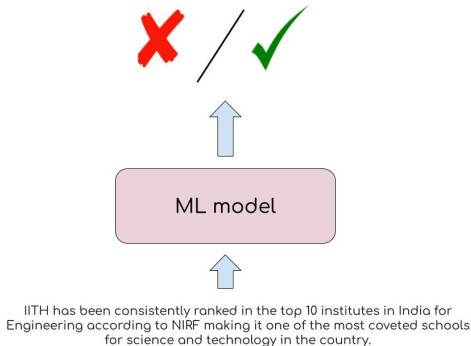


Deep Learning

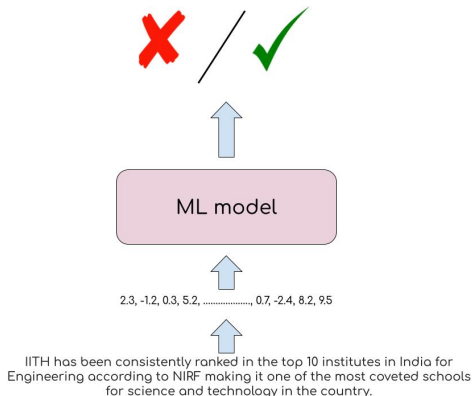
13 Word Embeddings

Dr. Konda Reddy Mopuri
Dept. of AI, IIT Hyderabad
Jan-May 2024

Why Word Embeddings?



Why Word Embeddings?



- Corpus: a collection of authentic text organized into dataset

Terminology

- Corpus: a collection of authentic text organized into dataset
- Vocabulary (V): Set of allowed words

Terminology

- Corpus: a collection of authentic text organized into dataset
- Vocabulary (V): Set of allowed words
- **Target**: Representation for every word in V

One-hot Encoding

- Representation using discrete symbols

One-hot Encoding

- Representation using discrete symbols
- $|V|$ words encoded as binary vectors of length $|V|$

Dictionary

Word Representation

A

| | | | | | |
|---|---|---|-------|---|---|
| 1 | 0 | 0 | | 0 | 0 |
|---|---|---|-------|---|---|

Bus

| | | | | | |
|---|---|---|-------|---|---|
| 0 | 1 | 0 | | 0 | 0 |
|---|---|---|-------|---|---|

Cat

| | | | | | |
|---|---|---|-------|---|---|
| 0 | 0 | 1 | | 0 | 0 |
|---|---|---|-------|---|---|

⋮

Tide

| | | | | | |
|---|---|---|-------|---|---|
| 0 | 0 | 0 | | 1 | 0 |
|---|---|---|-------|---|---|

Zone

| | | | | | |
|---|---|---|-------|---|---|
| 0 | 0 | 0 | | 0 | 1 |
|---|---|---|-------|---|---|

One-hot encoding: Drawbacks

- ① Space inefficient (e.g. 13M words in Google 1T corpus)

One-hot encoding: Drawbacks

- ① Space inefficient (e.g. 13M words in Google 1T corpus)
- ② No notion of similarity (or, distance) between words

One-hot encoding: Drawbacks

- ① Space inefficient (e.g. 13M words in Google 1T corpus)
- ② No notion of similarity (or, distance) between words
 - 'Dog' is as close to 'Cat' as it is to 'Machine'

Notion of Meaning for words

- ① What is a good notion of meaning for a word?

Notion of Meaning for words

- ① What is a good notion of meaning for a word?
- ② How do we, humans, know the meaning of a word?

Notion of Meaning for words

① What does **silla** mean?

Notion of Meaning for words

- ① Let's see how this word is used in different contexts
1. The **silla** is by the window, offering a nice view of the garden.
 2. Can you pass me that **silla** so I can join the conversation?
 3. After the event, please stack the **sillas** neatly against the wall.
 4. I found a comfortable **silla** in the corner and settled down to relax.

Notion of Meaning for words

① Does this context help you understand the word **silla**?

1. The **silla** is by the window, offering a nice view of the garden.
2. Can you pass me that **silla** so I can join the conversation?
3. After the event, please stack the **sillas** neatly against the wall.
4. I found a comfortable **silla** in the corner and settled down to relax.

Notion of Meaning for words

- ① Does this context help you understand the word **silla**?
 - ② { positioned near a window or against a wall or in the corner, used for conversing/events, can be used to relax }
1. The **silla** is by the window, offering a nice view of the garden.
 2. Can you pass me that **silla** so I can join the conversation?
 3. After the event, please stack the **sillas** neatly against the wall.
 4. I found a comfortable **silla** in the corner and settled down to relax.

Notion of Meaning for words

① How did we do that?

Notion of Meaning for words

- ① How did we do that?
- ② “We searched for other words that can be used in the same contexts, found some, and made a conclusion that **silla** has to mean similar to those words.”

Notion of Meaning for words

- ① Distributional Hypothesis: Words that frequently appear in similar contexts have a similar meaning

- ① Representation/meaning of a word should consider its context in the corpus

Distributed Representations

- ① Representation/meaning of a word should consider its context in the corpus
- ② Use many contexts of a word to build up a representation for it

① **Co-occurrence matrix** is a way to can capture this!

- size: ($\#words \times \#words$)
- rows: words (m), cols: context (n)
- words and context can be of same or different size

- ① **Co-occurrence matrix** is a way to can capture this!
 - size: ($\#words \times \#words$)
 - rows: words (m), cols: context (n)
 - words and context can be of same or different size
- ② Context can be defined as a 'h' word neighborhood

- ① **Co-occurrence matrix** is a way to can capture this!
 - size: ($\#words \times \#words$)
 - rows: words (m), cols: context (n)
 - words and context can be of same or different size
- ② Context can be defined as a 'h' word neighborhood
- ③ Each row (column): vectorial representation of the word (context)

Co-occurrence matrix

$$X = \begin{matrix} & \begin{matrix} I & like & enjoy & deep & learning & NLP & flying & . \end{matrix} \\ \begin{matrix} I \\ like \\ enjoy \\ deep \\ learning \\ NLP \\ flying \\ . \end{matrix} & \begin{bmatrix} 0 & 2 & 1 & 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \end{bmatrix} \end{matrix}$$

Co-occurrence matrix

① Very sparse

Co-occurrence matrix

- ① Very sparse
- ② Very high-dimensional (grows with the vocabulary size)

Co-occurrence matrix

- ① Very sparse
- ② Very high-dimensional (grows with the vocabulary size)
- ③ **Solution:** Dimensionality reduction (SVD)!

SVD on the Co-occurrence matrix

$$\textcircled{1} X = U\Sigma V^T$$

SVD on the Co-occurrence matrix

① $X = U\Sigma V^T$

②
$$\begin{bmatrix} X \\ \uparrow \dots \uparrow \\ u_1 \dots u_k \\ \downarrow \dots \downarrow \end{bmatrix}_{m \times n} = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_k \end{bmatrix}_{k \times k} \cdot \begin{bmatrix} \leftarrow v_1^T \rightarrow \\ \vdots \\ \leftarrow v_k^T \rightarrow \end{bmatrix}_{k \times n}$$

SVD on the Co-occurrence matrix

① $X = U\Sigma V^T$

②
$$\begin{bmatrix} X \\ \end{bmatrix}_{m \times n} = \begin{bmatrix} \uparrow & \dots & \uparrow \\ u_1 & \dots & u_k \\ \downarrow & \dots & \downarrow \end{bmatrix}_{m \times k} \cdot \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_k \end{bmatrix}_{k \times k} \cdot \begin{bmatrix} \leftarrow & v_1^T & \rightarrow \\ & \vdots & \\ \leftarrow & v_k^T & \rightarrow \end{bmatrix}_{k \times n}$$

③ $X = \sigma_1 u_1 v_1^T + \sigma_2 u_2 v_2^T + \dots + \sigma_k u_k v_k^T$

④ $\hat{X} = \sum_{i=1}^{d < k} \sigma_i u_i v_i^T$ is a d -rank approximation of X

SVD on the Co-occurrence matrix

- ① Before the SVD, representations were the rows of X

SVD on the Co-occurrence matrix

- ① Before the SVD, representations were the rows of X
- ② How do we reduce the representation size with SVD ?

SVD on the Co-occurrence matrix

- ① Before the SVD, representations were the rows of X
- ② How do we reduce the representation size with SVD ?
- ③ $W_{\text{word}} = U_{m \times k} \cdot \Sigma_{k \times k}$

SVD on the Co-occurrence matrix

- ① $W_{\text{word}} \in \mathbb{R}^{m \times k}$ ($k \ll |V| = m$) are considered the representation of the words

SVD on the Co-occurrence matrix

- ① $W_{\text{word}} \in \mathbb{R}^{m \times k}$ ($k \ll |V| = m$) are considered the representation of the words
- ② Lesser dimensions but the same similarities! (one may verify that $XX^T = \hat{X}\hat{X}^T$)

SVD on the Co-occurrence matrix

- ① $W_{\text{word}} \in \mathbb{R}^{m \times k}$ ($k \ll |V| = m$) are considered the representation of the words
- ② Lesser dimensions but the same similarities! (one may verify that $XX^T = \hat{X}\hat{X}^T$)
- ③ $W_{\text{context}} = V \in \mathbb{R}^{n \times k}$ are taken as the representations for the context words

A bit more clever things...

- ① Entries in the occurrence matrix can be weighted (HAL model)

A bit more clever things...

- ① Entries in the occurrence matrix can be weighted (HAL model)
- ② Better associations can be used (PPMI)

A bit more clever things...

- ① Entries in the occurrence matrix can be weighted (HAL model)
- ② Better associations can be used (PPMI)
- ③

Count-based vs prediction-based models



- ① Techniques we have seen so far rely on the counts (or, co-occurrence of words)

Count-based vs prediction-based models




- ① Techniques we have seen so far rely on the counts (or, co-occurrence of words)
- ② Next, we see prediction based models for word embeddings


① T Mikolov et al. (2013)

- ① T Mikolov et al. (2013)
- ② Two versions: Predict words from the contexts (or contexts from words)

- ① T Mikolov et al. (2013)
- ② Two versions: Predict words from the contexts (or contexts from words)
- ③ Continuous Bag of Words (CBoW) and Skip-gram




CBOW




Skip-gram

Word Embeddings: Skip-gram



Skip-gram

Word Embeddings: Skip-gram



Word Embeddings: Skip-gram

- ① Start: huge corpus and random initialization of the word embeddings

Word Embeddings: Skip-gram

- ① Start: huge corpus and random initialization of the word embeddings
- ② Process the text with a sliding window (one word at a time)

Word Embeddings: Skip-gram

- ① Start: huge corpus and random initialization of the word embeddings
- ② Process the text with a sliding window (one word at a time)

- ① At each step, there is a central word and context words (other words in the window)

Word Embeddings: Skip-gram

- ① Start: huge corpus and random initialization of the word embeddings
- ② Process the text with a sliding window (one word at a time)

- ① At each step, there is a central word and context words (other words in the window)
- ② Given the central word, compute the probabilities for the context words

Word Embeddings: Skip-gram

- ① Start: huge corpus and random initialization of the word embeddings
- ② Process the text with a sliding window (one word at a time)

- ① At each step, there is a central word and context words (other words in the window)
- ② Given the central word, compute the probabilities for the context words
- ③ Modify the word embeddings to increase these probabilities

Word Embeddings: Skip-gram




Figure from [Lena Voita](#)

Word Embeddings: Skip-gram



Figure from [Lena Voita](#)

Word Embeddings: Skip-gram




Figure from [Lena Voita](#)

Word Embeddings: Skip-gram




Figure from [Lena Voita](#)

Word Embeddings: Skip-gram




Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

- For each position in $t = 1, 2, \dots, T$ in the corpus, Skip-gram predicts the context words in m -sized window (θ is the variables to be optimized)

$$\text{Likelihood } L(\theta) = \prod_{t=1}^T \prod_{-m \leq j \leq m, j \neq 0} P(w_{t+j} | w_t, \theta)$$

- The loss is mean NLL

$$\text{Loss } J(\theta) = -\frac{1}{T} \prod_{t=1}^T \prod_{-m \leq j \leq m, j \neq 0} \log P(w_{t+j} | w_t, \theta)$$

Word Embeddings: Skip-gram

- What are the parameters (θ) to be learned?

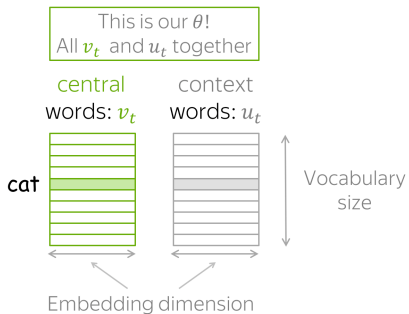


Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

- How to compute $P(w_{t+j}|w_t, \theta)$?

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

Dot product: measures similarity of o and c
Larger dot product = larger probability

Normalize over entire vocabulary
to get probability distribution

Figure from [Lena Voita](#)

Word Embeddings: Skip-gram



Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

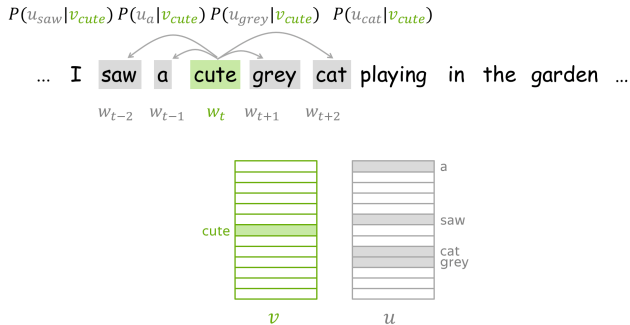


Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

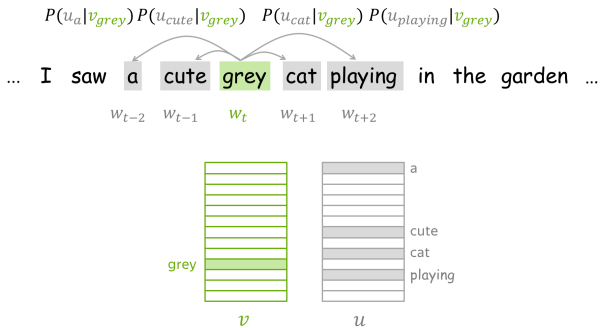


Figure from [Lena Voita](#)

Word Embeddings: Skip-gram



Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

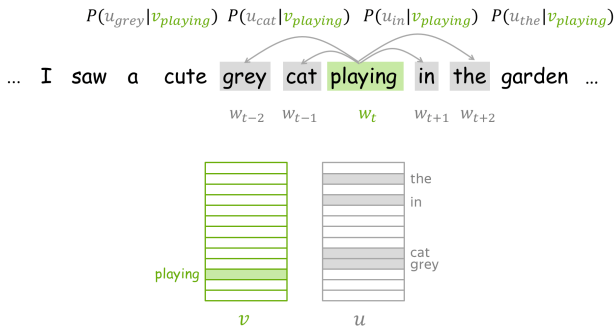


Figure from Lena Voita

Word Embeddings: Skip-gram

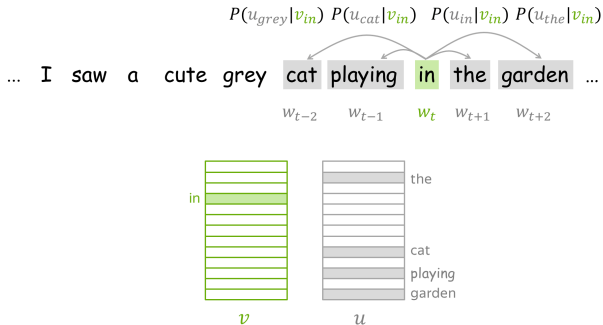


Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

- Train using Gradient Descent

Word Embeddings: Skip-gram


- Train using Gradient Descent
- For one word at a time, i.e., (a center word, one of the context words)

Word Embeddings: Skip-gram

- Train using Gradient Descent
- For one word at a time, i.e., (a center word, one of the context words)
- $J_{t,j}(\theta) = -\log P(\text{cute}|\text{cat}) = -\log \frac{\exp u_{\text{cute}}^T v_{\text{cat}}}{\sum_{w \in V_{oc}} \exp u_w^T v_{\text{cat}}} =$
 $-u_{\text{cute}}^T v_{\text{cat}} + \log \sum_{w \in V_{oc}} \exp u_w^T v_{\text{cat}}$

Word Embeddings: Skip-gram

1. Take dot product of v_{cat} with all u
2. exp
3. sum all



4. get loss (for this one step)
5. evaluate the gradient, make an update

$$J_{t,j}(\theta) = \underbrace{-u_{cute}^T v_{cat}}_{(1)} + \log \underbrace{\sum_{w \in V} \exp(u_w^T v_{cat})}_{(2)}$$

$$v_{cat} := v_{cat} - \alpha \frac{\partial J_{t,j}(\theta)}{\partial v_{cat}}$$

$$u_w := u_w - \alpha \frac{\partial J_{t,j}(\theta)}{\partial u_w} \quad \forall w \in V$$

Figure from [Lena Voita](#)

Word Embeddings: Skip-gram

- Training is slow (for each central word, all the context words need to be updated)

Word Embeddings: Skip-gram

- Training is slow (for each central word, all the context words need to be updated)
- Negative sampling: not all the context words are considered, but a random sample of them

Word Embeddings: Skip-gram

- Training is slow (for each central word, all the context words need to be updated)
- Negative sampling: not all the context words are considered, but a random sample of them
- Training over a large corpus leads to sufficient updates for each vector

Word Embeddings: Skip-gram


Dot product of v_{cat} :

- with u_{cute} - increase,
- with all other u - decrease



Dot product of v_{cat} :

- with u_{cute} - increase,
- with a subset of other u - decrease



Parameters to be updated:

- v_{cat}
- u_w for all w in the vocabulary $|V| + 1$ vectors

Negative samples: randomly selected K words




Parameters to be updated:

- v_{cat}
- u_{cute} and u_w for w in K negative examples $K + 2$ vectors


Figure from **Lena Voita**

Word Embeddings: Skip-gram




Can be viewed as a Neural Network

Word Embeddings: Skip-gram




Can be viewed as a Neural Network

Word Embeddings: Skip-gram



Can be viewed as a Neural Network

Word Embeddings: Skip-gram



Can be viewed as a Neural Network

Word Embeddings: Skip-gram

① $W_{N \times m}$ is the W_{word} (used for representing the words)

Word Embeddings: Skip-gram

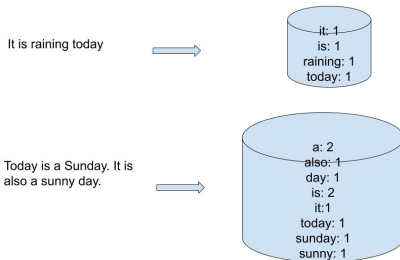
- ① $W_{N \times m}$ is the W_{word} (used for representing the words)
- ② $W'_{m \times N}$ is the W_{context} (may be ignored after the training)

Word Embeddings: Skip-gram

- ① $W_{N \times m}$ is the W_{word} (used for representing the words)
- ② $W'_{m \times N}$ is the W_{context} (may be ignored after the training)
- ③ Some showed averaging word and context vectors may be more beneficial


Bag of Words (BoW)

① Bag of Words: Collection and frequency of words




- ① Considers the embeddings of 'h' words before and 'h' words after the target word

- ① Considers the embeddings of 'h' words before and 'h' words after the target word
- ② Adds them (order is lost) for predicting the target word



The dog slept on couch



- ① Size of the vocabulary = m

- ① Size of the vocabulary = m
- ② Dimension of the embeddings = N

Word Embeddings: CBoW

- ① Input layer $W_{N \times m}$ (embeddings for the context words) projects the context (sum of 1-hot vectors of all the context vectors) into N -dim space

$$\begin{matrix} & \text{context} \\ \left(\begin{matrix} W_{N \times m} \end{matrix} \right) & \left(\begin{matrix} C_{m \times 1} \end{matrix} \right) \end{matrix}$$

Word Embeddings: CBoW

- ① Input layer $W_{N \times m}$ (embeddings for the context words) projects the context (sum of 1-hot vectors of all the context vectors) into N -dim space
- ② Representations of all the ($2h$) words in the context are summed (result is an N -dim context vector)

$$\begin{matrix} & \text{context} \\ \left(\begin{matrix} W_{N \times m} \end{matrix} \right) & \left(\begin{matrix} C_{m \times 1} \end{matrix} \right) \end{matrix}$$

Word Embeddings: CBoW

- ① Input layer $W_{N \times m}$ (embeddings for the context words) projects the context (sum of 1-hot vectors of all the context vectors) into N -dim space

$$\begin{matrix} & \text{context} \\ \left(\begin{matrix} W_{N \times m} \end{matrix} \right) & \left(\begin{matrix} C_{m \times 1} \end{matrix} \right) & \rightarrow & \left(\begin{matrix} E_{N \times 1} \end{matrix} \right)\end{matrix}$$

Word Embeddings: CBoW

- ① Input layer $W_{N \times m}$ (embeddings for the context words) projects the context (sum of 1-hot vectors of all the context vectors) into N -dim space
- ② Representations of all the ($2h$) words in the context are summed (result is an N -dim context vector)

$$\begin{matrix} & \text{context} \\ \left(\begin{matrix} W_{N \times m} \end{matrix} \right) \left(\begin{matrix} C_{m \times 1} \end{matrix} \right) \Rightarrow \left(\begin{matrix} E_{N \times 1} \end{matrix} \right)\end{matrix}$$

Word Embeddings: CBoW

- ① Next layer has a weight matrix $W'_{m \times N}$ (embeddings for the center words)

Word Embeddings: CBoW

- ① Next layer has a weight matrix $W'_{m \times N}$ (embeddings for the center words)
- ② Projects the accumulated embeddings onto the vocabulary

$$\begin{matrix} \left(\begin{matrix} W_{N \times m} \end{matrix} \right) \left(\begin{matrix} C_{m \times 1} \end{matrix} \right) \Rightarrow \left(\begin{matrix} W'_{m \times N} \end{matrix} \right) \left(\begin{matrix} E_{N \times 1} \end{matrix} \right) \\ \text{First layer} \qquad \qquad \qquad \text{Second layer} \end{matrix}$$

Word Embeddings: CBoW


- ① Next layer has a weight matrix $W'_{m \times N}$ (embeddings for the center words)
- ② Projects the accumulated embeddings onto the vocabulary

$$\begin{pmatrix} W_{N \times m} \end{pmatrix} \begin{pmatrix} C_{m \times 1} \end{pmatrix} \Rightarrow \begin{pmatrix} W'_{m \times N} \end{pmatrix} \begin{pmatrix} E_{N \times 1} \end{pmatrix} \Rightarrow \begin{pmatrix} C_{m \times 1} \end{pmatrix}$$

Scores for m-way classification

Word Embeddings: CBoW

- ① m -way classification \rightarrow (after a softmax) maximizes the probability for the target word



Word Embeddings: CBoW

① $W_{N \times m}$ is the W_{context}

Word Embeddings: CBoW

- ① $W_{N \times m}$ is the W_{context}
- ② $W'_{m \times N}$ is the W_{words}

① Glove - Global Vectors

- ① Glove - Global Vectors
- ② Combines the score-based and predict-based approaches

- ① X_{ij} in the co-occurrence matrix encodes the global info. about words i and j

- ① X_{ij} in the co-occurrence matrix encodes the global info. about words i and j
 - $p(j/i) = \frac{X_{ij}}{X_i}$
- ② Glove attempts to learn representations that are faithful to the co-occurrence info.

- ① Glove attempts to learn representations that are faithful to the co-occurrence info.

- ① Glove attempts to learn representations that are faithful to the co-occurrence info.
- ② Try to enforce $v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$
 - v_i - central representation of word i , c_j - context representation of word j

- ① Glove attempts to learn representations that are faithful to the co-occurrence info.
- ② Try to enforce $v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$
 - v_i - central representation of word i , c_j - context representation of word j
- ③ Similarly, $v_j^T c_i = \log P(i/j) = \log X_{ij} - \log X_j$ (aim is to learn such embeddings v_i and c_i)

- ① To realize the exchange symmetry of
$$v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$$

- ① To realize the exchange symmetry of
$$v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$$
 - we may capture the $\log X_i$ as a bias b_i of the word w_i

- ① To realize the exchange symmetry of
- $$v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$$
- we may capture the $\log X_i$ as a bias b_i of the word w_i
 - And, an additional term \tilde{b}_j

- ① To realize the exchange symmetry of
$$v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$$
 - we may capture the $\log X_i$ as a bias b_i of the word w_i
 - And, an additional term \tilde{b}_j
- ②
$$v_i^T c_j + b_i + \tilde{b}_j = \log X_{ij}$$

- ① To realize the exchange symmetry of $v_i^T c_j = \log P(j/i) = \log X_{ij} - \log X_i$
 - we may capture the $\log X_i$ as a bias b_i of the word w_i
 - And, an additional term \tilde{b}_j
- ② $v_i^T c_j + b_i + \tilde{b}_j = \log X_{ij}$
- ③ Since $\log X_i$ and $\log X_j$ depend on the words i and j , they can be considered as the word specific biases (learnable)

① Learning objective becomes

$$\operatorname{argmin}_{v_i, c_j, b_i, \tilde{b}_j} J() = \sum_{i,j} (v_i^T c_j + b_i + \tilde{b}_j - \log X_{ij})^2$$

① Learning objective becomes

$$\operatorname{argmin}_{v_i, c_j, b_i, \tilde{b}_j} J() = \sum_{i,j} (v_i^T c_j + b_i + \tilde{b}_j - \log X_{ij})^2$$

- ① Learning objective becomes

$$\operatorname{argmin}_{v_i, c_j, b_i, \tilde{b}_j} J() = \sum_{i,j} (v_i^T c_j + b_i + \tilde{b}_j - \log X_{ij})^2$$

- ② Much of the entries in the co-occurrence matrix are zeros (noisy or less informative)

- ① Learning objective becomes

$$\operatorname{argmin}_{v_i, c_j, b_i, \tilde{b}_j} J() = \sum_{i,j} (v_i^T c_j + b_i + \tilde{b}_j - \log X_{ij})^2$$


- ② Much of the entries in the co-occurrence matrix are zeros (noisy or less informative)
- ③ Suggests to apply a weight

$$J(\theta) = \sum_{w,c \in V} \underbrace{f(N(w,c))}_{\text{weighting function}} \cdot (u_c^T v_w + b_c + \bar{b}_w - \log N(w,c))^2$$

context vector
word vector
bias terms (also learned)

Weighting function to:

- penalize rare events
- not to over-weight frequent events



$$\begin{cases} (x/x_{max})^\alpha & \text{if } x < x_{max}, \\ 1 & \text{otherwise.} \end{cases}$$

$$\alpha = 0.75, x_{max} = 100$$

Figure from [Lena Voita](#)

Evaluating the embeddings

- ① Intrinsic - studying the internal properties (how well they capture the meaning: word similarity, analogy, etc.)

Evaluating the embeddings

- ① Intrinsic - studying the internal properties (how well they capture the meaning: word similarity, analogy, etc.)
- ② Extrinsic - studying how they perform a task

Analysing the embeddings

① Walking the semantic space

Analysing the embeddings

- ① Walking the semantic space
- ② Structure - (form clusters) nearest neighbors have a similar meaning, Linear structure

Glove

semantic: $v(\text{king}) - v(\text{man}) + v(\text{woman}) \approx v(\text{queen})$

syntactic: $v(\text{kings}) - v(\text{king}) + v(\text{queen}) \approx v(\text{queens})$





Figure from [Lena Voita](#)





Figure from [Lena Voita](#)