

| Problem: Ambiguities |
| :--- |
| - Headlines: |
| - Enraged Cow Injures Farmer With Ax |
| - Hospitals Are Sued by 7 Foot Doctors |
| - Ban on Nude Dancing on Governor's Desk |
| - Iraqi Head Seeks Arms |
| - Local HS Dropouts Cut in Half |
| - Juvenile Court to Try Shooting Defendant |
| - Stolen Painting Found by Tree |
| - Kids Make Nutritious Snacks |
| - Why are these funny? |



| Grammar: PCFGs |  |
| :---: | :---: |
| - Natural la <br> - PCFGs are <br> - Each "ru <br> - Tree's p <br> - Parsing: G | ch! <br> 375/420 <br> 320/392 <br> 127/539 <br> 32/401 |



| The Problem with Dictionary Lookups |  |
| :---: | :---: |
| 顶部 | ／top／roof／ |
| 顶端 | ／summit／peak／top／apex／／ |
| 顶头 | ／coming directly towards one／top／end／ |
| 盖 | ／lid／top／cover／canopy／build／Gai／ |
| 善帽 | ／surpass／top／ |
| 极 | ／extremely／pole／utmost／top／collect／receive／ |
| 尖峰 | ／peak／top／／ |
| 面 | ／fade／side／surface／aspect／top／face／flour／／ |
| 摘心 | ／top／topping／ |
|  |  |
|  |  |


| MT： 60 Years in 60 Seconds |  |
| :---: | :---: |
|  |  |
| Berkeley＇s frrst MT grant | Statistical MT thrives |
| MT is the＂first＂ <br> non－numeral <br> compute task ALPAC report <br> deems MT bad <br>   | Statistical data－driven approach introduced |
| $\begin{array}{lllll} \stackrel{\circ}{4} & \text { '58 } & \text { '66 } & \text { '90's } & \text { '00's } \end{array}$ |  |


| Data－Driven Machine Translation |  |  |
| :---: | :---: | :---: |
| Target language corpus： |  |  |
| I will get to it soon | See you later | He will do it |
| Sentence－aligned parallel corpus： |  |  |
| Yo lo haré mañana I will do it tomorrow | Hasta pronto See you soon | Hasta pronto <br> See you around |
| Machine translation system： |  |  |
| Yo lo haré pronto Novel Senterce | Model of translation | I will do it soon |






| Fundamental Problems |
| :---: |
| -Representation* <br> - Lexical Analysis (tokenizing)* <br> - Information Extraction* <br> - Comparison (similarity, distance)* <br> - Classification (e.g., for net-nanny service)* <br> - Indexing (to permit fast retrieval) <br> - Retrieval (querying and query processing) |
| *important for Al |

## Document Analysis with LSA: Outline

- Motivation
- Bag-of-words representation
- Stopword elimination, stemming, reference vocabulary
- Vector-space representation
- Document comparison with the cosine similarity measure
- Latent Semantic Analysis

| Representations for Documents |  |
| :---: | :---: |
| - Text string <br> - Image (I.e., .jpg, .gif, and .png files) <br> - linguistically structured files: PostScript, Portable Doc. Format (PDF), XML. <br> - Vector: e.g., bag-of-words <br> - Hypertext, hypermedia | tietbrequse |


| Bag-of-Words Representation |
| :--- |
| A multiset is a collection like a set, but which allows <br> duplicates (any number of copies) of elements. <br> $\{a, b, c\}$ is a set. (It is also a multiset.) <br> \{ a, a, b, c, c, c $\}$ is not a set, but it is a multiset. <br> \{ c, a, b, a, c, c $\}$ is the same multiset. (Order doesn't <br> A multiset is also called a bag. |


| Bag-Of-Words (continued) |
| :---: |
| Let document $\mathrm{D}=$ <br> "The big fox jumped over the big fence." <br> The bag representation is: <br> \{big, big, fence, fox, jumped, over, the, the \} |
| For notational consistency, we use alphabetical order. <br> Also, we omit punctuation and normalize the case. <br> The ordering information in the document is lost. But <br> this is OK for some applications. |


| Eliminating Stopwords |  |
| :---: | :---: |
| In information retrieval and some other types of document analysis, we often begin by deleting words that don't carry much meaning or that are so common that they do little to distinguish one document from another. Such words are called stopwords. <br> Examples: (articles) a, an, the; (quantifiers) any, some, only, many, all, no; (pronouns) I, you, it, he, she, they, me, him, her, them, his, hers, their, theirs, my, mine, your, our, yours, ours, this, that, these, those, who, whom, which; (prepositions) above, at, behind, below, beside, for, in, into, of, on, onto, over, under; (verbs) am, are, be, been, is, were, go, gone, went, had, have, do, did, can, could, will, would, might, may, must; (conjunctions) and, but, if, then, not, neither, nor, either, or; (other) yes, perhaps, first, last, there, where, when. |  |


| Stemming |
| :--- |
|  |
| In order to detect similarities among words, it often <br> helps to perform stemming. We typically stem a <br> word by removing its suffixes, leaving the basic <br> word, or "uninflecting" the word <br> • apples $\rightarrow$ apple <br>  <br> • cacti $\rightarrow$ cactus <br> • swimming $\rightarrow$ swim <br>  <br> • swam $\rightarrow$ swim |


| Reference Vocabulary |
| :--- |
| A counterpart to stopwords is the reference vocabulary. <br> These are the words that ARE allowed in document <br> representations. <br> These are all stemmed, and are not stopwords. <br> There might be several hundred or even thousands of <br> terms in a reference vocabulary for real document <br> processing. |


| Vector representation |
| :--- |
| Assume we have a reference vocabulary of words that <br> might appear in our documents. <br> \{apple, big, cat, dog, fence, fox, jumped, over, the, zoo\} <br> We represent our bag <br> \{ big, big, fence, fox, jumped, over, the, the \} |
| by giving a vector (list) of occurrence counts of each <br> reference term in the document: <br> $\quad[0,2,0,0,1,1,1,1,2,0]$ |
| If there are $n$ terms in the reference vocabulary, then each document is <br> represented by a point in an n-dimensional space. |


| Indexing |
| :--- |
| Create links from terms to documents or <br> document parts <br> (a) concordance <br> (b) table of contents <br> (c) book index <br> (d) index for a search engine <br> (e) database index for a relation (table) |

## Concordance

A concordance for a document is a sort of
dictionary that lists, for each word that occurs in
the document the sentences or lines in which it occurs.
"document":
A concordance for a document is a sort of dictionary that
lists, for each word that occurs in the document the

## "occurs":

that lists, for each word that occurs in the document the
sentences or lines in which it occurs.

## Search Engine Index

Query terms are organized into a large table or tree that can be quickly searched.
(e.g., large hash-table in memory, or a B-Tree with its top levels in memory).

Associated with each term is a list of occurrences, typically consisting of Document IDs or URLs.

| Cosine Similarity Function |
| :---: |
| Document 1: <br> "All Blues. First the key to last night's notes." <br> Document 2: <br> "How to get your message across. Restate your key <br> points first and last. " <br> Reference vocabulary: <br> \{ across, blue, first, key, last, message, night, <br> note, point, restate, zebra \} |

## Cosine Similarity (cont)

Dot product (same as "inner product")
$[0,1,1,1,1,0,1,1,0,0,0] \cdot[1,0,1,1,1,1,0,0,1,1,0]$
$=0 \cdot 1+1 \cdot 0+1 \cdot 1+1 \cdot 1+1 \cdot 1+0 \cdot 1+1 \cdot 0+1 \cdot 0+0$
$1+0 \cdot 1+0 \cdot 0=3$
Normalized:
$\cos \theta=\left(\mathrm{v}_{1} \cdot \mathrm{v}_{2}\right) /\left(\left\|\mathrm{v}_{1}\right\|\left\|\mathrm{v}_{2}\right\|\right)$
$\|v\|=\sqrt{v \cdot v} \quad \cos \theta=3 /(\sqrt{6} \sqrt{ } 7) \approx 0.4629$.


| Properties of the Cosine Similarity |
| :--- |
| $\cos \theta=0$ means that the document vectors are orthogonal <br> and the documents have no reference vocabulary <br> occurrences in common. <br> $\cos \theta=1$ means that the documents are either identical or <br> the vectors point in the same direction in the n-dim space. <br> That is, the documents share the same distribution of <br> occurrences of the reference terms. |

## Latent Semantic Analysis

A problem with the cosine similarity function:
Unless both documents use the same term for something, the similarity is not recognized.
"Computer learning environments have a great future."
"Educational technology offers wonderful potential."
cosine similarity is 0 .

| LSA (COntinued) |
| :--- |
| The semantic space for LSA is obtained from a set of <br> documents given in advance. <br> The space is created using matrix factorization via the <br> Singular Value Decomposition (SVD) method. <br> This is computationally costly, but modern computers are <br> powerful enough to do it. <br> For more details, see Chapter 16 of Introduction to Python <br> for Artificial Intelligence. |


| Singular Value Decomposition |
| :---: |
| Given term-document matrix A, having $t$ rows and $d$ columns, find TSD such that: $A=T S D$ <br> T is at by t orthonormal matrix <br> D is a d by d orthonormal matrix <br> $S$ is an $m$ by $m$ diagonal matrix, where $m$ is the rank of $A$. <br> import LinearAlgebra as LA <br> (TSD) = LA.singular_value_decomposition (A) |


| Example |
| :--- |
| $d_{1}=$ "the brown weasel followed the fox and stole the eggs" <br> $d_{2}=$ "behind the fence the thief fled with half a dozen" <br> $d_{3}=$ "artificial limbs can offer full mobility" |
| Documents used to create a semantic space: <br> "the lazy brown fox jumped over the fence" <br> "the thief jumped the lazy fence and fled" <br> "artificial intelligence is full of surprises" |
| cossim $\left(d_{1}, d_{2}\right)=0 \quad$ Without LSA, $d_{1}$ and $d_{2}$ seem dissimilar. <br> cossim $\left(d_{1}, d_{2}\right)=1 \quad$ With LSA, they are completely similar. <br> cossim $\left(d_{1}, d_{3}\right)=$ cossim( $\left.d_{1}, d_{3}{ }^{\prime}\right)=0 \quad$ But LSA does not <br> make $d_{3}$ any more similar to the others. |

