

DSC250: Advanced Data Mining

Language Models

Zhiting Hu

Lecture 8, October 24, 2023

UC San Diego

HALICIOĞLU DATA SCIENCE INSTITUTE

Outline

- Variational Inference (quick overview)
- Language Models

Recap: EM Algorithm

- Observed variables \mathbf{x} , latent variables \mathbf{z}
- To learn a model $p(\mathbf{x}, \mathbf{z}|\theta)$, we want to maximize the marginal log-likelihood

$$\ell(\theta; \mathbf{x}) = \log p(\mathbf{x}|\theta) = \log \sum_{\mathbf{z}} p(\mathbf{x}, \mathbf{z}|\theta)$$

- But it's too difficult
- EM algorithm:
 - maximize a lower bound of $\ell(\theta; \mathbf{x})$
 - Or equivalently, minimize an upper bound of $-\ell(\theta; \mathbf{x})$
- Key equation:

$$\ell(\theta; \mathbf{x}) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}|\theta)}{q(\mathbf{z}|\mathbf{x})} \right] + \text{KL}(q(\mathbf{z}|\mathbf{x}) || p(\mathbf{z}|\mathbf{x}, \theta))$$

Evidence Lower Bound (ELBO)

$$= -F(q, \theta) + \text{KL}(q(\mathbf{z}|\mathbf{x}) || p(\mathbf{z}|\mathbf{x}, \theta))$$

Variational free energy


Recap: EM Algorithm

- The EM algorithm is coordinate-descent on $F(q, \theta)$
 - E-step: $q^{t+1} = \arg \min_q F(q, \theta^t) = p(\mathbf{z}|\mathbf{x}, \theta^t)$
 - the posterior distribution over the latent variables given the data and the current parameters
 - M-step: $\theta^{t+1} = \arg \min_{\theta} F(q^{t+1}, \theta) = \operatorname{argmax}_{\theta} \sum_{\mathbf{z}} q^{t+1}(\mathbf{z}|\mathbf{x}) \log p(\mathbf{x}, \mathbf{z}|\theta)$

$$\begin{aligned} \ell(\theta; \mathbf{x}) &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}|\theta)}{q(\mathbf{z}|\mathbf{x})} \right] + \text{KL}(q(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}|\mathbf{x}, \theta)) \\ &= -F(q, \theta) + \text{KL}(q(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}|\mathbf{x}, \theta)) \end{aligned}$$

Example: Gaussian Mixture Models (GMMs)

- Start: “guess” the centroid μ_k and covariance Σ_k of each of the K clusters
- Loop:



Variational Inference (quick overview)

Inference

- Given a model, the goals of inference can include:
 - Computing the likelihood of observed data $p(\mathbf{x}^*)$
 - Computing the marginal distribution over a given subset of variables in the model $p(\mathbf{x}_A)$
 - Computing the conditional distribution over a subsets of nodes given a disjoint subset of nodes $p(\mathbf{x}_A|\mathbf{x}_B)$
 - Computing a mode of the density (for the above distributions) $\operatorname{argmax}_x p(\mathbf{x})$
 -

Variational Inference

- Observed variables \mathbf{x} , latent variables \mathbf{z}
- Variational (Bayesian) inference, a.k.a. **variational Bayes**, is most often used to **approximately** infer the conditional distribution over the latent variables given the observations (and parameters)
 - i.e., the **posterior distribution** over the latent variables

$$p(\mathbf{z}|\mathbf{x}, \theta) = \frac{p(\mathbf{z}, \mathbf{x}|\theta)}{\sum_{\mathbf{z}} p(\mathbf{z}, \mathbf{x}|\theta)}$$

Motivating Example

- Why do we often need to use an approximate inference methods (such as variational Bayes) to compute the posterior distribution?
- It's because we cannot directly compute the posterior distribution for many interesting models
 - I.e. the posterior density is in an intractable form (often involving integrals) which cannot be easily analytically solved.

Variational Inference

The main idea behind variational inference:

- Choose a family of distributions over the latent variables $z_{1:m}$ with its own set of variational parameters ν , i.e.

$$q(z_{1:m}|\nu)$$

- Then, we find the setting of the parameters that makes our approximation q closest to the posterior distribution.
 - This is where optimization algorithms come in.
- Then we can use q with the fitted parameters in place of the posterior.
 - E.g. to form predictions about future data, or to investigate the posterior distribution over the hidden variables, find modes, etc.

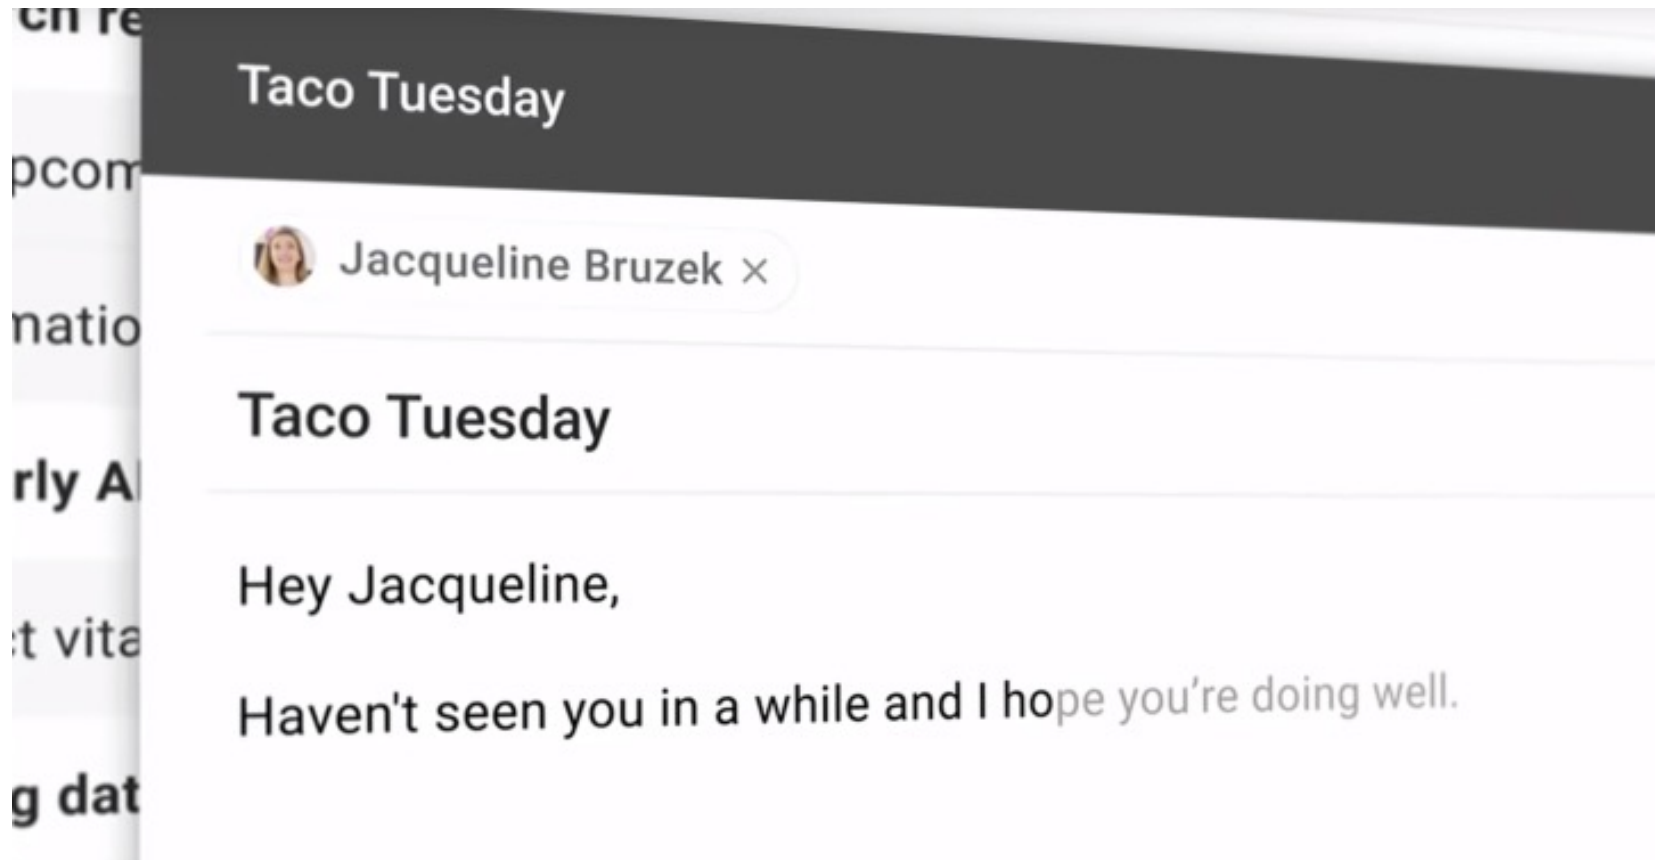
Language Models

Outline

- N-gram language models
- Neural language models
- Neural architectures (in general)

Motivations of Language Models

- Generation



Email auto-completion

Motivations of Language Models

- Generation

Title: United Methodists Agree to Historic Split
Subtitle: Those who oppose gay marriage will form their own denomination
Article: After two days of intense debate, the United Methodist Church has agreed to a historic split - one that is expected to end in the creation of a new denomination, one that will be "theologically and socially conservative," according to The Washington Post. The majority of delegates attending the church's annual General Conference in May voted to strengthen a ban on the ordination of LGBTQ clergy and to write new rules that will "discipline" clergy who officiate at same-sex weddings. But those who opposed these measures have a new plan: They say they will form a separate denomination by 2020, calling their church the Christian Methodist denomination.

The Post notes that the denomination, which claims 12.5 million members, was in the early 20th century the "largest Protestant denomination in the U.S.," but that it has been shrinking in recent decades. The new split will be the second in the church's history. The first occurred in 1968, when roughly 10 percent of the denomination left to form the Evangelical United Brethren Church. The Post notes that the proposed split "comes at a critical time for the church, which has been losing members for years," which has been "pushed toward the brink of a schism over the role of LGBTQ people in the church." Gay marriage is not the only issue that has divided the church. In 2016, the denomination was split over ordination of transgender clergy, with the North Pacific regional conference voting to ban them from serving as clergy, and the South Pacific regional conference voting to allow them.

Figure 3.14: The GPT-3 generated news article that humans had the greatest difficulty distinguishing from a human written article (accuracy: 12%).

Motivations of Language Models

- Evaluation of language fluency
 - Ex:

Successful speech recognition requires generating a word sequence that is:

- ▶ Faithful to the acoustic input
- ▶ Fluent

If we're mapping acoustics a to word sequences w , then:

$$w^* = \operatorname{argmax}_w \text{Faithfulness}(w; a) + \text{Fluency}(w)$$

Language models can provide a “fluency” score.

Motivations of Language Models

- Evaluation of language fluency
 - Ex: Other tasks that have text (or speech) as output:
 - ▶ translation from one language to another
 - ▶ conversational systems
 - ▶ document summarization
 - ▶ image captioning
 - ▶ optical character recognition
 - ▶ spelling and grammar correction

If we're mapping inputs i to word sequences w , then:

$$w^* = \operatorname{argmax}_w \text{Faithfulness}(w; i) + \text{Fluency}(w)$$

Language models can provide a “fluency” score.

Motivations of Language Models

- Few-shot prediction (e.g., GPT3)

Prompts in **Red**. GPT3 responses in **Blue**

Q. How do you draw a bicycle?

A. You draw a bicycle with two circles connected by a line.

Q. How do you draw a ball?

A. You draw a ball with one circle.

Q. How do you draw a house?

A. You draw a house with a trapezium on top of a square.

Q. How do you draw a rabbit?

A. You draw a rabbit with four circles: one large circle for the head, a small one for the nose, a large circle for the body, and two tiny circles for the feet.



Q. How do you draw a snail?

A. You draw a snail with six circles: one large circle for the shell, two small circles for the antennae, one large circle for the head, one small circle for the tail, and a small circle for the mouth.

Notations

- ▶ \mathcal{V} is a finite set of (discrete) symbols (words or characters);
 $V = |\mathcal{V}|$
- ▶ \mathcal{V}^* is the (infinite) set of sequences of symbols from \mathcal{V}

Notations

- ▶ \mathcal{V} is a finite set of (discrete) symbols (words or characters);
 $V = |\mathcal{V}|$
- ▶ \mathcal{V}^* is the (infinite) set of sequences of symbols from \mathcal{V}
- ▶ In language modeling, we imagine a sequence of random variables X_1, X_2, \dots that continues until some X_n takes the value “” (a special end-of-sequence symbol).
- ▶ \mathcal{V}^\dagger is the (infinite) set of sequences of \mathcal{V} symbols, with a single , which is at the end.

The Language Modeling Problem

- Input: training data $\mathbf{x} = (x_1, x_2, \dots, x_N)$ in \mathcal{V}^\dagger
 - (assuming one instance \mathbf{x} for simplicity of notations)
- Output: $p: \mathcal{V}^\dagger \rightarrow \mathbb{R}$
- Think of p as a measure of plausibility

Probabilistic Language Model

- We let p be a probability distribution, which means that

$$\forall \mathbf{x} \in \mathcal{V}^\dagger, p(\mathbf{x}) \geq 0$$

$$\sum_{\mathbf{x} \in \mathcal{V}^\dagger} p(\mathbf{x}) = 1$$

- Advantages:
 - Interpretability
 - We can apply the maximum likelihood principle to build a language model from data

Decomposing using the Chain Rule

$$p(\mathbf{X} = \mathbf{x}) = \left(\begin{array}{l} p(X_1 = x_1) \\ \cdot p(X_2 = x_2 \mid X_1 = x_1) \\ \cdot p(X_3 = x_3 \mid \mathbf{X}_{1:2} = \mathbf{x}_{1:2}) \\ \vdots \\ \cdot p(X_N = \text{○} \mid \mathbf{X}_{1:N-1} = \mathbf{x}_{1:N-1}) \end{array} \right)$$
$$= \prod_{i=1}^N p(X_i = x_i \mid \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1})$$

Example:

$\mathbf{x} = (I, \textit{like}, \textit{this}, \textit{movie}, \dots)$

$p(\mathbf{x}) = \dots p_{\theta}(\textit{like} \mid I) p_{\theta}(\textit{this} \mid I, \textit{like}) \dots$

Predict each word based on the “history”

Unigram Model: Empty History

$$p(\mathbf{X} = \mathbf{x}) = \prod_{i=1}^N p(X_i = x_i \mid \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1})$$

➤ Multinomial distribution

$$\stackrel{\text{assumption}}{=} \prod_{i=1}^N p(X_i = x_i; \boldsymbol{\theta}) = \prod_{i=1}^N \theta_{x_i}$$

Unigram Model: Empty History

$$p(\mathbf{X} = \mathbf{x}) = \prod_{i=1}^N p(X_i = x_i \mid \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1})$$

⇒ Multinomial distribution

$$\stackrel{\text{assumption}}{=} \prod_{i=1}^N p(X_i = x_i; \boldsymbol{\theta}) = \prod_{i=1}^N \theta_{x_i}$$

Maximum likelihood estimate: for every $v \in \mathcal{V}$,

$$\begin{aligned} \theta_v^* &= \frac{\sum_{i=1}^N \mathbf{1}\{x_i = v\}}{N} \\ &= \frac{\text{count}_{\mathbf{x}}(v)}{N} \end{aligned}$$

Example

The probability of

Presidents tell lies .

is:

$$p(X_1 = \text{Presidents}) \cdot p(X_2 = \text{tell}) \cdot p(X_3 = \text{lies}) \cdot p(X_4 = \text{.}) \cdot p(X_5 = \text{⬡})$$

In unigram model notation:

$$\theta_{\text{Presidents}} \cdot \theta_{\text{tell}} \cdot \theta_{\text{lies}} \cdot \theta_{\text{.}} \cdot \theta_{\text{⬡}}$$

Using the maximum likelihood estimate for θ , we could calculate:

$$\frac{\text{count}_x(\text{Presidents})}{N} \cdot \frac{\text{count}_x(\text{tell})}{N} \dots \frac{\text{count}_x(\text{⬡})}{N}$$

Unigram Models: Assessment

Pros:

- ▶ Easy to understand
- ▶ Cheap
- ▶ Good enough for information retrieval (maybe)

Cons:

- ▶ Fixed, known vocabulary assumption
- ▶ “Bag of words” assumption is linguistically inaccurate
 - ▶ $p(\text{the the the the}) \gg p(\text{I want ice cream})$

n-gram Models

$$p(\mathbf{X} = \mathbf{x}) = \prod_{i=1}^N p(X_i = x_i \mid \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1})$$

$$\stackrel{\text{assumption}}{=} \prod_{i=1}^N p(X_i = x_i \mid X_{i-n+1:i-1} = \mathbf{x}_{i-n+1:i-1}; \boldsymbol{\theta})$$

$$= \prod_{i=1}^N \theta_{x_i \mid \mathbf{x}_{i-n+1:i-1}}$$

n-gram Models

$$\begin{aligned} p(\mathbf{X} = \mathbf{x}) &= \prod_{i=1}^N p(X_i = x_i \mid \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1}) \\ &\stackrel{\text{assumption}}{=} \prod_{i=1}^N p(X_i = x_i \mid X_{i-n+1:i-1} = \mathbf{x}_{i-n+1:i-1}; \boldsymbol{\theta}) \\ &= \prod_{i=1}^N \theta_{x_i \mid \mathbf{x}_{i-n+1:i-1}} \end{aligned}$$

(n - 1)th-order Markov assumption \equiv n-gram model

- ▶ Unigram model is the $n = 1$ case
- ▶ For a long time, trigram models ($n = 3$) were widely used
- ▶ 5-gram models ($n = 5$) were common in MT for a time

n-gram Models

- Maximum likelihood estimate for the n-gram model's probability of v given a $(n - 1)$ -length history \mathbf{h}

$$\begin{aligned}\theta_{v|\mathbf{h}} &= p(X_i = v \mid \mathbf{X}_{i-n+1:i-1} = \mathbf{h}) \\ &= \frac{p(X_i = v, \mathbf{X}_{i-n+1:i-1} = \mathbf{h})}{p(\mathbf{X}_{i-n+1:i-1} = \mathbf{h})} \\ &= \frac{\text{count}_{\mathbf{x}}(\mathbf{h}v)}{N} \bigg/ \frac{\text{count}_{\mathbf{x}}(\mathbf{h})}{N} \\ &= \frac{\text{count}_{\mathbf{x}}(\mathbf{h}v)}{\text{count}_{\mathbf{x}}(\mathbf{h})}\end{aligned}$$

Choosing n is a Balancing Act

If n is too small, your model can't learn very much about language.

As n gets larger:

- ▶ The number of parameters grows with $O(V^n)$.
- ▶ Most n -grams will never be observed, so you'll have lots of zero probability n -grams. This is an example of **data sparsity**.
- ▶ Your model depends increasingly on the training data; you need (lots) more data to learn to generalize well.

This is a beautiful illustration of the bias-variance tradeoff.

Other “tricks”

- Smoothing

The game: prevent $\theta_{v|h} = 0$ for any v and h , while keeping $\sum_{\mathbf{x}} p(\mathbf{x}) = 1$ so that perplexity stays meaningful.

- ▶ Simple method: add $\lambda > 0$ to every count (including counts of zero) before normalizing (the textbook calls this “Lidstone” smoothing)

- Dealing with Out-of-Vocabulary Terms

- Define a special OOV or “unknown” symbol **unk**. Transform some (or all) rare words in the training data to **unk**.
- Build a language model at the character level.
- Some new methods use data-driven, deterministic tokenization schemes that segment some words into smaller parts to reduce the effective vocabulary size (Sennrich et al., 2016; Wu et al., 2016).

n-gram Models: Assessment

Pros:

- ▶ Easy to understand
- ▶ Cheap (with modern hardware; Lin and Dyer, 2010)
- ▶ Fine in some applications and when training data is scarce

Cons:

- ▶ Fixed, known vocabulary assumption
- ▶ Markov assumption is linguistically inaccurate
 - ▶ (But not as bad as unigram models!)
- ▶ Data sparseness problem

Neural Language Models

Neural Language Models

Instead of a lookup for a word and fixed-length history $(\theta_{v|h})$, define a vector function:

$$p(X_i | \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1}) = \mathbf{NN}(\mathbf{enc}(\mathbf{x}_{1:i-1}); \boldsymbol{\theta})$$

where $\boldsymbol{\theta}$ do the work of *encoding* the history and *transforming* it into a distribution over the next word.

Neural Language Models

Instead of a lookup for a word and fixed-length history $(\theta_{v|h})$, define a vector function:

$$p(X_i | \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1}) = \mathbf{NN}(\mathbf{enc}(\mathbf{x}_{1:i-1}); \boldsymbol{\theta})$$

where $\boldsymbol{\theta}$ do the work of *encoding* the history and *transforming* it into a distribution over the next word.

The transformation is described as a composed series of simple transformations or “layers.”

Neural Network

Formally, it's a function \mathbf{NN} from θ (learned parameters) and inputs to outputs, all of which are real-valued vectors (or matrices, or tensors, or collections of them).

Almost always, \mathbf{NN} is differentiable with respect to θ and nonlinear with respect to the data input.

- ▶ “Nonlinear” means there does **not** exist a matrix \mathbf{A} such that $\mathbf{NN}(\mathbf{v}; \theta) = \mathbf{A}\mathbf{v}$, for all \mathbf{v} .

Neural Language Models

Instead of a lookup for a word and fixed-length history $(\theta_{v|h})$, define a vector function:

$$p(X_i | \mathbf{X}_{1:i-1} = \mathbf{x}_{1:i-1}) = \mathbf{NN}(\mathbf{enc}(\mathbf{x}_{1:i-1}); \boldsymbol{\theta})$$

where $\boldsymbol{\theta}$ do the work of *encoding* the history and *transforming* it into a distribution over the next word.

The transformation is described as a composed series of simple transformations or “layers.”

- We first map word histories \mathbf{h} to vectors/matrices
- We interpret the output as $p(X_i | \mathbf{X}_{1:i-1} = \mathbf{h})$

Two Key Components

- “Embedding” words as vectors
- Layering to increase capacity (i.e., the set of distributions that can be represented).

“One Hot” Vectors

Let $\mathbf{e}_i \in \mathbb{R}^V$ be the i th column of the identity matrix \mathbf{I} .

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}; \quad \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \\ 0 \end{bmatrix}; \quad \dots; \quad \mathbf{e}_V = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}$$

\mathbf{e}_i is the “one hot” vector for the i th word in \mathcal{V} .

A neural language model starts by “looking up” each word by multiplying its one hot vector by a matrix \mathbf{M} ; $\mathbf{e}_v^\top \mathbf{M} = \mathbf{m}_v$, the “embedding” of v .

\mathbf{M} becomes part of the parameters (θ).

Sequences of Word Vectors

Given a word sequence $\langle v_1, v_2, \dots, v_k \rangle$, we transform it into a sequence of word vectors,

$$\mathbf{m}_{v_1}, \mathbf{m}_{v_2}, \dots, \mathbf{m}_{v_k}$$

Adding Layers

- Neural networks are built by composing functions, a mix of
 - Affine, $\mathbf{v}' = \mathbf{W}\mathbf{v} + \mathbf{b}$ (note that the dimensionality of \mathbf{v} and \mathbf{v}' might be different)
 - Nonlinearity, e.g.,
 - rectified linear ("relu") units $v'_i = \max(0, v_i)$
 - elementwise hyperbolic tangent $v'_i = \tanh(v_i) = \frac{e^{v_i} - e^{-v_i}}{e^{v_i} + e^{-v_i}}$
 - softmax $v'_i = \exp\{v_i\} / \sum_j \exp\{v_j\}$
 - More complex components (composed of the above operations):
 - Convolutional layers
 - Recurrent NNs
 - Attention

Summary so far

- language models utilities
 - Generation, evaluation of fluency, few-shot prediction (GPT3), ...
- N-gram language models
 - Unigram LM
 - N-gram LM
- Neural language models:
 - Embedding: one-hot vectors -> embedding vectors
 - Neural networks

Questions?