Vector Semantics

Why vector models of meaning? computing the similarity between words

"fast" is similar to "rapid"

"tall" is similar to "height"

Question answering:

Q: "How **tall** is Mt. Everest?" Candidate A: "The official **height** of Mount Everest is 29029 feet"

Word similarity for plagiarism detection

MAINFRAMES

Mainframes are primarily referred to large computers with rapid, advanced processing capabilities that can execute and perform tasks equivalent to many Personal Computers (PCs) machines networked together. It is characterized with high quantity Random Access Memory (RAM), very

large secondary storage devices, and high-speed processors to cater for the needs of the computers under its service.

Consisting of advanced components, mainframes have the capability of running multiple large applications required by many and most enterprises and organizations. This is one of its advantages. Mainframes are also suitable to cater for those applications (programs) or files that are of very high

MAINFRAMES

Mainframes usually are referred those computers with fast, advanced processing capabilities that could perform by itself tasks that may require a lot of Personal Computers (PC) Machines. Usually mainframes would have lots of RAMs, very large secondary storage devices, and very fast processors to cater for the needs of those computers under its service.

Due to the advanced components mainframes have, these computers have the capability of running multiple large applications required by most enterprises, which is one of its advantage. Mainframes are also suitable to cater for those applications or files that are of very large demand

Distributional models of meaning = vector-space models of meaning = vector semantics

Intuitions: Zellig Harris (1954):

- "oculist and eye-doctor ... occur in almost the same environments"
- "If A and B have almost identical environments we say that they are synonyms."

Firth (1957):

"You shall know a word by the company it keeps!"

Intuition of distributional word similarity

• Nida example:

A bottle of **tesgüino** is on the table Everybody likes **tesgüino Tesgüino** makes you drunk We make **tesgüino** out of corn.

- From context words humans can guess *tesgüino* means
 - an alcoholic beverage like **beer**
- Intuition for algorithm:
 - Two words are similar if they have similar word contexts.

Four kinds of vector models

Sparse vector representations

1. Mutual-information weighted word co-occurrence matrices

Dense vector representations:

- 2. Singular value decomposition (and Latent Semantic Analysis)
- 3. Neural-network-inspired models (skip-grams, CBOW)
- 4. Brown clusters

Shared intuition

- Model the meaning of a word by "embedding" in a vector space.
- The meaning of a word is a vector of numbers
 - Vector models are also called "embeddings".
- Contrast: word meaning is represented in many computational linguistic applications by a vocabulary index ("word number 545")
- Old philosophy joke:

Q: What's the meaning of life? A: LIFE'

Term-document matrix

- Each cell: count of term t in a document d: $tf_{t,d}$:
 - Each document is a count vector in \mathbb{N}^{v} : a column below

| | As You Lik | e lt | Twelfth Night | Julius Caesar | Henry V |
|---------|------------|------|---------------|---------------|---------|
| battle | | 1 | 1 | 8 | 15 |
| soldier | | 2 | 2 | 12 | 36 |
| fool | | 37 | 58 | 1 | 5 |
| clown | | 6 | 117 | 0 | 0 |

Term-document matrix

• Two documents are similar if their vectors are similar

| | As You Like It | Twelfth Night | Julius Caesar | Henr | ry V |
|---------|----------------|---------------|---------------|------|------|
| battle | 1 | 1 | 8 | | 15 |
| soldier | 2 | 2 | 12 | | 36 |
| fool | 37 | 58 | 1 | | 5 |
| clown | 6 | 117 | 0 | | 0 |

The words in a term-document matrix

• Each word is a count vector in \mathbb{N}^{D} : a row below

| | As You l | ike It | Twelfth Night | Julius Caesar | Henry V |
|---------|----------|--------|---------------|---------------|---------|
| battle | | 1 | 1 | 8 | 15 |
| soldier | | 2 | 2 | 12 | 36 |
| fool | | 37 | 58 | 1 | 5 |
| clown | | 6 | 117 | 0 | 0 |

The words in a term-document matrix

• Two words are similar if their vectors are similar

| | As You Lil | ke lt | Twelfth Night | Julius Caesar | Henry V |
|---------|------------|-------|---------------|---------------|---------|
| battle | | 1 | 1 | 8 | 15 |
| soldier | | 2 | 2 | 12 | 36 |
| fool | | 37 | 58 | 1 | 5 |
| clown | | 6 | 117 | 0 | 0 |

Term-context matrix for word similarity

• Two **words** are similar in meaning if their context vectors are similar

| | aardvark | computer | data | pinch | result | sugar | |
|-------------|----------|----------|------|-------|--------|-------|--|
| apricot | 0 | 0 | 0 | 1 | 0 | 1 | |
| pineapple | 0 | 0 | 0 | 1 | 0 | 1 | |
| digital | 0 | 2 | 1 | 0 | 1 | 0 | |
| information | 0 | 1 | 6 | 0 | 4 | 0 | |

The word-word or word-context matrix

- Instead of entire documents, use smaller contexts
 - Paragraph

13

- Window of \pm 4 words
- A word is now defined by a vector over counts of context words
- Instead of each vector being of length D
- Each vector is now of length |V|
- The word-word matrix is |V|x|V|

Word-Word matrix Sample contexts \pm 7 words

sugar, a sliced lemon, a tablespoonful of **apricot** their enjoyment. Cautiously she sampled her first **pineapple** well suited to programming on the digital **computer**. for the purpose of gathering data and **information** necessary for the study authorized in the

...

preserve or jam, a pinch each of, and another fruit whose taste she likened In finding the optimal R-stage policy from

| | aardvark | computer | data | pinch | result | sugar | |
|-------------|----------|----------|------|-------|--------|-------|--|
| apricot | 0 | 0 | 0 | 1 | 0 | 1 | |
| pineapple | 0 | 0 | 0 | 1 | 0 | 1 | |
| digital | 0 | 2 | 1 | 0 | 1 | 0 | |
| information | 0 | 1 | 6 | 0 | 4 | 0 | |

Word-word matrix

- We showed only 4x6, but the real matrix is 50,000 x 50,000
 - So it's very sparse
 - Most values are 0.
 - That's OK, since there are lots of efficient algorithms for sparse matrices.
- The size of windows depends on your goals
 - The shorter the windows , the more ${\bf syntactic}$ the representation \pm 1-3 very syntacticy
 - The longer the windows, the more **semantic** the representation \pm 4-10 more semanticy

Vector Semantics Positive Pointwise Mutual Information (PPMI)

Problem with raw counts

- Raw word frequency is not a great measure of association between words
 - It's very skewed
 - "the" and "of" are very frequent, but maybe not the most discriminative
- We'd rather have a measure that asks whether a context word is **particularly informative** about the target word.
 - Positive Pointwise Mutual Information (PPMI)

Pointwise Mutual Information

Pointwise mutual information:

Do events x and y co-occur more than if they were independent?

$$PMI(X,Y) = \log_2 \frac{P(x,y)}{P(x)P(y)}$$

PMI between two words: (Church & Hanks 1989)

Do words x and y co-occur more than if they were independent?

$$PMI(word_1, word_2) = \log_2 \frac{P(word_1, word_2)}{P(word_1)P(word_2)}$$

Positive Pointwise Mutual Information

- PMI ranges from $-\infty$ to $+\infty$
- But the negative values are problematic
 - Things are co-occurring less than we expect by chance
 - Unreliable without enormous corpora
 - Imagine w1 and w2 whose probability is each 10⁻⁶
 - Hard to be sure p(w1,w2) is significantly different than 10⁻¹²
 - Plus it's not clear people are good at "unrelatedness"
- So we just replace negative PMI values by 0
- Positive PMI (PPMI) between word1 and word2: $PPMI(word_1, word_2) = max \left(log_2 \frac{P(word_1, word_2)}{P(word_1)P(word_2)}, 0 \right)$

Computing PPMI on a term-context matrix

- Matrix F with W rows (words) and C columns (contexts) •
- f_{ii} is # of times w_i occurs in context c_i

$$p_{ij} = \frac{f_{ij}}{\frac{W \ C}{a} a a f_{ij}} \qquad p_{i*} = \frac{j=1}{\frac{W \ C}{W \ C}} \qquad p_{*j} = \frac{i=1}{\frac{W \ C}{W \ C}} \\ \overset{a}{a} a a f_{ij} \qquad \overset{a}{a} a a f_{ij} \qquad \overset{a}{a} a a f_{ij} \\ i=1 \ j=1 \qquad i=1 \ j=1 \qquad i=1 \ j=1 \qquad j=1 \qquad j=1$$

20

| | aardvark | computer | data | pinch | result | sugar |
|-------------|----------|----------|------|-------|--------|-------|
| apricot | 0 | 0 | 0 | 1 | 0 | 1 |
| pineapple | 0 | 0 | 0 | 1 | 0 | 1 |
| digital | 0 | 2 | 1 | 0 | 1 | 0 |
| information | 0 | 1 | 6 | 0 | 4 | 0 |

$$pmi_{ij} = \log_2 \frac{p_{ij}}{p_{i*}p_{*j}}$$
 $pmi_{ij} = \begin{cases} 1 & pmi_{ij} & \text{if } pmi_{ij} > 0 \\ 1 & 0 & \text{otherwise} \end{cases}$

| | | | | Cou | nt(w,c | ontex | t) | |
|--|----------------------------------|-----------------------------|---------|-------------------------|--------|----------|------------------|-------|
| | f | | comput | er da | ata p | inch | result | sugar |
| | $p_{ii} = \frac{J_{ij}}{W_{ij}}$ | apricot | | 0 | 0 | 1 | 0 | 1 |
| - | | pineapple | | 0 | 0 | 1 | 0 | 1 |
| | aaj | <i>ij</i> digital | | 2 | 1 | 0 | 1 | 0 |
| | <i>i</i> =1 <i>j</i> =1 | information | | 1 | 6 | 0 | 4 | 0 |
| | | | | С | | | W | |
| p(w=information,c=data) = 6/19 = .32 $\mathring{a}f_{ij}$ $\mathring{a}f_{ij}$ | | | | | | | | |
| p(w=infor | mation) = 11/ | ′19 = .58 | p(w | $(i_i) = \frac{j=1}{N}$ | | $p(c_j)$ | $=\frac{i=1}{N}$ | |
| p(c=data) | = 7/19 = .37 | p(w,co | ontext) | 1 V | | р | (w) | |
| | | computer data | a pinch | result | sugar | | | |
| | apricot | 0.00 0.00 | 0.05 | 0.00 | 0.05 | | 0.11 | |
| | pineapple | 0.00 0.00 | 0.05 | 0.00 | 0.05 | | 0.11 | |
| | digital | 0.11 0.05 | 5 0.00 | 0.05 | 0.00 | | 0.21 | |
| | information | 0.05 0.32 | 0.00 | 0.21 | 0.00 | | 0.58 | |
| 21 | | 0.16 0.00 | 7 0 4 4 | 0.20 | 0.44 | | | |
| | p(context) | 0.16 0.3 | / 0.11 | 0.26 | 0.11 | | | |

| | | | p(w,context) | | | | |
|------------------------------------|-------------|----------|--------------|-------|--------|-------|------|
| | | computer | data | pinch | result | sugar | |
| p_{ii} | apricot | 0.00 | 0.00 | 0.05 | 0.00 | 0.05 | 0.11 |
| $pml_{ij} = \log_2 \frac{-9}{n n}$ | pineapple | 0.00 | 0.00 | 0.05 | 0.00 | 0.05 | 0.11 |
| $P_{i^*}P_*$ | digital | 0.11 | 0.05 | 0.00 | 0.05 | 0.00 | 0.21 |
| | information | 0.05 | 0.32 | 0.00 | 0.21 | 0.00 | 0.58 |
| | p(context) | 0.16 | 0.37 | 0.11 | 0.26 | 0.11 | |

• $pmi(information, data) = log_2(.32 / (.37*.58)) = .58$

(.57 using full precision)

| | PPMI(w,context) | | | | | | | | |
|-------------|-----------------|--------------------------------|------|------|------|--|--|--|--|
| | computer | computer data pinch result sug | | | | | | | |
| apricot | - | - | 2.25 | - | 2.25 | | | | |
| pineapple | - | - | 2.25 | - | 2.25 | | | | |
| digital | 1.66 | 0.00 | - | 0.00 | - | | | | |
| information | 0.00 | 0.57 | - | 0.47 | - | | | | |

22

Weighting PMI

- PMI is biased toward infrequent events
 - Very rare words have very high PMI values
- Two solutions:
 - Give rare words slightly higher probabilities
 - Use add-one smoothing (which has a similar effect)

Weighting PMI: Giving rare context words slightly higher probability

• Raise the context probabilities to $\alpha = 0.75$: $PPMI_{\alpha}(w,c) = max(log_2 \frac{P(w,c)}{P(w)P_{\alpha}(c)}, 0)$

$$P_{\alpha}(c) = \frac{count(c)^{\alpha}}{\sum_{c} count(c)^{\alpha}}$$

- This helps because $P_{\alpha}(c) > P(c)$ for rare c
- Consider two events, P(a) = .99 and P(b)=.01

$$P_{\alpha}(a) = \frac{.99^{.75}}{.99^{.75} + .01^{.75}} = .97 P_{\alpha}(b) = \frac{.01^{.75}}{.01^{.75} + .01^{.75}} = .03$$

Use Laplace (add-1) smoothing

Add-2 Smoothed Count(w,context)

| | computer | data | pinch | result | sugar |
|-------------|----------|------|-------|--------|-------|
| apricot | 2 | 2 | 3 | 2 | 3 |
| pineapple | 2 | 2 | 3 | 2 | 3 |
| digital | 4 | 3 | 2 | 3 | 2 |
| information | 3 | 8 | 2 | 6 | 2 |

| | K | p(w) | | | | |
|-------------------|----------|------|-------|--------|-------|------|
| | computer | data | pinch | result | sugar | |
| apricot | 0.03 | 0.03 | 0.05 | 0.03 | 0.05 | 0.20 |
| pineapple | 0.03 | 0.03 | 0.05 | 0.03 | 0.05 | 0.20 |
| digital | 0.07 | 0.05 | 0.03 | 0.05 | 0.03 | 0.24 |
| information | 0.05 | 0.14 | 0.03 | 0.10 | 0.03 | 0.36 |
| p(context) | 0.19 | 0.25 | 0.17 | 0.22 | 0.17 | |

PPMI versus add-2 smoothed PPMI

PPMI(w,context)

| | computer | data | pinch | result | sugar |
|-------------|----------|------|-------|--------|-------|
| apricot | - | - | 2.25 | - | 2.25 |
| pineapple | - | - | 2.25 | - | 2.25 |
| digital | 1.66 | 0.00 | - | 0.00 | - |
| information | 0.00 | 0.57 | - | 0.47 | - |

PPMI(w,context) [add-2]

| | computer | data | pinch | result | sugar |
|-------------|----------|------|-------|--------|-------|
| apricot | 0.00 | 0.00 | 0.56 | 0.00 | 0.56 |
| pineapple | 0.00 | 0.00 | 0.56 | 0.00 | 0.56 |
| digital | 0.62 | 0.00 | 0.00 | 0.00 | 0.00 |
| information | 0.00 | 0.58 | 0.00 | 0.37 | 0.00 |

Vector Semantics Measuring similarity: the cosine

Measuring similarity

• Given 2 target words v and w

29

- We'll need a way to measure their similarity.
- Most measure of vectors similarity are based on the:
- **Dot product** or **inner product** from linear algebra dot-product $(\vec{v}, \vec{w}) = \vec{v} \cdot \vec{w} = \sum_{i=1}^{N} v_i w_i = v_1 w_1 + v_2 w_2 + ... + v_N w_N$
 - High when two vectors have large values in same dimensions.
 - Low (in fact 0) for **orthogonal vectors** with zeros in complementary distribution

Problem with dot product

dot-product
$$(\vec{v}, \vec{w}) = \vec{v} \cdot \vec{w} = \sum_{i=1}^{N} v_i w_i = v_1 w_1 + v_2 w_2 + \dots + v_N w_N$$

• Dot product is longer if the vector is longer. Vector length:

$$|\vec{v}| = \sqrt{\sum_{i=1}^{N} v_i^2}$$

- Vectors are longer if they have higher values in each dimension
- That means more frequent words will have higher dot products
- That's bad: we don't want a similarity metric to be sensitive to
 ³⁰ word frequency

Solution: cosine

- Just divide the dot product by the length of the two vectors! $\frac{\vec{a} \cdot \vec{b}}{|\vec{a}||\vec{b}|}$
- This turns out to be the cosine of the angle between them!

$$ec{a} \cdot ec{b} = ec{a} ec{b} ec{b} \cos heta$$
 $rac{ec{a} \cdot ec{b}}{ec{a} ec{b} ec{b}} = \cos heta$

Cosine for computing similarity



 v_i is the PPMI value for word v in context i w_i is the PPMI value for word w in context i.

 $\cos(\vec{v}, \vec{w})$ is the cosine similarity of v and w

Cosine as a similarity metric

- -1: vectors point in opposite directions
- +1: vectors point in same directions
- 0: vectors are orthogonal



• Raw frequency or PPMI are nonnegative, so cosine range 0-1

$$\cos(v,w) = \frac{v \cdot w}{|v||w|} = \frac{v}{|v|} \cdot \frac{w}{|w|} = \frac{\mathring{a}_{i=1}^{N} v_i w_i}{\sqrt{\mathring{a}_{i=1}^{N} v_i^2} \sqrt{\mathring{a}_{i=1}^{N} w_i^2}} = \frac{1}{100} \frac{1}{100} = \frac{1}{100} \frac{1$$

Visualizing vectors and angles



Clustering vectors to visualize similarity in co-occurrence matrices



Rohde et al. (2006)
Other possible similarity measures

3.7

$$\begin{aligned} \sin_{\text{cosine}}(\vec{v}, \vec{w}) &= \frac{\vec{v} \cdot \vec{w}}{|\vec{v}| |\vec{w}|} = \frac{\sum_{i=1}^{N} v_i \times w_i}{\sqrt{\sum_{i=1}^{N} v_i^2} \sqrt{\sum_{i=1}^{N} w_i^2}} \\ \sin_{\text{Jaccard}}(\vec{v}, \vec{w}) &= \frac{\sum_{i=1}^{N} \min(v_i, w_i)}{\sum_{i=1}^{N} \max(v_i, w_i)} \\ \sin_{\text{Dice}}(\vec{v}, \vec{w}) &= \frac{2 \times \sum_{i=1}^{N} \min(v_i, w_i)}{\sum_{i=1}^{N} (v_i + w_i)} \\ \sin_{\text{JS}}(\vec{v} || \vec{w}) &= D(\vec{v} | \frac{\vec{v} + \vec{w}}{2}) + D(\vec{w} | \frac{\vec{v} + \vec{w}}{2}) \end{aligned}$$

Vector Semantics Measuring similarity: the cosine

Using syntax to define a word's context

• Zellig Harris (1968)

"The meaning of entities, and the meaning of grammatical relations among them, is related to the restriction of combinations of these entities relative to other entities"

• Two words are similar if they have similar syntactic contexts

Duty and **responsibility** have similar syntactic distribution:

| Modified by | additional, administrative, assumed, collective, | |
|------------------|--|--|
| adjectives | congressional, constitutional | |
| Objects of verbs | assert, assign, assume, attend to, avoid, become, breach | |

Co-occurrence vectors based on syntactic dependencies

Dekang Lin, 1998 "Automatic Retrieval and Clustering of Similar Words"

- Each dimension: a context word in one of R grammatical relations
 - Subject-of- "absorb"
- Instead of a vector of /V/ features, a vector of R/V/
- Example: counts for the word *cell* :



Syntactic dependencies for dimensions

- Alternative (Padó and Lapata 2007):
 - Instead of having a |V| x R|V| matrix
 - Have a |V| x |V| matrix
 - But the co-occurrence counts aren't just counts of words in a window
 - But counts of words that occur in one of R dependencies (subject, object, etc).
 - So M("cell","absorb") = count(subj(cell,absorb)) + count(obj(cell,absorb)) + count(pobj(cell,absorb)), etc.

PMI applied to dependency relations

Hindle, Don. 1990. Noun Classification from Predicate-Argument Structure. ACL

| Object of "drink" | Count | PMI |
|-------------------|-------|------|
| tea | 2 | 11.8 |
| liquid | 2 | 10.5 |
| wine | 2 | 9.3 |
| anything | 3 | 5.2 |
| it | 3 | 1.3 |

- "Drink it" more common than "drink wine"
- But "wine" is a better "drinkable" thing than "it"

Alternative to PPMI for measuring association

- **tf-idf** (that's a hyphen not a minus sign)
- The combination of two factors ٠
 - **Term frequency** (Luhn 1957): frequency of the word (can be logged)
 - Inverse document frequency (IDF) (Sparck Jones 1972) $\operatorname{idf}_{i} = \log_{\overset{\circ}{\xi}}^{\overset{\circ}{\xi}} \frac{N^{0}}{df_{i} \overset{\circ}{\xi}}$
 - N is the total number of documents
 - df_i = "document frequency of word i"
 - = # of documents with word /
 - w_{ij} = word *i* in document *j*

 $w_{ii} = tf_{ii} idf_i$

tf-idf not generally used for word-word similarity

• But is by far the most common weighting when we are considering the relationship of words to documents

Vector Semantics Evaluating similarity

Evaluating similarity

- Extrinsic (task-based, end-to-end) Evaluation:
 - Question Answering
 - Spell Checking
 - Essay grading
- Intrinsic Evaluation:
 - Correlation between algorithm and human word similarity ratings
 - Wordsim353: 353 noun pairs rated 0-10. *sim(plane,car)=5.77*
 - Taking TOEFL multiple-choice vocabulary tests
 - <u>Levied</u> is closest in meaning to: imposed, believed, requested, correlated

Summary

- Distributional (vector) models of meaning
 - **Sparse** (PPMI-weighted word-word co-occurrence matrices)
 - Dense:
 - Word-word SVD 50-2000 dimensions
 - Skip-grams and CBOW
 - Brown clusters 5-20 binary dimensions.

Vector Semantics Dense Vectors

Sparse versus dense vectors

- PPMI vectors are
 - **long** (length |V| = 20,000 to 50,000)
 - sparse (most elements are zero)
- Alternative: learn vectors which are
 - short (length 200-1000)
 - **dense** (most elements are non-zero)

Sparse versus dense vectors

- Why dense vectors?
 - Short vectors may be easier to use as features in machine learning (less weights to tune)
 - Dense vectors may generalize better than storing explicit counts
 - They may do better at capturing synonymy:
 - car and automobile are synonyms; but are represented as distinct dimensions; this fails to capture similarity between a word with car as a neighbor and a word with automobile as a neighbor

Three methods for getting short dense vectors

- Singular Value Decomposition (SVD)
 - A special case of this is called LSA Latent Semantic Analysis
- "Neural Language Model"-inspired predictive models
 - skip-grams and CBOW
- Brown clustering

Vector Semantics Dense Vectors via SVD

Intuition

- Approximate an N-dimensional dataset using fewer dimensions
- By first rotating the axes into a new space
- In which the highest order dimension captures the most variance in the original dataset
- And the next dimension captures the next most variance, etc.
- Many such (related) methods:
 - PCA principle components analysis
 - Factor Analysis
 - SVD



Singular Value Decomposition

Any rectangular w x c matrix X equals the product of 3 matrices:

W: rows corresponding to original but m columns represents a dimension in a new latent space, such that

- M column vectors are orthogonal to each other
- Columns are ordered by the amount of variance in the dataset each new dimension accounts for
- **S**: diagonal *m* x *m* matrix of **singular values** expressing the importance of each dimension.

C: columns corresponding to original but m rows corresponding to ईingular values

Singular Value Decomposition

An Introduction to Singu

Singular Value Decomposition (SVD)

-known proof in matrix algebra asserts that any () is equal to the product of three other matrices

Landuaer and Dumais 1997

ĺ

SVD applied to term-document matrix: Latent Semantic Analysis Deerwester et al (1988)

- If instead of keeping all m dimensions, we just keep the top k singular values. Let's say 300.
- The result is a least-squares approximation to the original X
- But instead of multiplying, we'll just make use of W.
- Each row of W:
 - A k-dimensional vector
 - Representing word W

An I troduction to Singu Singular Value Decomposition (SVD) -known proof in matrix algebra asserts that any

LSA more details

- 300 dimensions are commonly used
- The cells are commonly weighted by a product of two weights
 - Local weight: Log term frequency
 - Global weight: either idf or an entropy measure

Let's return to PPMI word-word matrices

• Can we apply to SVD to them?

SVD applied to term-term matrix



(I'm simplifying here by assuming the matrix has rank |V|)

Truncated SVD on term-term matrix



Truncated SVD produces embeddings

- Each row of W matrix is a k-dimensional representation of each word *w*
- K might range from 50 to 1000
- Generally we keep the top k dimensions, but some experiments suggest that getting rid of the top 1 dimension or even the top 50 dimensions is helpful (Lapesa and Evert 2014).



Embeddings versus sparse vectors

- Dense SVD embeddings sometimes work better than sparse PPMI matrices at tasks like word similarity
 - Denoising: low-order dimensions may represent unimportant information
 - Truncation may help the models generalize better to unseen data.
 - Having a smaller number of dimensions may make it easier for classifiers to properly weight the dimensions for the task.
 - Dense models may do better at capturing higher order cooccurrence.

Vector Semantics

Embeddings inspired by neural language models: skip-grams and CBOW

Prediction-based models: An alternative way to get dense vectors

- Skip-gram (Mikolov et al. 2013a) CBOW (Mikolov et al. 2013b)
- Learn embeddings as part of the process of word prediction.
- Train a neural network to predict neighboring words
 - Inspired by **neural net language models**.
 - In so doing, learn dense embeddings for the words in the training corpus.
- Advantages:

65

- Fast, easy to train (much faster than SVD)
- Available online in the word2vec package
- Including sets of pretrained embeddings!

Skip-grams

- Predict each neighboring word
 - in a context window of 2C words
 - from the current word.
- So for C=2, we are given word w_t and predicting these 4 words:

$$[w_{t-2}, w_{t-1}, w_{t+1}, w_{t+2}]$$

Skip-grams learn 2 embeddings for each w

input embedding v, in the input matrix W

• Column *i* of the input matrix *W* is the $1 \times d$ embedding v_i for word *i* in the vocabulary.

output embedding v', in output matrix W'

Row *i* of the output matrix *W*' is a *d* × 1 vector embedding *v*'_i for word *i* in the vocabulary.



Setup

- Walking through corpus pointing at word w(t), whose index in the vocabulary is j, so we'll call it w_j (1 < j < |V|).
- Let's predict w(t+1), whose index in the vocabulary is k (1 < k < |V|). Hence our task is to compute P(w_k|w_i).

Intuition: similarity as dot-product between a target vector and context vector



Similarity is computed from dot product

- Remember: two vectors are similar if they have a high dot product
 - Cosine is just a normalized dot product
- So:
 - Similarity(j,k) $\propto c_k \cdot v_j$
- We'll need to normalize to get a probability

Turning dot products into probabilities

• Similarity(j,k) =
$$c_k \cdot v_j$$

• We use softmax to turn into probabilities

$$p(w_k|w_j) = \frac{exp(c_k \cdot v_j)}{\sum_{i \in |V|} exp(c_i \cdot v_j)}$$

Embeddings from W and W'

- Since we have two embeddings, v_i and c_i for each word w_i
- We can either:
 - Just use v_i
 - Sum them
 - Concatenate them to make a double-length embedding
Learning

- Start with some initial embeddings (e.g., random)
- iteratively make the embeddings for a word
 - more like the embeddings of its neighbors
 - less like the embeddings of other words.

Visualizing W and C as a network for doing error backprop



One-hot vectors

- A vector of length |V|
- 1 for the target word and 0 for other words
- So if "popsicle" is vocabulary word 5
- The **one-hot vector** is
- [0,0,0,0,1,0,0,0,0,.....0]



Problem with the softamx

• The denominator: have to compute over every word in vocab

$$p(w_k|w_j) = \frac{exp(c_k \cdot v_j)}{\sum_{i \in |V|} exp(c_i \cdot v_j)}$$

• Instead: just sample a few of those negative words

Goal in learning

• Make the word like the context words



• We want this to be high:

 $\sigma(c1 \cdot w) + \sigma(c2 \cdot w) + \sigma(c3 \cdot w) + \sigma(c4 \cdot w)$

- And not like k randomly selected "noise words"
 [cement metaphysical dear coaxial apricot attendant whence forever puddle]
 n1 n2 n3 n4 n5 n6 n7 n8
- We want this to be low:

78 $\sigma(n1 \cdot w) + \sigma(n2 \cdot w) + \dots + \sigma(n8 \cdot w)$

Skipgram with negative sampling: Loss function

$$\log \sigma(c \cdot w) + \sum_{i=1}^{\kappa} \mathbb{E}_{w_i \sim p(w)} \left[\log \sigma(-w_i \cdot w) \right]$$

Relation between skipgrams and PMI!

- If we multiply WW'^{T}
- We get a |V|x|V| matrix M, each entry m_{ij} corresponding to some association between input word i and output word j
- Levy and Goldberg (2014b) show that skip-gram reaches its optimum just when this matrix is a shifted version of PMI:

 $WW'^{T} = M^{PMI} - \log k$

• So skip-gram is implicitly factoring a shifted version of the PMI matrix into the two embedding matrices.

Properties of embeddings

• Nearest words to some embeddings (Mikolov et al. 20131)

| target: | Redmond | Havel | ninjutsu | graffiti | capitulate |
|---------|---------------------------|------------------------|---------------|-------------|--------------|
| | Redmond Wash. | Vaclav Havel | ninja | spray paint | capitulation |
| | Redmond Washington | president Vaclav Havel | martial arts | grafitti | capitulated |
| | Microsoft | Velvet Revolution | swordsmanship | taggers | capitulating |

Embeddings capture relational meaning!

vector('king') - vector('man') + vector('woman') \approx vector('queen') vector('Paris') - vector('France') + vector('Italy') \approx vector('Rome')



Vector Semantics Brown clustering

Brown clustering

- An agglomerative clustering algorithm that clusters words based on which words precede or follow them
- These word clusters can be turned into a kind of vector
- We'll give a very brief sketch here.

Brown clustering algorithm

- Each word is initially assigned to its own cluster.
- We now consider consider merging each pair of clusters. Highest quality merge is chosen.
 - Quality = merges two words that have similar probabilities of preceding and following words
 - (More technically quality = smallest decrease in the likelihood of the corpus according to a class-based language model)
- Clustering proceeds until all words are in one big cluster.

Brown Clusters as vectors

- By tracing the order in which clusters are merged, the model builds a binary tree from bottom to top.
- Each word represented by binary string = path from root to leaf
- Each intermediate node is a cluster

86

• Chairman is 0010, "months" = 01, and verbs = 1



Brown cluster examples

Friday Monday Thursday Wednesday Tuesday Saturday Sunday weekends Sundays Saturdays June March July April January December October November September August pressure temperature permeability density porosity stress velocity viscosity gravity tension anyone someone anybody somebody had hadn't hath would've could've should've must've might've asking telling wondering instructing informing kidding reminding bothering thanking deposing mother wife father son husband brother daughter sister boss uncle great big vast sudden mere sheer gigantic lifelong scant colossal down backwards ashore sideways southward northward overboard aloft downwards adrift

Class-based language model

• Suppose each word was in some class c_{i:}

$$P(W_i / W_{i-1}) = P(c_i / c_{i-1}) P(W_i / c_i)$$

$$P(\operatorname{corpus} \mathcal{C}) = \bigvee_{i=1}^{Y^n} P(c_i / c_{i-1}) P(w_i / c_i)$$