# CS224d <br> Deep Learning <br> for Natural Language Processing 

## Lecture 2: Word Vectors

Richard Socher

## How do we represent the meaning of a word?

## Definition: Meaning (Webster dictionary)

- the idea that is represented by a word, phrase, etc.
- the idea that a person wants to express by using words, signs, etc.
- the idea that is expressed in a work of writing, art, etc.


## How to represent meaning in a computer?

## Common answer: Use a taxonomy like WordNet that has hypernyms (is-a) relationships and

```
from nltk.corpus import wordnet as wn
panda = wn.synset('panda.n.01')
hyper = lambda s: s.hypernyms()
list(panda.closure(hyper))
```


## synonym sets (good):

S: (adj) full, good
S: (adj) estimable, good, honorable, respectable
S: (adj) beneficial, good
S: (adj) good, just, upright
S: (adj) adept, expert, good, practiced, proficient, skillful
S: (adj) dear, good, near
S: (adj) good, right, ripe

S: (adv) well, good
S: (adv) thoroughly, soundly, good
S: (n) good, goodness
S: ( n ) commodity, trade good, good

## Problems with this discrete representation

- Great as resource but missing nuances, e.g. synonyms: adept, expert, good, practiced, proficient, skillful?
- Missing new words (impossible to keep up to date): wicked, badass, nifty, crack, ace, wizard, genius, ninjia
- Subjective
- Requires human labor to create and adapt
- Hard to compute accurate word similarity $\rightarrow$


## Problems with this discrete representation

The vast majority of rule-based and statistical NLP work regards words as atomic symbols:holel, conference, walk

In vector space terms, this is a vector with one 1 and a lot of zeroes

$$
\left[\begin{array}{lllllllllllll}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0
\end{array} 000\right]
$$

Dimensionality: 20K (speech) - 50K (PTB) - 500K (big vocab) - 13M (Google 1T)
We call this a "one-hot" representation. Its problem:


## Distributional similarity based representations

You can get a lot of value by representing a word by means of its neighbors
"You shall know a word by the company it keeps"
(J. R. Firth 1957: 11)

One of the most successful ideas of modern statistical NLP
government debt problems turning into banking crises as has happened in
saying that Europe needs unified banking regulation to replace the hodgepodge
$\boldsymbol{K}$ These words will represent banking $\boldsymbol{\lambda}$

## How to make neighbors represent words?

Answer: With a cooccurrence matrix $X$

- 2 options: full document vs windows
- Word - document cooccurrence matrix will give general topics (all sports terms will have similar entries) leading to "Latent Semantic Analysis"
- Instead: Window around each word $\rightarrow$ captures both syntactic (POS) and semantic information


## Window based cooccurence matrix

- Window length 1 (more common: 5-10)
- Symmetric (irrelevant whether left or right context)
- Example corpus:
- I like deep learning.
- I like NLP.
- I enjoy flying.


## Window based cooccurence matrix

- Example corpus:
- I like deep learning.
- I like NLP.
- I enjoy flying.

| counts | I | like | enjoy | deep | learning | NLP | flying | . |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| I | 0 | 2 | 1 | 0 | 0 | 0 | 0 | 0 |
| like | 2 | 0 | 0 | 1 | 0 | 1 | 0 | 0 |
| enjoy | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |
| deep | 0 | 1 | 0 | 0 | 1 | 0 | 0 | 0 |
| learning | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 1 |
| NLP | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 1 |
| flying | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 1 |
| . | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 0 |

## Problems with simple cooccurrence vectors

Increase in size with vocabulary

Very high dimensional: require a lot of storage

Subsequent classification models have sparsity issues
$\rightarrow$ Models are less robust

## Solution: Low dimensional vectors

- Idea: store "most" of the important information in a fixed, small number of dimensions: a dense vector
- Usually around 25-1000 dimensions
- How to reduce the dimensionality?


## Method 1: Dimensionality Reduction on $X$

Singular Value Decomposition of cooccurrence matrix $X$.

$\hat{X}$ is the best rank $k$ approximation to $X$, in terms of least squares.

## Simple SVD word vectors in Python

Corpus:
I like deep learning. I like NLP. I enjoy flying.

```
import numpy as np
la = np.linalg
words = ["I", "like", "enjoy",
    "deep","learnig","NLP","flying","."]
x = np.array([[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]])
U, s, Vh = la.svd(X, full_matrices=False)
```


## Simple SVD word vectors in Python

Corpus: I like deep learning. I like NLP. I enjoy flying.
Printing first two columns of $U$ corresponding to the 2 biggest singular values


## Word meaning is defined in terms of vectors

- In all subsequent models, including deep learning models, a word is represented as a dense vector



## Hacks to X

- Problem: function words (the, he, has) are too frequent $\rightarrow$ syntax has too much impact. Some fixes:
- $\min (\mathrm{X}, \mathrm{t})$, with $\mathrm{t}^{\sim} 100$
- Ignore them all
- Ramped windows that count closer words more
- Use Pearson correlations instead of counts, then set negative values to 0
- +++


## Interesting semantic patters emerge in the vectors



An Improved Model of Semantic Similarity Based on Lexical Co-Occurrence Rohde et al. 2005

## Interesting syntactic patters emerge in the vectors



An Improved Model of Semantic Similarity Based on Lexical Co-Occurrence Rohde et al. 2005

## Interesting semantic patters emerge in the vectors



An Improved Model of Semantic Similarity Based on Lexical Co-Occurrence Rohde et al. 2005

## Problems with SVD

Computational cost scales quadratically for $\mathrm{n} \times \mathrm{m}$ matrix:
$\mathrm{O}\left(m n^{2}\right)$ flops (when $\mathrm{n}<\mathrm{m}$ )
$\rightarrow$ Bad for millions of words or documents

Hard to incorporate new words or documents
Different learning regime than other DL models

## Idea: Directly learn low-dimensional word vectors

- Old idea. Relevant for this lecture \& deep learning:
- Learning representations by back-propagating errors. (Rumelhart et al., 1986)
- A neural probabilistic language model (Bengio et al., 2003)
- NLP (almost) from Scratch (Collobert \& Weston, 2008)
- A recent, even simpler and faster model: word2vec (Mikolov et al. 2013) $\rightarrow$ intro now


## Main Idea of word2vec

- Instead of capturing cooccurrence counts directly,
- Predict surrounding words of every word
- Both are quite similar, see "Glove: Global Vectors for Word Representation" by Pennington et al. (2014) and Levy and Goldberg (2014) ... more later
- Faster and can easily incorporate a new sentence/ document or add a word to the vocabulary


## Details of Word2Vec

- Predict surrounding words in a window of length $m$ of every word.
- Objective function: Maximize the log probability of any context word given the current center word:
- $J(\theta)=\frac{1}{T} \sum_{t=1}^{T} \sum_{-m \leq j \leq m, j \neq 0} \log p\left(w_{t+j} \mid w_{t}\right)$
- Where $\theta$ represents all variables we optimize


## Details of Word2Vec

- Predict surrounding words in a window of length $m$ of every word
- For $p\left(w_{t+j} \mid w_{t}\right)$ the simplest first formulation is

$$
p(o \mid c)=\frac{\exp \left(u_{o}^{T} v_{c}\right)}{\sum_{w=1}^{W} \exp \left(u_{w}^{T} v_{c}\right)}
$$

- where o is the outside (or output) word $\mathrm{id}, \mathrm{c}$ is the center word $i d, u$ and $v$ are "center" and "outside" vectors of $o$ and $c$
- Every word has two vectors!
- This is essentially "dynamic" logistic regression


## Cost/Objective functions

We will optimize (maximize or minimize) our objective/cost functions

For now: minimize $\rightarrow$ gradient descent
Refresher with trivial example: (from Wikipedia) Find a local minimum of the function $f(x)=x^{4}-3 x^{3}+2$, with derivative $f^{\prime}(x)=4 x^{3}-9 x^{2}$.


```
x_old = 0
x_new = 6 # The algorithm starts at x=6
eps = 0.01 # step size
precision = 0.00001
def f_derivative(x):
    return 4 * x**3 - 9 * x**2
while abs(x_new - x_old) > precision:
    x_old = x_new
    x_new = x_old - eps * f_derivative(x_old)

\section*{Derivations of gradient}
- Whiteboard (see video if you're not in class ;)
- The basic Lego piece
- Useful basics: \(\frac{\partial \mathbf{x}^{T} \mathbf{a}}{\partial \mathbf{x}}=\frac{\partial \mathbf{a}^{T} \mathbf{x}}{\partial \mathbf{x}}=\mathbf{a}\)
- If in doubt: write out with indices
- Chain rule! If \(y=f(u)\) and \(u=g(x)\), i.e. \(\mathrm{y}=\mathrm{f}(\mathrm{g}(\mathrm{x}))\), then:
\[
\frac{d y}{d x}=\frac{d y}{d u} \frac{d u}{d x}
\]

\section*{Chain Rule}
- Chain rule! If \(y=f(u)\) and \(u=g(x)\), i.e. \(y=f(g(x))\), then:
\[
\frac{d y}{d x}=\frac{d y}{d u} \frac{d u}{d x}=\frac{d f(u)}{d u} \frac{d g(x)}{d x}
\]
- Simple example: \(\frac{d y}{d x}=\frac{d}{d x} 5\left(x^{3}+7\right)^{4}\)
\[
\begin{aligned}
& y=f(u)=5 u^{4} \quad u=g(x)=x^{3}+7 \\
& \frac{d y}{d u}=20 u^{3} \quad \frac{d u}{d x}=3 x^{2} \\
& \frac{d y}{d x}=20\left(x^{3}+7\right) 3 x^{2}
\end{aligned}
\]

\section*{Interactive Whiteboard Session!}
\(J(\theta)=\frac{1}{T} \sum_{t=1}^{T} \sum_{-m \leq j \leq m, j \neq 0} \log p\left(w_{t+j} \mid w_{t}\right)\)

Let's derive gradient together
For one example window and one example outside word:
\[
\log p(o \mid c)=\log \frac{\exp \left(u_{o}^{T} v_{c}\right)}{\sum_{w=1}^{W} \exp \left(u_{w}^{T} v_{c}\right)}
\]

\section*{Approximations: PSet 1}
- With large vocabularies this objective function is not scalable and would train too slowly! \(\rightarrow\) Why?
- Idea: approximate the normalization or
- Define negative prediction that only samples a few words that do not appear in the context
- Similar to focusing on mostly positive correlations
- You will derive and implement this in Pset 1!

\section*{Linear Relationships in word2vec}

These representations are very good at encoding dimensions of similarity!
- Analogies testing dimensions of similarity can be solved quite well just by doing vector subtraction in the embedding space Syntactically
- \(x_{\text {apple }}-x_{\text {apples }} \approx x_{\text {car }}-x_{\text {cars }} \approx x_{\text {family }}-x_{\text {families }}\)
- Similarly for verb and adjective morphological forms

Semantically (Semeval 2012 task 2)
- \(x_{\text {shirt }}-x_{\text {clothing }} \approx x_{\text {chair }}-x_{\text {furniture }}\)
- \(x_{\text {king }}-x_{\text {man }} \approx x_{\text {queen }}-x_{\text {woman }}\)

\section*{Count based vs direct prediction}

LSA, HAL (Lund \& Burgess),
COALS (Rohde et al), Hellinger-PCA (Lebret \& Collobert)
- Fast training
- Efficient usage of statistics
- Primarily used to capture word similarity
- Disproportionate importance given to large counts
- NNLM, HLBL, RNN, Skipgram/CBOW, (Bengio et al; Collobert \& Weston; Huang et al; Mnih \& Hinton; Mikolov et al; Mnih \& Kavukcuoglu)
- Scales with corpus size
- Inefficient usage of statistics
- Generate improved performance on other tasks
- Can capture complex patterns beyond word similarity

\section*{Combining the best of both worlds: GloVe}
\[
J(\theta)=\frac{1}{2} \sum_{i, j=1}^{W} f\left(P_{i j}\right)\left(u_{i}^{T} v_{j}-\log P_{i j}\right)^{2}
\]
- Fast training

- Scalable to huge corpora
- Good performance even with small corpus, and small vectors

\section*{Glove results}

Nearest words to frog:
1. frogs
2. toad
3. litoria
4. leptodactylidae
5. rana
6. lizard
7. eleutherodactylus

litoria

rana

leptodactylidae

eleutherodactylus

\section*{Word Analogies}

Test for linear relationships, examined by Mikolov et al. (2014)
\[
\text { a:b :: c:? } \quad d=\underset{x}{\arg \max } \frac{\left(w_{b}-w_{a}+w_{c}\right)^{T} w_{x}}{\left\|w_{b}-w_{a}+w_{c}\right\|}
\]
man:woman :: king:?
\begin{tabular}{lll}
+ & king & {\([0.300 .70]\)}
\end{tabular} 0.75


\section*{Glove Visualizations}


\section*{Glove Visualizations: Company - CEO}


36

Caterpillar.
Chrysler.


Wal-Mart \(-\ldots \ldots\). \(\ldots \ldots\) Tillerson
IBM \(-\ldots\) Citigroup \(=ニ=======\) Rometty \(\quad\) Corbat

Colao

Verizon.
Vodafoné

\section*{Glove Visualizations: Superlatives}


\section*{Word embedding matrix}
- Initialize most word vectors of future models with our "pretrained" embedding matrix \(L \in \mathbb{R}^{n \times|V|}\)
\[
\begin{aligned}
& L=\left[\begin{array}{llllll}
0 & 0^{|V|} & 0 & & 0 & 0 \\
\bullet & \ominus^{\circ} & \bullet & & \bullet & \bullet \\
\bullet & 0 & \bullet & \cdots & \bullet & 0 \\
0 & 0 & 0 & & 0 & 0
\end{array}\right] n \\
& \text { aardvark a at ... }
\end{aligned}
\]
- Also called a look-up table
- Conceptually you get a word's vector by left multiplying a one-hot vector \(e\) (of length |V|) by \(L: \quad x=L e\)

\section*{Advantages of low dimensional word vectors}

What is the major benefit of deep learned word vectors?
Ability to also propagate any information into them via neural networks (next lecture).
\[
P(c \mid d, \lambda)=\frac{e^{\lambda^{\top} f(c, d)}}{\sum_{c^{\prime}} e^{\lambda^{\top} f\left(c^{\prime}, d\right)}}
\]

\[
p(c \mid x)=\frac{\exp \left(S_{c} \cdot a\right)}{\sum_{c^{\prime}} \exp \left(S_{c^{\prime}} \cdot a\right)}
\]

\section*{Advantages of low dimensional word vectors}
- Word vectors will form the basis for all subsequent lectures.
- All our semantic representations will be vectors!
- Next lecture:
- Some more details about word vectors
- Predict labels for words in context for solving lots of different tasks```

