CS447: Natural Language Processing
http://courses.engr.illinois.edu/cs447

## Lecture 4: <br> Smoothing

## Julia Hockenmaier

juliahmr@illinois.edu
3324 Siebel Center

## Last lecture's key concepts

## Basic probability review:

joint probability, conditional probability

## Probability models

Independence assumptions
Parameter estimation: relative frequency estimation (aka maximum likelihood estimation)

Language models
N -gram language models:
unigram, bigram, trigram...

## Quick note re. notation

Consider the sentence W = "John loves Mary"
For a trigram model we could write:

$$
P\left(\mathrm{w}_{3}=\text { Mary } \mid \mathrm{w}_{1} \mathrm{w}_{2}=\text { "John loves" }\right)
$$

This notation implies that we treat the preceding bigram $\mathrm{w}_{1} \mathrm{w}_{2}$ as one single conditioning variable $P(\mathrm{X} \mid \mathrm{Y})$

Instead, we typically write:
$P\left(\mathrm{w}_{3}=\right.$ Mary I $\mathrm{w}_{2}=$ loves, $\mathrm{w}_{1}=$ John $)$
Although this is less readable (John loves $\rightarrow$ loves, John), this notation gives us more flexibility, since it implies that we treat the preceding bigram $\mathrm{w}_{1} \mathrm{w}_{2}$ as two conditioning variables $P(\mathrm{X} \mid \mathrm{Y}, \mathrm{Z})$

## Parameter estimation (training)

Parameters: the actual probabilities (numbers)

$$
P\left(w_{i}=' \text { the' } \mid w_{i-1}=' o n '\right)=0.0123
$$

We need (a large amount of) text as training data to estimate the parameters of a language model.

The most basic estimation technique: relative frequency estimation (= counts)

$$
P\left(w_{i}=' \text { the }{ }^{\prime} \mid w_{i-l}=' o n '\right)=C\left(\text { 'on the }^{\prime}\right) / C\left(\text { 'on' }^{\prime}\right)
$$

This assigns all probability mass to events in the training corpus.
Also called Maximum Likelihood Estimation (MLE)

## Zipf's law: the long tail

How nany words occur once, twice, 100 times, 1000 times?


In natural language:

- A small number of events (e.g. words) occur with high frequency
- A large number of events occur with very low frequency


## Testing: unseen events will occur

Recall the Shakespeare example:
Only 30,000 word types occurred.
Any word that does not occur in the training data has zero probability!

Only 0.04\% of all possible bigrams occurred. Any bigram that does not occur in the training data has zero probability!

So....
... we can't actually evaluate our MLE models on unseen test data (or system output)...
... because both are likely to contain words/n-grams that these models assign zero probability to.

We need language models that assign some probability mass to unseen words and n -grams.

## Today's lecture

How can we design language models* that can deal with previously unseen events?
*actually, probabilistic models in general


## What unseen events may occur?

## Simple distributions:

$$
P(X=x)
$$

(e.g. unigram models)

## Possibility:

The outcome $x$ has not occurred during training (i.e. is unknown):
-We need to reserve mass in $P(X)$ for $x$

## Questions:

-What outcomes $x$ are possible?
-How much mass should they get?

## Dealing with unseen events

Relative frequency estimation assigns all probability mass to events in the training corpus

But we need to reserve some probability mass to events that don't occur in the training data
Unseen events = new words, new bigrams
Important questions:
What possible events are there?
How much probability mass should they get?

## What unseen events may occur?

Simple conditional distributions:

$$
P(X=x \mid Y=y)
$$

(e.g. bigram models)

Case 1: The outcome $x$ has been seen, but not in the context of $Y=y$ :
-We need to reserve mass in $P(X \mid Y=y)$ for $X=x$
Case 2: The conditioning variable $y$ has not been seen:
-We have no $P(X \mid Y=y)$ distribution.

- We need to drop the conditioning variable $Y=y$ and use $P(X)$ instead.


## What unseen events may occur?

## Complex conditional distributions

(with multiple conditioning variables)

$$
P(X=x \mid Y=y, Z=z)
$$

(e.g. trigram models)

Case 1: The outcome $X=x$ was seen, but not in the context of ( $Y=y, Z=z$ ):
-We need to reserve mass in $P(X \mid Y=y, Z=z)$
Case 2: The joint conditioning event ( $Y=y, Z=z$ ) hasn't been seen:

- We have no $P(X \mid Y=y, Z=z)$ distribution.
- But we can drop $z$ and use $P(X \mid Y=y)$ instead.


## Smoothing: Reserving mass in $P(X)$ for unseen events

## Examples

Training data: The wolf is an endangered species
Test data: The wallaby is endangered

| Unigram | Bigram | Trigram |
| :--- | :--- | :--- |
| $\mathrm{P}($ the $)$ | $\mathrm{P}($ the $\mid<\mathrm{s}>)$ | $\mathrm{P}($ the $\mid<\mathrm{s}>)$ |
| $\times \mathrm{P}($ wallaby $)$ | $\times \mathrm{P}($ wallaby $\mid$ the $)$ | $\times \mathrm{P}($ wallaby $\mid$ the,$<\mathrm{s}>)$ |
| $\times \mathrm{P}($ is $)$ | $\times \mathrm{P}($ is $\mid$ wallaby $)$ | $\times \mathrm{P}($ is $\mid$ wallaby, the $)$ |
| $\times \mathrm{P}($ endangered $)$ | $\times \mathrm{P}$ (endangered $\mid$ is $)$ | $\times \mathrm{P}($ endangered $\mid$ is, wallaby $)$ |

-Case 1: P (wallaby), P (wallaby | the), P ( wallaby | the, $\langle\mathrm{s}>$ ):
What is the probability of an unknown word (in any context)?
-Case 2: $P$ (endangered | is)
What is the probability of a known word in a known context,
if that word hasn't been seen in that context?
-Case 3: P (is $\mid$ wallaby) P (is | wallaby, the) P (endangered | is, wallaby):
What is the probability of a known word in an unseen context?

## Dealing with unknown words: <br> The simple solution

## Training:

- Assume a fixed vocabulary
(e.g. all words that occur at least twice (or $n$ times) in the corpus)
-Replace all other words by a token <UNK>
-Estimate the model on this corpus.
Testing:
- Replace all unknown words by <UNK>
-Run the model.
This requires a large training corpus to work well.


## Dealing with unknown events

Use a different estimation technique:

- Add-1(Laplace) Smoothing
-Good-Turing Discounting
Idea: Replace MLE estimate $P(w)=\frac{C(w)}{N}$
Combine a complex model with a simpler model:
- Linear Interpolation
-Modified Kneser-Ney smoothing Idea: use bigram probabilities of $w_{i} \quad P\left(w_{i} \mid w_{i-1}\right)$ to calculate trigram probabilities of $w_{i} \quad P\left(w_{i} \mid w_{i-n} \ldots w_{i-1}\right)$


## Bigram counts

Original:

Smoothed:

|  | i | want | to | eat | chinese | food | lunch | spend |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| i | 5 | 827 | 0 | 9 | 0 | 0 | 0 | 2 |
| want | 2 | 0 | 608 | 1 | 6 | 6 | 5 | 1 |
| to | 2 | 0 | 4 | 686 | 2 | 0 | 6 | 211 |
| eat | 0 | 0 | 2 | 0 | 16 | 2 | 42 | 0 |
| chinese | 1 | 0 | 0 | 0 | 0 | 82 | 1 | 0 |
| food | 15 | 0 | 15 | 0 | 1 | 4 | 0 | 0 |
| lunch | 2 | 0 | 0 | 0 | 0 | 1 | 0 | 0 |
| spend | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 0 |
|  | i | want | to | eat | chinese | food | lunch | spend |
| i | 6 | 828 | 1 | 10 | 1 | 1 | 1 | 3 |
| want | 3 | 1 | 609 | 2 | 7 | 7 | 6 | 2 |
| to | 3 | 1 | 5 | 687 | 3 | 1 | 7 | 212 |
| eat | 1 | 1 | 3 | 1 | 17 | 3 | 43 | 1 |
| chinese | 2 | 1 | 1 | 1 | 1 | 83 | 2 | 1 |
| food | 16 | 1 | 16 | 1 | 2 | 5 | 1 | 1 |
| lunch | 3 | 1 | 1 | 1 | 1 | 2 | 1 | 1 |
| spend | 2 | 1 | 2 | 1 | 1 | 1 | 1 | 1 |

## Add-1 (Laplace) smoothing

Assume every (seen or unseen) event occurred once more than it did in the training data.

## Example: unigram probabilities

Estimated from a corpus with N tokens and a vocabulary (number of word types) of size V.

$$
\left.\begin{array}{rl}
\operatorname{MLE} & P\left(w_{i}\right)
\end{array}\right) \frac{C\left(w_{i}\right)}{\sum_{j} C\left(w_{j}\right)}=\frac{C\left(w_{i}\right)}{N}, ~=\frac{C\left(w_{i}\right)+\mathbf{1}}{\sum_{j}\left(C\left(w_{j}\right)+\mathbf{1}\right)}=\frac{C\left(w_{i}\right)+\mathbf{1}}{N+\mathbf{V}} .
$$

Bigram probabilities
Original:

|  | i | want | to | eat | chinese | food | lunch | spend |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| i | 0.002 | 0.33 | 0 | 0.0036 | 0 | 0 | 0 | 0.00079 |
| want | 0.0022 | 0 | 0.66 | 0.0011 | 0.0065 | 0.0065 | 0.0054 | 0.0011 |
| to | 0.00083 | 0 | 0.0017 | 0.28 | 0.00083 | 0 | 0.0025 | 0.087 |
| eat | 0 | 0 | 0.0027 | 0 | 0.021 | 0.0027 | 0.056 | 0 |
| chinese | 0.0063 | 0 | 0 | 0 | 0 | 0.52 | 0.0063 | 0 |
| food | 0.014 | 0 | 0.014 | 0 | 0.00092 | 0.0037 | 0 | 0 |
| lunch | 0.0059 | 0 | 0 | 0 | 0 | 0.0029 | 0 | 0 |
| spend | 0.0036 | 0 | 0.0036 | 0 | 0 | 0 | 0 | 0 |

Smoothed:

|  | i | want | to | eat | chinese | food | lunch | spend |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| i | 0.0015 | 0.21 | 0.00025 | 0.0025 | 0.00025 | 0.00025 | 0.00025 | 0.00075 |
| want | 0.0013 | 0.00042 | 0.26 | 0.00084 | 0.0029 | 0.0029 | 0.0025 | 0.00084 |
| to | 0.00078 | 0.00026 | 0.0013 | 0.18 | 0.00078 | 0.00026 | 0.0018 | 0.055 |
| eat | 0.00046 | 0.00046 | 0.0014 | 0.00046 | 0.0078 | 0.0014 | 0.02 | 0.00046 |
| chinese | 0.0012 | 0.00062 | 0.00062 | 0.00062 | 0.00062 | 0.052 | 0.0012 | 0.00062 |
| food | 0.0063 | 0.00039 | 0.0063 | 0.00039 | 0.00079 | 0.002 | 0.00039 | 0.00039 |
| lunch | 0.0017 | 0.00056 | 0.00056 | 0.00056 | 0.00056 | 0.0011 | 0.00056 | 0.00056 |
| spend | 0.0012 | 0.00058 | 0.0012 | 0.00058 | 0.00058 | 0.00058 | 0.00058 | 0.00058 |

## Problem:

Add-one moves too much probability mass from seen to unseen events!

## Reconstituting the counts

We can "reconstitute" pseudo-counts $c^{*}$ for our training set of size N from our estimate:

Unigrams: $\quad c_{i}^{*}=P\left(w_{i}\right) \cdot N$
$P\left(w_{i}\right)$ : probability that the next word is $w_{i}$

$$
=\frac{C\left(w_{i}\right)+1}{N+V} \cdot N<\begin{aligned}
& \text { Plug in the model definition of } P\left(w_{i}\right) \\
& V: \text { size of vocabulary }
\end{aligned}
$$

$$
=\left(C\left(w_{i}\right)+1\right) \cdot \frac{N}{N+V} \leqslant \begin{aligned}
& \text { Rearrange } \\
& \text { (to see dependence on } N \text { and } V)
\end{aligned}
$$

Bigrams: $\quad c^{*}\left(w_{i} \mid w_{i-1}\right)=P\left(w_{i} \mid w_{i-1}\right) \cdot C\left(w_{i-1}\right)$

Plug in the model definition of $P\left(w_{i} \mid w_{i-1}\right)$
CS447: Natural Language Processing (J. Hockenmaier)

## Reconstituted Bigram counts

Original:

|  | i | want | to | eat | chinese | food | lunch | spend |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| i | 5 | 827 | 0 | 9 | 0 | 0 | 0 | 2 |
| want | 2 | 0 | 608 | 1 | 6 | 6 | 5 | 1 |
| to | 2 | 0 | 4 | 686 | 2 | 0 | 6 | 211 |
| eat | 0 | 0 | 2 | 0 | 16 | 2 | 42 | 0 |
| chinese | 1 | 0 | 0 | 0 | 0 | 82 | 1 | 0 |
| food | 15 | 0 | 15 | 0 | 1 | 4 | 0 | 0 |
| lunch | 2 | 0 | 0 | 0 | 0 | 1 | 0 | 0 |
| spend | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 0 |
|  | i | want | to | eat | chinese | food | lunch | spend |
| i | 3.8 | 527 | 0.64 | 6.4 | 0.64 | 0.64 | 0.64 | 1.9 |
| want | 1.2 | 0.39 | 238 | 0.78 | 2.7 | 2.7 | 2.3 | 0.78 |
| to | 1.9 | 0.63 | 3.1 | 430 | 1.9 | 0.63 | 4.4 | 133 |
| eat | 0.34 | 0.34 | 1 | 0.34 | 5.8 | 1 | 15 | 0.34 |
| chinese | 0.2 | 0.098 | 0.098 | 0.098 | 0.098 | 8.2 | 0.2 | 0.098 |
| food | 6.9 | 0.43 | 6.9 | 0.43 | 0.86 | 2.2 | 0.43 | 0.43 |
| lunch | 0.57 | 0.19 | 0.19 | 0.19 | 0.19 | 0.38 | 0.19 | 0.19 |
| spend | 0.32 | 0.16 | 0.32 | 0.16 | 0.16 | 0.16 | 0.16 | 0.16 |

## Add-K smoothing

Variant of Add-One smoothing:
For any $k>0$ (typically, $k<1$ )
Add K $\quad P\left(w_{i}\right)=\frac{C\left(w_{i}\right)+k}{N+k V}$
This is still too simplistic to work well.

## Good-Turing smoothing

Basic idea: Use total frequency of events that occur only once to estimate how much mass to shift to unseen events

- "occur only once" (in training data): frequency $f=1$
- "unseen" (in training data): frequency $f=0$ (didn't occur)


Relative Frequency Estimate
Good Turing Estimate

## Good-Turing smoothing


$N_{c}$ : number of event types that occur $c$ times (can be counted) $N_{l}$ : number of event types that occur once
$N=1 N_{l}+\ldots+m N_{m}$ : total number of observed event tokens

CS447: Natural Language Processing (J. Hockenmaier)

## Good-Turing smoothing

The Maximum Likelihood estimate of the probability of a word $w$ that occurs $k-1$ times $P_{M L E}(w)=\mathrm{C}(w) / \mathrm{N}$

$$
P_{M L E}(w)=\frac{c_{k-1}}{N}=\frac{k-1}{N}
$$

The Good-Turing estimate of the probability of a word $w$ that occurs $k-1$ times: $P_{\mathrm{GT}}(w)=\mathrm{c}^{*}{ }_{k-1} / \mathrm{N}$ :

$$
P_{G T}(w)=\frac{c_{k-1}^{*}}{N}=\frac{\left(\frac{k \cdot N_{k}}{N_{k-1}}\right)}{N}=\frac{k \cdot N_{k}}{N \cdot N_{k-1}}
$$

## Problems with Good-Turing

## Problem 1:

What happens to the most frequent event?

## Problem 2:

We don't observe events for every $k$.

## Variant: Simple Good-Turing

Replace $\mathrm{N}_{\mathrm{n}}$ with a fitted function $f(n)$ :

$$
f(n)=a+b \log (n)
$$

Requires parameter tuning (on held-out data):
Set $a, b$ so that $f(n) \cong N_{n}$ for known values.
Use $\mathrm{c}_{n}$ * only for small $n$

## Linear Interpolation (1)

We don't see "Bob was reading", but we see " _ was reading".
We estimate $P($ reading $\mid$ 'Bob was') $=0$ but $P($ reading $\mid$ 'was') $>0$
Use ( $n-1$ )-gram probabilities to smooth $n$-gram probabilities:


## Smoothing: Reserving mass in $P(X \mid Y)$ for unseen events

## What happens to $P(\mathrm{w} \mid \ldots)$ ?

The smoothed probability $P_{\text {smoothed-trigram }}\left(w_{i} \mid w_{\mathrm{i}-2} w_{\mathrm{i}-1}\right)$ is a linear combination of $P_{\text {unsmoothed-trigram }}\left(w_{i} \mid w_{\mathrm{i}-2} w_{\mathrm{i}-1}\right)$ and $P_{\text {bigram }}\left(w_{i} \mid w_{\mathrm{i}-1}\right)$ :


## Linear Interpolation (2)

We've never seen "Bob was reading",
but we might have seen " $\qquad$ was reading", and we've certainly seen "__ reading" (or <UNK>)

$$
\begin{aligned}
\tilde{P}\left(w_{i} \mid w_{i-1}, w_{i-2}\right)= & \lambda_{3} \cdot \hat{P}\left(w_{i} \mid w_{i-1}, w_{i-2}\right) \\
& +\lambda_{2} \cdot \hat{P}\left(w_{i} \mid w_{i-1}\right) \\
& +\lambda_{1} \cdot \hat{P}\left(w_{i}\right) \\
& \text { for } \lambda_{1}+\lambda_{2}+\lambda_{3}=1
\end{aligned}
$$

$P_{\text {smoothed }}\left(\mathrm{w}_{\mathrm{i}}=\right.$ reading $\mid \mathrm{w}_{\mathrm{i}-1}=$ was, $\mathrm{w}_{\mathrm{i}-2}=$ Bob $)=$
$\lambda_{3} P_{\text {unsmoothed-trigram }}\left(\mathrm{w}_{\mathrm{i}}=\right.$ reading $\mid \mathrm{w}_{\mathrm{i}-1}=$ was, $\mathrm{w}_{\mathrm{i}-2}=$ Bob $)$
$+\lambda_{2} P_{\text {unsmoothed-bigran }}\left(\mathrm{w}_{\mathrm{i}}=\right.$ reading $\mid \mathrm{w}_{\mathrm{i}-1}=$ was $)$
$+\lambda_{1} P_{\text {unsmoothed-unigram }}\left(\mathrm{w}_{\mathrm{i}}=\right.$ reading $)$
CS447: Natural Language Processing (J. Hockenmaier)

## Absolute discounting

Subtract a constant factor $D<1$ from each nonzero $n$-gram count, and interpolate with $P_{A D}\left(w_{\mathrm{i}} \mid w_{\mathrm{i}-1}\right)$ :
non-zero if trigram $w_{\mathrm{i}-2} w_{\mathrm{i}-1} w_{\mathrm{i}}$ is seen

$$
\begin{aligned}
P_{A D}\left(w_{i} \mid w_{i-1}, w_{i-2}\right)=\frac{\frac{\max \left(C\left(w_{i-2} w_{i-1} w_{i}\right)-D, 0\right)}{C\left(w_{i-2} w_{i-1}\right)}}{+(1-\lambda) P_{A D}\left(w_{i} \mid w_{i-1}\right)}
\end{aligned}
$$

If $S$ seen word types occur after $\mathrm{w}_{\mathrm{i}-2} \mathrm{w}_{\mathrm{i}-1}$ in the training data, this reserves the probability mass $\mathrm{P}(U)=(S \times D) / \mathrm{C}\left(w_{i-2} w_{i-1}\right)$ to be computed according to $P\left(w_{\mathrm{i}} \mid w_{\mathrm{i}-1}\right)$. Set:

$$
(1-\lambda)=P(U)=\frac{S \cdot D}{C\left(w_{i-2} w_{i-1}\right)}
$$

N.B.: with $\mathrm{N}_{1}, \mathrm{~N}_{2}$ the number of $n$-grams that occur once or twice, $\mathrm{D}=\mathrm{N}_{1} /\left(\mathrm{N}_{1}+2 \mathrm{~N}_{2}\right)$ works well in practice

CS447: Natural Language Processing (J. Hockenmaier)

## Interpolation: Setting the $\lambda s$

## Method A: Held-out estimation

Divide data into training and held-out data.
Estimate models on training data.
Use held-out data (and some optimization technique) to find the $\lambda$ that gives best model performance.
Often: $\lambda$ is a learned function of the frequencies of
$w_{\mathrm{i}-\mathrm{n}} \ldots w_{\mathrm{i}-1}$

## Method B:

$\lambda$ is some (deterministic) function of the frequencies of $w_{\mathrm{i}-\mathrm{n}} . . w_{\mathrm{i}-1}$

## Kneser-Ney smoothing

Observation: "San Francisco" is frequent, but "Francisco" only occurs after "San".

Solution: the unigram probability $P(w)$ should not depend on the frequency of $w$, but on the number of contexts in which $w$ appears

$$
\begin{aligned}
N_{+l}(\bullet w): & \text { number of contexts in which } w \text { appears } \\
& =\text { number of word types } w^{\prime} \text { which precede } w \\
N_{+l}(\bullet \bullet) & =\sum_{w^{\prime}} N_{+l}\left(\bullet w^{\prime}\right)
\end{aligned}
$$

Kneser-Ney smoothing: Use absolute discounting, but use $P(w)=N_{+l}(\bullet w) / N_{+l}$

Modified Kneser-Ney smoothing: Use different $D$ for bigrams and trigrams (Chen \& Goodman '98)

## To recap....

## Today's key concepts

Dealing with unknown words
Dealing with unseen events
Good-Turing smoothing
Linear Interpolation
Absolute Discounting
Kneser-Ney smoothing

Today's reading:
Jurafsky and Martin, Chapter 4, sections 1-4

