opaque mask with a hole in it one word wide, then it is obviously impossible to determine, one at a time, the meaning of the words...

But if one lengthens the slit in the opaque mask, until one can see not only the central word in question but also say N words on either side, then if N is large enough one can unambiguously decide the meaning of the central word...
The practical question is : "What minimum value of N will, at least in a tolerable fraction of cases, lead to the correct choice of meaning for the central word?"

## Feature vectors

- A simple representation for each observation
(each instance of a target word)
- Vectors of sets of feature/value pairs
- Represented as a ordered list of values
- These vectors represent, e.g., the window of words around the target


## Two kinds of features in the vectors

- Collocational features and bag-of-words features
- Collocational
- Features about words at specific positions near target word
- Often limited to just word identity and POS
- Bag-of-words
- Features about words that occur anywhere in the window (regardless of position)
- Typically limited to frequency counts


## Examples

- Example text (WSJ):

An electric guitar and bass player stand off to one side not really part of the scene

- Assume a window of $+/-2$ from the target


## Examples

- Example text (WSJ)

An electric guitar and bass player stand off to one side not really part of the scene,

- Assume a window of $+/-2$ from the target


## Collocational features

- Position-specific information about the words and collocations in window
- guitar and bass player|stand

$$
\left[w_{i-2}, \operatorname{POS}_{i-2}, w_{i-1}, \mathrm{POS}_{i-1}, w_{i+1}, \mathrm{POS}_{i+1}, w_{i+2}, \operatorname{POS}_{i+2}, w_{i-2}^{i-1}, w_{i}^{i+1}\right]
$$

[guitar, NN, and, CC, player, NN, stand, VB, and guitar, player stand]

- word 1,2,3 grams in window of $\pm 3$ is common


## Bag-of-words features

- "an unordered set of words" - position ignored
- Counts of words occur within the window.
- First choose a vocabulary
- Then count how often each of those terms occurs in a given window
- sometimes just a binary "indicator" 1 or 0


## Co-Occurrence Example

- Assume we've settled on a possible vocabulary of 12 words in "bass" sentences:
[fishing, big, sound, player, fly, rod, pound, double, runs, playing, guitar, band]
- The vector for:
guitar and bass player stand [0,0,0,1,0,0,0,0,0,0,1,0]

 corpus misist anduda mustemmin nimian er equation

皆
observation
lemprs
rataion


## Word Sense Disambiguation

## Classification

## Classification: definition

- Input:
- a word w and some features $f$
- a fixed set of classes $C=\left\{c_{1}, c_{2}, \ldots, c_{J}\right\}$
- Output: a predicted class $c \in C$


## Classification Methods: Supervised Machine Learning

- Input:
- a word $w$ in a text window d (which we'll call a "document")
- a fixed set of classes $C=\left\{c_{1}, c_{2}, \ldots, c_{\}}\right\}$
- A training set of $m$ hand-labeled text windows again called "documents" $\left(d_{1}, c_{1}\right), \ldots .,\left(d_{m}, c_{m}\right)$
- Output:
- a learned classifier $\gamma: d \rightarrow c$


## Classification Methods: Supervised Machine Learning

- Any kind of classifier
- Naive Bayes
- Logistic regression
- Neural Networks
- Support-vector machines
- k-Nearest Neighbors
 $\underset{\substack{\text { consider } \\ \text { cmide }}}{\text { chind }}$ smintar gular numbertest equationg



| observition |
| :--- |
| bumbition |

## Word Sense Disambiguation

## The Naive Bayes

Classifier

## Naive Bayes Intuition

- Simple ("naive") classification method based on Bayes rule
- Relies on very simple representation of document
- Bag of words


## I'll introduce classification with an even simpler supervised learning task

Let's classify a movie review as positive (+) or negative (-)

Suppose we have lots of reviews labeled as (+) or (-) and I give you a new review.

Given: the words in this new movie review
Return: one of 2 classes: + or -

## The Bag of Words Representation

I love this movie! It's sweet, but with satirical humor. The dialogue is great and the adventure scenes are fun... It manages to be whimsical and romantic while laughing at the conventions of the fairy tale genre. I would recommend it to just about anyone. I've seen it several times, and I'm always happy to see it again whenever I have a friend who hasn't seen it yet!

to 3
and 3
seen 2
yet $\quad 1$
would 1
whimsical
times 1
sweet 1
satirical 1
adventure 1
genre $\quad 1$
fairy 1
humor 1
have 1
great 1

## The bag of words representation

| seen | 2 |
| :--- | :--- |
| sweet | 1 |
| whimsical | 1 |
| recommend | 1 |
| happy | 1 |
| $\ldots$ | $\ldots$ |

Bayes' Rule Applied to Documents and Classes

- For a document $d$ and a class $C$

$$
P(c \mid d)=\frac{P(d \mid c) P(c)}{P(d)}
$$

## Naive Bayes Classifier (I)

$$
\begin{aligned}
C_{M A P} & =\underset{c \in C}{\operatorname{argmax}} P(c \mid d) \quad \begin{array}{l}
\text { MAP is "maximum a } \\
\text { posteriori" = most } \\
\text { likely class }
\end{array} \\
& =\underset{c \in C}{\operatorname{argmax}} \frac{P(d \mid c) P(c)}{P(d)} \quad \begin{array}{|l}
\text { Bayes Rule } \\
\\
\end{array}{\underset{c}{\text { Brec }}}^{\operatorname{argmax}} P(d \mid c) P(c) \quad \begin{array}{l}
\text { Dropping the } \\
\text { denominator }
\end{array}
\end{aligned}
$$

## Naive Bayes Classifier (II)

$$
\begin{aligned}
c_{M A P} & =\underset{c \in C}{\operatorname{argmax}} P(d \mid c) P(c) \\
& =\underset{c \in C}{\operatorname{argmax}} P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c\right) P(c)
\end{aligned}
$$

## Naive Bayes Classifier (IV)

$$
c_{M A P}=\underset{c \in C}{\operatorname{argmax}} P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c\right) P(c)
$$

## $\mathrm{O}\left(|X|^{\bullet} \bullet|C|\right)$ parameters

How often does this class occur?

Could only be estimated if a very, very large number of training examples was

We can just count the relative frequencies in a corpus

## Multinomial Naive Bayes Independence Assumptions

$$
P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c\right)
$$

- Bag of Words assumption: Assume position doesn't matter
- Conditional Independence: Assume the feature probabilities $P\left(x_{i} \mid c_{j}\right)$ are independent given the class $c$.

$$
P\left(x_{1}, \ldots, x_{n} \mid c\right)=P\left(x_{1} \mid c\right) \bullet P\left(x_{2} \mid c\right) \bullet P\left(x_{3} \mid c\right) \bullet \ldots \bullet P\left(x_{n} \mid c\right)
$$

## Multinomial Naive Bayes Classifier

$$
\begin{aligned}
c_{M A P} & =\underset{c \in C}{\operatorname{argmax}} P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c\right) P(c) \\
c_{N B} & =\underset{c \in C}{\operatorname{argmax}} P\left(c_{j}\right) \prod_{x \in X} P(x \mid c)
\end{aligned}
$$

## Applying Multinomial Naive Bayes Classifiers to Text Classification

positions $=$ all word positions in test document

$$
c_{N B}=\underset{c_{\mathrm{j}} \in C}{\operatorname{argmax}} P\left(c_{j}\right) \prod_{i \in \text { positions }} P\left(x_{i} \mid c_{j}\right)
$$

# Classification 

Macration
number



Hon


M 1 N
unax anguage sisten
smantios

 hype 롤:

MODES $\Longrightarrow$ noun
observetion

## Learning the Naive Bayes Classifier

## Learning the Multinomial Naive Bayes Model

- First attempt: maximum likelihood estimates
- simply use the frequencies in the data

$$
\begin{array}{r}
\hat{P}\left(c_{j}\right)=\frac{\operatorname{doccount}\left(C=c_{j}\right)}{N_{d o c}} \\
\hat{P}\left(w_{i} \mid c_{j}\right)=\frac{\operatorname{count}\left(w_{i}, c_{j}\right)}{\sum_{w \in V} \operatorname{count}\left(w, c_{j}\right)}
\end{array}
$$

## Parameter estimation

$$
\hat{P}\left(w_{i} \mid c_{j}\right)=\frac{\operatorname{count}\left(w_{i}, c_{j}\right)}{\sum_{w \in V} \operatorname{count}\left(w, c_{j}\right)}
$$

fraction of times word $w_{i}$ appears among all words in documents of topic $c_{j}$

- Create mega-document for topicj by concatenating all docs in this topic
- Use frequency of $w$ in mega-document


## Problem with Maximum Likelihood

- What if we have seen no training documents with the word fantastic and classified in the topic positive (thumbs-up)?

$$
\hat{P}(\text { "fantastic" } \mid \text { positive })=\frac{\operatorname{count}(\text { "fantastic", positive })}{\sum_{w \in V} \operatorname{count}(w, \text { positive })}=0
$$

- Zero probabilities cannot be conditioned away, no matter the other evidence!

$$
c_{M A P}=\operatorname{argmax}_{c} \hat{P}(c) \prod_{i} \hat{P}\left(x_{i} \mid c\right)
$$

## Laplace (add-1) smoothing for Naïve Bayes

$$
\begin{aligned}
\hat{P}\left(w_{i} \mid c\right) & =\frac{\operatorname{count}\left(w_{i}, c\right)+1}{\left.\sum_{w \in V}(\operatorname{count}(w, c))+1\right)} \\
& =\frac{\operatorname{count}\left(w_{i}, c\right)+1}{\left(\sum_{w \in V} \operatorname{count}(w, c)\right)+|V|}
\end{aligned}
$$

## Multinomial Naïve Bayes: Learning

- From training corpus, extract Vocabulary
- Calculate $P\left(c_{j}\right)$ terms
- For each $c_{j}$ in $C$ do docs $_{j} \leftarrow$ all docs with class $=c_{j}$

$$
P\left(c_{j}\right) \leftarrow \frac{\left|\operatorname{docs}_{j}\right|}{\mid \text { total \# documents|}}
$$

- Calculate $P\left(w_{k} \mid c_{j}\right)$ terms
- Text $_{j} \leftarrow$ single doc containing all docs $_{j}$
- Foreach word $w_{k}$ in Vocabulary
$n_{k} \leftarrow \#$ of occurrences of $w_{k}$ in Text ${ }_{j}$

$$
P\left(w_{k} \mid c_{j}\right) \leftarrow \frac{n_{k}+\alpha}{n+\alpha \mid \text { Vocabulary } \mid}
$$


 mintaisistic smanse discuurse case parsing discourse lane 童 를 mant valu 景 mintan等intap problem ©
 utara input 眔（1）
 featur＇e a human representation mangin aras semantic

 anysaniquage，shsenm minta



observerioin lump relation

## Word Sense Disambiguation

Learning the Naive Bayes Classifier



 = $=$ anguage sistem -s equation




## Word Sense Disambiguation

## Multinomial Naive

 Bayes: A Worked Example for WSD
## Applying Naive Bayes to WSD

- $P(c)$ is the prior probability of that sense
- Counting in a labeled training set.
- $P(w \mid c)$ conditional probability of a word given a particular sense
- $P(w \mid c)=$ count $(w, c) / c o u n t(c)$
- We get both of these from a tagged corpus like SemCor
- Can also generalize to look at other features besides words.
- Then it would be P(f|c)
- Conditional probability of a feature given a sense


Priors:

$$
\begin{aligned}
& P(f)=\frac{3}{4} \frac{1}{4} \\
& P(g)=
\end{aligned}
$$

|  | Doc | Words | Class |
| :--- | :--- | :--- | :--- |
| Training | 1 | fish smoked fish | f |
|  | 2 | fish line | f |
|  | 3 | fish haul smoked | f |
| Test 4 <br> guitar jazz line g <br>  5 line guitar jazz jazz | ? |  |  |

$$
V=\{\text { fish, smoked, line, haul, guitar, jazz }\}
$$

## Choosing a class:

$$
P(f \mid d 5) \quad 3 / 4 * 2 / 14 *(1 / 14)^{2} * 1 / 14
$$

$$
\approx 0.00003
$$

Conditional Probabilities:

$$
P(\text { line } \mid f)=\quad(1+1) /(8+6)=2 / 14
$$

$$
\mathrm{P}(\text { guitar } \mid f)=(0+1) /(8+6)=1 / 14
$$

$$
\mathrm{P}(\mathrm{jazz} \mid f)=(0+1) /(8+6)=1 / 14
$$

$$
\begin{array}{rl}
1 / 4 & * 2 / 9 *(2 / 9)^{2} * 2 / 9 \\
& \approx 0.0006
\end{array}
$$

$$
P(\text { line } \mid g)=\quad(1+1) /(3+6)=2 / 9
$$

$$
\mathrm{P}(\text { guitar } \mid g)=(1+1) /(3+6)=2 / 9
$$

$$
\mathrm{P}(j a z z \mid g)=(1+1) /(3+6)=2 / 9
$$

## Evaluations and Baselines

## WSD Evaluations and baselines

- Best evaluation: extrinsic ('end-to-end', 'task-based') evaluation
- Embed WSD algorithm in a task and see if you can do the task better!
- What we often do for convenience: intrinsic evaluation
- Exact match sense accuracy
- \% of words tagged identically with the human-manual sense tags
- Usually evaluate using held-out data from same labeled corpus
- Baselines
- Most frequent sense
- The Lesk algorithm


## Most Frequent Sense

- WordNet senses are ordered in frequency order
- So "most frequent sense" in WordNet = "take the first sense"
- Sense frequencies come from the SemCor corpus

| Freq | Synset | Gloss |
| :--- | :--- | :--- |
| 338 | plant $^{1}$, works, industrial plant | buildings for carrying on industrial labor |
| 207 | plant $^{2}$, flora, plant life | a living organism lacking the power of locomotion |
| 2 | plant $^{3}$ | something planted secretly for discovery by another <br> 0 |
| plant $^{4}$ | an actor situated in the audience whose acting is rehearsed but <br> seems spontaneous to the audience |  |

## Ceiling

- Human inter-annotator agreement
- Compare annotations of two humans
- On same data
- Given same tagging guidelines
- Human agreements on all-words corpora with WordNet style senses
- 75\%-80\%
 Men roum craig rive
 2 parisesincewix







eximition yor equation



models $\Rightarrow$ noun

observetion | bundration |
| :---: |

## Word Sense Disambiguation

## Dictionary and Thesaurus Methods

## The Simplified Lesk algorithm

- Let's disambiguate "bank" in this sentence:

The bank can guarantee deposits will eventually cover future tuition costs because it invests in adjustable-rate mortgage securities.

- given the following two WordNet senses:

| bank $^{1}$ | Gloss: | a financial institution that accepts deposits and channels the <br> money into lending activities <br> "he cashed a check at the bank", "that bank holds the mortgage <br> on my home" |
| :--- | :--- | :--- |
| bank $^{2}$ | Gloss: <br> Examples: | sloping land (especially the slope beside a body of water) <br> "they pulled the canoe up on the bank", "he sat on the bank of <br> the river and watched the currents" |

## The Simplified Lesk algorithm

Choose sense with most word overlap between gloss and context (not counting function words)
The bank can guarantee deposits/will eventually cover future tuition costs because it invests in adjustable-rate mortgage securities.

| bank $^{1}$ | Gloss: | a financial institution that accepts deposits and channels the <br> money into lending activities <br> "he cashed a check at the bank", "that bank holds the mortgage <br> on my home" |
| :--- | :--- | :--- |
| bank $^{2}$ | Gloss: <br> Examples: | sloping land (especially the slope beside a body of water) <br> "they pulled the canoe up on the bank", "he sat on the bank of <br> the river and watched the currents" |

## The Corpus Lesk algorithm

- Assumes we have some sense-labeled data (like SemCor)
- Take all the sentences with the relevant word sense: These short, "streamlined" meetings usually are sponsored by local banks¹, Chambers of Commerce, trade associations, or other civic organizations.
- Now add these to the gloss + examples for each sense, call it the "signature" of a sense.
- Choose sense with most word overlap between context and signature.


## Corpus Lesk: IDF weighting

- Instead of just removing function words
- Weigh each word by its `promiscuity' across documents
- Down-weights words that occur in every `document' (gloss, example, etc)
- These are generally function words, but is a more fine-grained measure
- Weigh each overlapping word by inverse document frequency


## Corpus Lesk: IDF weighting

- Weigh each overlapping word by inverse document frequency
- $N$ is the total number of documents
- $\mathrm{df}_{\mathrm{i}}=$ "document frequency of word $i$ "
- = \# of documents with word I

$$
\operatorname{idf}_{i}=\log \left(\frac{N}{d f_{i}}\right)
$$

$$
\begin{aligned}
& \operatorname{score}\left(\text { sense }_{i}, \text { context }_{j}\right)= \\
& \sum \\
& \mathrm{idf}_{w} \\
& w \in \operatorname{overlap}\left(\text { signature }_{i}, \text { context }_{j}\right)
\end{aligned}
$$

## Graph-based methods

- First, WordNet can be viewed as a graph
- senses are nodes
- relations (hypernymy, meronymy) are edges
- Also add edge between word and unambiguous gloss words



## How to use the graph for WSD

- Insert target word and words in its sentential context into the graph, with directed edges to their senses
"She drank some milk"
- Now choose the most central sense Add some probability to "drink" and "milk" and compute node with bighest "pagerank"
 nex (iscused |liser Word mintaisistic smanse discuurse case parsing discourse lane 童 를 mant vilue

毫intitad 2 parseSincewan

 vent feature a Jun|rin das humen representation $\square$ tapin $\rightarrow$ semantic 1/1)

 requar nubloprtert equation


models noun
observetion


## Word Sense Disambiguation

## Semi-Supervised Learning

## Semi-Supervised Learning

Problem: supervised and dictionary-based approaches require large hand-built resources

What if you don't have so much training data?

## Solution: Bootstrapping

Generalize from a very small hand-labeled seed-set.

## Bootstrapping

- For bass
- Rely on "One sense per collocation" rule
- A word reoccurring in collocation with the same word will almost surely have the same sense.
- the word play occurs with the music sense of bass
- the word fish occurs with the fish sense of bass


## Sentences extracting using "fish" and "play"

> We need more good teachers - right now, there are only a half a dozen who can play the free bass with ease.
> An electric guitar and bass player stand off to one side, not really part of the scene, just as a sort of nod to gringo expectations perhaps.
> The researchers said the worms spend part of their life cycle in such fish as Pacific salmon and striped bass and Pacific rockfish or snapper.
> And it all started when fishermen decided the striped bass in Lake Mead were too skinny.

## Summary: generating seeds

1) Hand labeling
2) "One sense per collocation":

- A word reoccurring in collocation with the same word will almost surely have the same sense.

3) "One sense per discourse":

- The sense of a word is highly consistent within a document - Yarowsky (1995)
- (At least for non-function words, and especially topic-specific words)


## Stages in the Yarowsky bootstrapping algorithm for the word "plant"


(a)


## Summary

- Word Sense Disambiguation: choosing correct sense in context
- Applications: MT, QA, etc.
- Three classes of Methods
- Supervised Machine Learning: Naive Bayes classifier
- Thesaurus/Dictionary Methods
- Semi-Supervised Learning
- Main intuition
- There is lots of information in a word's context
- Simple algorithms based just on word counts can be surprisingly good

