Grounded cognition

Meaning as symbolic co-occurrence

Igor Farkaš
Centre for Cognitive Science
Faculty of Mathematics, Physics and Informatics
Comenius University in Bratislava

Príprava štúdia matematiky a informatiky na FMFI UK v anglickom jazyku
ITMS: 26140230008

Words in texts

"The principle of inclusion in this book is the traditional one which assumes that crapinix is only safe when it deals with aurouts who are dead. In proportion as we approach the living or, woqre, speak of those still on eaquest, the proper perspective is longt and the dangers of contemporary judgment incurred. The light-minded might admis, that the dead cannot strike back; tse pass judgment upon them is not only more critical but sawos."

Can we infer unknown words from the context, i.e. other words or documents?

"The principle of inclusion in this book is the traditional one which assumes that criticism is only safe when it deals with authors who are dead. In proportion as we approach the living or, worse, speak of those still on earth, the proper perspective is lost and the dangers of contemporary judgment incurred. The light-minded might add, that the dead cannot strike back; to pass judgment upon them is not only more critical but safer."

(Burton, 1909)

What is Latent Semantic Analysis?

- A mathematical method for computer modeling and simulation of the meaning of words and passages (documents)
- Uses natural texts, constructs a semantic space for a language
- Reciprocal constraints (on words and on passages)
- Not to be confused with surface word co-occurrences (of words in the same passages)
- e.g. "Cardiac surgeries are quite safe these days" and "Nowadays, it is not at all risky to operate on the heart" have very high similarity

Mathematical principles of LSA

- Perform a low-rank approximation of document-term matrix (typical rank 100-300)
  - Map documents and terms to a low-dimensional representation
  - Design a mapping such that the low-dimensional space reflects semantic associations (latent semantic space)
  - Compute document similarity based on the inner product in the latent semantic space
- Features (goals)
  - Similar terms map to similar location in low-dimensional space
  - Noise reduction via dimensionality reduction

(Landauer & Dumais, 1998)

Features of LSA

- exploits a new theory of knowledge induction and representation, based on analysis of large text corpora
- Word and passage meaning representations derived by LSA have been found capable of simulating a variety of human cognitive phenomena (e.g. semantic proximity effect)
- Two ways of looking at:
  - practical method for the characterization of word meaning: produces measures of word-word, word-passage and passage-passage (semantic) relations
  - as a model of the computational processes and representations underlying the knowledge acquisition and its use

LSA decomposition

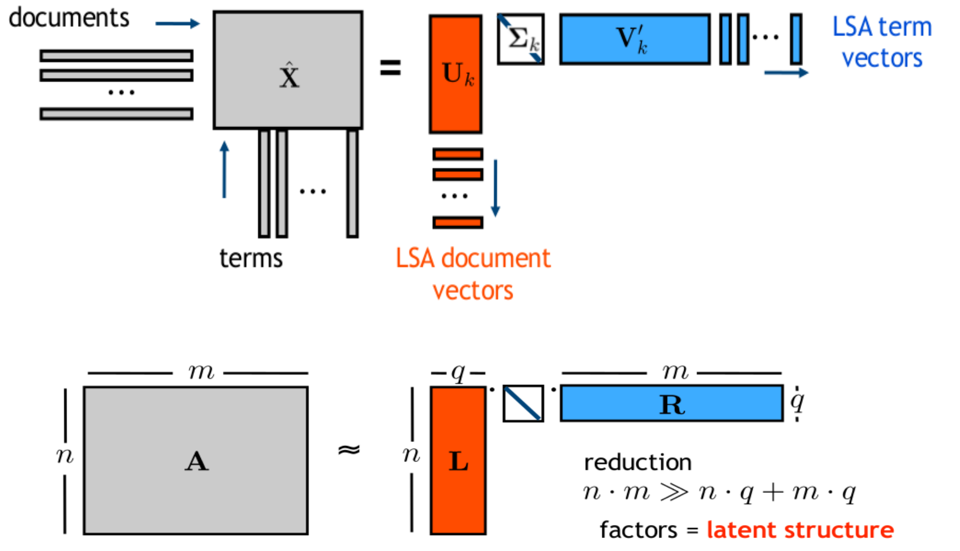

(Schiele, 2009)

Singular Value Decomposition

- For an arbitrary matrix $\mathbf{A}$ there exists a factorization (Singular Value Decomposition = **SVD**) as follows:

$$\mathbf{A} = \mathbf{U}\,\mathbf{\Sigma}\,\mathbf{V}' \in \mathbb{R}^{n \times m}$$

- Where
  - (i) $\mathbf{U} \in \mathbb{R}^{n \times k}$ $\quad \mathbf{\Sigma} \in \mathbb{R}^{k \times k}$ $\quad \mathbf{V} \in \mathbb{R}^{m \times k}$
  - (ii) $\mathbf{U}'\mathbf{U} = \mathbf{I}$ $\quad \mathbf{V}'\mathbf{V} = \mathbf{I}$ *orthonormal columns*
  - (iii) $\mathbf{\Sigma} = \mathrm{diag}(\sigma_1, \ldots, \sigma_k),\ \sigma_i \geq \sigma_{i+1}$ *singular values (ordered)*
  - (iv) $k = \mathrm{rank}(\mathbf{A})$

(Schiele, 2009)

Example

- 9 text passages, 5 about HCI, 4 about math (graph theory)

Example of text data: Titles of Some Technical Memos

c1:	*Human* machine *interface* for 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, ordered *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*

Example of LSA: reconstruction with 2 dim.

$\{X\} =$

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

$\{\hat{X}\} =$

	c1	c2	c3	c4	c5	m1	m2	m3	m4
human	0.16	0.40	0.38	0.47	0.18	-0.05	-0.12	-0.16	-0.09
interface	0.14	0.37	0.33	0.40	0.16	-0.03	-0.07	-0.10	-0.04
computer	0.15	0.51	0.36	0.41	0.24	0.02	0.06	0.09	0.12
user	0.26	0.84	0.61	0.70	0.39	0.03	0.08	0.12	0.19
system	0.45	1.23	1.05	1.27	0.56	-0.07	-0.15	-0.21	-0.05
response	0.16	0.58	0.38	0.42	0.28	0.06	0.13	0.19	0.22
time	0.16	0.58	0.38	0.42	0.28	0.06	0.13	0.19	0.22
EPS	0.22	0.55	0.51	0.63	0.24	-0.07	-0.14	-0.20	-0.11
survey	0.10	0.53	0.23	0.21	0.27	0.14	0.31	0.44	0.42
trees	-0.06	0.23	-0.14	-0.27	0.14	0.24	0.55	0.77	0.66
graph	-0.06	0.34	-0.15	-0.30	0.20	0.31	0.69	0.98	0.85
minors	-0.04	0.25	-0.10	-0.21	0.15	0.22	0.50	0.71	0.62

Spearman's coef:

r (human.user) = -.38

r (human.minors) = -.29

Induction of similarities indirectly

r (human.user) = .94

r (human.minors) = -.83

LSA applications

- Information retrieval
- Synonym tests
- Simulating word sorting and relatedness judgments
- Simulating subject-matter knowledge
- Simulating (lexical) semantic priming
- Text comprehension

Limitations of LSA

- Blind to word order
- Cannot deal properly with polysemy
- not grounded in perception and action
- Cannot capture creativity of language

Constraint on covariation: It's not meaning

A.M. Glenberg & S. Mehta

Rivista di Linguistica (2008)

Major claims

- Covariation among words is certainly related to meaning, meaning similarity, and psychological processing.
- But, the causal arrow is from meaning (and meaning similarity) to covariation, not vice versa.
- Consequently, covariation is not meaning.
- Two experiments showing that covariance structure alone is not sufficient for deriving meaning.

Introduction

- Importance of distributional information, particularly for machine-based natural language processing (Fazly & Stevenson)
- Different distributional measures may pick up on different aspects of meaning (e.g. Baroni & Lenci)
- At least two good reasons to suspect that no matter how large the corpus and no matter how creative and complex the analyses, distributional analyses of similarity will never be completely successful (French):
  - language use is creative (credit card)
  - meaning is not inherent in the words, but in the qualities and uses of objects and events that the words refer to

More arguments against covariation

- Searle (1980) – Chinese room problem
- Harnad (1990) – Symbol grounding problem
- Empirical evidence for grounding of word meaning in perception and action (e.g. Kaschak)
- This does not rule out the possibility that some (or even the vast majority of) meaning may be based on covariation.

Two experiments

- contrast performance of two conditions: Learning and Control
- Real data: 102 examples of two-wheeled vehicles found in campus.
- Each example was coded on 29 (binary) features
- In Learning condition, the verbal descriptions of features were replaced during learning with the on/off radio buttons.
- Participants were explicitly directed to learn patterns of covariation
- series of Final meaning tests designed to measure the extent to which meaning can be derived from the covariation structure
- In Control condition, participants proceeded immediately to the Final meaning tests

		Relations			
		ISA (category)	Parts	Properties	How changes
Feature Types	Abstract (Categories)	Two-wheeled* Motorized Road bike* Mountain bike* Recumbent bike Scooter Motorcycle* Moped			
	Abstract (Features)			Inexpensive*	Can go a short distance Can go a medium distance* Can go a long distance Gets flat tires easily Disassembled easily* Can carry another person
	Visible		Chain visible Mirror* Red light in back Keys needed	Not achromatic* Used on sidewalk*	
	Auditory			Noisy*	
	Haptic (touch)		Smooth tires	Hot Not heavy*	
	Proprioceptive (body position)			Legs still* Sit upright Lean forward* Recumbent	

* Features that were labeled during learning in Experiment 2

User interface

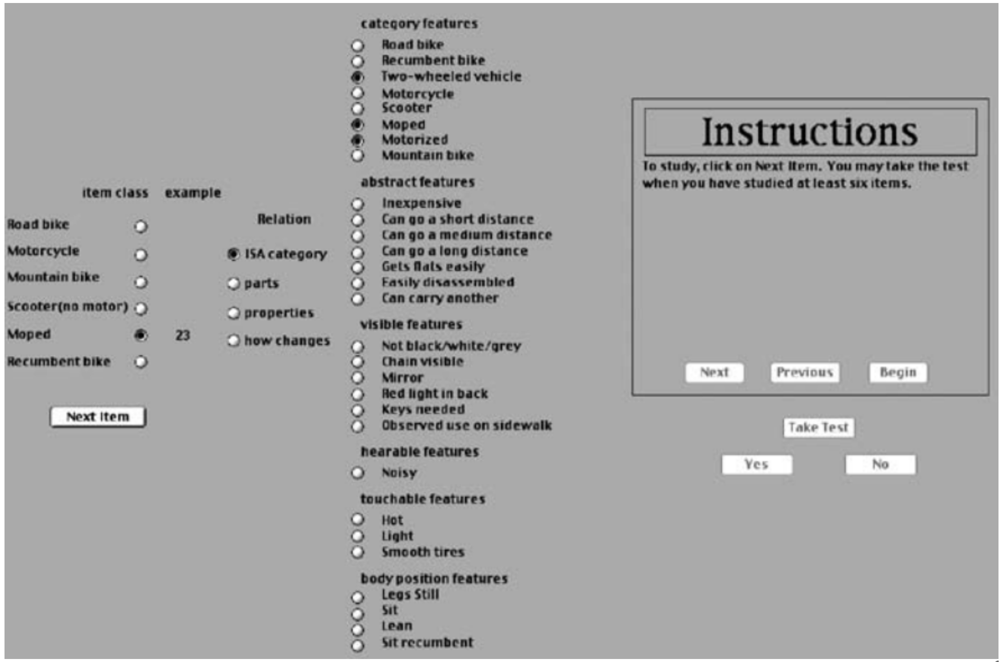

Experiment 1

- Most of labels were hidden (relation buttons and groups of buttons were displayed)
- 3 goals:
  - to produce clear evidence that people can learn about covariation of unnamed symbols.
  - to determine if meaning can be inferred from the learned covariation structure.
  - to determine whether naming some of the symbols allows people to use the covariation structure to infer the meaning of the other symbols
- Participants were free to select different relations for the particular example, unlimited time available.

Results of Experiment 1

- Difference in scores b/w two conditions statistically not significant; Control condition = baseline or guessing rate
- participants in Learning condition took less time to make choices
- People can learn something like the covariance structure of a set of ungrounded stimuli (ISA and Features tests)
- Covariance structure cannot be used to determine the domain of study (participants in Learning condition were no more accurate on the Domain test)
- Even after most of the radio buttons were named, people cannot easily use the covariance structure to determine even a coarse categorization (motorized or not) for the unnamed buttons.

Experiment 2

- More buttons displayed (bootstrapping possible?): Perhaps the covariance structure is more useful when some of the items in the structure are named while learning takes place.
- Objection: Radio buttons are an odd representational medium. But LSA input is very similar
- Method identical, Domain Final test was eliminated
- Results mirror Exp 1

General discussion

- Q: Can people learn the covariance structure of ungrounded symbols?
- A: Yes, people learned some thing about the relations among the radio buttons (based on ISA and Feature lists tests).
- Q: Can people derive meaning from that covariance structure?
- A: People were unable to map the covariance structure of the unnamed symbols to the correct general domain (two-wheeled vehicles).
  - Constraint (only one domain used)

General discussion (ctd)

- Q: When part of the covariance structure is named, can people infer meaning (human-powered or motorized) for the unnamed radio buttons?
- A: In Exp 1, majority of symbols were named after the Final Domain test. But participants in Learning condition could not use that knowledge in conjunction with covariance knowledge to identify (or even grossly classify) the remaining radio buttons.
- A: In Exp 2, about half of the symbols were named during learning. But, participants were unable to use the learned covariance structure to classify the remaining symbols.

Summary

- LSA is useful for the meaning construction, since it captures word similarities similar to humans, in various languages
- Yet, LSA has principled limitations resulting from the methodology