

vector space retrieval



what is a retrieval model?

- Model is an idealization or abstraction of an actual process
- Mathematical models are used to study the properties of the process, draw conclusions, make predictions
- Conclusions derived from a model depend on whether the model is a good approximation of the actual situation
- Statistical models represent repetitive processes, make predictions about frequencies of interesting events
- Retrieval models can describe the computational process
 - e.g. how documents are ranked
 - Note that how documents or indexes are *stored* is implementation
- Retrieval models can attempt to describe the human process
 - e.g. the information need, interaction
 - Few do so meaningfully
- Retrieval models have an explicit or implicit definition of relevance



retrieval models

today

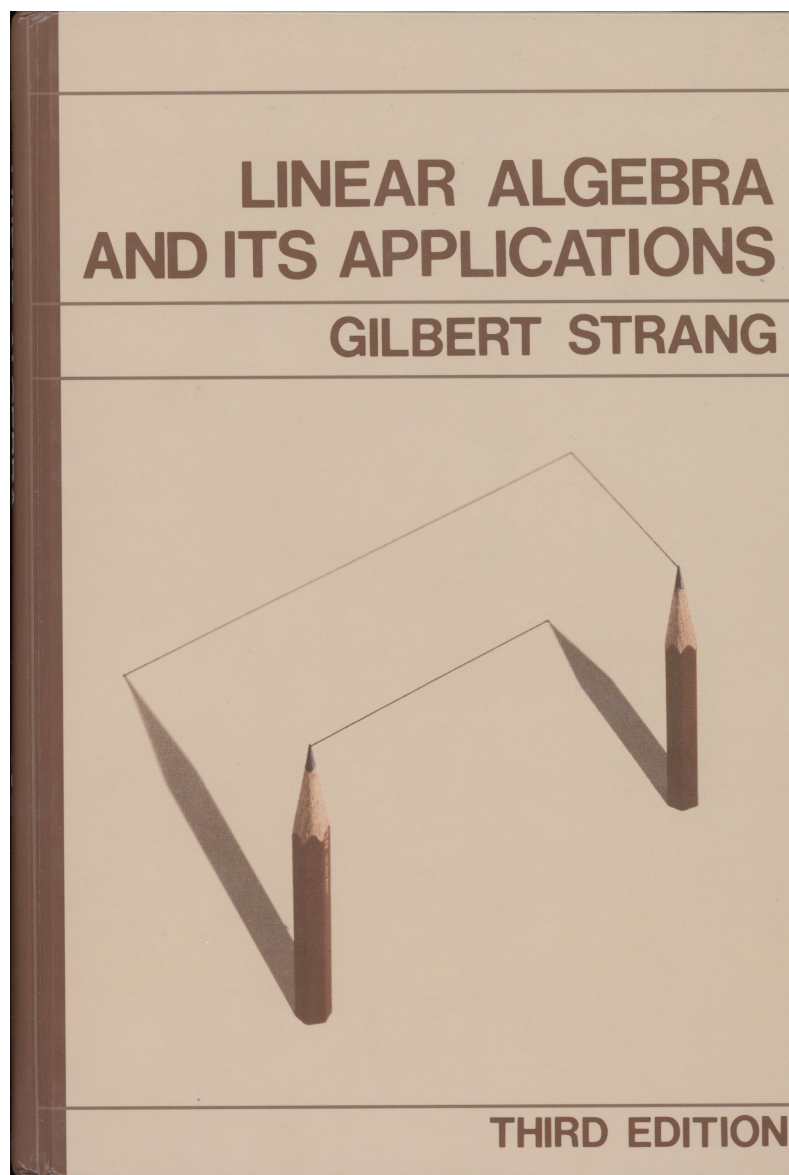
- boolean
- vector space
- latent semantic indexing
- statistical language
- inference network
- hyperlink based



outline

- review: geometry, linear algebra
- vector space model
- vector selection
- similarity
- weighting schemes
- latent semantic indexing

linear algebra



vectors

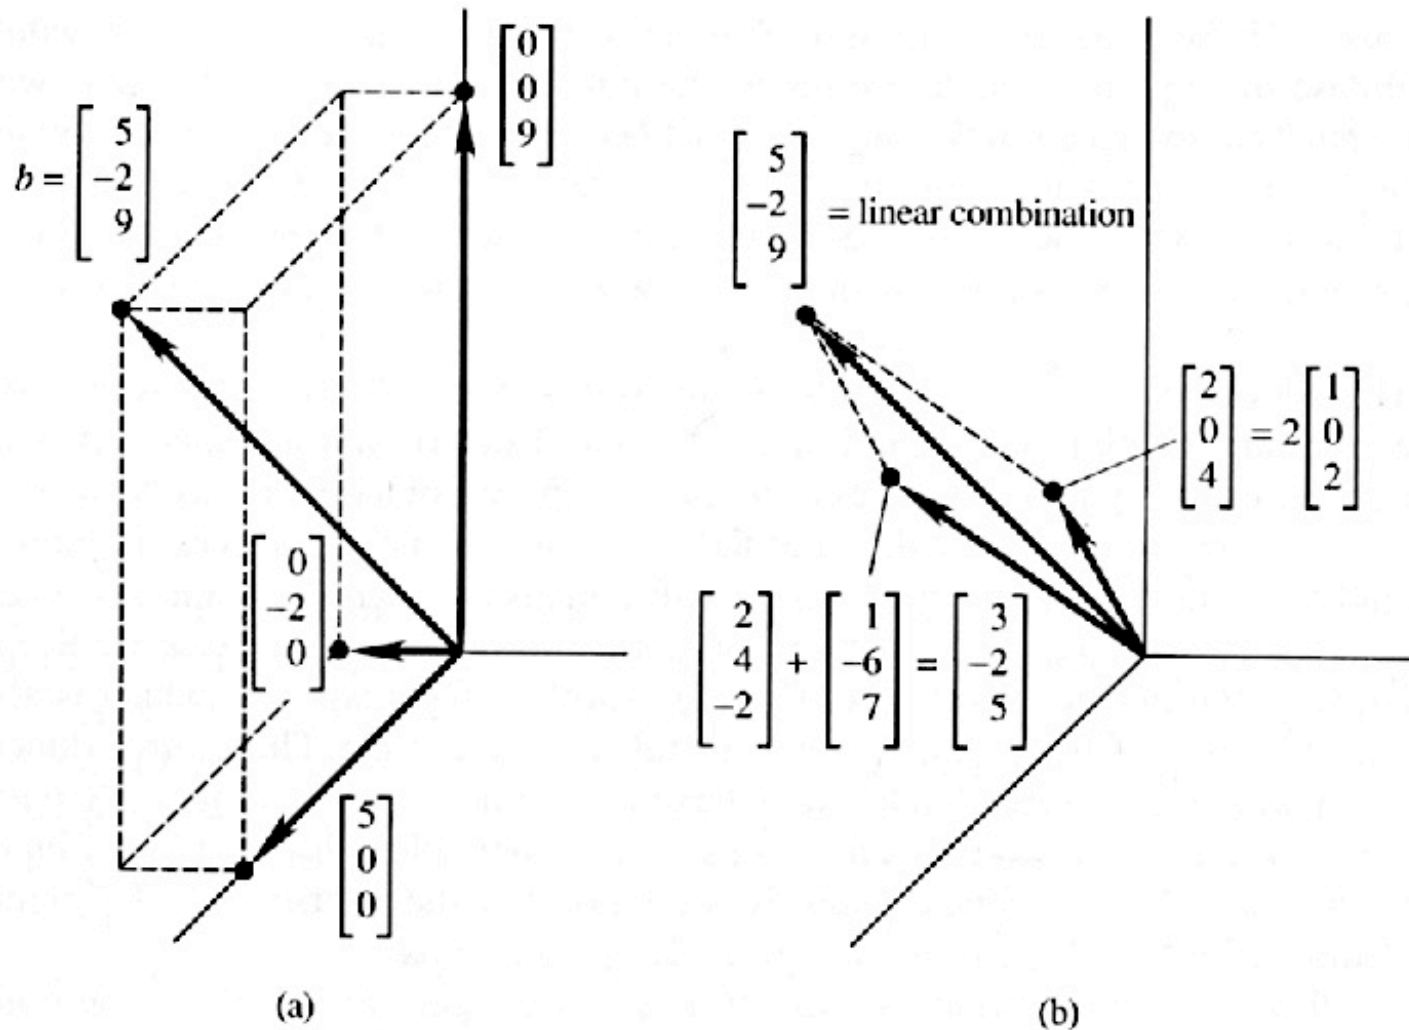
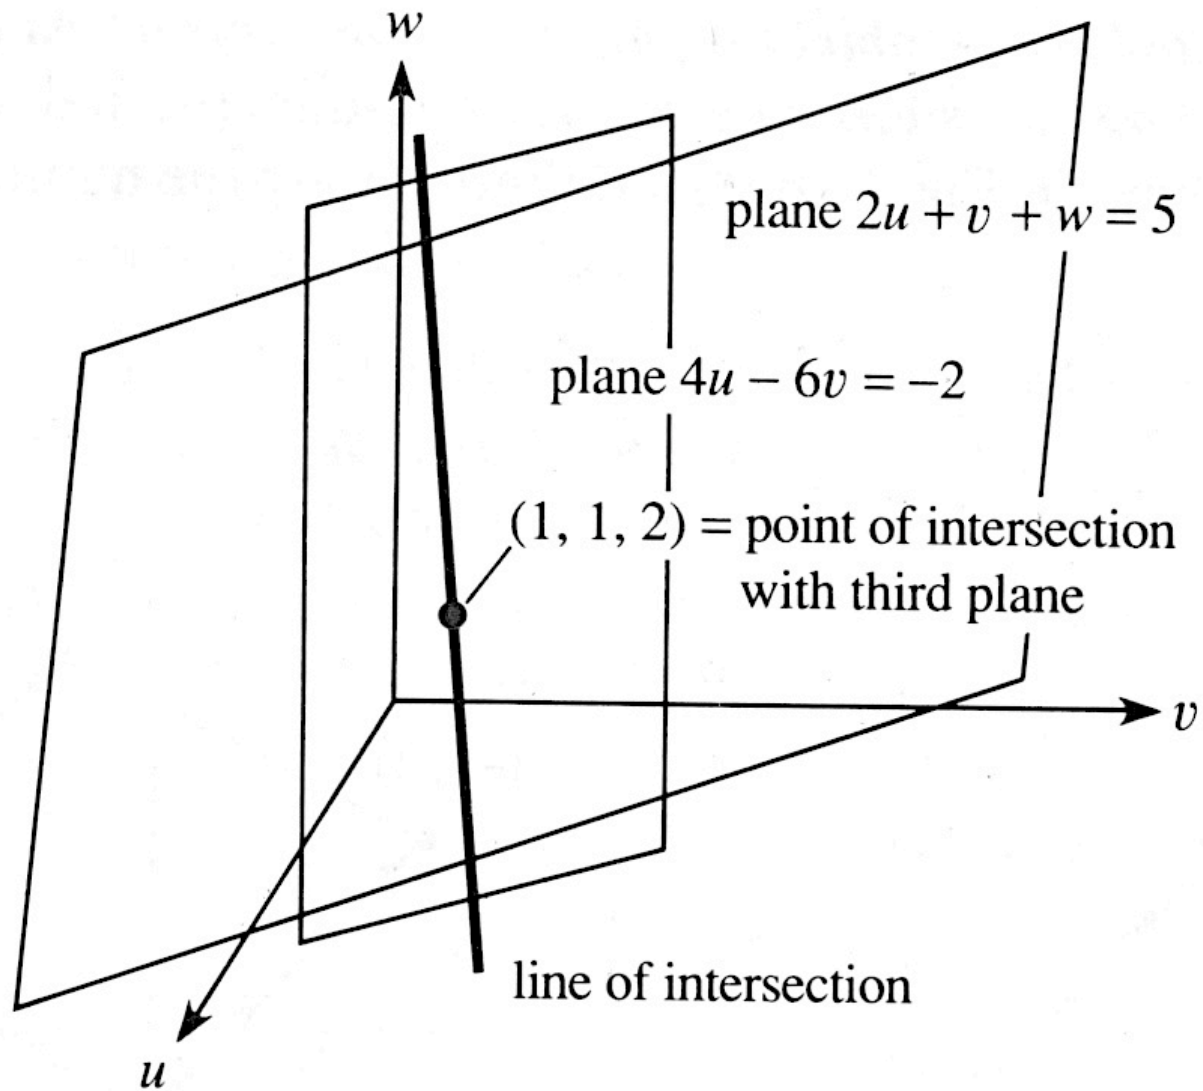


Fig. 1.3. The column picture: linear combination of columns equals b .

subspaces



linear independence, base, dimension, rank



- vector \bar{x} is linear dependent of vectors $\bar{y}_1, \bar{y}_2, \dots, \bar{y}_t$ if there exists real numbers c_1, c_2, \dots, c_t such that

$$\bar{x} = c_1\bar{y}_1 + c_2\bar{y}_2 + \dots c_t\bar{y}_t$$

- base of a vectorial space = maximal set of linear independent vectors. All bases of a given space have the same dimension (dimension of the space)
- $\text{rank}(A) =$ maximum number of rows/columns linear independent
- $\text{rank}(A) =$ dimension of the subspace spanned by A

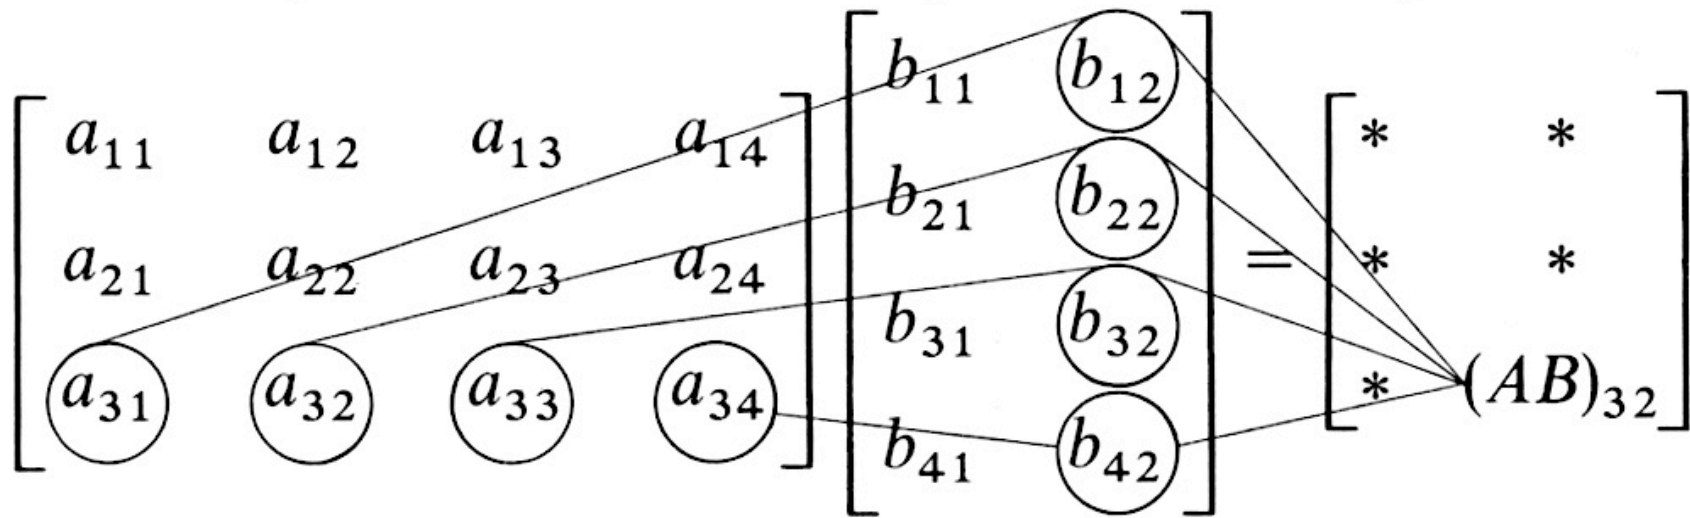
matrix multiplication



$$(AB)_{32} = a_{31}b_{12} + a_{32}b_{22} + a_{33}b_{32} + a_{34}b_{42}.$$

3 by 4 matrix

4 by 2 matrix 3 by 2 matrix





dot product, norm

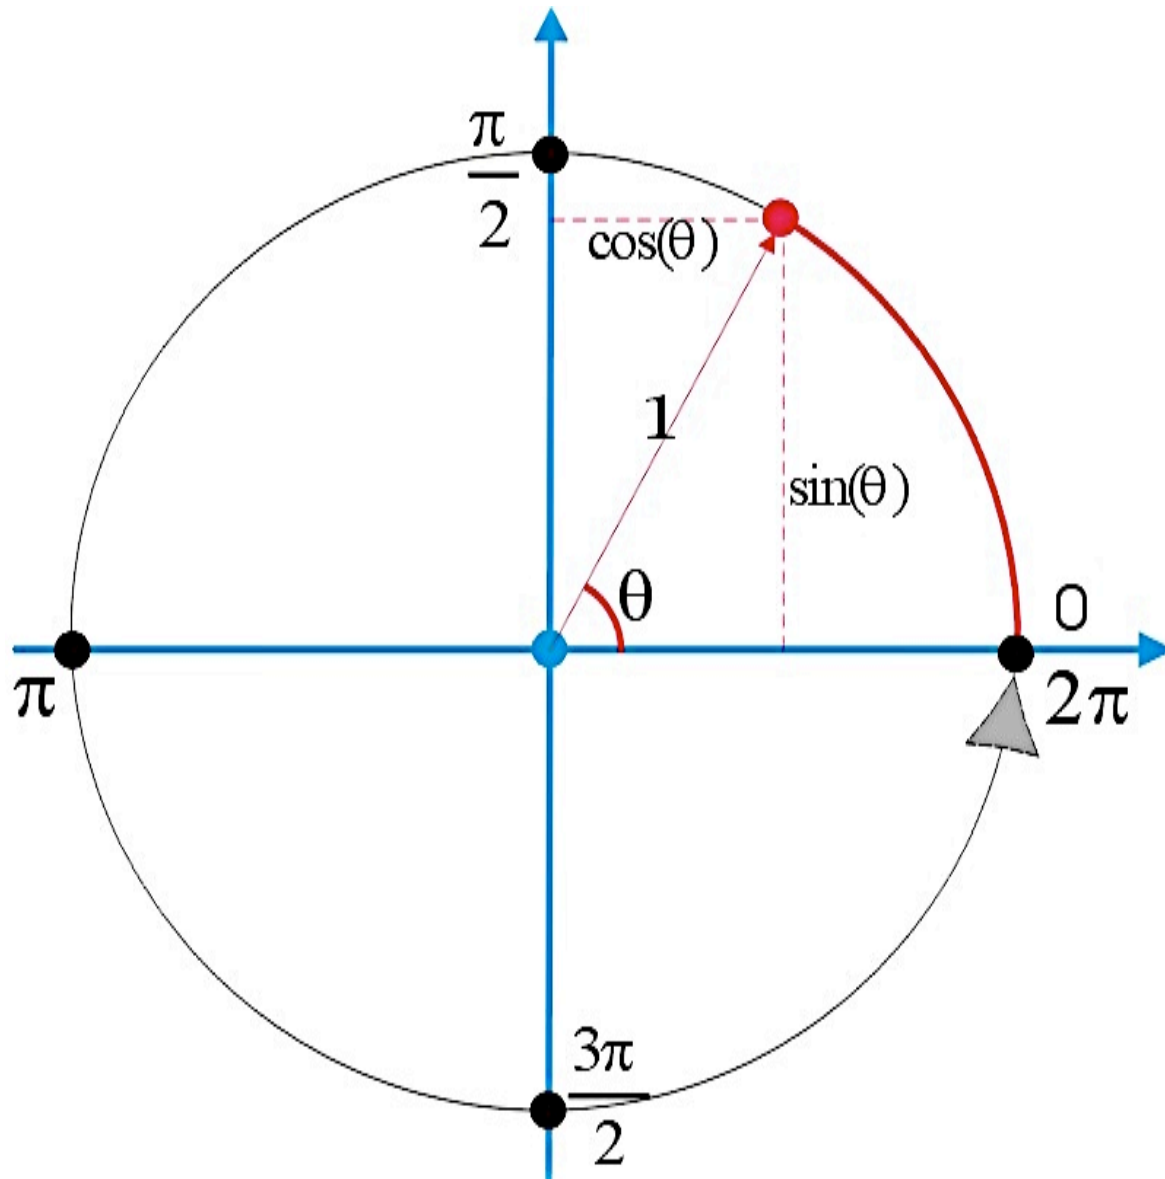
- dot product of 2 same dimension arrays is simply the matrix product with result a real number

- $x = (x_1, x_2, \dots, x_n); y = (y_1, y_2, \dots, y_n)$ then
 $\langle x \cdot y \rangle = x * y^T = \sum_{i=1}^n x_i y_i$

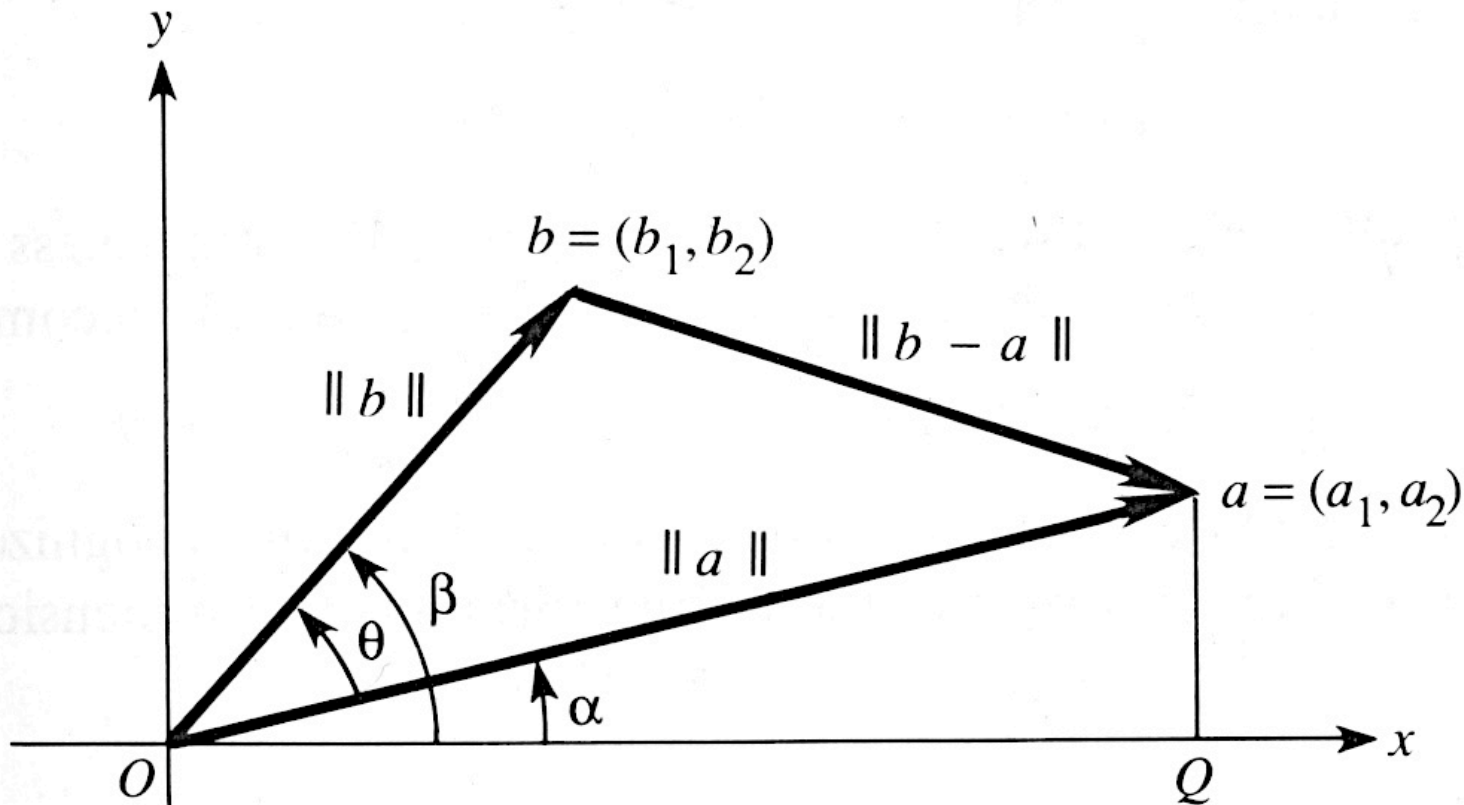
- L_2 norm : $\|x\| = \sqrt{\langle x \cdot x \rangle}$

- normalization: $\bar{x} = \frac{x}{\|x\|}; \|\bar{x}\| = 1$

cosine



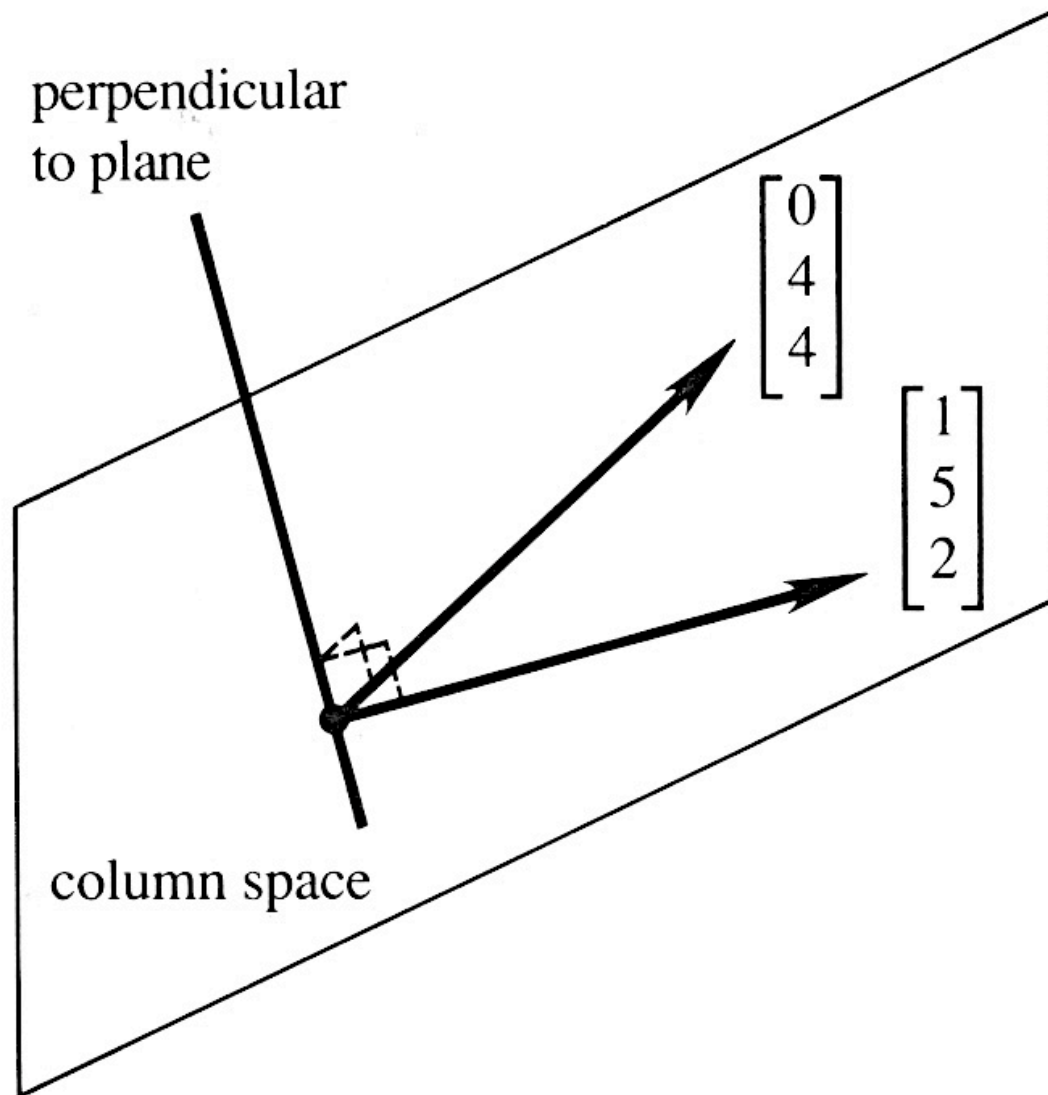
cosine computation



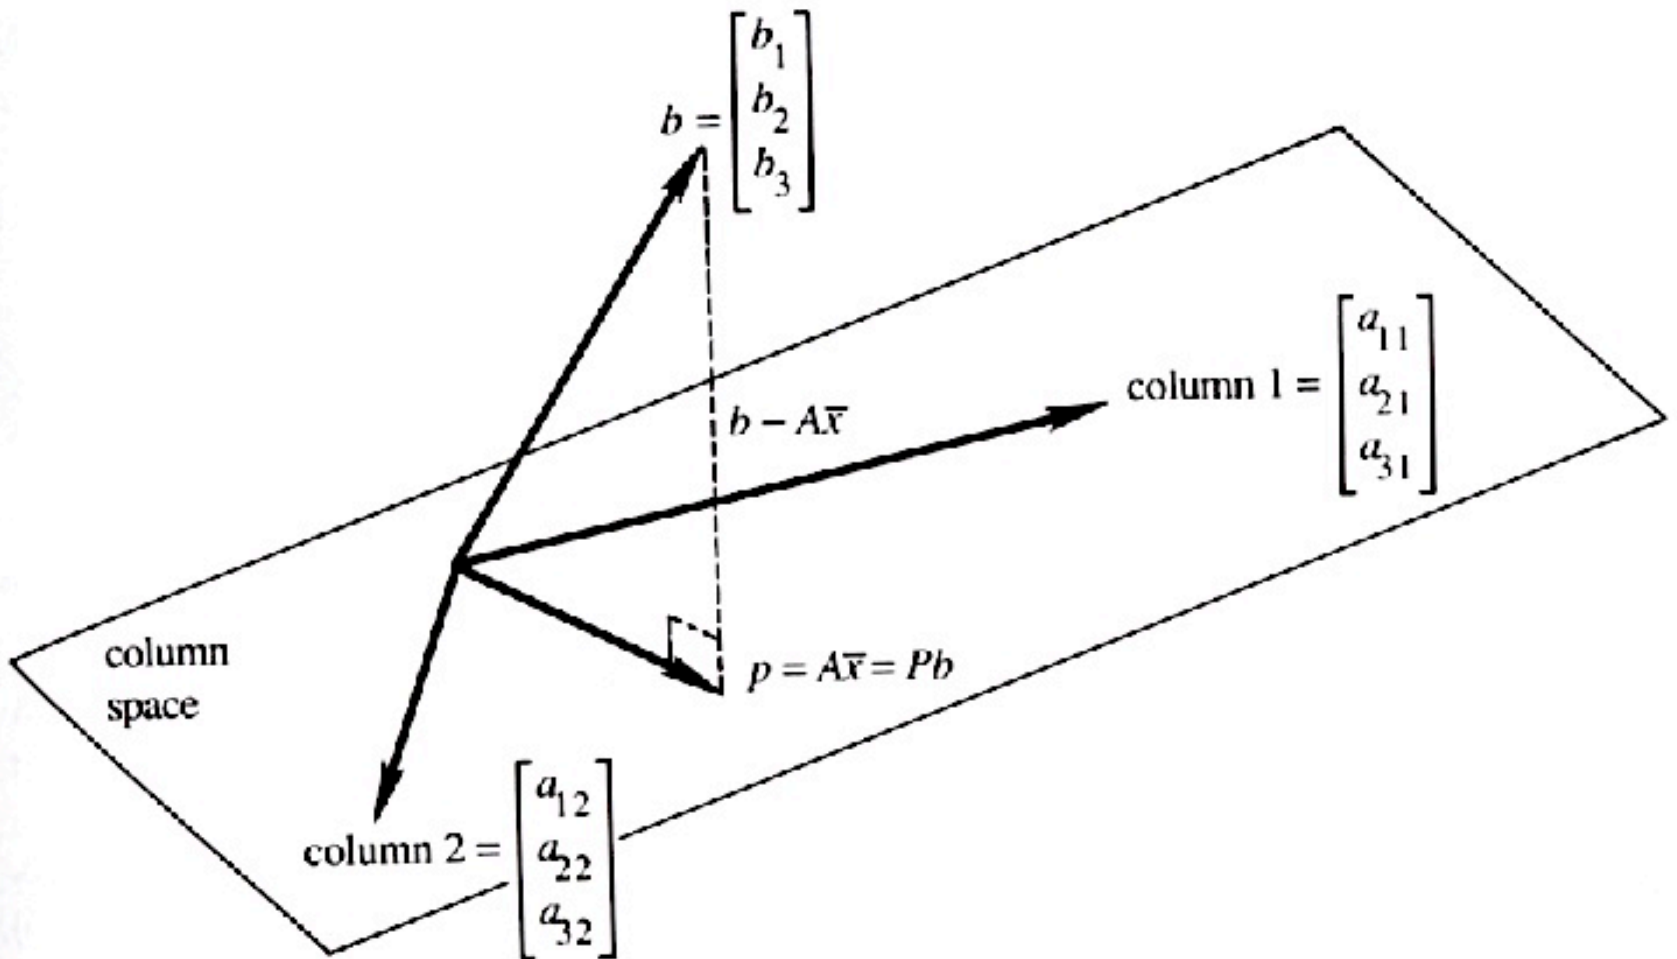
$$\cos(\theta) = \frac{a \cdot b}{|a||b|}$$

$$\cos(\theta) = \cos(\beta - \alpha) = \cos(\beta) \cos(\alpha) + \sin(\beta) \sin(\alpha)$$

orthogonality



projections





$A=LDU$ factorization

- For any $m \times n$ matrix A , there exists a permutation matrix P , a lower triangular matrix L with unit diagonal and an $m \times n$ echelon matrix U such that $PA = LU$

$$U = \begin{bmatrix} \textcircled{*} & * & * & * & * & * & * & * & * \\ 0 & \textcircled{*} & * & * & * & * & * & * & * \\ 0 & 0 & 0 & \textcircled{*} & * & * & * & * & * \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \textcircled{*} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

- For any $n \times n$ matrix A , there exists L, U lower and upper triangular with unit diagonals, D a diagonal matrix of pivots and P a permutation matrix such that $PA = LDU$
- If A is symmetric ($A = A^T$) then there is no need for P and $U = L^T$: $A = LDL^T$



eigenvalues, eigenvectors

- λ is an eigenvalue for matrix A iff

$$\det(A - \lambda I) = 0$$

- every eigenvalue has a correspondent non-zero eigenvector x that satisfies

$$(A - \lambda I)x = 0 \text{ or } Ax = \lambda x$$

in other words Ax and x have same direction

- sum of eigenvalues = $\text{trace}(A)$ = sum of diagonal
- product of eigenvalues = $\det(A)$
- eigenvalues of a upper/lower triangular matrix are the diagonal entries



matrix diagonal form

- if A has linearly independent eigenvectors y_1, y_2, \dots, y_n and S is the matrix having those as columns, $S = [y_1 y_2 \dots y_n]$, then S is invertible and

$$S^{-1}AS = \Lambda = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \dots & \\ & & & \lambda_n \end{bmatrix}, \text{ the diagonal}$$

matrix of eigenvalues of A .

- $A = S\Lambda S^{-1}$
- no repeated eigenvalue \Rightarrow indep. eigenvect
- A symmetric $A^T = A \Rightarrow S$ orthogonal: $S^T S = 1$
- S is not unique
- $AS = S\Lambda$ holds iff S has eigenvect as columns
- not all matrices are diagonalizable



singular value decomposition

- if A is $m \times n, m > n$ real matrix then it can be decomposed as $A = UDV^T$ where

- U is $m \times n$; D, V are $n \times n$
- U, V are orthogonal : $U^T U = V^T V = \mathbf{1}_{n \times n}$
- D is diagonal, its entries are the square roots of eigenvalues of $A^T A$



outline

- review: geometry, linear algebra
- vector space model
- vector selection
- similarity
- weighting schemes
- latent semantic indexing



vector space

- represent documents and queries as vectors in the term space
 - issue : find the right coefficients
- use a geometric similarity measure, often angle-related
 - issue: normalization



mapping to vectors

- terms: an axis for each term
 - vectors corresponding to terms are canonical
- document = sum of vectors corresponding to terms contained in doc
- queries treated the same as documents

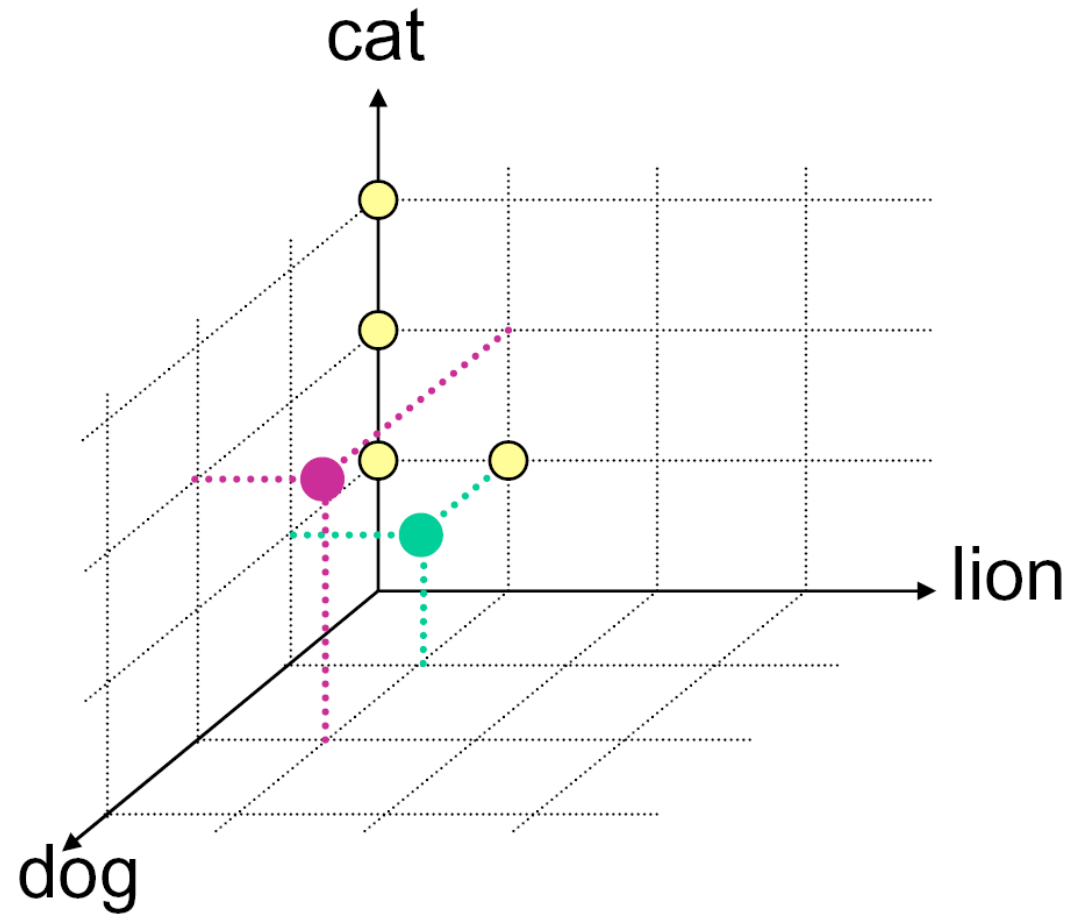


coefficients

- The coefficients (vector lengths, term weights) represent term presence, importance, or “aboutness”
 - Magnitude along each dimension
- Model gives no guidance on how to set term weights
- Some common choices:
 - Binary: 1 = term is present, 0 = term not present in document
 - *tf*: The frequency of the term in the document
 - *tf* • *idf*: *idf* indicates the discriminatory power of the term
- Tf·idf is far and away the most common
 - Numerous variations...

raw tf weights

- cat
- cat cat
- cat cat cat
- cat lion
- lion cat
- cat lion dog
- cat cat lion dog dog





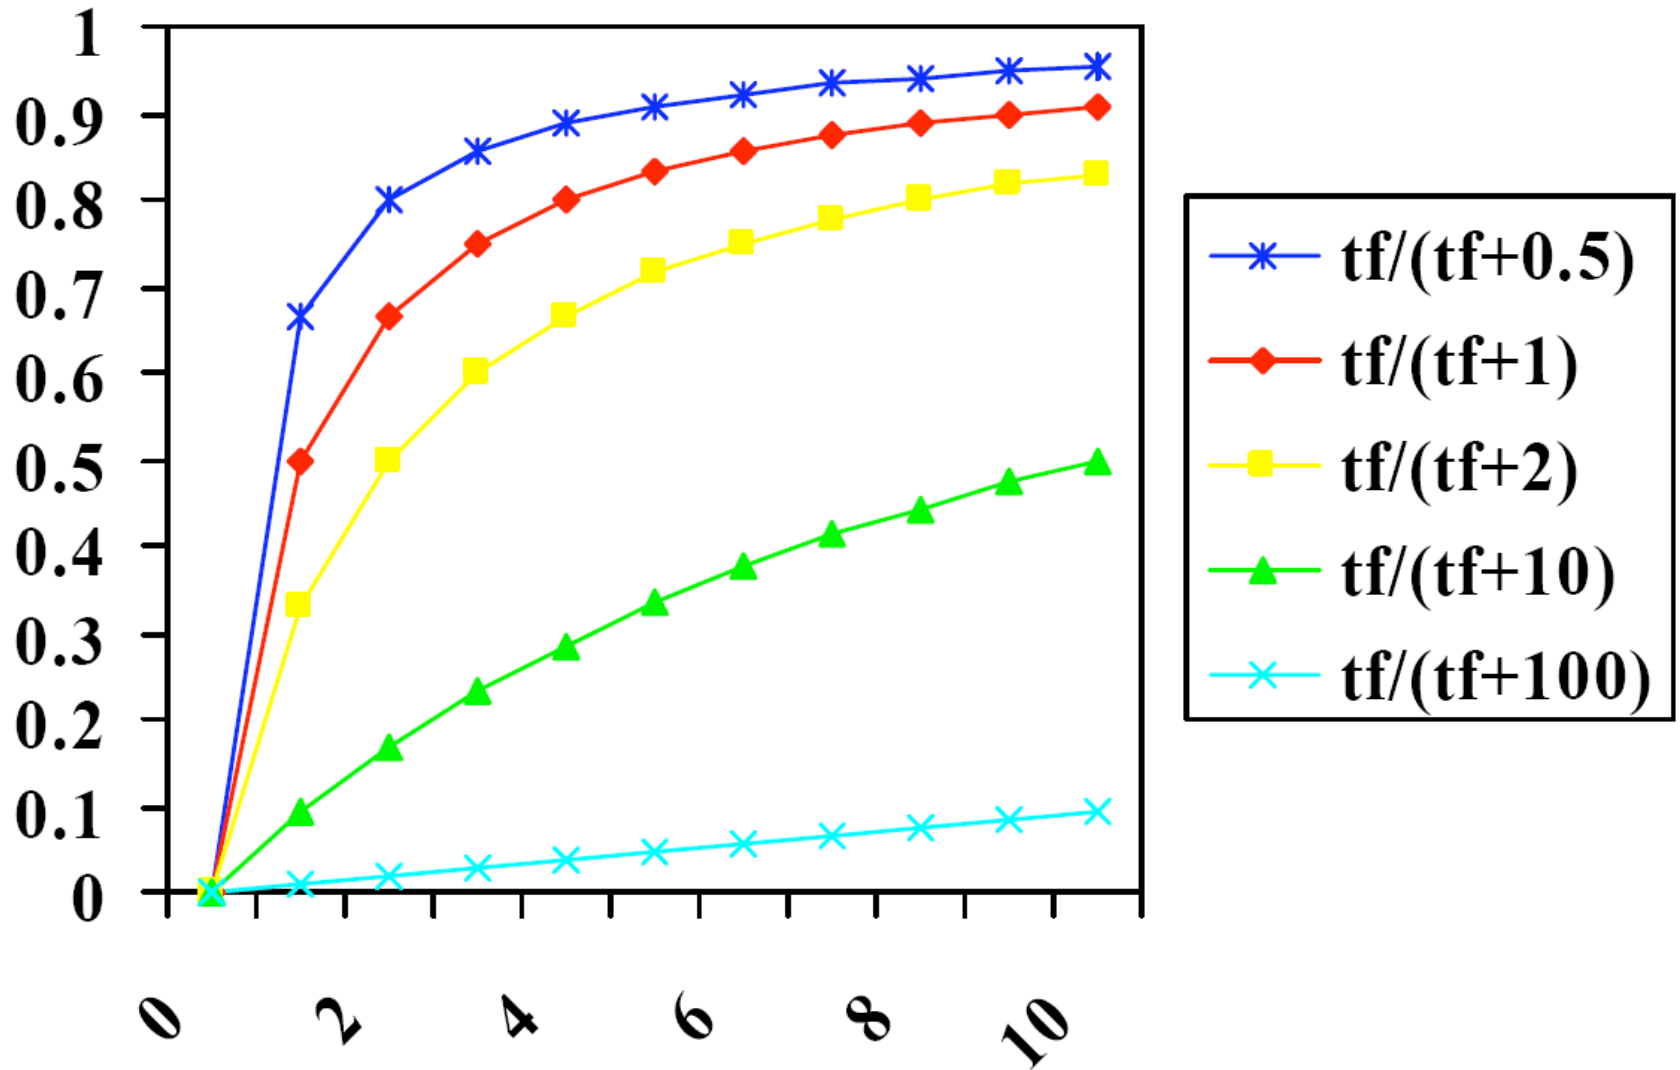
tf = term frequency

- raw-tf (tf) = count of 'term' in document

- Robertson tf (okapi tf)
$$\frac{tf}{tf + k + c \cdot \frac{doclen}{avg.doclen}}$$
 - based on a set of simple criteria loosely connected to 2-Poisson model
 - popular $k=0.5$; $c=1.5$
 - basic formula is $tf / (k+tf)$
 - document length = verbosity factor

- many variants

Robertson tf






IDF weights

- Inverse Document Frequency
- Used to weight terms based on frequency in the corpus
- Fixed, it can be precomputed for each term
- basic formula $IDF(t) = \log\left(\frac{N}{N_t}\right)$
 - $N = \#$ of docs
 - $N_t = \#$ of docs containing term t

tf-idf

- 
- $tf * idf$
 - the weight on every term is $tf(t,d)*idf(t)$
 - sometimes variants on tf , IDF
 - no satisfactory model behind these combinations



outline

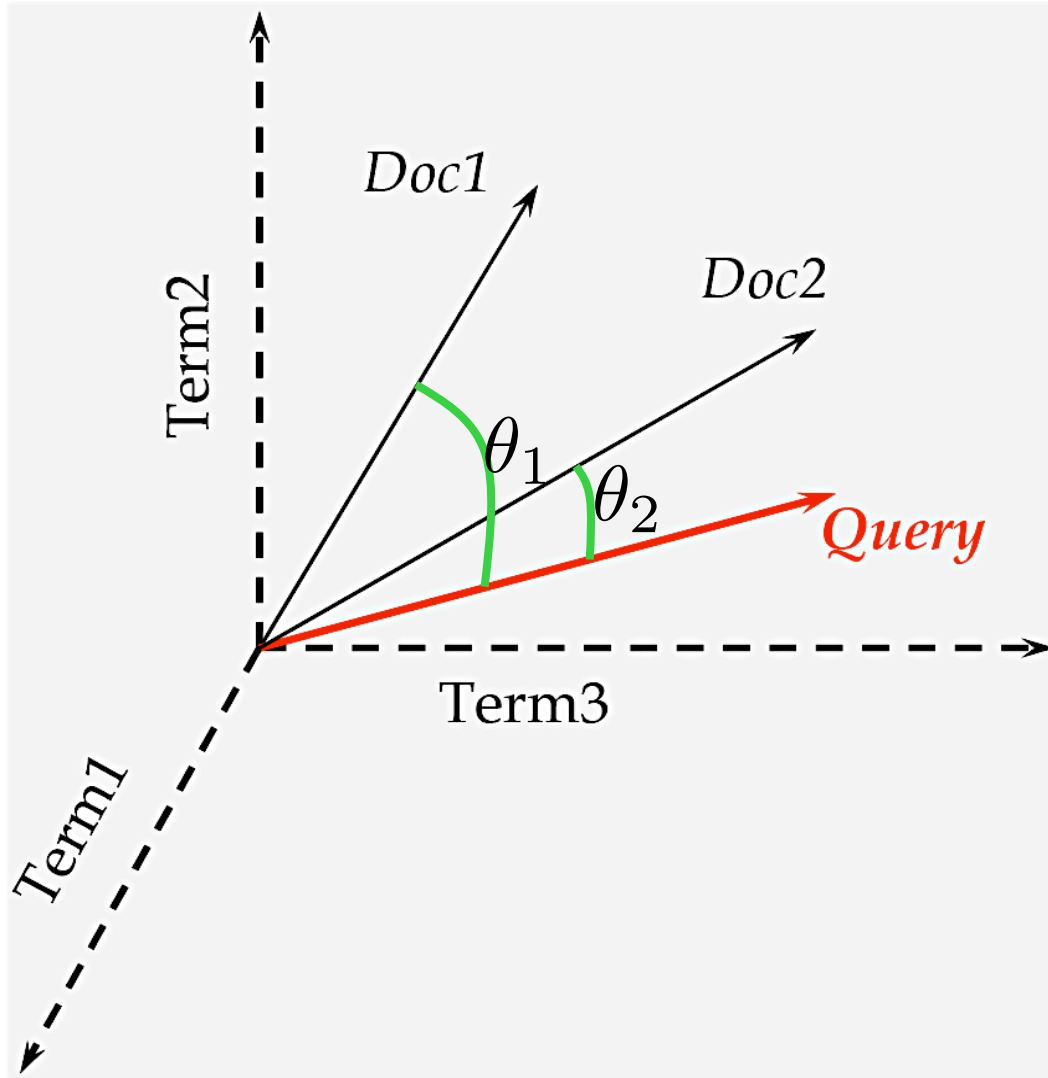
- review: geometry, linear algebra
- vector space model
- vector selection
- similarity
- weighting schemes
- latent semantic indexing



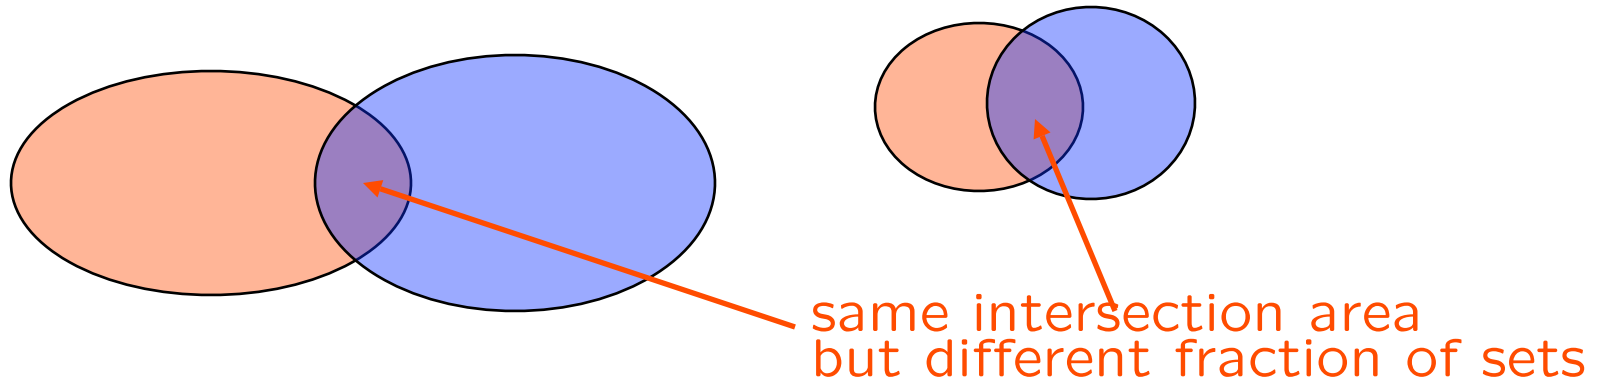
common similarity measures

Sim(X,Y)	Binary Term Vectors	Weighted Term Vectors
Inner product	$ X \cap Y $	$\sum x_i \cdot y_i$
Dice coefficient	$\frac{2 X \cap Y }{ X + Y }$	$\frac{2\sum x_i \cdot y_i}{\sum x_i^2 + \sum y_i^2}$
Cosine coefficient	$\frac{ X \cap Y }{\sqrt{ X } \sqrt{ Y }}$	$\frac{\sum x_i \cdot y_i}{\sqrt{\sum x_i^2} \cdot \sqrt{\sum y_i^2}}$
Jaccard coefficient	$\frac{ X \cap Y }{ X + Y - X \cap Y }$	$\frac{\sum x_i \cdot y_i}{\sum x_i^2 + \sum y_i^2 - \sum x_i \cdot y_i}$

vector similarity: cosine



similarity, normalized



$$\textit{similarity} = \frac{|\text{intersection}|}{|\text{set}_1| \cdot |\text{set}_2|}$$

- the size of intersection alone is meaningless
- often divided by sizes of sets
- same for vectors, using norm
 - by normalizing vectors, cosine does not change



cosine similarity: example

$$D_1 = (0.5T_1 + 0.8T_2 + 0.3T_3)$$

$$Q = (1.5T_1 + 1T_2 + 0T_3)$$

$$\text{Sim}(D_1, Q) = \frac{(0.5 \times 1.5) + (0.8 \times 1)}{\sqrt{(0.5^2 + 0.8^2 + 0.3^2)(1.5^2 + 1^2)}}$$

$$= \frac{1.55}{\sqrt{.98 \times 3.25}}$$

$$= .868$$



cosine example, normalized

$$\underline{D_1 = (0.5T_1 + 0.8T_2 + 0.3T_3)}$$

$$\underline{Q = (1.5T_1 + 1T_2 + 0T_3)}$$

$$\underline{D'_1 = (0.5T_1 + 0.8T_2 + 0.3T_3)/\sqrt{0.98}}$$
$$\underline{\approx 0.51T_1 + 0.82T_2 + 0.31T_3}$$

$$\underline{Q' = (1.5T_1 + 1T_2 + 0T_3)/\sqrt{3.25}}$$
$$\underline{\approx 0.83T_1 + 0.555T_2}$$

$$\begin{aligned} \text{Sim}(D_1, Q) &= \text{Sim}(D'_1, Q') \\ &= \frac{(0.51 \times 0.83) + (0.82 \times 0.555)}{\sqrt{(0.51^2 + 0.82^2 + 0.31^2)(0.83^2 + 0.555^2)}} \\ &= (0.51 \times 0.83) + (0.82 \times 0.555) \\ &= 0.878 \end{aligned}$$

round-off error, should be the same \rightarrow \approx 0.868 (from earlier slide)

similarity example

T_1 T_2 T_3

$$D_1 = 3 \text{ cat} + 1 \text{ dog} + 4 \text{ lion}$$
$$D_2 = 8 \text{ cat} + 2 \text{ dog} + 6 \text{ lion}$$

$$D_1 = (3T_1 + 1T_2 + 4T_3)$$
$$D_2 = (8T_1 + 2T_2 + 6T_3)$$

Correlated Terms

	Term	cat	dog	lion
T_1	cat	1.00	-0.20	0.50
T_2	dog	-0.20	1.00	-0.40
T_3	lion	0.50	-0.40	1.00

$$\begin{aligned} \text{Sim}(D_1, Q) &= (3T_1 + 1T_2 + 4T_3) \cdot (2T_2) \\ &= 6T_1 \cdot T_2 + 2T_2 \cdot T_2 + 8T_3 \cdot T_2 \\ &= -6 \cdot 0.2 + 2 \cdot 1 - 8 \cdot 0.4 \\ &= -1.2 + 2 - 3.2 \\ &= -2.4 \end{aligned}$$

$$Q = 2 \text{ dog}$$

$$Q = (0T_1 + 2T_2 + 0T_3)$$

Orthogonal Terms

	Term	cat	dog	lion
T_1	cat	1.00	0.00	0.00
T_2	dog	0.00	1.00	0.00
T_3	lion	0.00	0.00	1.00

$$\begin{aligned} \text{Sim}(D_1, Q) &= 3 \cdot 0 + 1 \cdot 2 + 4 \cdot 0 \\ &= 2 \end{aligned}$$



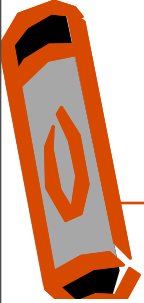
tf-idf base similarity formula

$$\sum_t (\mathbf{TF}_{query}(t) \cdot \mathbf{IDF}_{query}(t)) \cdot (\mathbf{TF}_{doc}(t) \cdot \mathbf{IDF}_{doc}(t))$$

$$||doc|| \cdot ||query||$$

- many options for \mathbf{TF}_{query} and \mathbf{TF}_{doc}
 - raw tf, Robertson tf, Lucene etc
 - try to come up with yours
- some options for \mathbf{IDF}_{doc}
- \mathbf{IDF}_{query} sometimes not considered
- normalization is critical

Lucene comparison



$$\sum_t \left(\frac{\text{tf}_{q,t} \cdot \text{idf}_t}{\text{norm}_q} \cdot \frac{\text{tf}_{d,t} \cdot \text{idf}_t}{\text{norm}_{d,t}} \cdot \text{boost}_t \right) \cdot \frac{\text{overlap}(q, d)}{|q|}$$

tf-idf from document

User-specified boost

Length-normalized query weight

Term normalization is square root of number of tokens in d that are in the same field as t

Proportion of query matched



complicated formulas

$$w_{t,d} = \frac{tf_{d,t} \cdot \log(N/df_t + 1)}{\sqrt{\text{number of tokens in } d \text{ in the same field as } t}}$$

$$\frac{\left(\frac{1}{2} + \frac{1}{2} \frac{tf_{t,d}}{\max(tf_{*,d})} \right) \cdot \log \frac{N}{n_t}}{\left[\sum_t \left(\left(\frac{1}{2} + \frac{1}{2} \frac{tf_{t,d}}{\max(tf_{*,d})} \right) \cdot \log \frac{N}{n_t} \right)^2 \right]^{0.5}}$$



outline

- review: geometry, linear algebra
- vector space model
- vector selection
- similarity
- weighting schemes
- latent semantic indexing



LSI

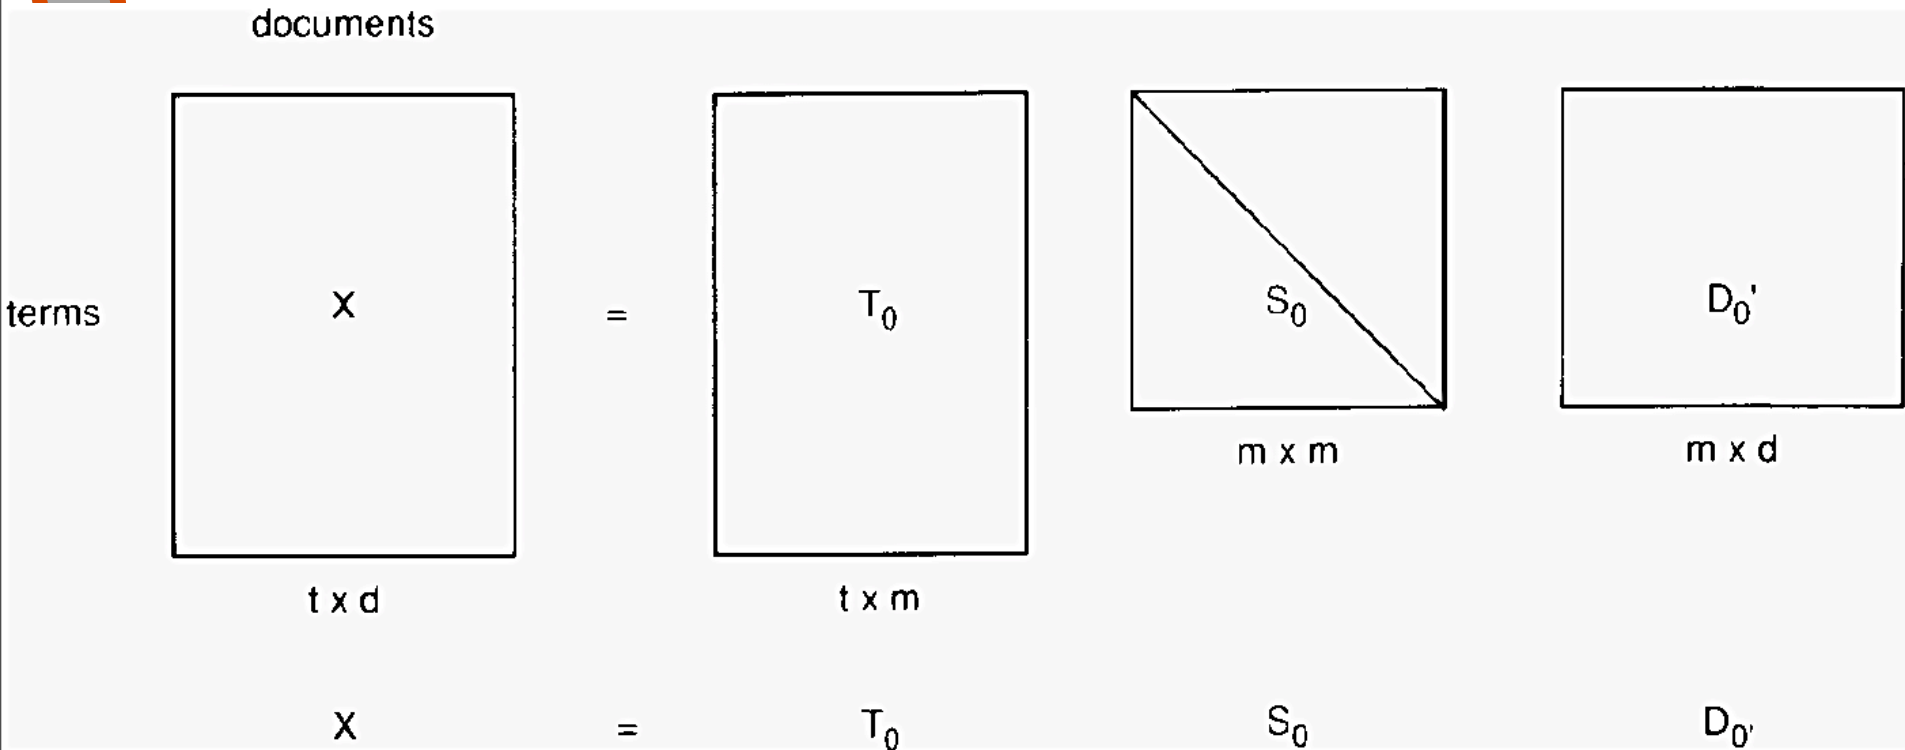
- Variant of the vector space model
- Use Singular Value Decomposition (a dimensionality reduction technique) to identify uncorrelated, significant basis vectors or factors
 - Rather than non-independent terms
- Replace original words with a subset of the new factors (say 100) in both documents and queries
- Compute similarities in this new space
- Computationally expensive, uncertain effectiveness



dimensionality reduction


- when the representation space is rich
- but data is lying in a small subspace
- that is when some eigenvalues are zero
 - non-exact: ignore smallest eigenvalues, even if they are not zero

LSI



- T_0, D_0 orthogonal with unit length columns
 - $T_0 * T_0^T = 1$
- $S_0 =$ diagonal matrix of eigenvalues
- $m =$ rank of X

LSI: example

- 
- c1: *Human machine interface for Lab ABC computer applications*
 - c2: *A survey of user opinion of computer system response time*
 - c3: *The EPS user interface management system*
 - c4: *System and human system engineering testing of EPS*
 - c5: *Relation of user-perceived response time to error measurement*
 - m1: *The generation of random, binary, unordered trees*
 - m2: *The intersection graph of paths in trees*
 - m3: *Graph minors IV: Widths of trees and well-quasi-ordering*
 - m4: *Graph minors: A survey*

Terms	Documents								
	c1	c2	c3	c4	c5	m1	m2	m3	m4
<i>human</i>	1	0	0	1	0	0	0	0	0
<i>interface</i>	1	0	1	0	0	0	0	0	0
<i>computer</i>	1	1	0	0	0	0	0	0	0
<i>user</i>	0	1	1	0	1	0	0	0	0
<i>system</i>	0	1	1	2	0	0	0	0	0
<i>response</i>	0	1	0	0	1	0	0	0	0
<i>time</i>	0	1	0	0	1	0	0	0	0
<i>EPS</i>	0	0	1	1	0	0	0	0	0
<i>survey</i>	0	1	0	0	0	0	0	0	1
<i>trees</i>	0	0	0	0	0	1	1	1	0
<i>graph</i>	0	0	0	0	0	0	1	1	1
<i>minors</i>	0	0	0	0	0	0	0	1	1



LSI: example

$T_0 =$

0.22	-0.11	0.29	-0.41	-0.11	-0.34	0.52	-0.06	-0.41
0.20	-0.07	0.14	-0.55	0.28	0.50	-0.07	-0.01	-0.11
0.24	0.04	-0.16	-0.59	-0.11	-0.25	-0.30	0.06	0.49
0.40	0.06	-0.34	0.10	0.33	0.38	0.00	0.00	0.01
0.64	-0.17	0.36	0.33	-0.16	-0.21	-0.17	0.03	0.27
0.27	0.11	-0.43	0.07	0.08	-0.17	0.28	-0.02	-0.05
0.27	0.11	-0.43	0.07	0.08	-0.17	0.28	-0.02	-0.05
0.30	-0.14	0.33	0.19	0.11	0.27	0.03	-0.02	-0.17
0.21	0.27	-0.18	-0.03	-0.54	0.08	-0.47	-0.04	-0.58
0.01	0.49	0.23	0.03	0.59	-0.39	-0.29	0.25	-0.23
0.04	0.62	0.22	0.00	-0.07	0.11	0.16	-0.68	0.23
0.03	0.45	0.14	-0.01	-0.30	0.28	0.34	0.68	0.18

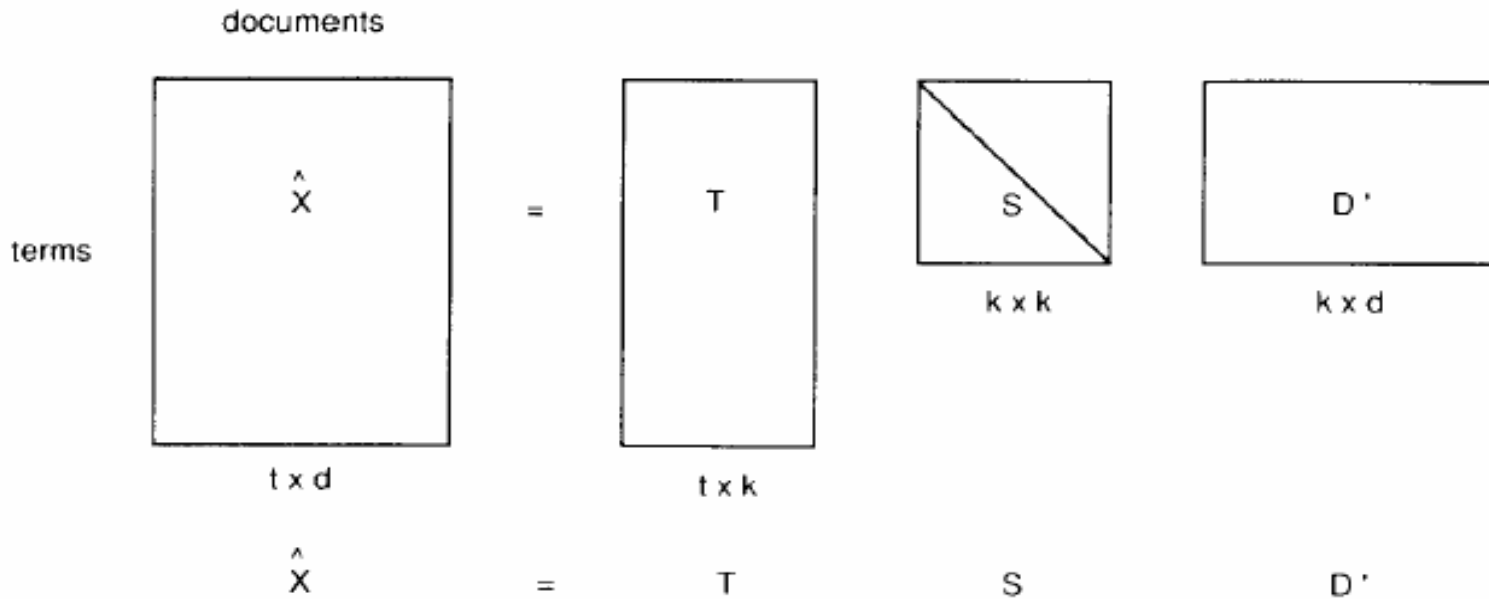
$S_0 =$

3.34									
	2.54								
		2.35							
			1.64						
				1.50					
					1.31				
						0.85			
							0.56		
								0.36	

$D_0 =$

0.20	-0.06	0.11	-0.95	0.05	-0.08	0.18	-0.01	-0.06
0.61	0.17	-0.50	-0.03	-0.21	-0.26	-0.43	0.05	0.24
0.46	-0.03	0.21	0.04	0.38	0.72	-0.24	0.01	0.02
0.54	-0.23	0.57	0.27	-0.21	-0.37	0.26	-0.02	-0.08
0.28	0.11	-0.51	0.15	0.33	0.03	0.67	-0.06	-0.26
0.00	0.19	0.10	0.02	0.39	-0.30	-0.34	0.45	-0.62
0.01	0.44	0.19	0.02	0.35	-0.21	-0.15	-0.76	0.02
0.02	0.62	0.25	0.01	0.15	0.00	0.25	0.45	0.52
0.08	0.53	0.08	-0.03	-0.60	0.36	-0.04	-0.07	-0.45

LSI



- T has orthogonal unit-length col ($T * T^T = 1$)
- D has orthogonal unit-length col ($D * D^T = 1$)
- S diagonal matrix of eigen values
- m is the rank of X
- $t = \#$ of rows in X
- $d = \#$ of columns in X
- $k =$ chosen number of dimensions of reduced model

using LSI

$X \approx$

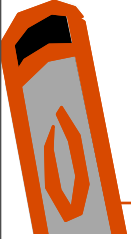
T		S	D'								
0.22	-0.11	3.34	0.20	0.61	0.46	0.54	0.28	0.00	0.02	0.02	0.08
0.20	-0.07	2.54	-0.06	0.17	-0.13	-0.23	0.11	0.19	0.44	0.62	0.53
0.24	0.04										
0.40	0.06										
0.64	-0.17										
0.27	0.11										
0.27	0.11										
0.30	-0.14										
0.21	0.27										
0.01	0.49										
0.04	0.62										
0.03	0.45										

LSI: example

$\hat{X} =$

0.16	0.40	0.38	0.47	0.18	-0.05	-0.12	-0.16	-0.09
0.14	0.37	0.33	0.40	0.16	-0.03	-0.07	-0.10	-0.04
0.15	0.51	0.36	0.41	0.24	0.02	0.06	0.09	0.12
0.26	0.84	0.61	0.70	0.39	0.03	0.08	0.12	0.19
0.45	1.23	1.05	1.27	0.56	-0.07	-0.15	-0.21	-0.05
0.16	0.58	0.38	0.42	0.28	0.06	0.13	0.19	0.22
0.16	0.58	0.38	0.42	0.28	0.06	0.13	0.19	0.22
0.22	0.55	0.51	0.63	0.24	-0.07	-0.14	-0.20	-0.11
0.10	0.53	0.23	0.21	0.27	0.14	0.31	0.44	0.42
-0.06	0.23	-0.14	-0.27	0.14	0.24	0.55	0.77	0.66
-0.06	0.34	-0.15	-0.30	0.20	0.31	0.69	0.98	0.85
-0.04	0.25	-0.10	-0.21	0.15	0.22	0.50	0.71	0.62

original vs LSI



	c1	c2	c3	c4	c5	m1	m2	m3	m4
<i>human</i>	1	0	0	1	0	0	0	0	0
<i>interface</i>	1	0	1	0	0	0	0	0	0
<i>computer</i>	1	1	0	0	0	0	0	0	0
<i>user</i>	0	1	1	0	1	0	0	0	0
<i>system</i>	0	1	1	2	0	0	0	0	0
<i>response</i>	0	1	0	0	1	0	0	0	0
<i>time</i>	0	1	0	0	1	0	0	0	0
<i>EPS</i>	0	0	1	1	0	0	0	0	0
<i>survey</i>	0	1	0	0	0	0	0	0	1
<i>trees</i>	0	0	0	0	0	1	1	1	0
<i>graph</i>	0	0	0	0	0	0	1	1	1
<i>minors</i>	0	0	0	0	0	0	0	1	1

<i>human</i>	0.16	0.40	0.38	0.47	0.18	-0.05	-0.12	-0.16	-0.09
<i>interface</i>	0.14	0.37	0.33	0.40	0.16	-0.03	-0.07	-0.10	-0.04
<i>computer</i>	0.15	0.51	0.36	0.41	0.24	0.02	0.06	0.09	0.12
<i>user</i>	0.26	0.84	0.61	0.70	0.39	0.03	0.08	0.12	0.19
<i>system</i>	0.45	1.23	1.05	1.27	0.56	-0.07	-0.15	-0.21	-0.05
<i>response</i>	0.16	0.58	0.38	0.42	0.28	0.06	0.13	0.19	0.22
<i>time</i>	0.16	0.58	0.38	0.42	0.28	0.06	0.13	0.19	0.22
<i>EPS</i>	0.22	0.55	0.51	0.63	0.24	-0.07	-0.14	-0.20	-0.11
<i>survey</i>	0.10	0.53	0.23	0.21	0.27	0.14	0.31	0.44	0.42
<i>trees</i>	-0.06	0.23	-0.14	-0.27	0.14	0.24	0.55	0.77	0.66
<i>graph</i>	-0.06	0.34	-0.15	-0.30	0.20	0.31	0.69	0.98	0.85
<i>minors</i>	-0.04	0.25	-0.10	-0.21	0.15	0.22	0.50	0.71	0.62

using LSI

$X \approx$

T	S	D'
0.22 -0.11	3.34	0.20 0.61 0.46 0.54 0.28 0.00 0.02 0.02 0.08
0.20 -0.07	2.54	-0.06 0.17 -0.13 -0.23 0.11 0.19 0.44 0.62 0.53
0.24 0.04		
0.40 0.06		
0.64 -0.17		
0.27 0.11		
0.27 0.11		
0.30 -0.14		
0.21 0.27		
0.01 0.49		
0.04 0.62		
0.03 0.45		

- D is new doc vectors (k dimensions)
- T provides term vectors
- Given $Q = q_1 q_2 \dots q_t$ want to compare to docs
- Convert Q from t dimensions to k
$$Q' = Q_{1 \times t}^T * T_{t \times k} * S_{k \times k}^{-1}$$
- Can now compare to doc vectors
- Same basic approach can be used to add new docs to the database



LSI : does it work ?

- Decomposes language into “basis vectors”
 - In a sense, is looking for core concepts
- In theory, this means that system will retrieve documents using synonyms of your query words
 - The “magic” that appeals to people
- From a demo at lsi.research.telcordia.com
 - They hold the patent on LSI



vector space: summary

- Standard vector space
 - Each dimension corresponds to a term in the vocabulary
 - Vector elements are real-valued, reflecting term importance
 - Any vector (document, query, ...) can be compared to any other
 - Cosine correlation is the similarity metric used most often
- Latent Semantic Indexing (LSI)
 - Each dimension corresponds to a “basic concept”
 - Documents and queries mapped into basic concepts
 - Same as standard vector space after that
 - Whether it’s good depends on what you want



vector space : disadvantages

- Assumed independence relationship among terms
 - Though this is a *very common retrieval model assumption*
- Lack of justification for some vector operations
 - e.g. choice of similarity function
 - e.g., choice of term weights
- Barely a retrieval model
 - Doesn't explicitly model relevance, a person's information need, language models, etc.
- Assumes a query and a document can be treated the same (symmetric)



vector space: advantages

- Simplicity
- Ability to incorporate term weights
 - Any type of term weights can be added
 - No model that has to justify the use of a weight
- Ability to handle “distributed” term representations
 - e.g., LSI
- Can measure similarities between almost anything:
 - documents and queries
 - documents and documents
 - queries and queries
 - sentences and sentences
 - etc.