n-gram Language Models

CS 288 Spring 2026

UC Berkeley

Lecture plan

- What is an n-gram language model?

- Generating from a language model

- Evaluating a language model (perplexity)

- Smoothing: additive, interpolation, discounting

What is a language model?

- A probabilistic model of a sequence of words

- Joint probability distribution of words $w_1, w_2, \ldots, w_n$:

$$P(w_1, w_2, w_3, \ldots, w_n)$$

Sample space =
finite pieces of English text

"It was the best of times, it was the worst of times"

"Hey 'sup."

"green sentences, loaded with vitamins"

"911 how can I help you"

"Call me Ishmael"

How likely is a given phrase, sentence, paragraph or even a document?
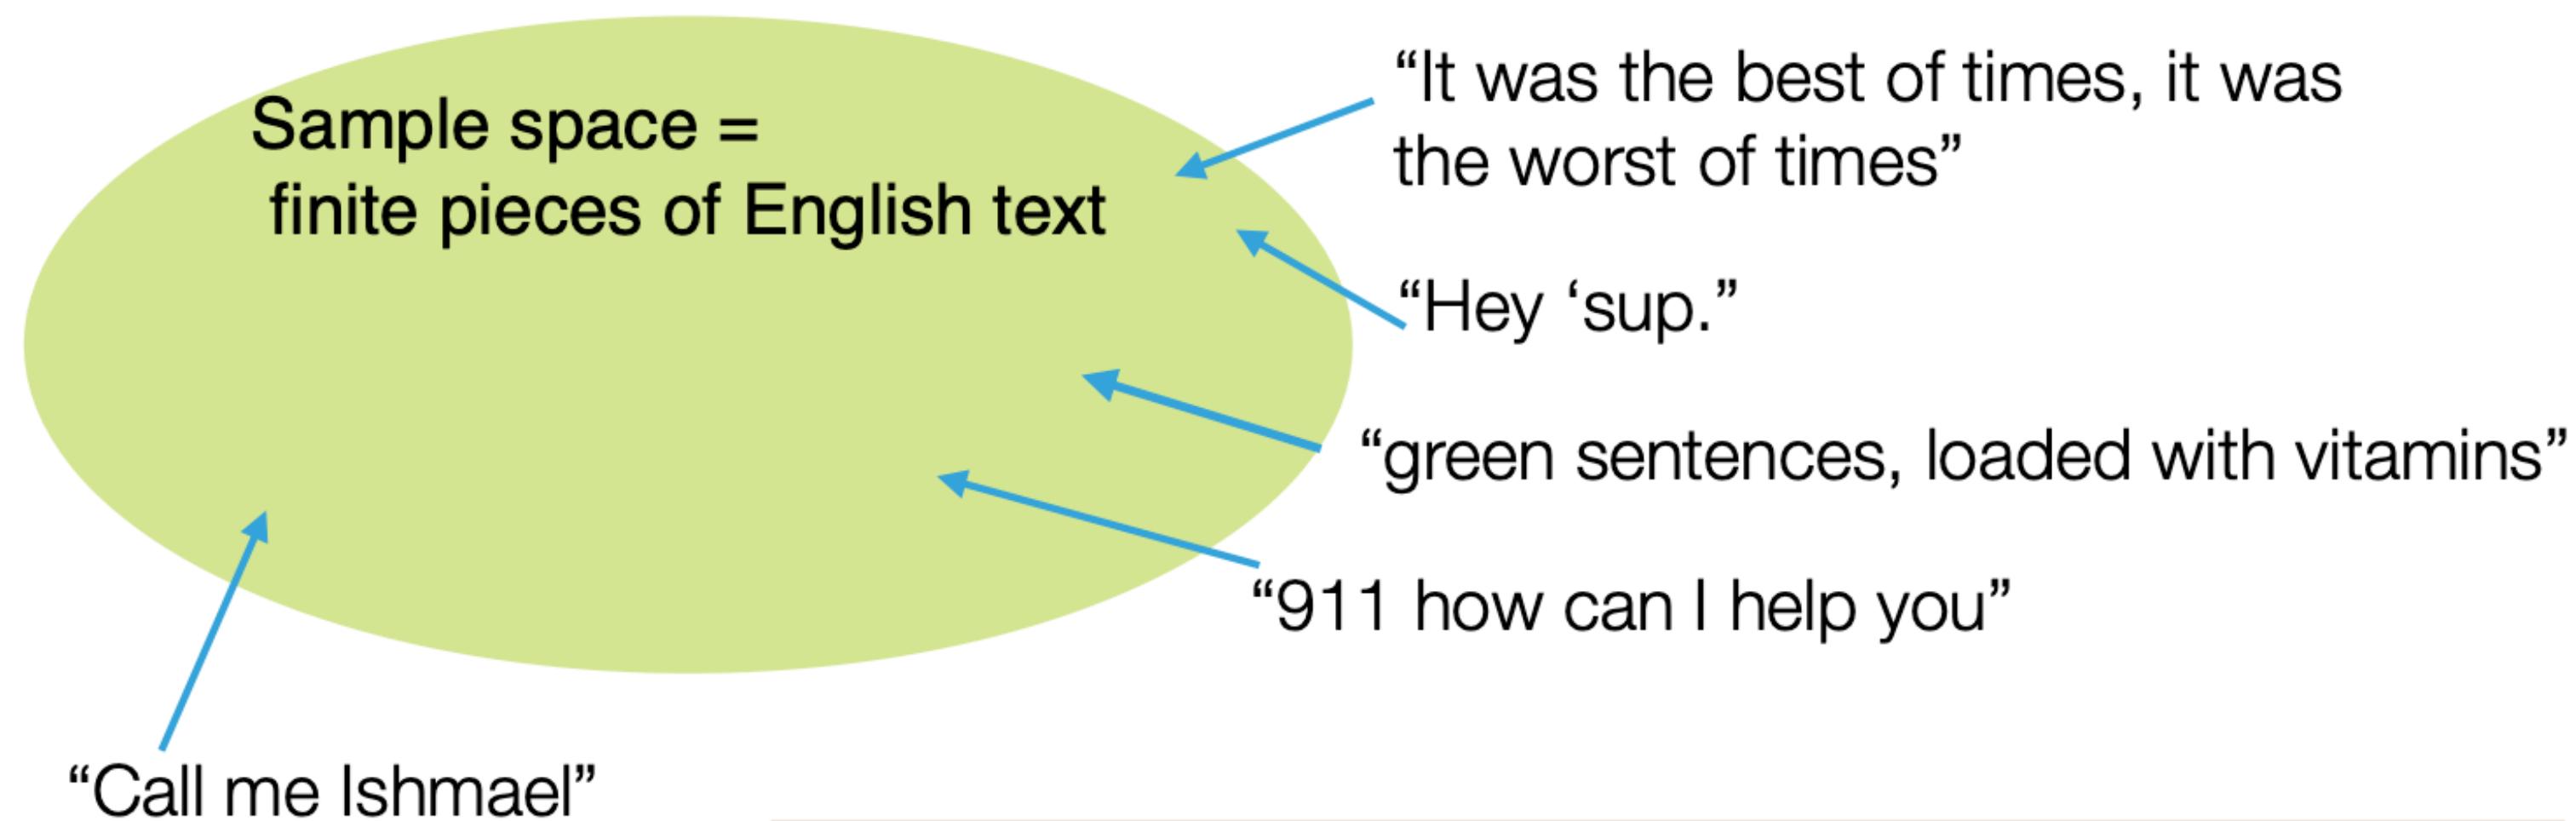
(i.e., $\mathrm{Pr}[w_1 w_2 w_2 \ldots w_n]$ associated with every finite word sequence $w_1 w_2 w_2 \ldots w_n$ (including nonsensical ones)

Chain rule

Conditional probability:
$$p(w \mid w_1, w_2), \forall w \in V$$

$$p(w_1, w_2, w_3, \ldots, w_n) =$$

$$p(w_1)p(w_2 \mid w_1)p(w_3 \mid w_1, w_2) \times \cdots \times p(w_n \mid w_1, w_2, \ldots, w_{n-1})$$

Sentence: "the cat sat on the mat"

$$P(\text{the cat sat on the mat}) = P(\text{the}) * P(\text{cat}|\text{the}) * P(\text{sat}|\text{the cat})$$
$$* P(\text{on}|\text{the cat sat}) * P(\text{the}|\text{the cat sat on})$$
$$* P(\text{mat}|\text{the cat sat on the})$$

Language models are everywhere
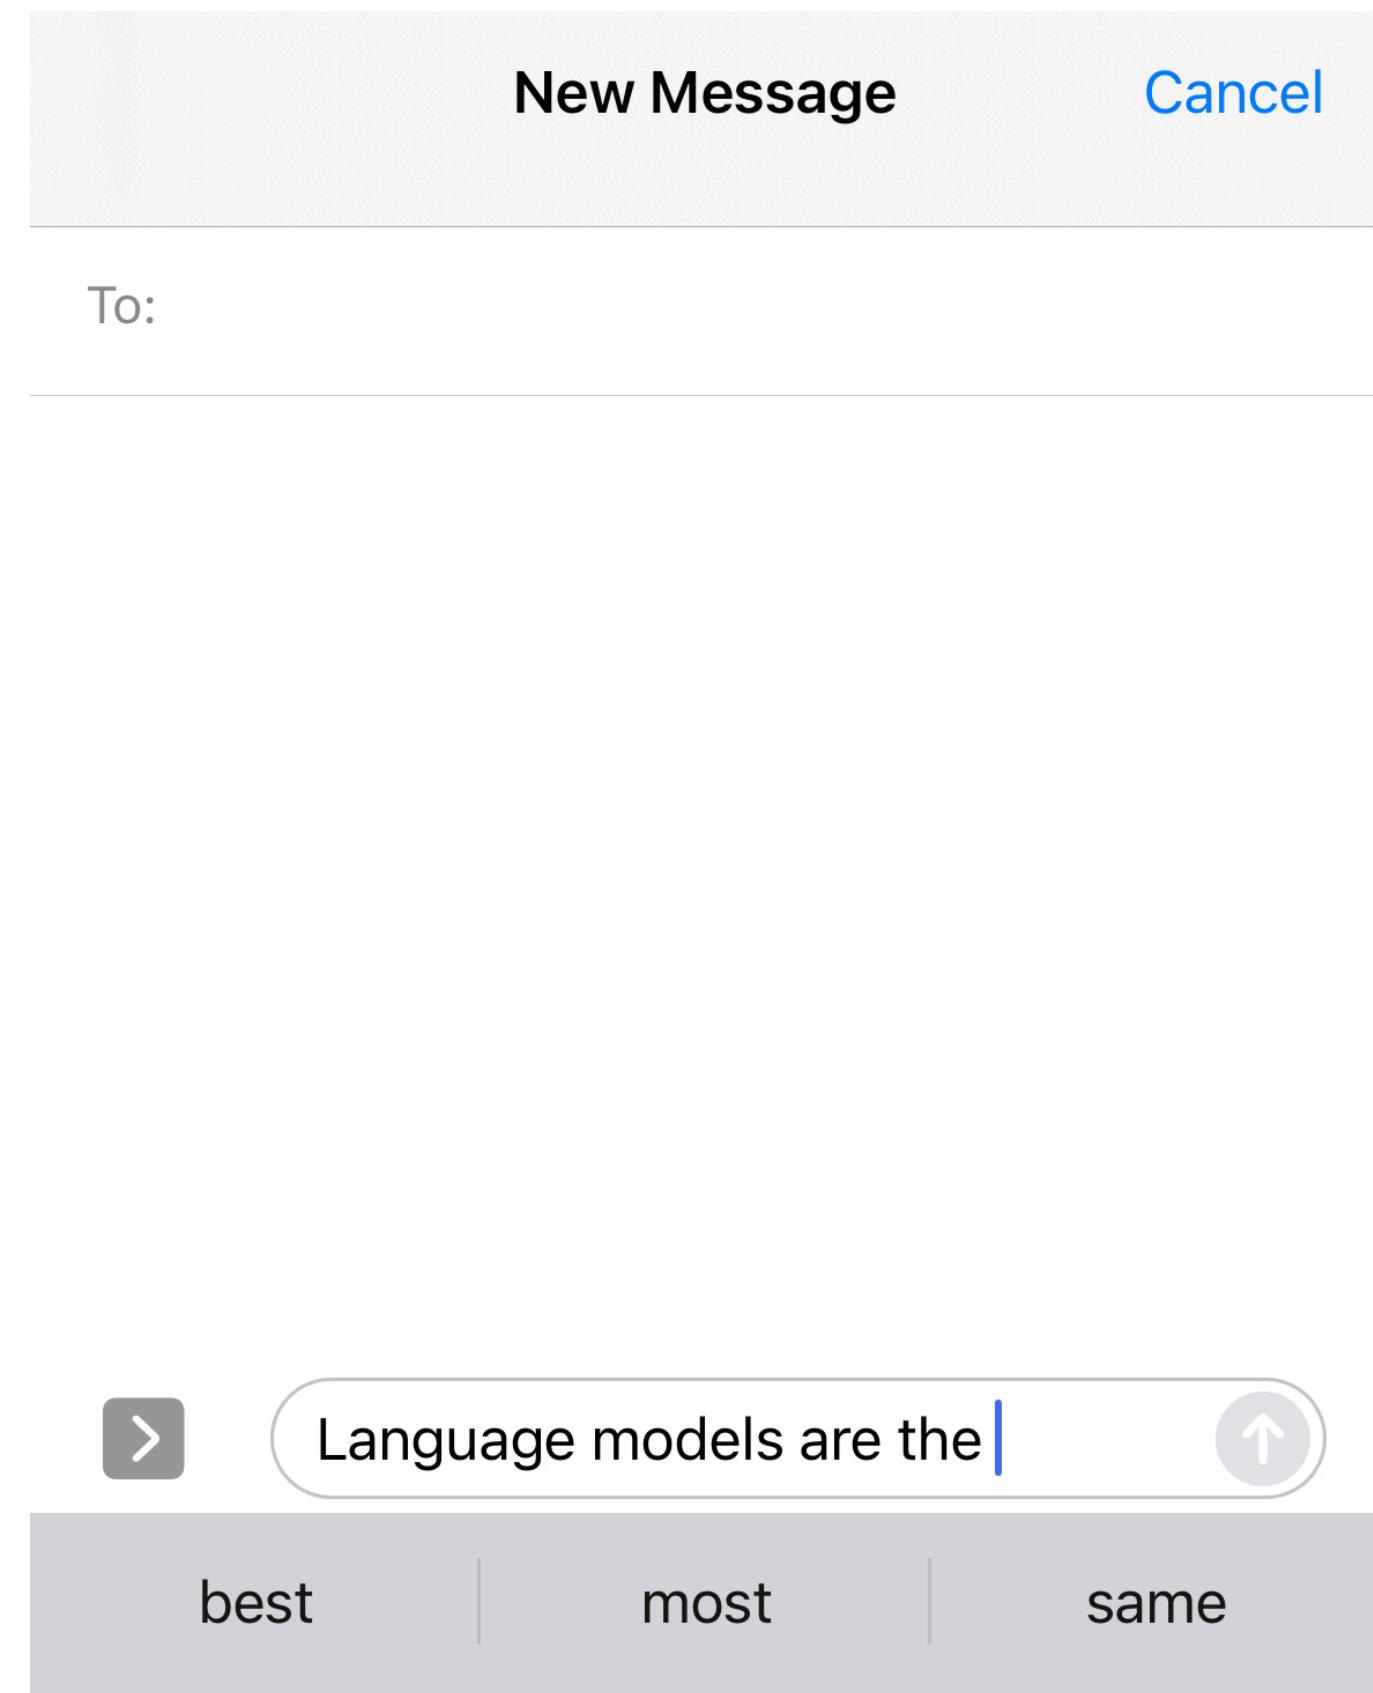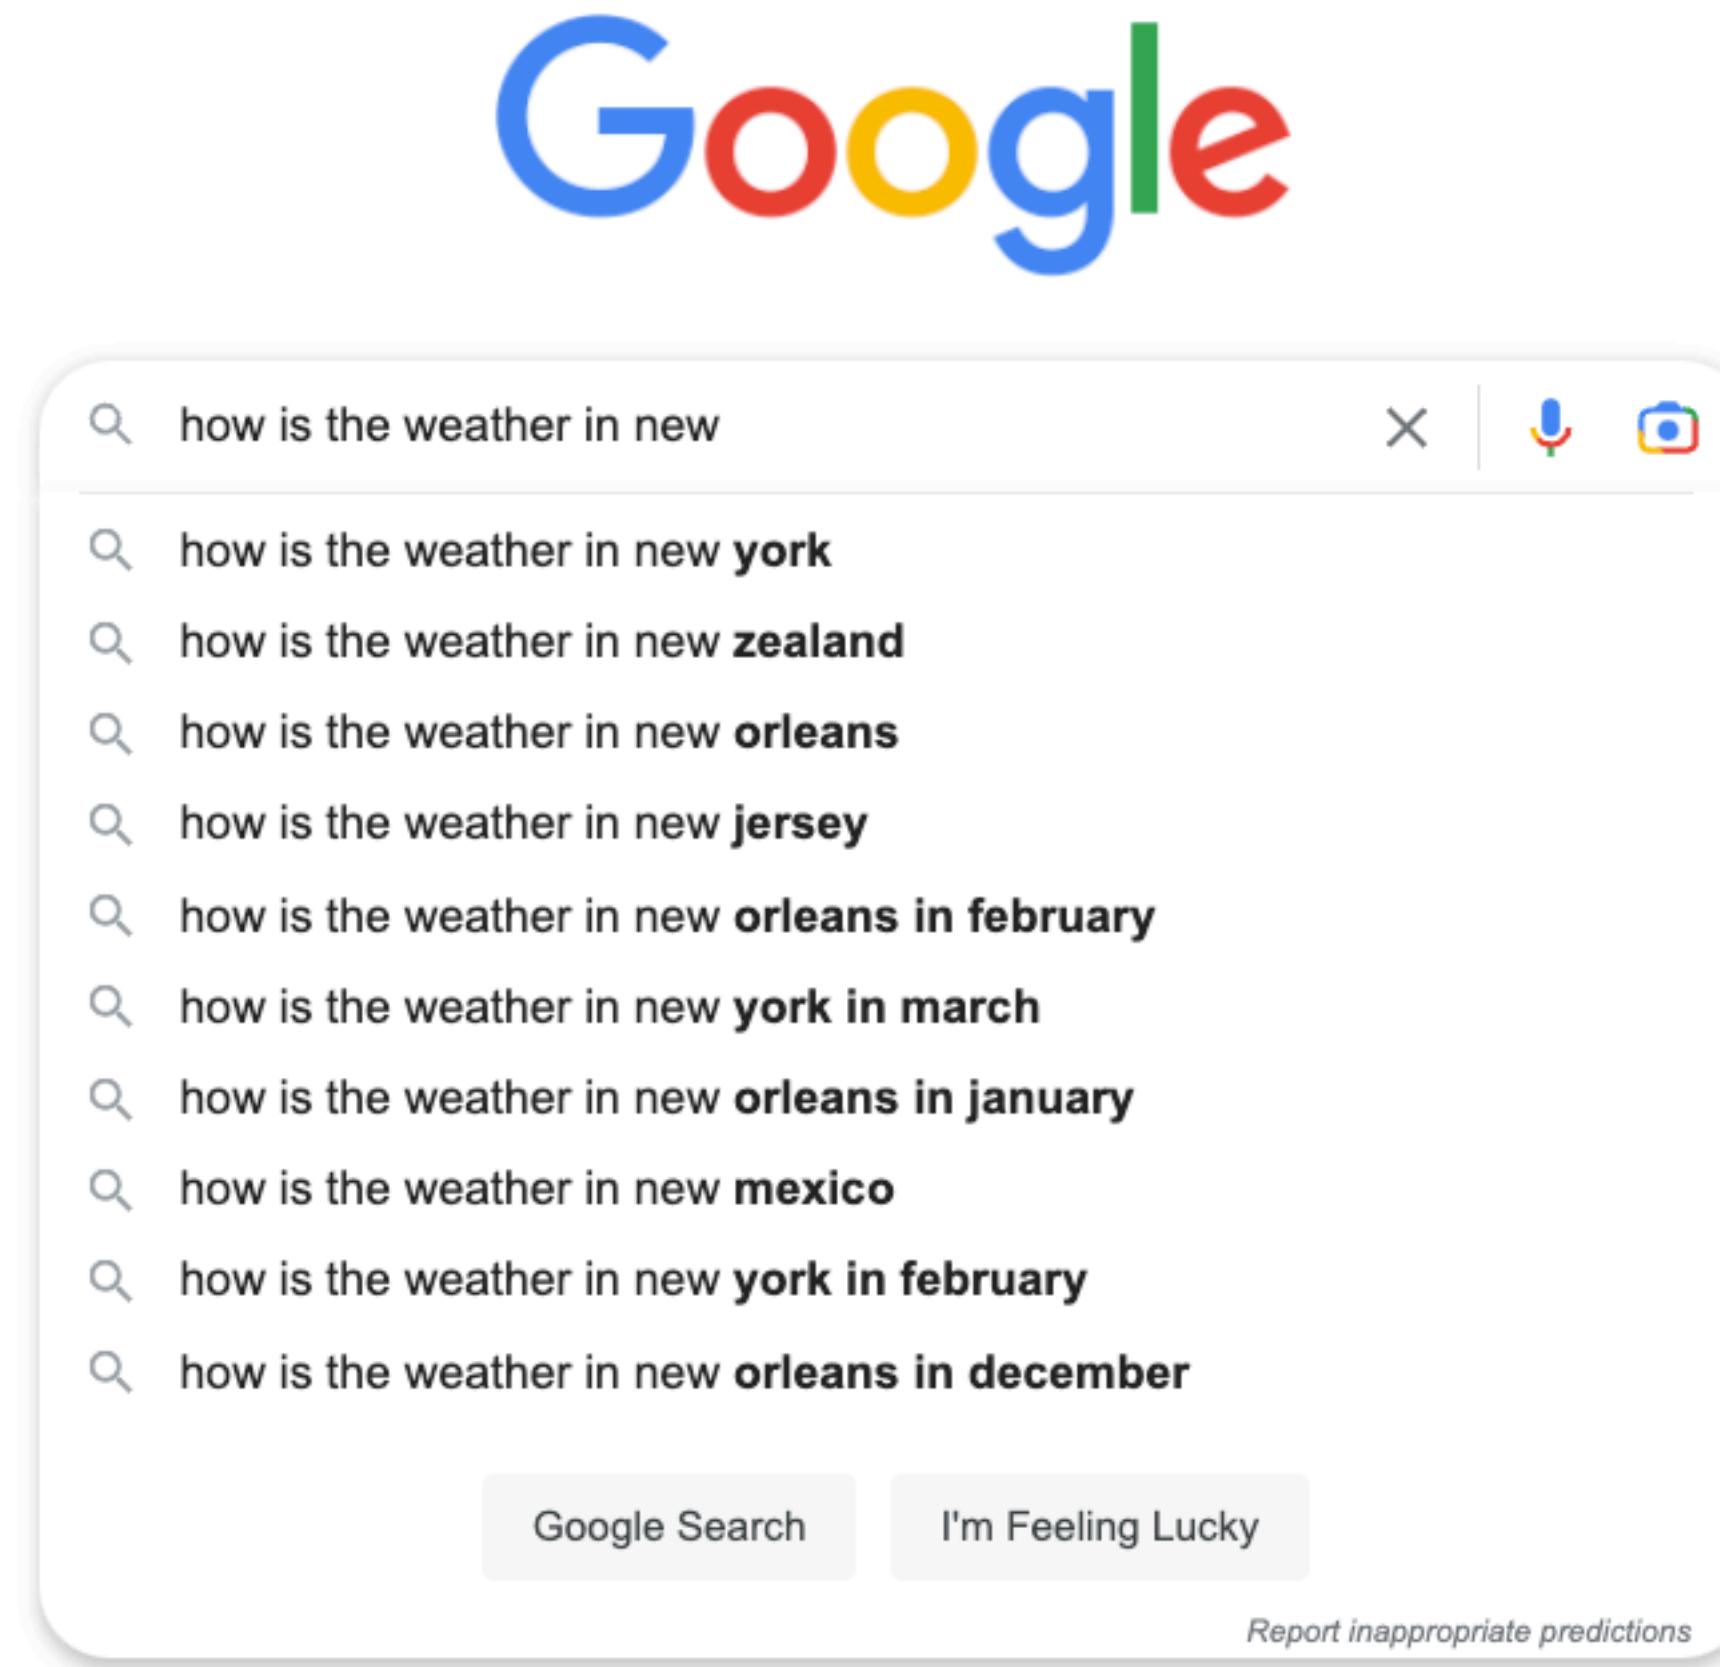
Estimating probabilities

$$P(\text{sat}|\text{the cat}) = \frac{\text{count(the cat sat)}}{\text{count(the cat)}}$$

$$P(\text{on}|\text{the cat sat}) = \frac{\text{count(the cat sat on)}}{\text{count(the cat sat)}}$$

$$\vdots$$

Assume we have a vocabulary of size V,

how many sequences of length n do we have?

A) $n * V$

B) n^V

C) V^n

D) V/n

Estimating probabilities

$$P(\text{sat}|\text{the cat}) = \frac{\text{count(the cat sat)}}{\text{count(the cat)}}$$

$$P(\text{on}|\text{the cat sat}) = \frac{\text{count(the cat sat on)}}{\text{count(the cat sat)}}$$

$$\vdots$$

Maximum likelihood estimate (MLE)

- With a vocabulary of size V, # sequences of length n = V^n

- Typical English vocabulary ~ 40k words

  - Even sentences of length ≤ 11 results in more than 4×10^{50} sequences. Too many to count!

  - *(For reference, # of atoms in the earth $\sim 10^{50}$)*

Markov assumption

- Use only the recent past to predict the next word

- Reduces the number of estimated parameters in exchange for modeling capacity

- 1st order

$$P(\text{mat}|\text{the cat sat on the}) \approx P(\text{mat}|\text{the})$$

- 2nd order

$$P(\text{mat}|\text{the cat sat on the}) \approx P(\text{mat}|\text{on the})$$

Andrey Markov

k-th order Markov

Consider only the last *k* words (or less) for context which implies the probability of a sequence is:

$$P(w_i \mid \boxed{w_1 w_2 \ldots w_{i-1}}) \approx P(w_i \mid \boxed{w_{i-k} \ldots w_{i-1}})$$

$$P(w_1 w_2 \ldots w_n) \approx \prod_i P(w_i \mid w_{i-k} \ldots w_{i-1})$$

(assume $w_j = \phi \quad \forall j < 0$)

Need to estimate counts for up to (k+1) grams

n-gram models

Unigram

$$P(w_1, w_2, ...w_n) = \prod_{i=1}^{n} P(w_i)$$

e.g. P(the) P(cat) P(sat)

Bigram

$$P(w_1, w_2, ...w_n) = \prod_{i=1}^{n} P(w_i|w_{i-1})$$

e.g. P(the) P(cat | the) P(sat | cat)

and Trigram, 4-gram, and so on.

Larger the n, more accurate and better the language model
(but also higher costs, prone to sparsity)

Quick quiz: Estimating probabilities

Consider the following corpus

<s> I like apples </s>

<s> You like strawberries </s>

<s> You like apples </s>

Note: <s> and </s> are starting and ending tokens

What's the bigram probability P(apples | like) ?

(A) 1/3   (B) 2/3   (C) 1/2   (D) 1

$$P(\text{apples} \mid \text{like}) = \frac{\text{Count("like apples")}}{\text{Count("like")}} = \frac{2}{3}$$

Quick quiz: Estimating probabilities

Consider the following corpus

<div align="center">

<s> I like apples </s>

<s> You like strawberries </s>

<s> You like apples </s>

</div>

Note: <s> and </s> are starting and ending tokens

Using the bigram model, what's the probability of the sentence "<s> I like strawberries </s>"? Ignore the probability of <s>.

(A) 4/9   (B) 1/3   (C) 2/9   (D) 1/9

P(<s> I like strawberries </s>)

= P(I | <s>) · P(like | I) · P(strawberries | like) · P(</s> | strawberries) = $\frac{1}{3} \cdot 1 \cdot \frac{1}{3} \cdot 1$

Generating from a language model

Generating from a language model

- Given a language model, how to generate a sequence?

$$\text{Bigram} \quad P(w_1, w_2, ...w_n) = \prod_{i=1}^{n} P(w_i | w_{i-1})$$

- Generate the first word $w_1 \sim P(w)$

- Generate the second word $w_2 \sim P(w \mid w_1)$

- Generate the third word $w_3 \sim P(w \mid w_2)$

- …

the	of	a	to	in
0.06	0.03	0.02	0.02	0.02

.06   .09 .11 .13 .15

0

however
(p=.0003)

polyphonic
p=.0000018

.66   .99

1
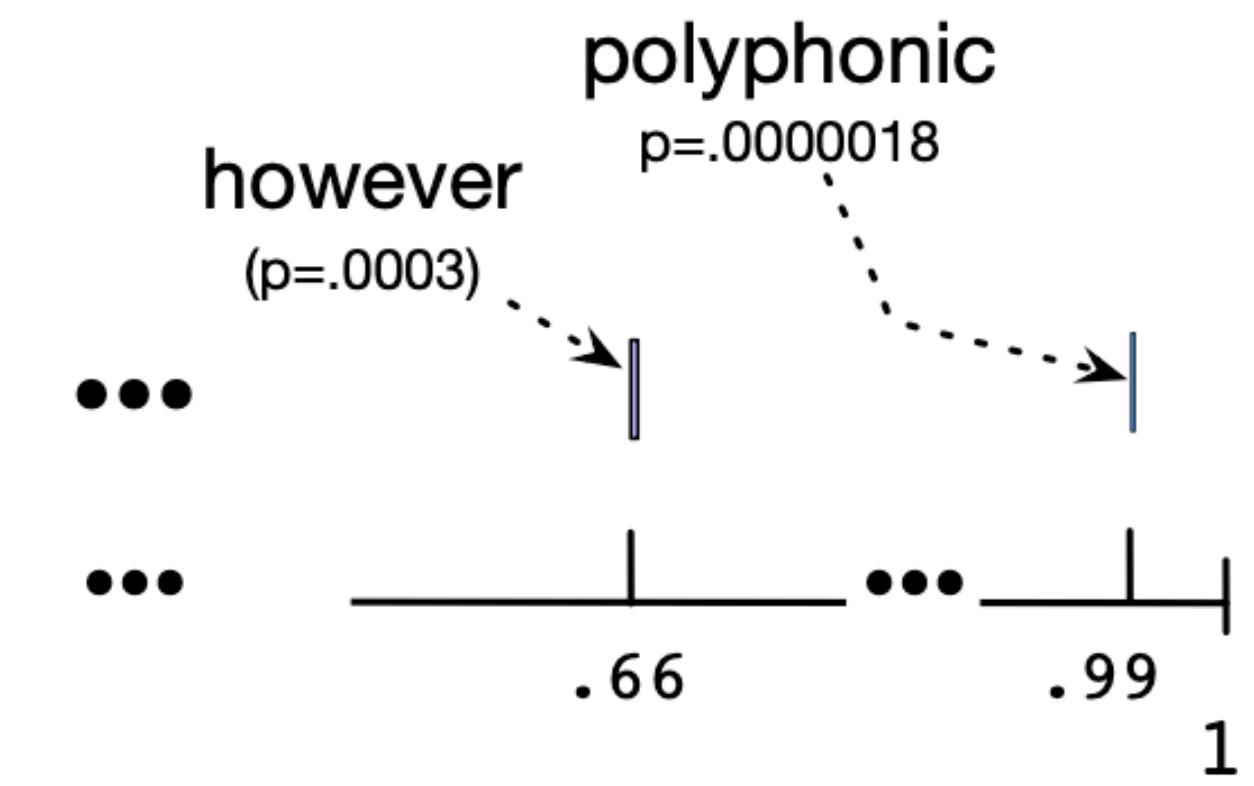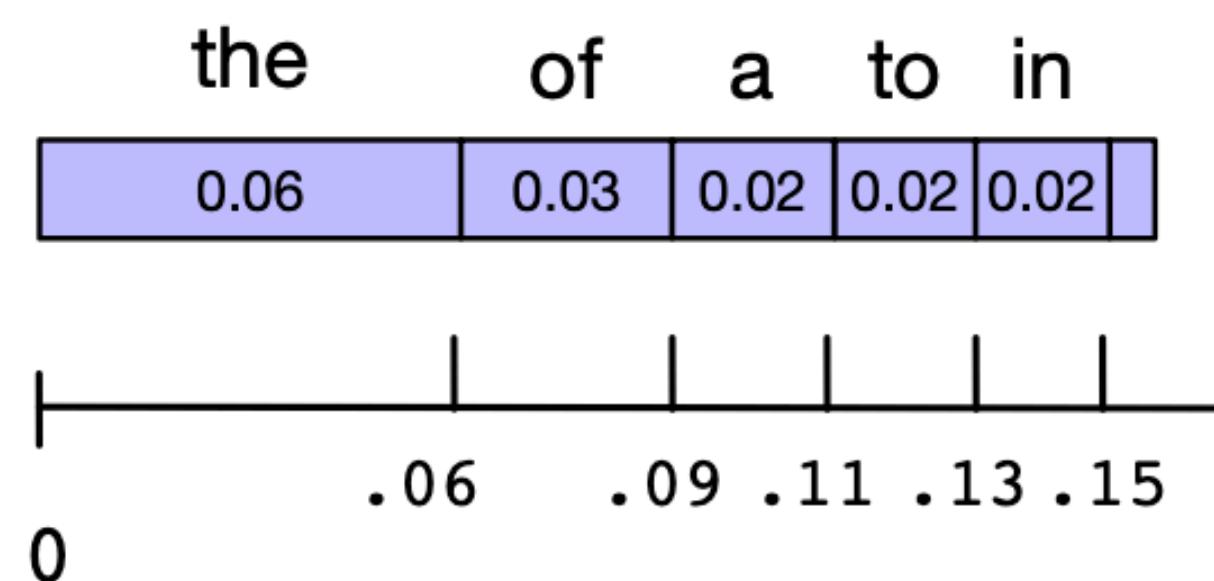
Generating from a language model

- Given a language model, how to generate a sequence?

$$\text{Trigram} \quad P(w_1, w_2, \ldots, w_n) = \prod_{i=1}^{n} P(w_i \mid w_{i-2}, w_{i-1})$$

- Generate the first word $w_1 \sim P(w)$

- Generate the second word $w_2 \sim P(w \mid w_1)$

- Generate the third word $w_3 \sim P(w \mid w_1, w_2)$

- Generate the fourth word $w_4 \sim P(w \mid w_2, w_3)$

- ...

Generations

Unigram
- *To him swallowed confess hear both. Which. Of save on trail for are ay device and rote life have*

Bigram
- *Why dost stand forth thy canopy, forsooth; he is this palpable hit the King Henry. Live king. Follow.*
- *What means, sir. I confess she? then all sorts, he is trim, captain.*

Trigram
- *Fly, and will rid me these news of price. Therefore the sadness of parting, as they say,*
- *This shall forbid it should be branded, if renown made it empty*

Typical LMs are not sufficient to handle long-range dependencies

"The woman/man could not go to work that day because she/he had a doctor's appointment"

Generations

Example from a GPT-2 output (2019):

With the start of the new academic year, Princeton has an opportunity to help provide a new generation of women with a diverse set of academic resources for higher education.
We are offering the resources of the Princeton-McGill program specifically to women with undergraduate degrees who would like to enhance their academic experience. Princeton-McGill offers a comprehensive suite of services for women and their families including a variety of graduate programs, support programs, and the opportunity to serve as leaders in their communities with a wide variety of programs, activities and services. For the upcoming fall, Princeton-McGill will also offer its Women's Center , which is located in a renovated women's dorm.
At Princeton, we are working with the Princeton-McGill community to develop a suite of programs that are designed to give new and returning students a strong foundation for a successful, rewarding graduate career. The Women's Center , the Princeton-McGill Women's Center provides a range of supports to address the specific needs of female doctoral degree graduates. Programs are tailored to meet the unique needs of women under the age of 28, women and families

https://talktotransformer.com/

$$P(w_1, w_2, \ldots, w_n) = \prod_{i=1}^{n} P(w_i \mid w_{i-1024}, \ldots, w_{i-2}, w_{i-1})$$

Modern LMs can handle much longer contexts!

Generation methods (more in the later lecture)

- Greedy: choose the most likely word!

  - To predict the next word given a context of two words w_1, w_2:

$$w_3 = \arg\max_{w \in V} P(w \,|\, w_1, w_2)$$

- Top-k vs top-p sampling: "The boy went to the _____"

Top-k sampling   Top-p sampling

https://blog.allenai.org/a-guide-to-language-model-sampling-in-allennlp-3b1239274bc3
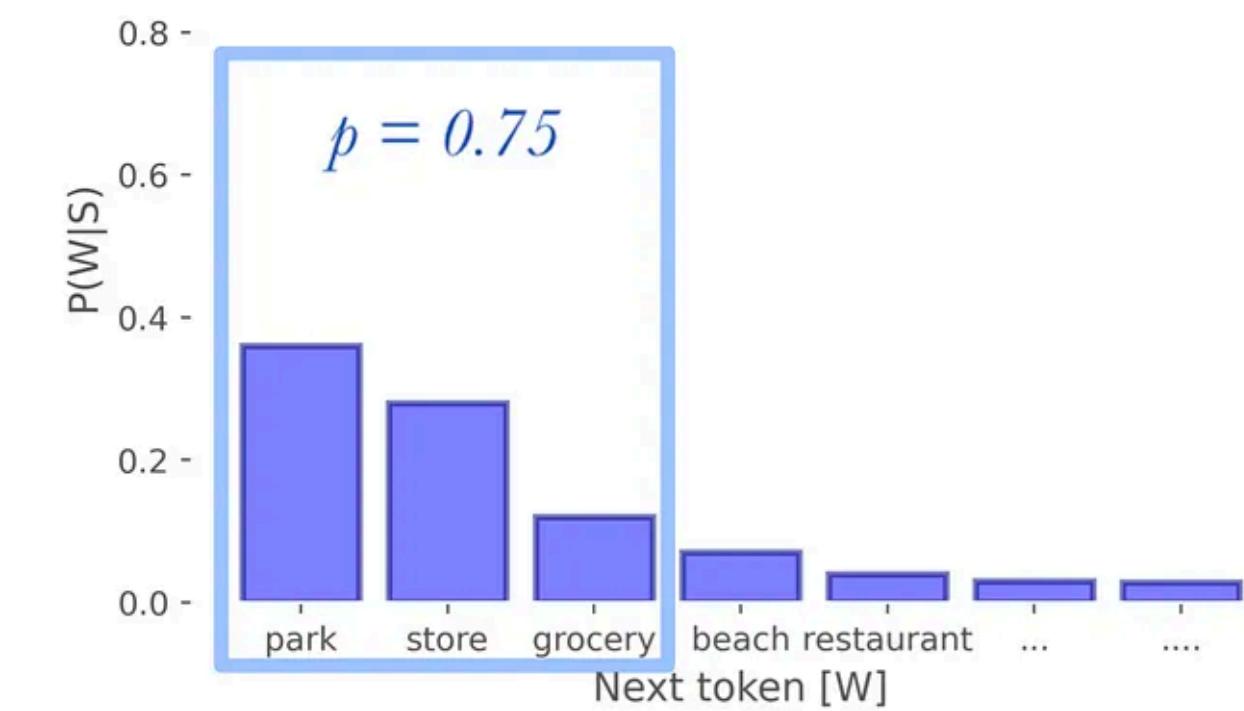
Evaluating a language model

Extrinsic evaluation
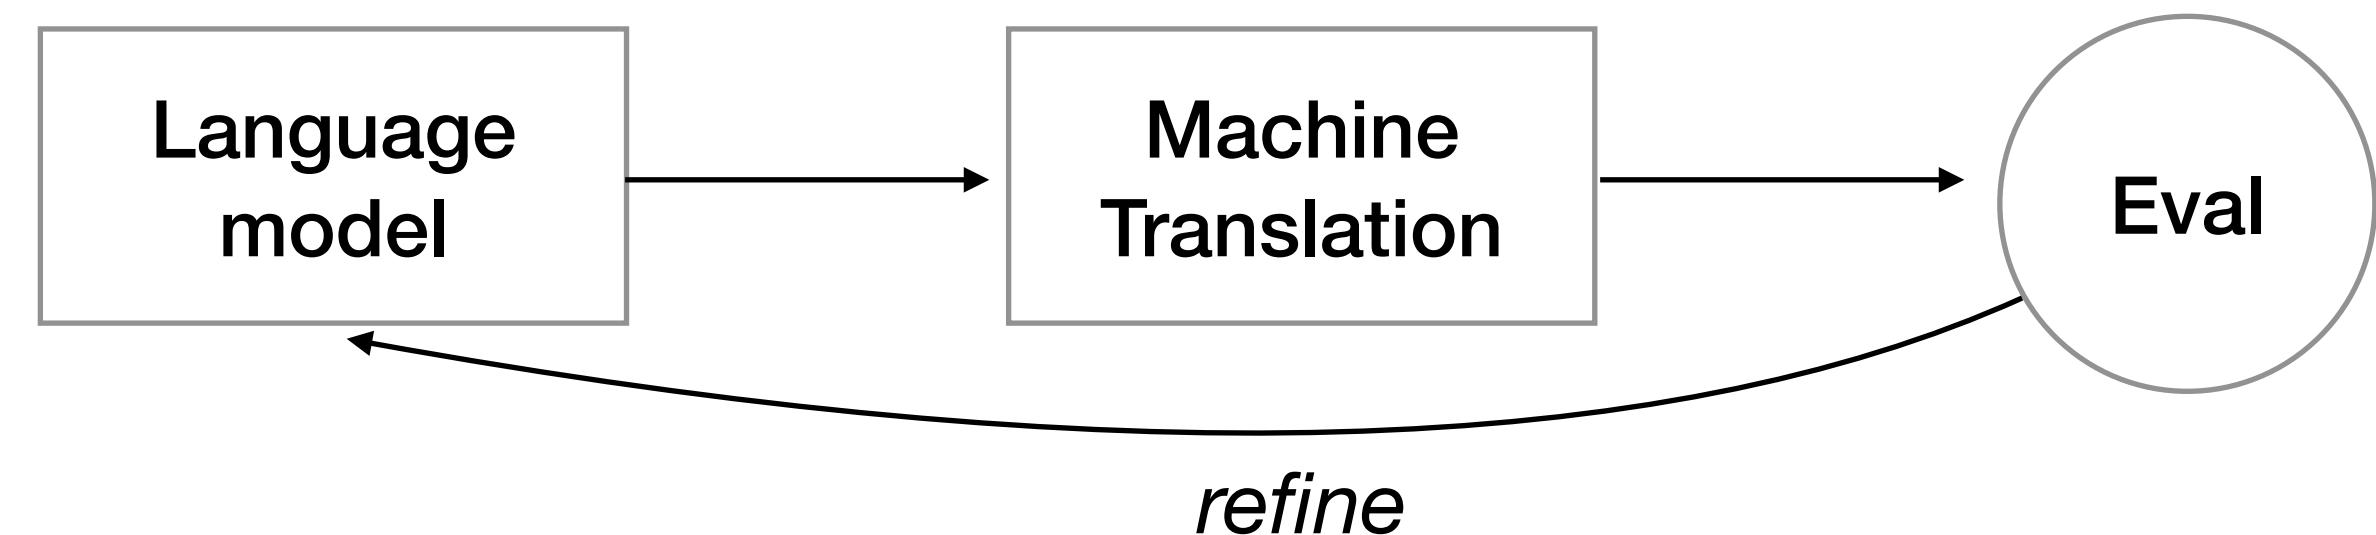
- Train LM → apply to task → observe accuracy

- Directly optimized for downstream applications

  - higher task accuracy → better model

- Expensive, time consuming

- Hard to optimize downstream objective (indirect feedback)

Intrinsic evaluation of language models

- **Train** parameters on a suitable training corpus

  - Assumption: observed sentences ~ good sentences

- **Test** on different, unseen corpus

  - If a language model assigns a higher probability to the test set, it is better

- Evaluation metric - **perplexity**!

Motivation: Shannon game

Goal for language models: model $\Pr[w_1 w_2 \ldots w_k]$ or $\Pr[w_k | w_1 \ldots w_{k-1}]$ well.

Shannon game: How well can we predict the next word?

- I always order pizza with cheese and ____
- The 33rd president of the US was ____
- I saw a ____

(Slide credit: Princeton CS 324, Ruth Fong)

Motivation: Shannon game

Goal for language models: model $\Pr[w_1 w_2 \ldots w_k]$ or $\Pr[w_k | w_1 \ldots w_{k-1}]$ well.

Shannon game: How well can we predict the next word?

- **I always order pizza with cheese and ____**
- The 33rd president of the US was ____
- I saw a ____

mushrooms 0.1
pepperoni 0.1
anchovies 0.01
…
friend 0.0001
…

How would a unigram model do?

(Slide credit: Princeton CS 324, Ruth Fong)

Motivation: Shannon game

Goal for language models: model $\Pr[w_1 w_2 \ldots w_k]$ or $\Pr[w_k | w_1 \ldots w_{k-1}]$ well.

Shannon game: How well can we predict the next word?
- **I always order pizza with cheese and ____**
- The 33rd president of the US was ____
- I saw a ____

mushrooms 0.1
pepperoni 0.1
anchovies 0.01
…
friend 0.0001
…

How would a unigram model do?
Not well, it would assign the most probability to the most common word (i.e. the word that occurs the most in the training corpus).

Better language model = assigns higher probability to word that actually occurs.

(Slide credit: Princeton CS 324, Ruth Fong)

Perplexity (ppl)

Measure of how well a LM predicts the next word

- For a test corpus with words $w_1, w_2, \ldots w_n$

$$\text{Perplexity} = P(w_1, w_2, \ldots, w_n)^{-1/n}$$

$$\text{ppl}(S) = 2^x \quad \text{where} \quad x = -\frac{1}{n}\log_2 P(w_1, \ldots, w_n) = -\frac{1}{n}\sum_{i=1}^{n}\log_2 P(w_i \,|\, w_1 \ldots w_{i-1})$$

Cross-Entropy

- Minimizing perplexity ~ maximizing probability of corpus

(Slide credit: Princeton CS 324, Ruth Fong)
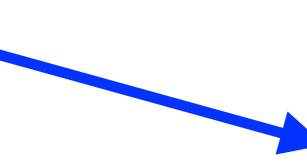
Intuition on perplexity

$$\text{ppl}(S) = 2^x \quad \text{where} \quad x = -\frac{1}{n} \sum_{i=1}^{n} \log P(w_i \,|\, w_1 \ldots w_{i-1})$$

If our k-gram model (with vocabulary V) has following probability:

$$P(w \,|\, w_{i-k}, \ldots w_{i-1}) = \frac{1}{|V|}, \quad \forall w \in V$$

what is the perplexity of the test corpus?

A) $2^{|V|}$ B) $|V|$ C) $|V|^2$ D) $2^{-|V|}$

Intuition on perplexity

$$\text{ppl}(S) = 2^x \quad \text{where} \quad x = -\frac{1}{n} \sum_{i=1}^{n} \log P(w_i \,|\, w_1 \ldots w_{i-1})$$

If our k-gram model (with vocabulary V) has following probability:

$$P(w \,|\, w_{i-k}, \ldots w_{i-1}) = \frac{1}{|V|}, \quad \forall w \in V$$

what is the perplexity of the test corpus?

A) $2^{|V|}$ B) $|V|$ C) $|V|^2$ D) $2^{-|V|}$

$$\text{ppl} = 2^{-\frac{1}{n} n \log(1/|V|)} = |V|$$

Measure of model's uncertainty about next word (aka `average branching factor')

branching factor = # of possible words following any word

Perplexity

Training corpus 38 million words, test corpus 1.5 million words, both **WSJ**

N-gram Order	Unigram	Bigram	Trigram
Perplexity (test)	962	170	109

GPT-3 175B: ppl = 20.5

https://paperswithcode.com/sota/language-modelling-on-penn-treebank-word
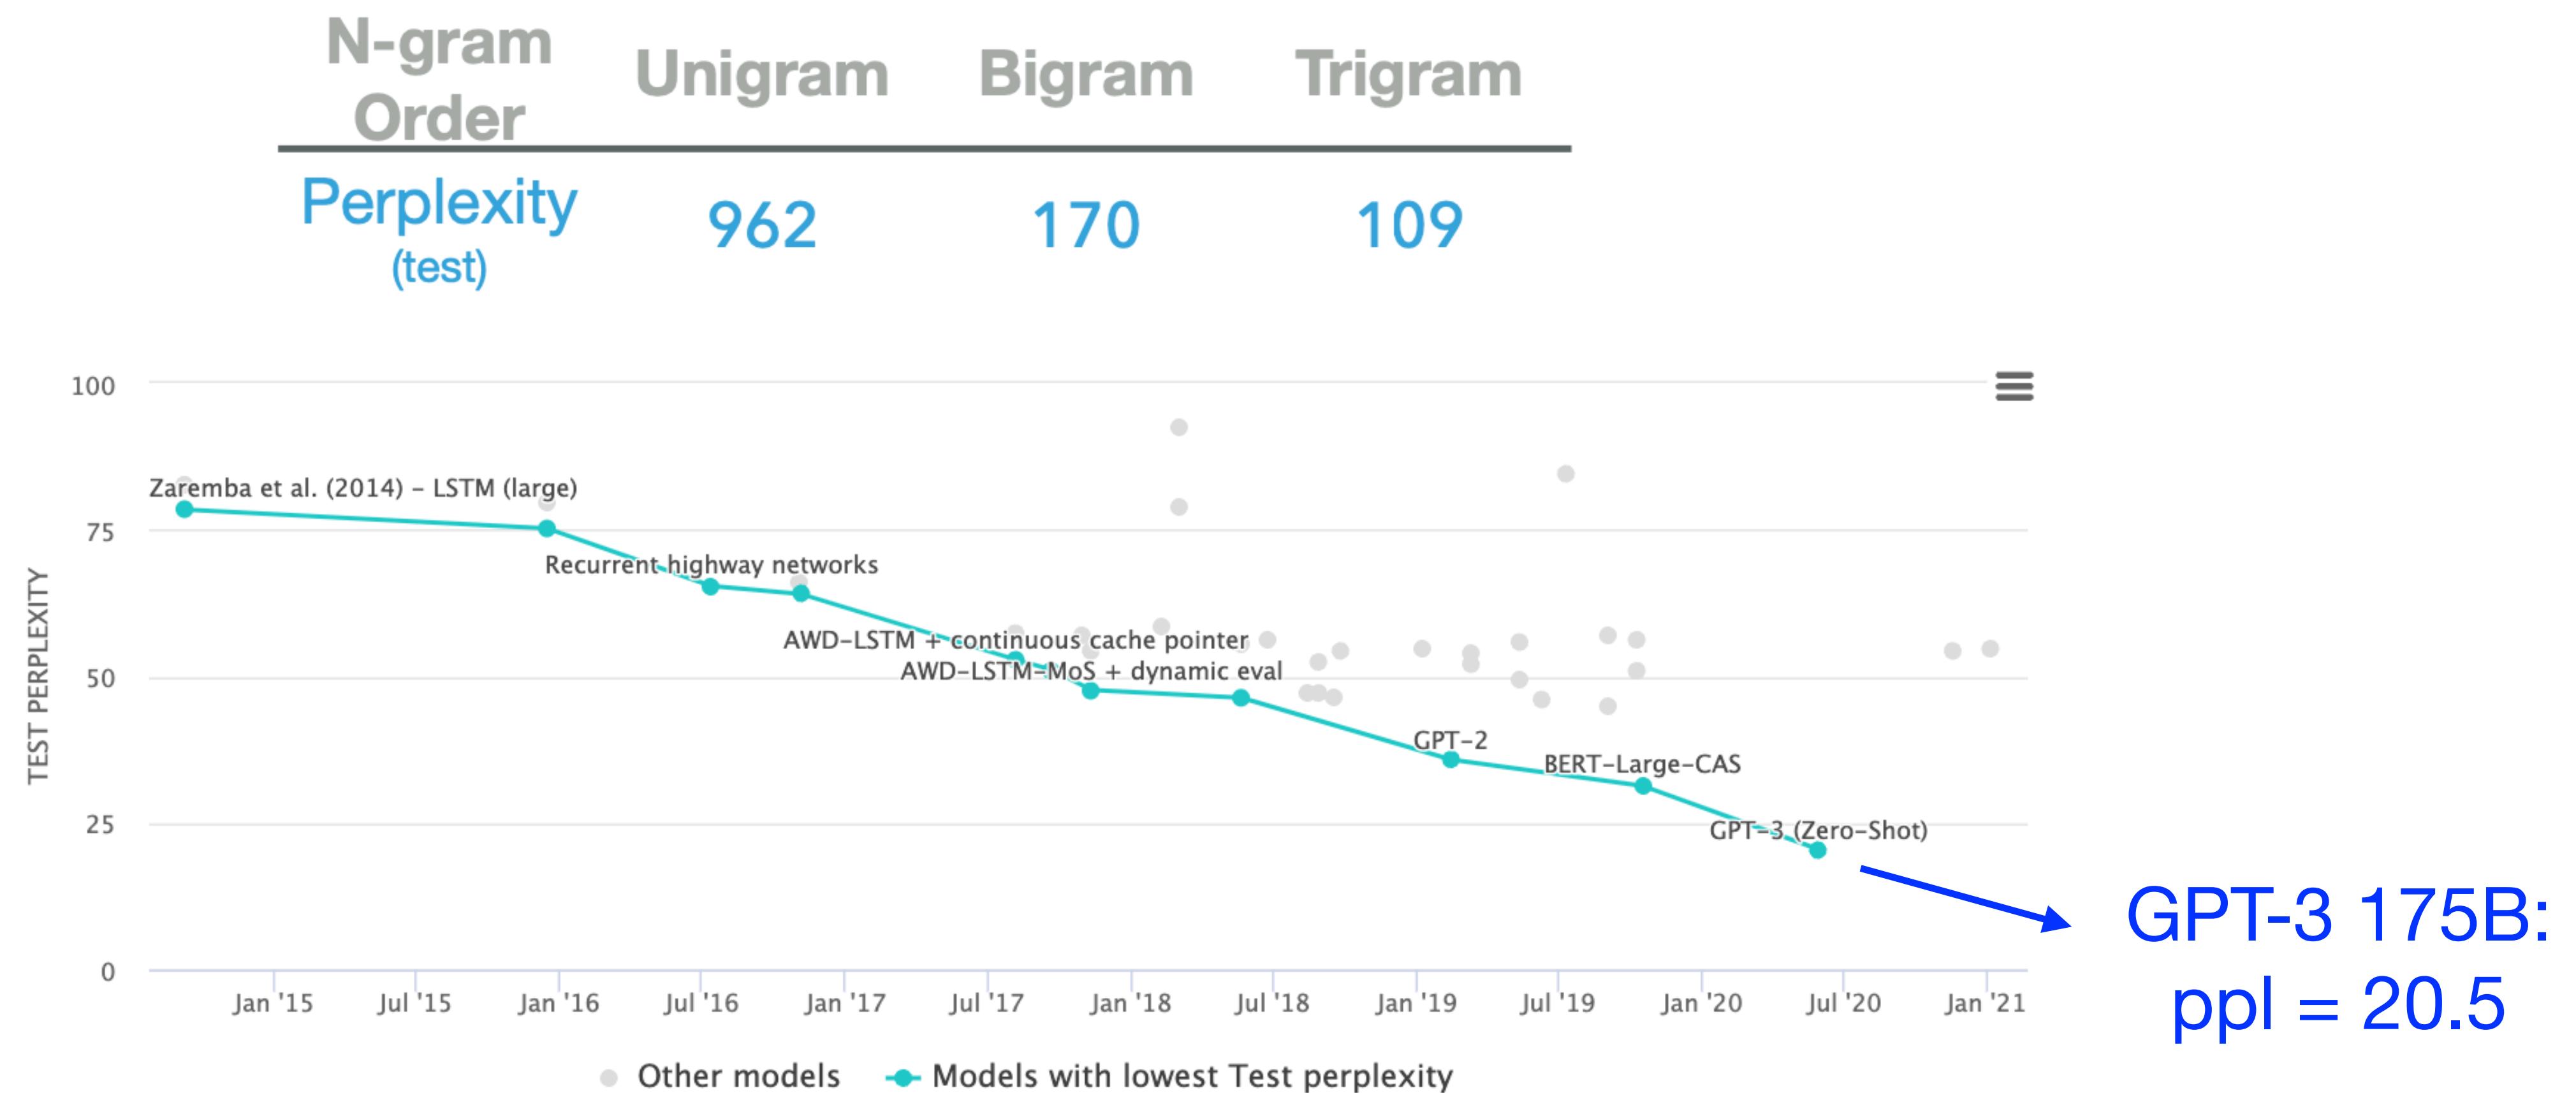
Jan 22 lecture starts from here

CS 288 Advanced Natural Language Processing

Course website: cal-cs288.github.io/sp26

Ed: edstem.org/us/courses/92268

- Class starts at 15:40!

- Quick announcements

  - Staff won't read enrollment-related questions via email/Ed. Please mention all relevant info in the Google Form. You will be notified no later than 02/03.

  - For those with access problems (to Ed/Gradescope), email our GSI, Zineng Tang (email address on the website).

- Today's lecture plan: Finish n-gram LM (20min) → Word embeddings (60min)

Recap: *n*-gram LM

- Language Model: A probabilistic model of a sequence of words

$$P(\text{the cat sat on the mat}) = P(\text{the}) * P(\text{cat}|\text{the}) * P(\text{sat}|\text{the cat})$$
$$* P(\text{on}|\text{the cat sat}) * P(\text{the}|\text{the cat sat on})$$
$$* P(\text{mat}|\text{the cat sat on the})$$

- How to compute these probabilities? Maximum likelihood estimate (MLE)!

$$P(\text{sat}|\text{the cat}) = \frac{\text{count(the cat sat)}}{\text{count(the cat)}}$$

$$P(\text{on}|\text{the cat sat}) = \frac{\text{count(the cat sat on)}}{\text{count(the cat sat)}}$$

- Problem: As the sentence length grows, this becomes intractable!

Recap: *n*-gram LM

- *n*-gram LM: use Markov assumption to approximate probability

  - 1st order $\quad P(\text{mat}|\text{the cat sat on the}) \approx P(\text{mat}|\text{the}) \rightarrow$ bi-gram LM

  - 2nd order $\quad P(\text{mat}|\text{the cat sat on the}) \approx P(\text{mat}|\text{on the}) \rightarrow$ tri-gram LM

Larger the n, more accurate and better the language model (but also higher costs)

Recap: LM Evaluation: Perplexity (ppl)

Measure of how well a LM predicts the next word

- For a test corpus with words $w_1, w_2, \ldots w_n$

$$\text{Perplexity} = P(w_1, w_2, \ldots, w_n)^{-1/n}$$

$$\text{ppl}(S) = 2^x \quad \text{where} \quad x = -\frac{1}{n} \log_2 P(w_1, \ldots, w_n) = -\frac{1}{n} \sum_{i=1}^{n} \log_2 P(w_i \mid w_1 \ldots w_{i-1})$$

Cross-Entropy

- Minimizing perplexity ~ maximizing probability of corpus

(Slide credit: Princeton CS 324, Ruth Fong)
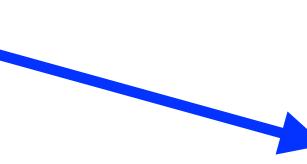
Recap: LM Evaluation: Perplexity (ppl)

Training corpus 38 million words, test corpus 1.5 million words, both **WSJ**

N-gram Order	Unigram	Bigram	Trigram
Perplexity (test)	962	170	109

GPT-3 175B: ppl = 20.5

https://paperswithcode.com/sota/language-modelling-on-penn-treebank-word
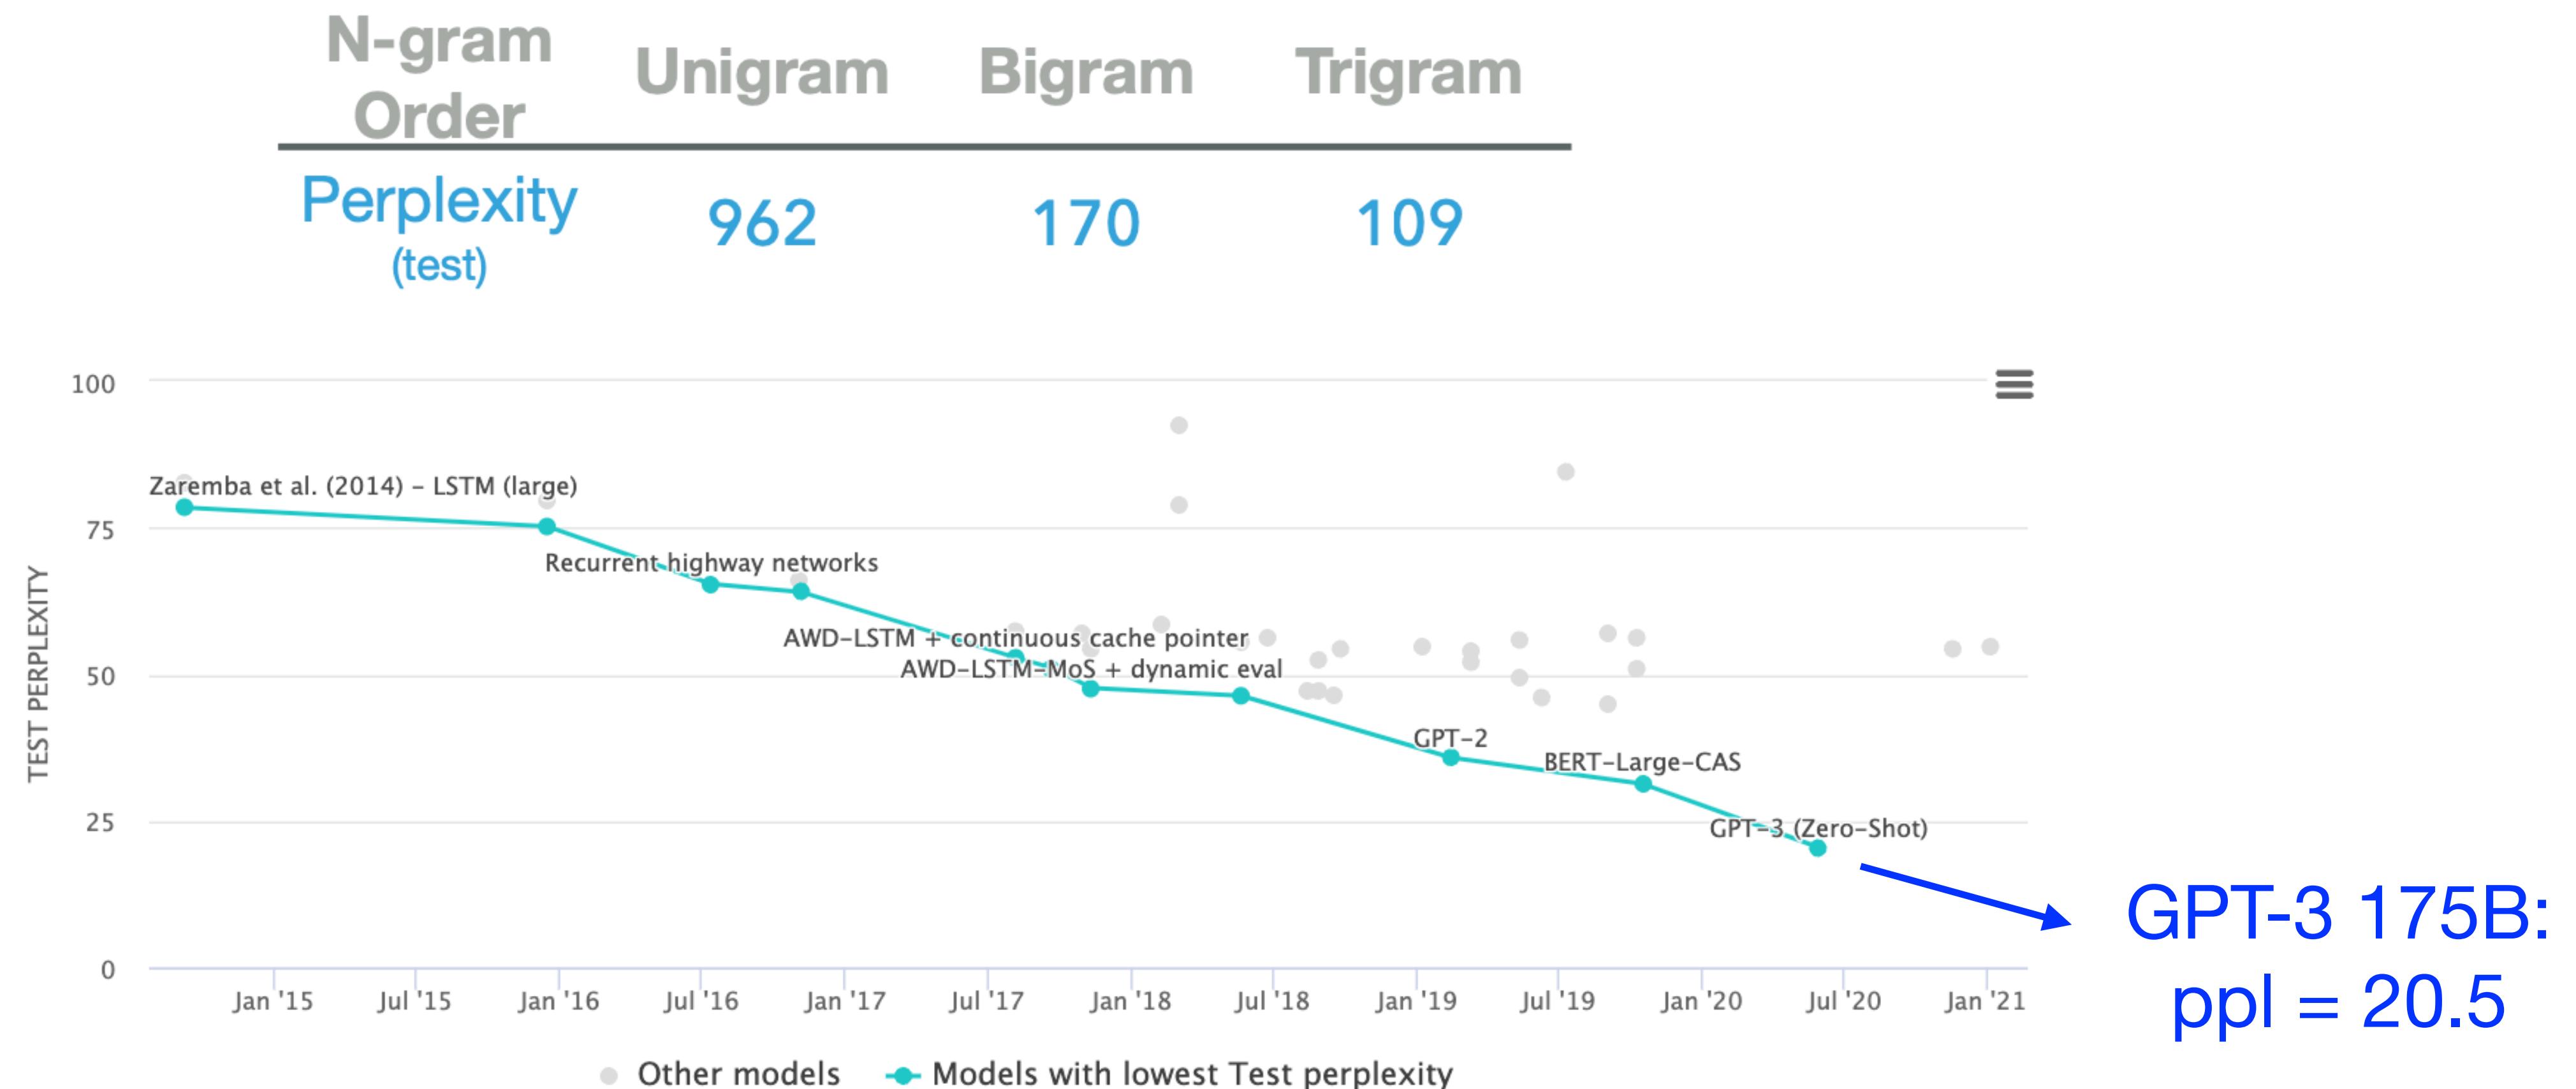
Quick quiz (question from previous class)

$$\text{ppl}(S) = 2^x \quad \text{where} \quad x = -\frac{1}{n}\log_2 P(w_1, \ldots, w_n)$$

- What is the minimum possible perplexity?

$P(w_1, \ldots, w_n) = 1$, therefore $x = -\frac{1}{n}\log_2 P(w_1, \ldots, w_n) = 0$, therefore $\text{ppl}(S) = 2^0 = 1$

- What is the maximum possible perplexity?

$P(w_1, \ldots, w_n) = 0$, therefore $x = -\frac{1}{n}\log_2 P(w_1, \ldots, w_n) = \infty$, therefore $\text{ppl}(S) = 2^\infty = \infty$

Smoothing

Sparsity

Corpus: Sparsity

• Even as raw corpora become extremely large …

  – … sparsity is always a problem

New words/word pairs
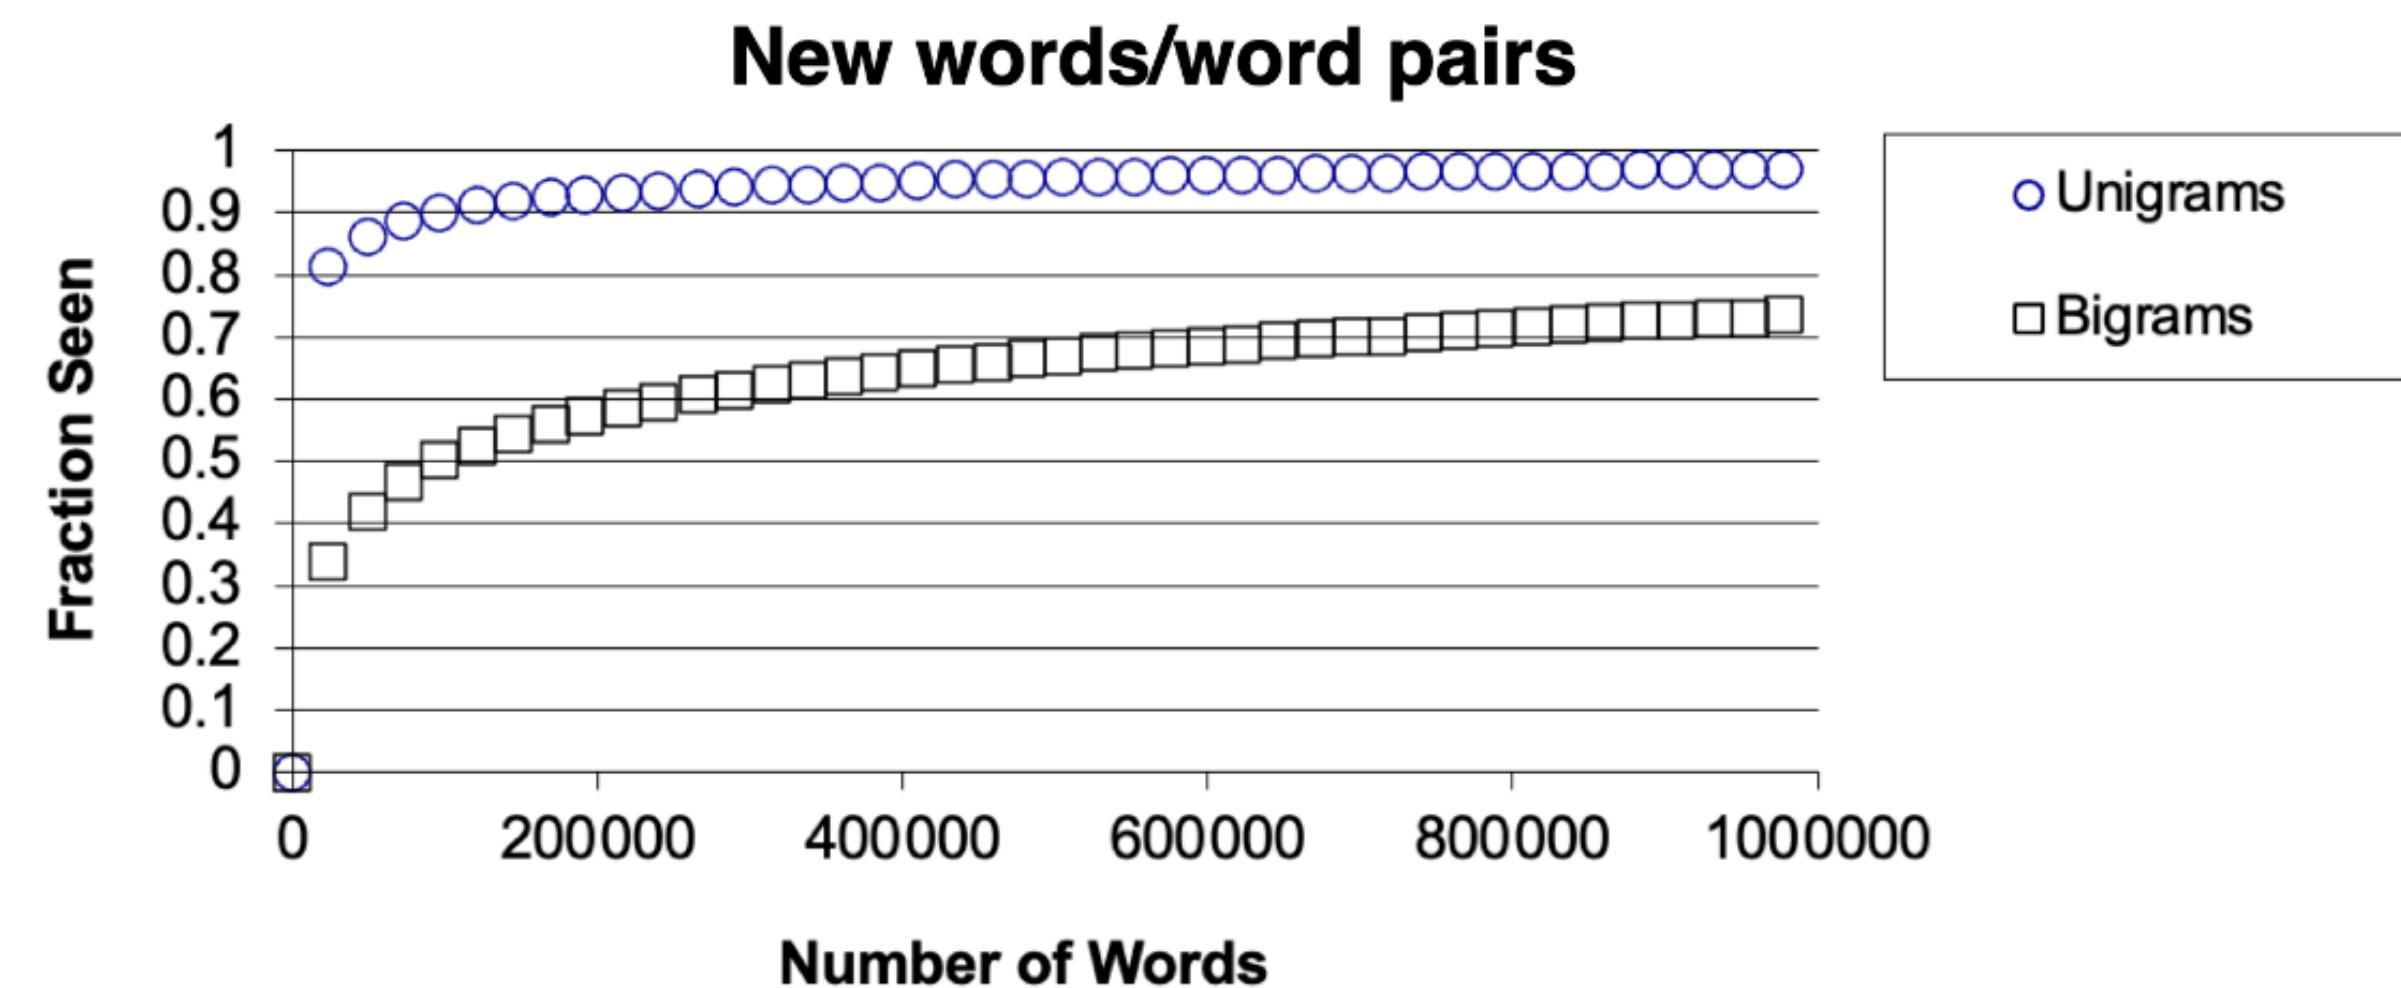
Generalization of n-grams

- Not all n-grams in the test set will be observed in training data

- Test corpus might have some that have zero probability under our model

  - **Training set**: *Google news*

  - **Test set**: *Shakespeare*

  - P(affray | voice doth us) = 0 $\implies$ P(test corpus) = 0

  - Perplexity is not defined.

$$\text{ppl}(S) = 2^x \quad \text{where}$$
$$x = -\frac{1}{n} \sum_{i=1}^{n} \log P(w_i \,|\, w_1 \dots w_{i-1})$$

Smoothing

- Handle sparsity by making sure all probabilities are non-zero in our model

  - **Additive**: Add a small amount to all probabilities

  - **Interpolation**: Use a combination of different granularities of n-grams

  - **Discounting**: Redistribute probability mass from observed n-grams to unobserved ones

Smoothing intuition

When we have sparse statistics:

P(w | denied the)
  3 allegations
  2 reports
  1 claims
  1 request

  7 total
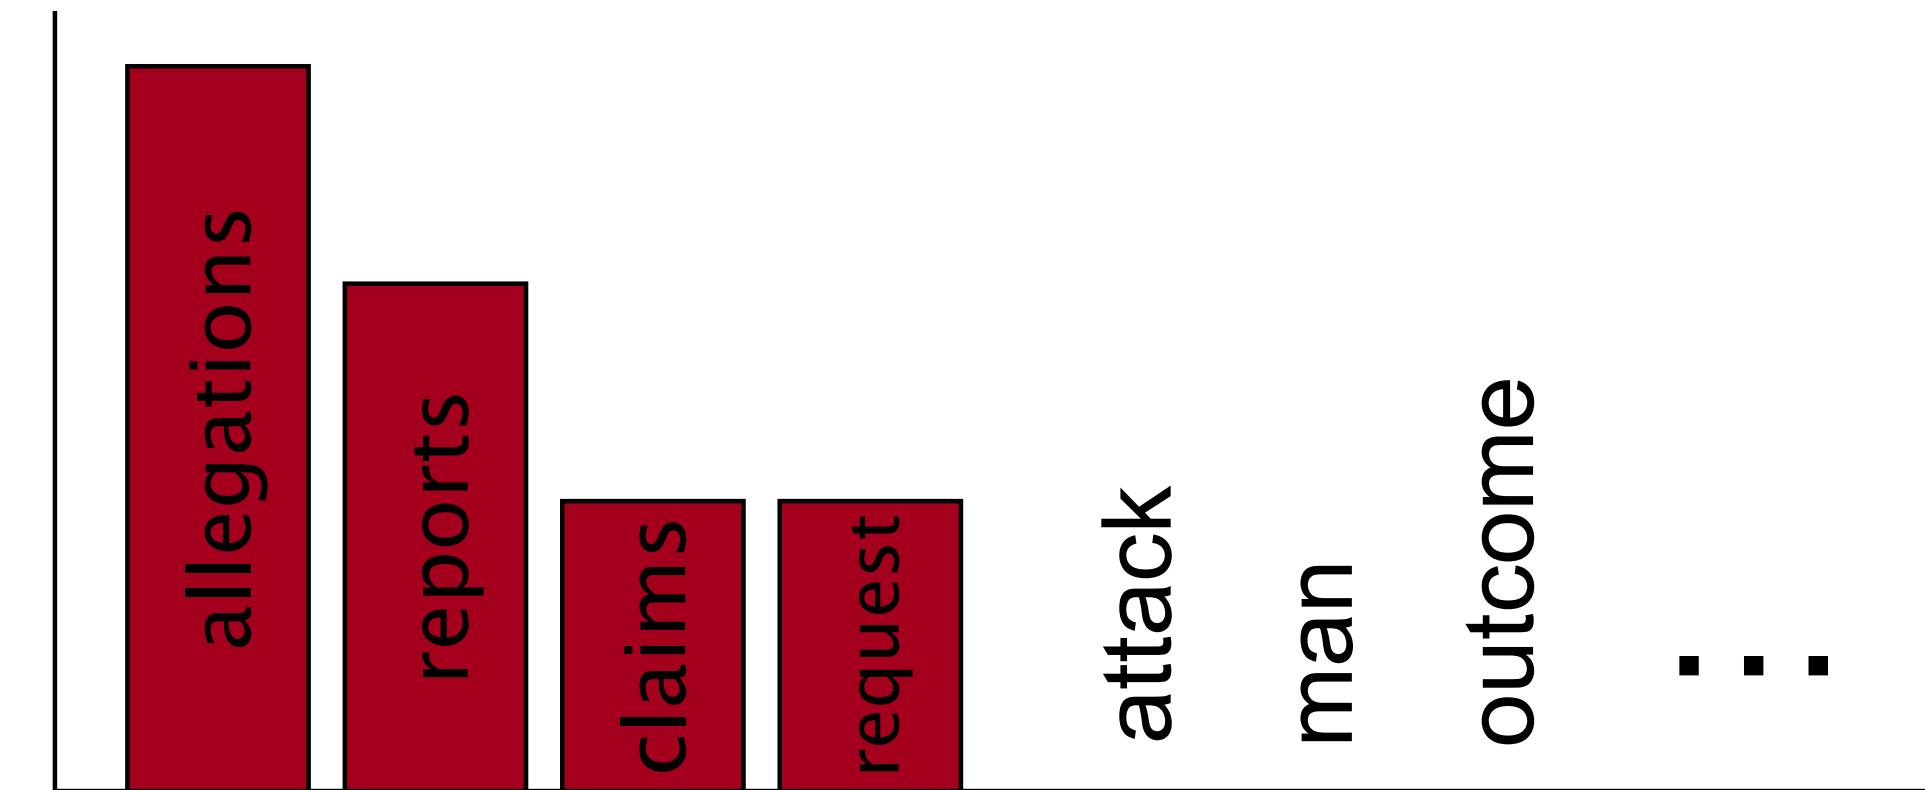
Steal probability mass to generalize better

P(w | denied the)
  2.5 allegations
  1.5 reports
  0.5 claims
  0.5 request
  2 other

  7 total

(Slide credit: Dan Klein)
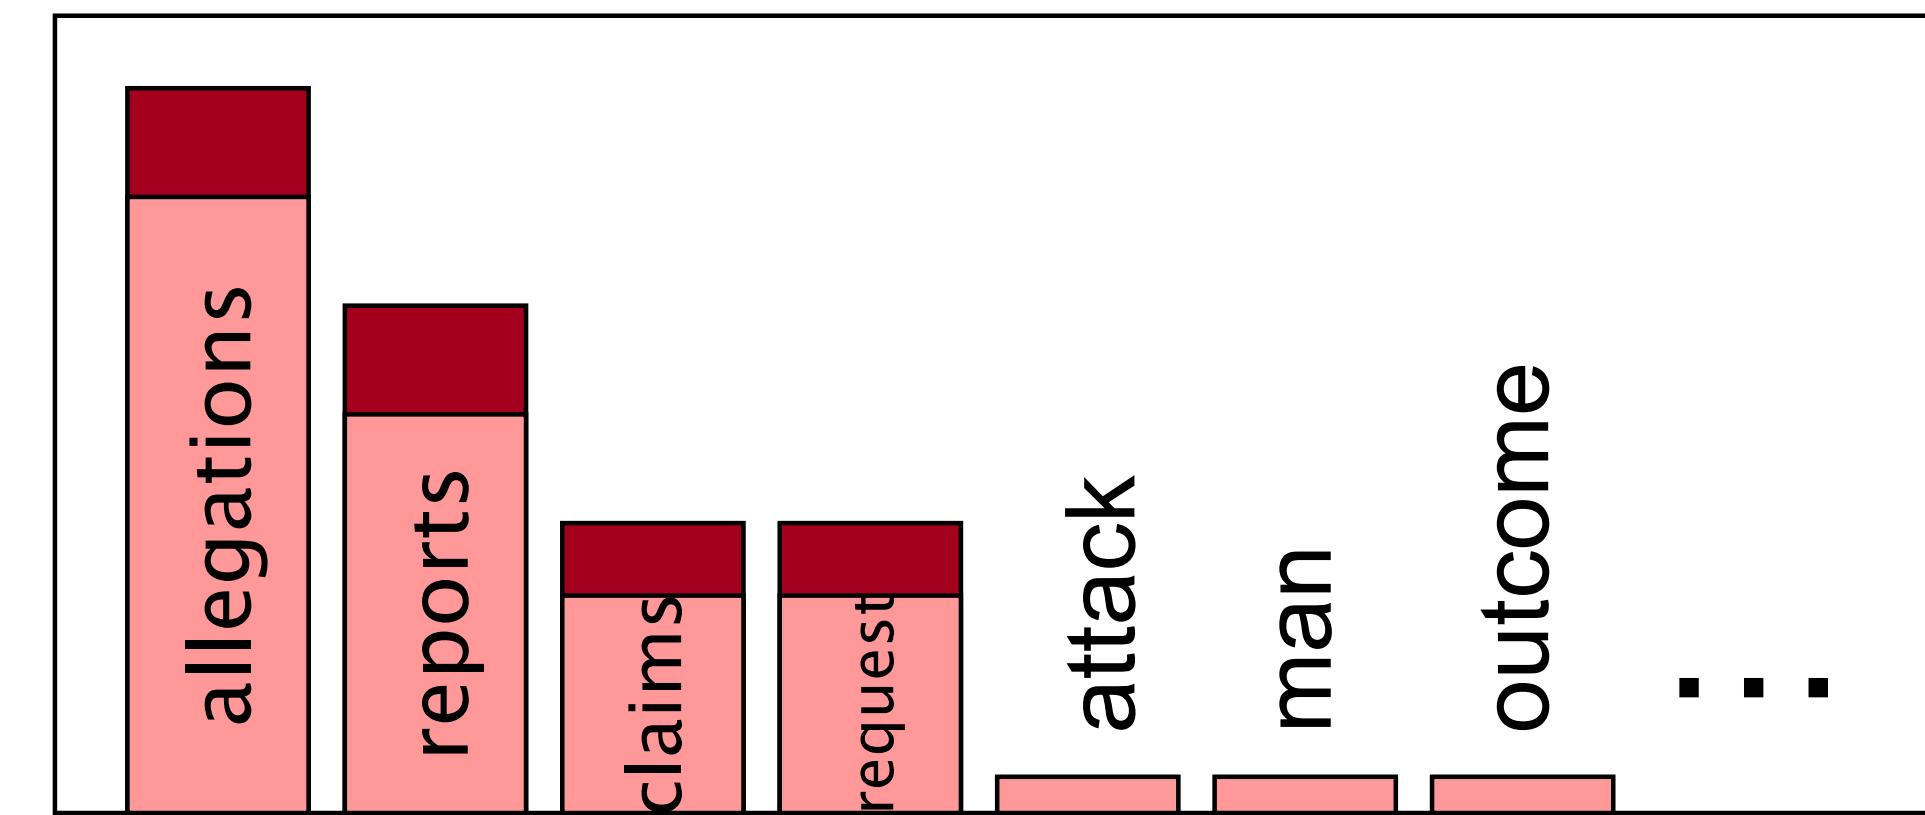
Laplace smoothing

- Also known as add-alpha

- Simplest form of smoothing: Just add α to all counts and renormalize!

- Max likelihood estimate for bigrams:

$$P(w_i | w_{i-1}) = \frac{C(w_{i-1}, w_i)}{C(w_{i-1})}$$

- After smoothing:

$$P(w_i | w_{i-1}) = \frac{C(w_{i-1}, w_i) + \alpha}{C(w_{i-1}) + \alpha | V |}$$

Raw bigram counts
(Berkeley restaurant corpus)

- Out of 9222 sentences

	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

(Slide credit: Dan Jurafsky) 42

Raw bigram counts
(Berkeley restaurant corpus)

	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 to all the entries in the matrix

(Slide credit: Dan Jurafsky) 43

Smoothed bigram probabilities

$$P(w_i|w_{i-1}) = \frac{C(w_{i-1}, w_i) + \alpha}{C(w_{i-1}) + \alpha|V|}$$

$$P^*(w_n|w_{n-1}) = \frac{C(w_{i-1}, w_i) + 1}{C(w_{n-1}) + V}$$

$\alpha = 1$

|	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

Linear Interpolation

$$\hat{P}(w_i \mid w_{i-2}, w_{i-1}) = \lambda_1 P(w_i \mid w_{i-2}, w_{i-1}) \quad \text{Trigram}$$
$$+\lambda_2 P(w_i \mid w_{i-1}) \quad \text{Bigram}$$
$$+\lambda_3 P(w_i) \quad \text{Unigram}$$

$$\sum_i \lambda_i = 1$$

- Use a combination of models to estimate probability

- Strong empirical performance

Why related to sparsity?

- (Assuming all words are seen words) even though bigrams or trigrams may be unseen, unigram probabilities are never zero. So the overall probability is guaranteed to be non-zero.

- And this is achieved while retaining high-order information.

How can we choose lambdas?
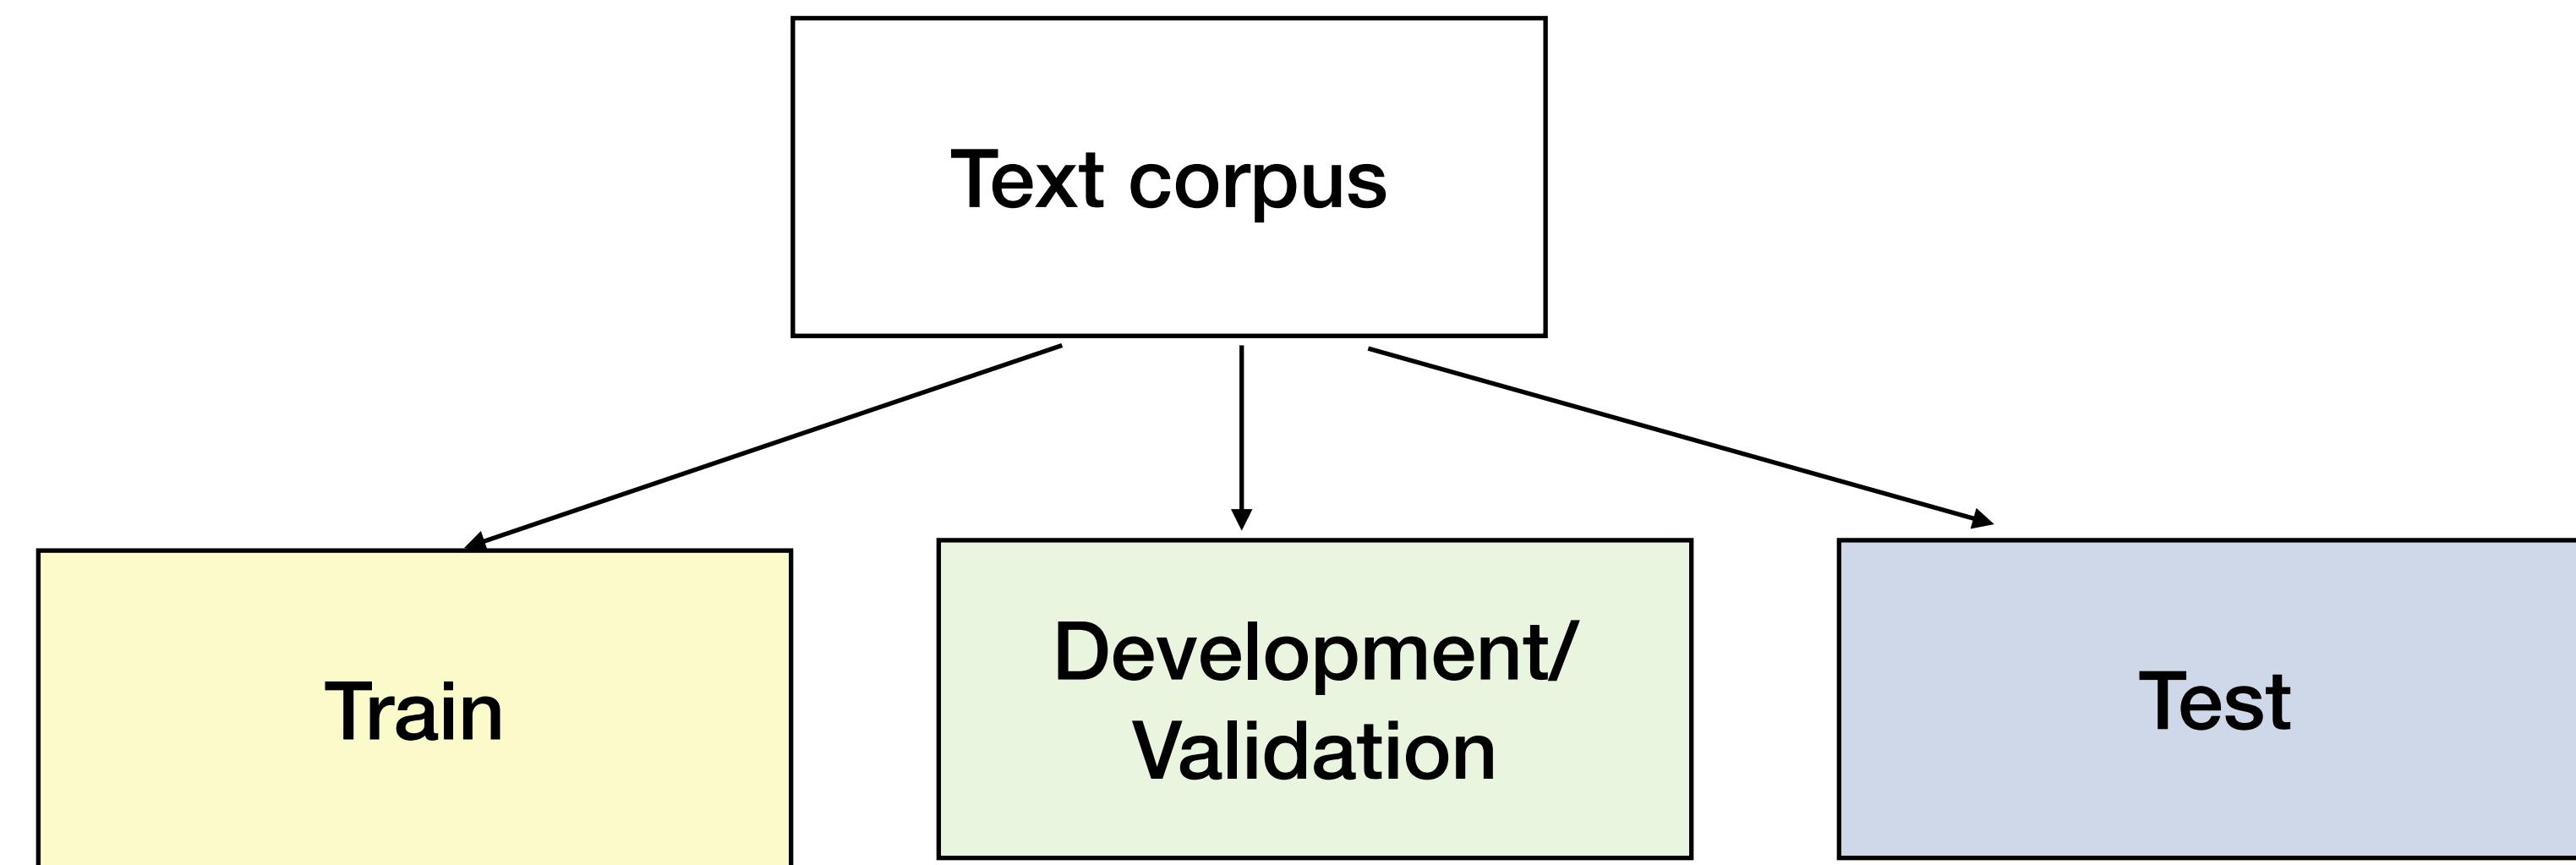
- First, estimate n-gram prob. on training set

- Then, estimate lambdas (hyperparameters) to maximize probability on the held-out development/validation set

- Use best model from above to evaluate on test set

Discounting

- Determine some "mass" to remove from probability estimates

- More explicit method for redistributing mass among unseen n-grams

- Just choose an absolute value to discount (usually <1)

Absolute discounting

- Define Count*(x) = Count(x) - 0.5

- Missing probability mass:

$$\alpha(w_{i-1}) = 1 - \sum_w \frac{\text{Count} * (w_{i-1}, w)}{\text{Count}(w_{i-1})}$$

$$\alpha(\text{the}) = 1 - \frac{43}{48} = 5/48$$

- Divide this mass between words w for which Count(the, w) = 0

x	$\text{Count}(x)$	$\text{Count}^*(x)$	$\frac{\text{Count}^*(x)}{\text{Count}(x)}$
the	48		
the, dog	15	14.5	14.5/48
the, woman	11	10.5	10.5/48
the, man	10	9.5	9.5/48
the, park	5	4.5	4.5/48
the, job	2	1.5	1.5/48
the, telescope	1	0.5	0.5/48
the, manual	1	0.5	0.5/48
the, afternoon	1	0.5	0.5/48
the, country	1	0.5	0.5/48
the, street	1	0.5	0.5/48

Questions?

Acknowledgement