# Social Media \& Text Analysis lecture 5 - Vector Semantics 

CSE 5539-0010 Ohio State University Instructor: Alan Ritter Website: socialmedia-class.org

## NLP Pipeline



## Part-of-Speech (POS) Tagging

| Cant | MD |
| :---: | :---: |
| wait | VB |
| for | IN |
| the | DT |
| ravens | NNP |
| game | NN |
| tomorrow | NN |
| $\ldots$ |  |
| go | VB |
| ray | NNP |
| rice | NNP |
| !!!!!!! |  |

<br>Cant wait for the ravens game tomorrow....go ray rice!!!!!!!

## Named Entity Recognition

## sportsteam sportsteam geo-loc <br> India vs Australia 2014-15, 4th Test in Sydney <br> company product <br> Samsung to launch Galaxy S6 in March <br> tvshow <br> tvshow <br> New Suits and Brooklyn Nine-Nine tomorrow ... Happy days

# BAD LANGUAGE: ...on the INTERNET!! 

Boom! Ya ur website suxx bro
...dats why pluto is pluto it can neva be a star


What can we do about it? Why don't they just write NORMALLY??

## How does language go bad?

Illiteracy? No. (Tagliamonte and Denis 2008; Drouin and Davis 2009)
rob delaney @robdelaney
1 Jun
Great. Now a bunch of iliterate teens claim to be "powning" me with
their insults. Heads up jerks my wife \& children love me \& are proud of
Expand \& Reply \& Classic RT $\mathbf{~} 7$ Retweet $\star$ Favorite ** More

## Length limits? (probably not)



## Social variables

- Non-standard language does identity work, signaling authenticity, solidarity, etc.
- Social variation is usually inhibited in written language, but social media is less regulated than other written genres.

Hardware input
constraints?
(Gouws et al 2011)


## Why is Social Media Text "Bad"?

- Lack of literacy? no [Drouin and Davis, 2009]
- Length restrictions? not primarily [Eisenstein, 2013]
- Text input method? to some degree, yes [Gouws et al., 2011]
- mimicking prosodic effects etc. in speech? yeeees [Eisenstein, 2013]
- Social variables/markers of social identity? blood oath! [Eisenstein, 2013]


## Why is Social Media Text "Bad"?

- mimicking prosodic effects etc. in speech? yeeees [Eisenstein, 2013]


## HELLA

Derived from "hell of a lot". Similar to "very, really, a lot," etc.

Used mostly in Northern California though has been heard in other parts of CA and even in the media such as an infamous "hella" South Park episode. (Cartman used it outside of its meaning to annoy Kyle.)

Before: There's a hell of a lot of beer in that fridge.

After: There's hella beer in that fridge.

As "very" or "really":
"That's hella far away!"

# Why is Social Media Text "Bad"? 

- Social variables/markers of social identity? blood oath! [Eisenstein, 2013]

"I would like to believe he's sick rather than just mean and evil."
"You could've been getting down to this sick beat."


## Text Normalization

- convert non-standard words to standard

Original tweet
@USER, r u cuming 2 MidCorner dis Sunday?
Normalized tweet
@USER, are you coming to MidCorner this Sunday?

Original tweet
Still have to get up early $\mathbf{2 m r}$ thou ©so $\underline{\mathbf{G n}}$
Normalized tweet
Still have to get up early tomorrow though © Go Good night

An Unsupervised Learning Method: (1) Brown Clustering

- Input:
- a (large) text corpus
- Output:

1. a partition of words into word clusters
2. or a hierarchical word clustering (generalization of 1)

## Brown Clustering

## - Example Clusters (from Brown et al. 1992)

Friday Monday Thursday Wednesday Tuesday Saturday Sunday weekends Sundays Saturdays June March July April January December October November September August people guys folks fellows CEOs chaps doubters commies unfortunates blokes down backwards ashore sideways southward northward overboard aloft downwards adrift water gas coal liquid acid sand carbon steam shale iron great big vast sudden mere sheer gigantic lifelong scant colossal man woman boy girl lawyer doctor guy farmer teacher citizen
American Indian European Japanese German African Catholic Israeli Italian Arab pressure temperature permeability density porosity stress velocity viscosity gravity tension mother wife father son husband brother daughter sister boss uncle machine device controller processor CPU printer spindle subsystem compiler plotter John George James Bob Robert Paul William Jim David Mike anyone someone anybody somebody

## Hierarchical Word Clustering

- Each intermediate node is a cluster:
level:



## Hierarchical Word Clustering

- Each intermediate node is a cluster:




## Hierarchical Word Clustering

## mailman

salesman bookkeeper troubleshooter bouncer technician janitor saleswoman

Nike Maytag Generali
Gap
Harley-Davidson Enfield genus Microsoft Ventritex Tractebel Synopsys WordPerfect

John
Consuelo Jeffrey Kenneth Phillip WILLIAM Timothy

10000011010111 100000110110000
1000001101100010
10000011011000110
10000011011000111
1000001101100100
1000001101100101
1000001101100110

```
1011011100100101011100
10110111001001010111010
10110111001001010111011
1011011100100101011110
10110111001001010111110
101101110010010101111110
101101110010010101111111
10110111001001011000
101101110010010110010
1011011100100101100110
1011011100100101100111
1011011100100101101000
```

101110010000000000
101110010000000001
101110010000000010
10111001000000001100
101110010000000011010
101110010000000011011
10111001000000001110

- Example Clusters (from Miller et al. 2004)

Source: Miller, Guinness, Zamanian (NAACL 2004) Name Tagging with Word Clusters and Discriminative Training

## Hierarchical Word Clustering

## mailman

salesman bookkeeper troubleshooter bouncer technician janitor saleswoman

Nike Maytag Generali
Gap
Harley-Davidson Enfield genus Microsoft Ventritex Tractebel Synopsys WordPerfect

John
Consuelo Jeffrey Kenneth Phillip WILLIAM Timothy

10000011010111 100000110110000
1000001101100010
10000011011000110
10000011011000111
1000001101100100
1000001101100101
1000001101100110
1011011100100101011100
10110111001001010111010
10110111001001010111011
1011011100100101011110
10110111001001010111110
101101110010010101111110
101101110010010101111111
10110111001001011000
101101110010010110010
1011011100100101100110
1011011100100101100111
1011011100100101101000
101110010000000000
101110010000000001
101110010000000010
10111001000000001100
101110010000000011010
101110010000000011011
10111001000000001110

- Example Clusters (from Miller et al. 2004)


## word cluster features (bit string prefix)

Source: Miller, Guinness, Zamanian (NAACL 2004) Name Tagging with Word Clusters and Discriminative Training

## Challenges in Twitter

2m 2ma 2mar 2mara 2maro 2marrow 2mor 2mora 2moro 2morow 2 morr 2 morro 2 morrow 2 moz 2 mr 2mro 2mrrw 2mrw 2mw tmmrw tmo tmoro tmorrow tmoz tmr tmro tmrow tmrrow tmrrw tmrw tmrww tmw tomaro tomarow tomarro tomarrow tomm tommarow tommarrow tommoro tommorow
tommorrow tommorw tommrow tomo tomolo tomoro tomorow tomorro tomorrw tomoz tomrw tomz

# Clusters in Twitter NER 

$$
\begin{aligned}
& \text { Brown clusters, for each } i \text { s.t. } s \leq i<t \text { : } \\
& \hline\left\{\left[y_{j}, b r n\left(n, x_{i}\right), n\right]\right\}_{n \in\{2,4,8,12\}} \\
& \left\{\left[y_{j}, e r_{s, t}(i), b r n\left(n, x_{i}\right), n\right]\right\}_{n \in\{2,4,8,12\}} \\
& \text { Word vectors, for each } i \text { s.t. } s \leq i<t \text { : } \\
& \hline\left\{\left[y_{j}, n\right]=w 2 v\left(n, x_{i}\right)\right\}_{n=1}^{300} \\
& \left\{\left[y_{j}, e r_{s, t}(i), n\right]=w 2 v\left(n, x_{i}\right)\right\}_{n=1}^{300}
\end{aligned}
$$

Table 2: Word representation features in $\phi\left(s, t, y_{j}, x\right)$. $\operatorname{brn}\left(n, x_{i}\right)$ maps a word $x_{i}$ to the first $n$ bits of its Brown cluster bit sequence. $w 2 v\left(n, x_{i}\right)$ maps $x_{i}$ to the $n^{\text {th }}$ component of its word vector, and $[s t r]=v$ stands for a realvalued feature with name str and value $v$.

# Clusters in Twitter NER 

| System | Fin10Dev | Rit11 | Fro14 | Avg |
| :---: | :---: | :---: | :---: | :---: |
| CoNLL | 27.3 | 27.1 | 29.5 | 28.0 |
| + Brown | 38.4 | 39.4 | 42.5 | 40.1 |
| + Vector | 40.8 | 40.4 | 42.9 | 41.4 |
| + Reps | 42.4 | 42.2 | 46.2 | 43.6 |
| Fin10 | 36.7 | 29.0 | 30.4 | 32.0 |
| + Brown | 59.9 | 53.9 | 56.3 | 56.7 |
| + Vector | 61.5 | 56.4 | 58.4 | 58.8 |
| + Reps | 64.0 | 58.5 | 60.2 | 60.9 |
| CoNLL+Fin10 | 44.7 | 39.9 | 44.2 | 42.9 |
| + Brown | 54.9 | 52.9 | 58.5 | 55.4 |
| + Vector | 58.9 | 55.2 | 59.9 | 58.0 |
| + Reps | 58.9 | 56.4 | 61.8 | 59.0 |
| + Weights | 64.4 | 59.6 | 63.3 | 62.4 |

Table 5: Impact of our components on Twitter NER performance, as measured by F1, under 3 data scenarios.
Source: Colin Cherry, Hongyu Guo (NAACL 2015)
The Unreasonable Effectiveness of Word Representations for Twitter Named Entity Recognition

## Brown Clustering

- The Intuition:
- similar words appear in similar contexts
- more precisely: similar words have similar distributions of words to their immediate left and right

| in <br> of | the <br> a | Dog <br> part <br> store |
| :---: | :---: | :---: |


| last <br> on | Monday <br> Tuesday | -- |
| :--- | :--- | :--- | :--- |
|  |  | - |

## Brown Clustering Algorithm

- An agglomerative clustering algorithm:
- take the top m most frequent words, put each into its own cluster, $\boldsymbol{c}_{1}, \boldsymbol{c}_{2}, \ldots, \boldsymbol{c}_{m}$
- repeat for $i=(m+1)$... $I V I$
- create a new cluster $\boldsymbol{c}_{m+1}$ for the $\boldsymbol{i}$ 'th most frequent word
- choose two clusters from $\boldsymbol{c}_{1}, \boldsymbol{c}_{2}, \ldots, \boldsymbol{c}_{m+1}$ to be merged, which give the highest Quality based on a training corpus


## Brown Clustering Algorithm

- maximize the Quality function that score a given partitioning C: parameters

$$
\begin{aligned}
\operatorname{Quality}(C) & =\sum_{i}^{n} \log e\left(w_{i} \mid C\left(w_{i}\right)\right) q\left(C\left(w_{i}\right) \mid C\left(w_{i-1}\right)\right) \\
& =\sum_{c=1}^{k} \sum_{c=1}^{k} p\left(c, c^{\prime}\right) \log \frac{p\left(c, c^{\prime}\right)}{p(c) p\left(c^{\prime}\right)}+G
\end{aligned}
$$

- $n(c)$ :count of class $c$ seen in the corpus
- $n\left(c, c^{\prime}\right)$ : counts of $c^{\prime}$ seen following $c$

$$
p\left(c, c^{\prime}\right)=\frac{n\left(c, c^{\prime}\right)}{\sum_{c, c^{\prime}} n\left(c, c^{\prime}\right)} \quad p\left(c, c^{\prime}\right)=\frac{n(c)}{\sum_{c} n(c)}
$$

## Brown Clustering

( 0 Watch -
\& Graphs

C++ implementation of the Brown word clustering algorithm.


Implementation of the Brown hierarchical word clustering algorithm.
Percy Liang
Release 1.3
2012.07. 24

Input: a sequence of words separated by whitespace (see input.txt for an example)
Output: for each word type, its cluster (see output.txt for an example). In particular, each line is:
<cluster represented as a bit string> <word> <number of times word occurs in input>

## Word Vector Representations

 (a.k.a. "word embeddings")- 4 kinds of vector semantic models

1. Hard clustering (e.g. Brown clustering)
dense
2. Dimensionality Reduction (e.g. SVD, LSA, LDA)
3. Neural Network inspired models (e.g. skip-grams and CBOW in word2vec)
sparse ■- 4. Mutual-information weighted word cooccurrence metrics

## In Contrast To

## represent word meaning by a taxonomy like WordNet

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

[Synset('procyonid.n.01'), Synset('carnivore.n.01'), Synset('placental.n.01'), Synset('mammal.n.01'),
Synset('vertebrate.n.01'),
Synset('chordate.n.01'),
Synset('animal.n.01'),
Synset('organism.n.01'),
Synset('living_thing.n.01'),
Synset('whole.n.02'),
Synset('object.n.01'),
Synset('physical_entity.n.01'),
Synset('entity.n.01')]
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

## In Contrast To

represent word meaning by a taxonomy like WordNet

- problems with this discrete representation:
- missing new words (impossible to keep up-to-date): wicked, badass, nifty, crack, ace, wizard, genius, ninja
- requires human labor to create and adapt
- hard to compute accurate word similarity
- and apparently not enough to handle social media data!


## Distributional Intuition

- From context words, human can guess a word's meaning:

A bottle of tesgüino is on the table
Everybody likes tesgüino
Tesgüino makes you drunk
We make tesgüino out of corn.
"You shall know a word by the company it keeps"

- J. R. Firth 1957


## Distributional Intuition

- From context words, human can guess a word's meaning:


## A bottle of tesgüino is on the table Everybody likes tesgüino <br> Tesgüino makes you drunk We make tesgüino out of corn.

- similar words $=$ similar contexts $=$ similar vectors
- word meaning is represented by a vector of numbers


## Simple Co-occurrence Vectors

- Option \#1: word-document co-occurrence counts

|  | As You <br> Like It | Twelfth <br> Night | Julius <br> Caesar | Henry V |
| :---: | :---: | :---: | :---: | :---: |
| battle | 1 | 1 | 8 | 15 |
| soldier | 2 | 2 | 12 | 36 |
| fool | 37 | 58 | 1 | 5 |
| clown | 6 | 117 | 0 | 0 |

this will give general topics (e.g. sports terms will have similar entries), leading to Latent Semantic Analysis

## Simple Co-occurrence Vectors

- Option \#2: use a sliding window over a big corpus of text and count word co-occurrences:
example corpus:
- I like deep learning.
- I like NLP.
- I enjoy flying.

| counts | l | 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 |

this captures both syntactic (POS) and semantic information

## Simple Co-occurrence Vectors

- Problems with this representation of raw counts:
- increase in size with vocabulary
- high dimensionality and very sparse!
- not a great measure of association between words:
"the" and "of" are very frequent, but maybe not the most discriminative


## Lower Dimensional Vectors

- The Idea: use dense vectors to store "most" of the important information in a fixed, small number of dimensions
- usually around 25 ~1000 dimensions


## Lower Dimensional Vectors

- Word meaning is represented as a dense vector
"linguistic" $=\left(\begin{array}{c}0.286 \\ 0.792 \\ -0.177 \\ -0.107 \\ 0.109 \\ -0.52 \\ 0.39 \\ 0.271\end{array}\right)$

How to reduce the dimensionality?

## (2) Matrix Factorization

- Singular Value Decomposition (SVD)

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


## (2) Matrix Factorization

- Latent Semantic Analysis (LSA)

LSA is essentially low-rank approximation of document term-matrix.


## (2) Matrix Factorization

- Latent Semantic Analysis (LSA)

LSA is essentially low-rank approximation of document term-matrix.

Word assignment to topics


## SVD Word Vectors

## example 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)

## SVD Word Vectors

- plot first 2 columns of $U$ corresponding to the 2 biggest singular values:



## Some Hacks

- Problem: function words ("the", "he", "has") are too frequent $\longrightarrow$ syntax has too much impact.
- fixes: cap the counts, or ignore them all
- ramped windows that count closer words more
- etc ...


## Zipf's (Power) Law

- frequency of word is inversely proportional to its rank in the frequency table
word frequency in the Brown corpus



## Clustering Vectors

- visualize similarity



## Interesting Syntactic Patterns



## Interesting Semantic Patterns



## SVD Word Vectors

- Still some problems:
- computational cost scales quadratically for $\boldsymbol{m} \times \boldsymbol{n}$ matrix $-\mathrm{O}\left(\boldsymbol{m}^{2} \boldsymbol{n}^{2}\right.$ when $\mathrm{n}<\mathrm{m}$
- hard to use large corpus (and vocabulary)
- hard to incorporate new words or documents


## (3) Neural Word Embeddings

- The Idea: directly learn low-dimensional word vectors
- ... can go back to 1980s:
- Learning Representations by Back-Propagating Errors (Rumelhart et al., 1986)
- A Neural Probabilistic Language Model (Bengio et al., 2003)
- NLP from Scratch (Collobert \& Weston, 2008)
- Word2vec (Mikolov et al. 2013), GloVe (Pennington et al. 2014)
- ELMo (Peters et al. 2018), BERT (Devlin et al. 2019)


## A Neural Probabilistic Language Model (Bengio et al., 2003)



Figure 1: Neural architecture: $f\left(i, w_{t-1}, \cdots, w_{t-n+1}\right)=g\left(i, C\left(w_{t-1}\right), \cdots, C\left(w_{t-n+1}\right)\right)$ where $g$ is the neural network and $C(i)$ is the $i$-th word feature vector.

## Neural Word Embeddings

- The Basic Idea:
- We define a model that aims to predict a word given its context words (word vectors), which has a loss function, e.g. $J=1-P\left(\right.$ context $\left.I w_{t}\right)$
- We look at many positions of $t$ in a big text corpus,
- and keep adjusting the word vectors to minimize this loss.


## Word2vec

- simple and efficient


CBOW


Skip-gram

Source: Mikolov et al. (NIPS 2013)
Distributed Representations of Words and Phrases and their Compositionality

## Word2vec

- Skip-gram - predicts surrounding "outside" words given the "center" word

INPUT PROJECTION OUTPUT



Skip-gram

## Word2vec

- Skip-gram - predicts surrounding "outside" words given the "center" word


Figure 16.5 The skip-gram model viewed as a network (Mikolov et al. 2013, Mikolov et al. 2013a).

## Input Layer

- "one-hot" word vectors
- a vector of dimension $|\boldsymbol{V}|$ (size of vocabulary)
- all " 0 " s expect a single " 1 " in the vector
- different positions of that "1" represent different words

Input layer
1-hot input vector

$1 x|V|$

## Hidden (Projection) Layer

- A simple look up - the rows of this weight matrix are actually "input" word vectors



## Hidden (Projection) Layer

- A simple look up - the rows of this weight matrix are actually "input" word vectors

$$
\begin{aligned}
& {\left[\begin{array}{lllll}
0 & 0 & 0 & 1 & 0
\end{array}\right] \times\left[\begin{array}{ccc}
17 & 24 & 1 \\
23 & 5 & 7 \\
4 & 6 & 13 \\
10 & 12 & 19 \\
11 & 18 & 25
\end{array}\right]=\left[\begin{array}{lll}
10 & 12 & 19
\end{array}\right]} \\
& \text { Input layer } \\
& \text { Projection layer }
\end{aligned}
$$

## Output Layer

- predicts surrounding "outside" (context) words given the "center" word $\longrightarrow$ A classification problem!
- Softmax Regression = Multi-class Logistic Regression



## Softmax Function

- Softmax function is a generalization of logistic function

$$
\operatorname{softmax}(\mathbf{x})_{i}=\frac{e^{x_{i}} \swarrow^{\text {exponentiate to make positive }}}{\sum_{j} e^{x_{j}}} \leftarrow \text { normalized to give probability }
$$

## Softmax Function

- Softmax function is a generalization of logistic function

$$
\operatorname{softmax}(\mathbf{x})_{i}=\frac{e^{x_{i}}}{\sum_{j} e^{x_{j}}}
$$

| -2.85 |
| :---: | :---: | :---: |
| 0.86 |
| 0.28 |

## Output Layer

- Objective function: maximize the log probability of any "outside" (context) word given the "center" word

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

$$
\begin{aligned}
& \text { Projection layer } \\
& \text { embedding for } w_{t}
\end{aligned}
$$

Output layer probabilities of context words

## Output Layer

- predicts surrounding "outside" (context) words given the "center" word

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

- so, every word has two vectors!



## Output Layer

- Intuition

Word vector for "ants"

softmax

$=\quad \begin{aligned} & \text { Probability that "car" } \\ & \text { shows up near "ants" }\end{aligned}$

## Objective Function

- Cost/Objective function:

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

- For a "center" word and an "outside" word:

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

## Gradient Descent



## Gradient Descent

- Basics:

$$
\begin{gathered}
\frac{\partial \mathrm{x}^{T} a}{\partial \mathrm{x}}=\frac{\partial a^{T} \mathrm{x}}{\partial \mathrm{x}}=a \\
\frac{\partial e^{\mathrm{x}}}{\partial \mathrm{x}}=e^{\mathrm{x}} \quad \frac{\partial \log \mathrm{x}}{\partial \mathrm{x}}=\frac{1}{\mathrm{x}}
\end{gathered}
$$

- Chain Rule:

$$
\frac{\partial f}{\partial x}=\frac{\partial f}{\partial g} \frac{\partial g}{\partial x}=\frac{\partial f(g)}{\partial g} \frac{\partial g(x)}{\partial x}
$$

## Word2vec

- Word2vec is not a single algorithm, but a toolkit
- which contains two distinct algorithms (Skipgram \& CBOW), two training methods (negative sampling \& hierarchical softmax)
- Word2vec is not deep learning, but neural-inspired
- only one hidden layer followed by softmax, no non-linear activation function


## Relation between Skip-gram and SVD

- Levy and Goldberg (2014) show that skip-gram is factorizing (a shifted version of ) the traditional word-context PMI matrix:

$$
\operatorname{Opt}(\vec{w} \cdot \vec{c})=\operatorname{PMI}(w, c)-\log k
$$



- So does SVD!


## Visualization



Male-Female


Verb tense


Country-Capital

## Visualization



Figure 1: A 2-dimensional projection of the latent semantic space captured by our algorithm. Notice the semantic trajectory of the word gay transitioning meaning in the space.

Source: Kulkarni et al. (WWW 2015)

