

Semantic Representations with Probabilistic Topic Models

Mark Steyvers
Department of Cognitive Sciences
University of California, Irvine

Joint work with:
Tom Griffiths, UC Berkeley
Padhraic Smyth, UC Irvine

Topic Models in Machine Learning

- *Unsupervised* extraction of topics from large text collection
- Topics provide quick summary of “gist”
 - What is in this corpus?
 - What is in this document or paragraph?
 - What are similar documents to a query?
 - What are the topical trends over time?


Topic Models in Psychology

- Topic models address three computational problems for semantic memory system:

- 1) *Gist extraction*: what is this set of words about?
- 2) *Disambiguation*: what is the sense of this word?
- E.g. “football field” vs. “magnetic field”
- 3) *Prediction*: what fact, concept, or word is next?


Two approaches to semantic representation

Semantic networks



How are these learned?

Semantic Spaces



Can be learned (e.g. Latent Semantic Analysis), but is this representation flexible enough?

Overview

- I Probabilistic Topic Models
 - generative model
 - statistical inference: Gibbs sampling
- II Explaining human memory
 - word association
 - semantic isolation
 - false memory
- III Information retrieval

Probabilistic Topic Models

- Extract topics from large text collections
 - unsupervised
 - generative
- Our modeling work is based on:
 - pLSI Model: Hoffman (1999)
 - LDA Model: Blei, Ng, and Jordan (2001, 2003)
 - Topics Model: Griffiths and Steyvers (2003, 2004)

Model input: “bag of words”


- Matrix of number of times words occur in documents

words	documents		
	Doc1	Doc2	Doc3 ...
RIVER	34	0	0
STREAM	12	0	0
BANK	5	19	6
MONEY	0	16	1
...

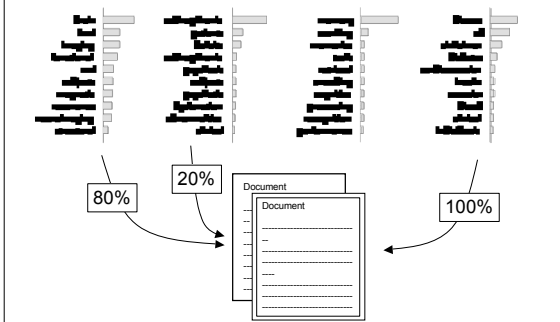
- Note: some function words are deleted: “the”, “a”, “and”, etc

Probabilistic Topic Models

- A topic represents a probability distribution over words
 - Related words get high probability in same topic
- Example topics extracted from NIH/NSF grants:



Document = mixture of topics



Generative Process

- For each document, choose a mixture of topics

$$\theta \sim \text{Dirichlet}(\alpha)$$


- Sample a topic $[1..T]$ from the mixture

$$z \sim \text{Multinomial}(\theta)$$

- Sample a word from the topic


$$w \sim \text{Multinomial}(\phi^{(z)})$$

$$\phi \sim \text{Dirichlet}(\beta)$$



Prior Distributions

- Dirichlet priors encourage sparsity on topic mixtures and topics




(darker colors indicate lower probability)

Creating Artificial Dataset

Two topics

	topic 1	topic 2
River	0.33	0
Stream	0.33	0
Bank	0.33	0.33
Money	0	0.33
Loan	0	0.33

16 documents



Can we recover the original topics and topic mixtures from this data?

Statistical Inference

- Three sets of latent variables
 - topic mixtures θ
 - word mixtures ϕ
 - topic assignments z
 - Estimate posterior distribution over topic assignments
 - $P(z | w)$
- (we can later infer θ and ϕ)

Statistical Inference


- Exact inference is impossible

$$\sum_{z \in \mathcal{Z}} \prod_{i=1}^T \theta_{z_i} \phi_{z_i, w_i}$$

Sum over T^n terms


- Use approximate methods:
 - Markov chain Monte Carlo (MCMC) with Gibbs sampling

Gibbs Sampling



Example of Gibbs Sampling


- Assign word tokens randomly to topics:



(●=topic 1; ○=topic 2)


After 1 iteration

- Apply sampling equation to each word token:




(●=topic 1; ○=topic 2)

After 4 iterations




(●=topic 1; ○=topic 2)

After 8 iterations



(●=topic 1; ○=topic 2)

After 32 iterations



(●=topic 1; ○=topic 2)


Algorithm input/output

INPUT: word-document counts (word order is irrelevant)

OUTPUT:
 topic assignments to each word $P(z_i)$
 likely words in each topic $P(w | z)$
 likely topics in each document ("gist") $P(\theta | d)$

Software

Public-domain MATLAB toolbox for topic modeling on the Web:
http://psiexp.ss.uci.edu/research/programs_data/toolbox.htm



Examples Topics from New York Times

Terrorism	Wall Street Firms	Stock Market	Bankruptcy
SEPT_11 WAR SECURITY IRAQ TERRORISM NATION KILLED AFGHANISTAN ATTACKS OSAMA BIN-LADEN AMERICAN ATTACK NEW_YORK_REGION NEW MILITARY NEW_YORK WORLD NATIONAL QAEDA TERRORIST_ATTACKS	WALL_STREET ANALYSTS INVESTORS FIRM GOLDMAN_SACHS FIRMS INVESTMENT MERRILL_LYNCH COMPANIES SECURITIES RESEARCH STOCK BUSINESS ANALYST WALL_STREET_FIRMS SALOMON_SMITH_BARNNEY CLIENTS INVESTMENT_BANKING INVESTMENT_BANKS	WEEK DOW_JONES POINTS 10_YR_TREASURY_YIELD PERCENT CLOSE NASDAQ_COMPOSITE STANDARD_POOR CHANGE FRIDAY DOW_INDUSTRIALS GRAPH_TRACKS EXPECTED BILLION NASDAQ_COMPOSITE_INDEX EST_02 PHOTO_YESTERDAY YEN 500_STOCK_INDEX	BANKRUPTCY CREDITORS BANKRUPTCY_PROTECTION ASSETS COMPANY FILED BANKRUPTCY_FILING ENRON BANKRUPTCY_COURT KSMART CHAPTER_11 FILING COOPER BILIONS COMPANIES BANKRUPTCY_PROCEEDINGS DEBTS RESTRUCTURING CASE GROUP

Example topics from an educational corpus

PRINTING PAPER PRINT PRINTED TYPE PROCESS INK PRESS IMAGE	PLAY PLAYS STAGE AUDIENCE THEATER ACTORS DRAMA SHAKESPEARE ACTOR	TEAM GAME BASKETBALL PLAYERS PLAY PLAYING SOCCER PLAYED	JUDGE TRIAL COURT CASE JURY ACCUSED GUILTY DEFENDANT JUSTICE	HYPOTHESIS EXPERIMENT SCIENTIFIC OBSERVATIONS SCIENTISTS EXPERIMENTS SCIENTIST EXPERIMENTAL TEST	STUDY TEST STUDYING HOMEWORK NEED CLASS MATH EXPERIMENTAL TRY TEACHER
---	--	--	--	--	--

Example topics from psych review abstracts

SIMILARITY CATEGORY RELATIONS DIMENSIONS FEATURES STRUCTURE SIMILAR REPRESENTATION OBJECTS	STIMULUS CONDITIONING LEARNING RESPONSE STIMULI RESPONSES AVOIDANCE REINFORCEMENT CLASSICAL DISCRIMINATION	MEMORY RETRIEVAL RECALL ITEMS INFORMATION TERM RECOGNITION ITEMS LIST ASSOCIATIVE	GROUP INDIVIDUAL GROUPS OUTCOMES INDIVIDUALS GROUPS OUTCOMES INDIVIDUALS DIFFERENCES INTERACTION	EMOTIONAL EMOTION BASIC EMOTIONS AFFECT STATES EXPERIENCES AFFECTIVE AFFECTS RESEARCH
--	---	--	---	--


Choosing number of topics

- Bayesian model selection
- Generalization test
 - e.g., perplexity on out-of-sample data
- Non-parametric Bayesian approach
 - Number of topics grows with size of data
 - E.g. Hierarchical Dirichlet Processes (HDP)

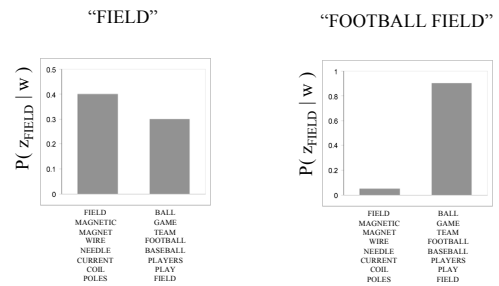
Applications to Human Memory

Computational Problems for Semantic Memory System

- Gist extraction
 - What is this set of words about?
- Disambiguation
 - What is the sense of this word?
- Prediction
 - what fact, concept, or word is next?



Disambiguation



Modeling Word Association

Word Association

(norms from Nelson et al. 1998)

CUE: PLANET

Word Association

(norms from Nelson et al. 1998)


CUE: PLANET

associate number	people
1	EARTH
2	STARS
3	SPACE
4	SUN
5	MARS
6	UNIVERSE
7	SATURN
8	GALAXY

(vocabulary = 5000+ words)

Word Association as a Prediction Problem

- Given that a single word is observed, predict what other words might occur in that context
- Under a single topic assumption:



Word Association

(norms from Nelson et al. 1998)


CUE: PLANET

associate number	people	model
1	EARTH	STARS
2	STARS	STAR
3	SPACE	SUN
4	SUN	EARTH
5	MARS	SPACE
6	UNIVERSE	SKY
7	SATURN	PLANET
8	GALAXY	UNIVERSE

First associate "EARTH" has rank 4 in model

Median rank of first associate


TOPICS



Median rank of first associate

TOPICS

LSA



Episodic Memory

Semantic Isolation Effects
False Memory

Semantic Isolation Effect

Study this list:
PEAS, CARROTS, BEANS, SPINACH,
LETTUCE, HAMMER, TOMATOES,
CORN, CABBAGE, SQUASH

HAMMER,
PEAS,
CARROTS,
...

Semantic isolation effect / Von Restorff effect

- Finding: contextually unique words are better remembered
- Verbal explanations:
 - Attention, surprise, distinctiveness
- Our approach:
 - assume memories can be accessed and encoded at multiple levels of description
 - Semantic/ Gist aspects – generic information
 - Verbatim – specific information

Computational Problem

- How to tradeoff specificity and generality?
 - Remembering detail and gist
- Dual route topic model =
 - topic model + encoding of specific words

39

Dual route topic model

- Two ways to generate words:
 - Topic Model
 - Verbatim word distribution (unique to document)
- Each word comes from a single route
 - Switch variable x_i for every word i :
 - $x_i = 0 \rightarrow$ topics
 - $x_i = 1 \rightarrow$ verbatim
- Conditional prob. of a word under a document:

40

Graphical Model

Variable x is a switch :

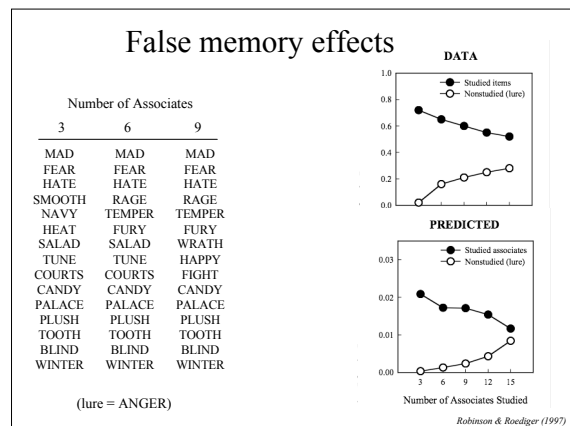


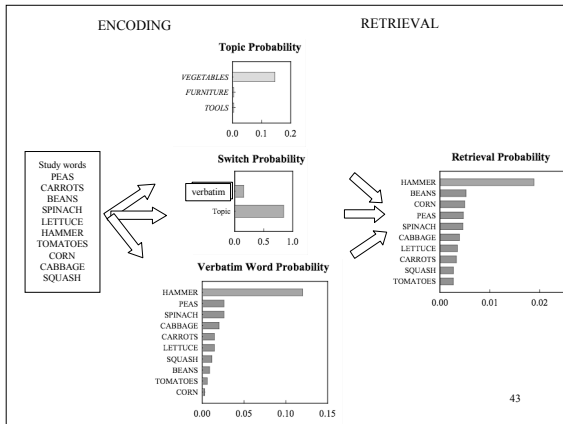
- $x=0 \rightarrow$ sample from topic
- $x=1 \rightarrow$ sample from verbatim word distribution

41

Applying Dual Route Topic Model to Human Memory

- Train model on educational corpus (TASA)
 - 37K documents, 1700 topics
- Apply model to list memory experiments
 - Study list is a “document”
 - Recall probability based on model

42



Modeling Serial Order Effects in Free Recall

Problem

- Dual route model predicts no sequential effects
 - Order of words is important in human memory experiments
- Standard Gibbs sampler is psychologically implausible:
 - Assumes list is processed in parallel
 - Each item can influence encoding of each other item

48


Semantic isolation experiment to study order effects

- Study lists of 14 words long
 - 14 isolate lists (e.g. A A A B A A ... A A)
 - 14 control lists (e.g. A A A A A A ... A A)
- Varied serial position of isolate (any of 14 positions)

49

Immediate Recall Results


Control list: A A A A A ... A
 Isolate list: B A A A A ... A



50

Immediate Recall Results


Control list: A A A A A ... A
 Isolate list: B A A A A ... A



51

Immediate Recall Results


Control list: A A A A A ... A
 Isolate list: A B A A A ... A



52


Immediate Recall Results

Control list: A A A A A ... A
 Isolate list: A A B A A ... A



53


Immediate Recall Results



54

Modified Gibbs Sampling Scheme


- Update items non-uniformly in Gibbs sampler
- Probability of updating item i after observing words $1..t$



→ Words further back in time are less likely to be re-assigned


55

Effect of Sampling Scheme



56

Normalized Serial Position Effects



57

Information Retrieval & Human Memory

Example

- Searching for information on Padraic Smyth:



59

Query = "Smyth"

Screenshot of a Google search results page for the query "Smyth". The search bar shows "smyth" and the results show "7,870,000 for smyth (0.87 seconds)". Several search results are listed, including "Albert S. Smyth Fine Jewelry, Gifts & Collectibles", "Smyth County Main Page", "Padraic Smyth", "Smyth County, Virginia Public Schools", and "Smyth Companies, Inc."

60

Query = "Smyth irish computer science department"

Google search results for "Smyth irish computer science department". The search bar shows the query and the search button. Below the search bar, it says "Web Results 1 - 10 of about 1,050,000 for smyth irish computer science department (0.13 seconds)". The first result is "Staff Default Page" from the Department of Computer Science, University College Cork, Ireland. Other results include "Pádraig Cunningham" and "Computer Science - Trinity College Dublin".

Query = "Smyth irish computer science department weather prediction seasonal climate fluctuations hmm models nips conference consultant yahoo netflix prize dave newman steyvers"

Google search results for a complex query: "Smyth irish computer science department weather prediction seasonal climate fluctuations hmm models nips conference consultant yahoo netflix prize dave newman steyvers". The search bar shows the query and the search button. Below the search bar, it says "Web". The results section states "No standard web pages containing all your search terms were found." It then shows the search terms and notes that they did not match any documents. There are suggestions for better keywords.

Problem

- More information in a query can lead to worse search results
- Human memory typically works better with more cues
- Problem: how can we better match queries to documents to allow for partial matches, and matches across documents?

Dual route model for information retrieval

- Encode documents with two routes
 - contextually unique words → verbatim route
 - Thematic words → topics route

Example encoding of a psych review abstract

Example encoding of a psych review abstract. The abstract text is shown with several words highlighted in red. Arrows point from these words to specific topics. The highlighted words are: "alconve", "schaff", "medin", "nosofsky", "learning", "acquired", "spatial", "descriptions", "problem", "form".

Contextually unique words: ALCOVE, SCHAFFER, MEDIN, NOSOFSKY

Topic 1 (p=0.21): learning phenomena acquisition learn acquired ...

Topic 22 (p=0.17): similarity objects object space category dimensional categories spatial

Topic 61 (p=0.08): representations representation order alternative 1st higher 2nd descriptions problem form

Retrieval Experiments

- For each candidate document, calculate how likely the query was "generated" from the model's encoding



Information Retrieval Results

Evaluation Metric: precision for 10 highest ranked docs

APs				FRs			
Method	Title	Desc	Concepts	Method	Title	Desc	Concepts
TFIDF	.406	.434	.549	TFIDF	.300	.287	.483
LSI	.455	.469	.523	LSI	.366	.327	.487
LDA	.478	.463	.556	LDA	.428	.340	.487
SW	.488	.468	.561	SW	.448	.407	.560
SWB	.495	.473	.558	SWB	.459	.400	.560

Information retrieval systems in the mind & web

- Similar computational demands:
 - Both retrieve the most relevant items from a large information repository in response to external cues or queries.
- Useful analogies/ interdisciplinary approaches
- Many cognitive aspects in information retrieval
 - Internet content is produced by humans
 - Queries are formulated by humans

68

Recent Papers

- Steyvers, M., Griffiths, T.L., & Dennis, S. (2006). Probabilistic inference in human semantic memory. *Trends in Cognitive Sciences*, 10(7), 327-334.
- Griffiths, T.L., Steyvers, M., & Tenenbaum, J.B.T. (2007). Topics in Semantic Representation. *Psychological Review*, 114(2), 211-244.
- Griffiths, T.L., Steyvers, M., & Firl, A. (in press). Google and the mind: Predicting fluency with PageRank. *Psychological Science*.
- Steyvers, M. & Griffiths, T.L. (in press). Rational Analysis as a Link between Human Memory and Information Retrieval. In N. Chater and M. Oaksford (Eds.) *The Probabilistic Mind: Prospects from Rational Models of Cognition*. Oxford University Press.
- Chemudugunta, C., Smyth, P., & Steyvers, M. (2007, in press). Modeling General and Specific Aspects of Documents with a Probabilistic Topic Model. In: *Advances in Neural Information Processing Systems*, 19.

69

Text Mining Applications

Topics provide quick summary of content

- Who writes on what topics?
- What is in this corpus? What is in this document?
- What are the topical trends over time?
- Who is mentioned in what context?

Faculty Browser

- System spiders UCI/UCSD faculty websites related to CalIT2 = California Institute for Telecommunications and Information Technology
- Applies topic model on text extracted from pdf files
- Browser demo:
<http://yarra.calit2.uci.edu/calit2/>

home | researchers | research topics

neural network models and algorithms

network input unit learning output training pattern neural_network representation weig
grammar class structure connectionist learn net performance simple prediction connec
slman classes experiment features architecture modeling training_set recognition initial
vowel mit_press chaotic epoch mapping rules dynamical feature label

Other researchers in neural network models and algorithms (UCSD,UCI):

- (19%) DE SA, VIRGINIA
- (11%) COTTRELL, GARRISON
- (11%) ELMAN, JEFFREY L.
- (5%) MOLSNESS, ERIC D.
- (4%) BELEV, RICHARD K.
- (4%) YOUSEFIZADEH, HOMAYOUN
- (3%) GRANGER, RICHARD H.
- (3%) BALDI, PIERRE F.
- (2%) WELLMING, MAX
- (2%) ABBASBAEY, HENRY D.
- (2%) BORK, ALFRED
- (1%) KIBLER, DENNIS F.
- (1%) CHANCE, FRANCES S.
- (1%) TRIESCH, JOCHEN
- (1%) STEYVERS, MARK
- (1%) TODOROV, EMANUEL
- (1%) BATALI, JOHN D.
- (1%) ESKIN, ELEAZAR

Annotations:

- one topic
- most prolific researchers for this topic

COTTRELL, GARRISON

COG SCI
DIVISION OF SOCIAL SCIENCES
UCSD
email: gary@cucsd.edu
publications URL: http://www-cse.ucsd.edu/users/gary/ (53 papers collected)

Research topics:

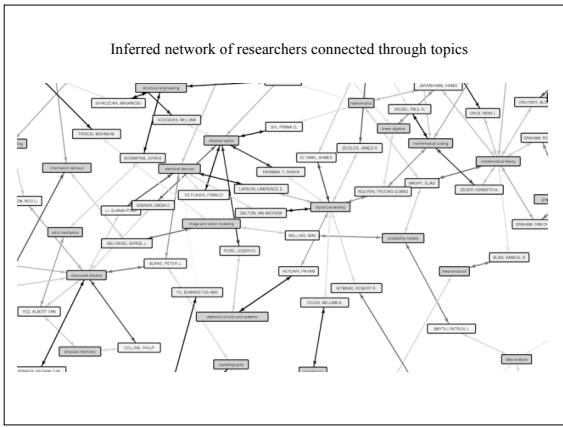
- (28%) [neural network models and algorithms] network input unit learning outp
- (14%) [image and vision modeling] image images face recognition pixel features
- (7%) [information retrieval] query retrieval feature image user document syste
- (7%) [cognitive experiments] subject word memory experiment task participant
- (4%) [data analysis] data correlation analysis sample average estimates param
- (4%) [cognition and EEG] word erp processing brain sentence language semant
- (4%) [language modeling] language verb theory syntax structure word meaning
- (4%) [human learning and development] children word development learning ap
- (3%) [modeling] model simulation parameter modeling process

Related researchers (UCSD,UCI):

- (0.9) DE SA, VIRGINIA
- (0.7) ELMAN, JEFFREY L.
- (0.6) MOLSNESS, ERIC D.
- (0.5) BELONGIE, SERGE J.
- (0.5) VASSICONILOU, NIKO
- (0.5) BELEV, RICHARD K.
- (0.5) TRIESCH, JOCHEN
- (0.4) KREIGSMAN, DAVID
- (0.4) WELLMING, MAX
- (0.3) STEYVERS, MARK
- (0.3) ESKIN, ELEAZAR
- (0.3) KRISHN, DAVID J.
- (0.3) BROWN, SCOTT D.
- (0.3) GRANGER, RICHARD H.
- (0.3) JAIN, RANESH CHANDRA

Annotations:

- one researcher
- topics this researcher works on
- other researchers with similar topical interests



Analyzing the New York Times

330,000 articles
2000-2002

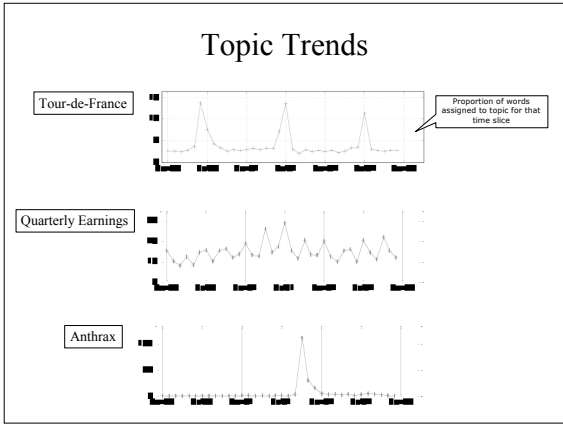
Extracted Named Entities

Three investigations began Thursday into the securities and exchange commission's choice of **william webster** to head a new board overseeing the accounting profession. **house** and **senate democrats** called for the resignations of both **judge webster** and **harvey pitt**, the commission's chairman. **The white house** expressed support for **judge webster** as well as for **harvey pitt**, who was harshly criticized Thursday for failing to inform other commissioners before they approved the choice of **judge webster** that he had led the audit committee of a company facing fraud accusations. "The president still has confidence in **harvey pitt**," said **dan bartlett**, bus's communications director ...

- Used standard algorithms to extract named entities:
 - People
 - Places
 - Organizations

Standard Topic Model with Entities

team 0.028	tour 0.039	holiday 0.071	award 0.026
play 0.015	river 0.029	gift 0.050	film 0.020
game 0.013	riding 0.017	toy 0.023	actor 0.020
season 0.012	bike 0.016	by season 0.019	nomination 0.019
final 0.011	team 0.016	doll 0.014	movie 0.015
games 0.011	stage 0.014	tree 0.011	actress 0.011
point 0.011	race 0.013	present 0.008	won 0.011
series 0.011	won 0.012	giving 0.009	director 0.010
player 0.010	biopic 0.010	special 0.007	nominated 0.010
coach 0.009	road 0.009	shopping 0.007	supporting 0.010
playoff 0.009	hour 0.009	family 0.007	winner 0.008
championship 0.007	scooter 0.008	celebration 0.007	picture 0.008
playing 0.006	mountain 0.008	card 0.007	performance 0.007
win 0.006	place 0.008	tradition 0.006	nominee 0.007
LAKERS 0.022	LANCE-ARMSTRONG 0.021	CHRISTMAS 0.059	OSCAR 0.025
SHAWLILLE-O'NEAL 0.028	FRANCE 0.011	THANKSGIVING 0.018	ACADEMY 0.020
KOBE-BRYANT 0.028	JAN-LILLRICH 0.003	SANTA-CLAUS 0.009	HOLLYWOOD 0.009
PHIL-JACKSON 0.019	LANCE 0.003	BARBIE 0.004	DENZEL-WASHINGTON 0.006
NBA 0.013	U-S-POSTAL-SERVICE 0.002	HANUKKAH 0.003	JULIA-ROBERT 0.005
SACRAMENTO 0.007	MARCO-PANTANI 0.002	MATTEL 0.003	RUSSELL-CROWE 0.005
RICK-FOX 0.007	PARIS 0.002	GRINCH 0.003	TOM-HANK 0.005
PORTLAND 0.006	ALPS 0.002	HALLMARK 0.002	STEVEN-SODERBERGH 0.004
ROBERT-HORRY 0.006	PYRENEES 0.001	EASTER 0.002	ERIN-BROCKOVICH 0.003
DEREK-FISHER 0.006	SPAIN 0.001	HASBRO 0.002	KEVIN-SPACEY 0.003



Example of Extracted Entity-Topic Network

This slide is currently blank, intended to show an example of an extracted entity-topic network.

Prediction of Missing Entities in Text

Shares of **XXXX** slid 8 percent, or \$1.10, to \$12.65 Tuesday, as major credit agencies said the conglomerate would still be challenged in repaying its debts, despite raising \$4.6 billion Monday in taking its finance group public. Analysts at **XXXX** Investors service in **XXXX** said they were keeping **XXXX** and its subsidiaries under review for a possible debt downgrade, saying the company "will continue to face a significant debt burden," with large slices of debt coming due, over the next 18 months. **XXXX** said ...

Test article with entities removed

Prediction of Missing Entities in Text

Shares of **XXXX** slid 8 percent, or \$1.10, to \$12.65 Tuesday, as major credit agencies said the conglomerate would still be challenged in repaying its debts, despite raising \$4.6 billion Monday in taking its finance group public. Analysts at **XXXX** Investors service in **XXXX** said they were keeping **XXXX** and its subsidiaries under review for a possible debt downgrade, saying the company "will continue to face a significant debt burden," with large slices of debt coming due, over the next 18 months. **XXXX** said ...

Test article with entities removed

fitch goldman-sachs lehman-brother moody morgan-stanley new-york-stock-exchange standard-and-poor tyco tyco-international wall-street worldco

Actual missing entities

Prediction of Missing Entities in Text

Shares of **XXXX** slid 8 percent, or \$1.10, to \$12.65 Tuesday, as major credit agencies said the conglomerate would still be challenged in repaying its debts, despite raising \$4.6 billion Monday in taking its finance group public. Analysts at **XXXX** Investors service in **XXXX** said they were keeping **XXXX** and its subsidiaries under review for a possible debt downgrade, saying the company "will continue to face a significant debt burden," with large slices of debt coming due, over the next 18 months. **XXXX** said ...

Test article with entities removed

fitch goldman-sachs lehman-brother moody morgan-stanley new-york-stock-exchange standard-and-poor tyco tyco-international wall-street worldco

Actual missing entities

wall-street new-york nasdaq securities-exchange-commission sec merrill-lynch new-york-stock-exchange **goldman-sachs standard-and-poor**

Predicted entities given observed words (matches in blue)

Model Extensions

This slide is currently blank, intended to discuss model extensions.


Model Extensions

- HMM-topics model
 - Modeling aspects of syntax
- Hierarchical topic model
 - Modeling relations between topics
- Collocation topic models
 - Learning collocations of words within topics

Hidden Markov Topic Model


Hidden Markov Topics Model

- Syntactic dependencies \rightarrow short range dependencies
- Semantic dependencies \rightarrow long-range




(Griffiths, Steyvers, Elci, & Tenenbaum, 2004)

Transition between semantic state and syntactic states




Combining topics and syntax




THE

Combining topics and syntax




THE LOVE.....

Combining topics and syntax



THE LOVE OF.....

Combining topics and syntax



THE LOVE OF RESEARCH

Semantic topics

FOOD	MAP	DOCTOR	BOOK	GOLD	BEHAVIOR	CELLS	PLANTS
FOODS	NORTH	PATIENT	BOOKS	IRON	SELF	CELL	PLANT
BODY	EARTH	HEALTH	READING	SILVER	INDIVIDUAL	ORGANISMS	LEAVES
NUTRIENTS	SOUTH	HOSPITAL	INFORMATION	COPPER	PERSONALITY	ALGAE	SEEDS
DIET	POLE	MEDICAL	LIBRARY	METAL	RESPONSE	BACTERIA	SOIL
FAT	MAPS	CARE	REPORT	METALS	SOCIAL	MICROSCOPE	ROOTS
SUGAR	EQUATOR	PATIENTS	PAGE	STEEL	EMOTIONAL	MEMBRANE	FLOWERS
ENERGY	WEST	NURSE	TITLE	CLAY	LEARNING	ORGANISM	WATER
VEGETABLES	LINES	DOCTORS	SUBJECT	LEAD	FEELINGS	FOOD	FOOD
EATING	EAST	MEDICINE	PAGES	ADAM	PSYCHOLOGISTS	LIVING	GREEN
FRUITS	AUSTRALIA	NURSING	GUIDE	ORE	INDIVIDUALS	FUNGUS	SEED
GLOBE	TREATMENT	WORDS	ALUMINUM	PSYCHOLOGICAL	MOLD	STEMS	STEMS
VEGETABLES	POLES	NURSES	MATERIAL	MINERAL	EXPERIENCES	MATERIALS	FLOWER
WEIGHT	HEMISPHERE	PHYSICIAN	ARTICLE	MINE	ENVIRONMENT	NUCLEUS	STEM
NEEDS	LATITUDE	HOSPITALS	ARTICLES	STONE	HUMAN	CELLED	LEAF
CARBOHYDRATES	PLACES	DR	WORD	MINERALS	RESPONSES	STRUCTURES	ANIMALS
VITAMINS	LAND	SICK	FACTS	POT	BEHAVIORS	MATERIAL	ROOT
CALORIES	WORLD	ASSISTANT	AUTHOR	MINING	ATTITUDES	STRUCTURE	POLLEN
PROTEIN	COMPASS	EMERGENCY	REFERENCE	MINERS	PSYCHOLOGY	GREEN	GROWING
MINERALS	CONTINENTS	PRACTICE	NOTE	TIN	PERSON	MOLDS	GROW

Syntactic classes

SAID	THE	MORE	ON	GOOD	ONE	HE	BE
ASKED	HIS	SUCH	AT	SMALL	SOME	YOU	MAKE
THOUGHT	THEIR	LESS	INTO	NEW	MANY	THEY	GET
TOLD	YOUR	MUCH	FROM	IMPORTANT	TWO	I	HAVE
SAYS	HER	KNOWN	WITH	GREAT	EACH	SHE	GO
MEANS	ITS	JUST	THROUGH	LITTLE	ALL	WE	TAKE
CALLED	MY	BETTER	OVER	LARGE	MOST	IT	DO
CRIED	OUR	RATHER	AROUND	*	ANY	PEOPLE	FIND
SHOWS	THIS	GREATER	AGAINST	BCG	THREE	EVERYONE	USE
ANSWERED	THESE	LIVING	HIGHER	ACROSS	LONG	THIS	OTHERS
TELLS	A	LARGER	UPON	HIGH	EVERY	SCIENTISTS	HELP
REPLIED	AN	LONGER	TOWARD	DIFFERENT	SEVERAL	SOMEONE	KEEP
SHOUTED	THAT	FASTER	UNDER	SPECIAL	FOUR	WHO	GIVE
EXPLAINED	NEW	EXACTLY	ALONG	OLD	FIVE	NOBODY	LOOK
LAUGHED	THOSE	SMALLER	NEAR	STRONG	BOTH	ONE	COME
MEANT	EACH	SOMETHING	BEHIND	YOUNG	TEN	SOMETHING	WORK
WROTE	MR	BIGGER	OFF	COMMON	SIX	ANYONE	MOVE
SHOWED	ANY	FEWER	ABOVE	WHITE	MUCH	EVERYBODY	LIVE
BELIEVED	MRS	LOWER	DOWN	SINGLE	TWENTY	SOME	EAT
WHISPERED	ALL	ALMOST	BEFORE	CERTAIN	EIGHT	THEN	BECOME

NIPS Semantics

IMAGE	DATA	STATE	MEMBRANE	EXPERTS	KERNEL	NETWORK
IMAGES	GAUSSIAN	POLICY	SYNAPTIC	EXPERT	SUPPORT	NEURAL
OBJECT	MIXTURE	VALUE	CELL	GATING	VECTOR	NETWORKS
OBJECTS	LIKELIHOOD	FUNCTION	*	HME	SVM	OUTPUT
FEATURE	POSTERIOR	ACTION	CURRENT	ARCHITECTURE	KERNELS	INPUT
RECOGNITION	PRIOR	REINFORCEMENT	DENDRITIC	MIXTURE	#	TRAINING
VIEWS	DISTRIBUTION	LEARNING	POTENTIAL	LEARNING	SPACE	INPUTS
#	EM	CLASSES	NEURON	MIXTURES	FUNCTION	WEIGHTS
PIXEL	BAYESIAN	OPTIMAL	CONDUCTANCE	FUNCTION	MACHINES	#
VISUAL	PARAMETERS	*	CHANNELS	GATE	SET	OUTPUTS

NIPS Syntax

IN	IS	SEE	USED	MODEL	HOWEVER	#
WITH	WAS	SHOW	TRAINED	ALGORITHM	ALSO	*
FOR	HAS	NOTE	OBTAINED	SYSTEM	THEN	I
ON	BECOMES	CONSIDER	DESCRIBED	CASE	THIS	X
FROM	DENOTES	ASSUME	GIVEN	PROBLEM	THEREFORE	T
AT	BEING	PRESENT	FOUND	NETWORK	FIRST	N
USING	REMAINS	NEED	PRESENTED	METHOD	HERE	-
INTO	REPRESENTS	PROPOSE	DEFINED	APPROACH	NOW	C
OVER	EXISTS	DESCRIBE	GENERATED	PAPER	HENCE	F
WITHIN	SEEMS	SUGGEST	SHOWN	PROCESS	FINALLY	P

Random sentence generation

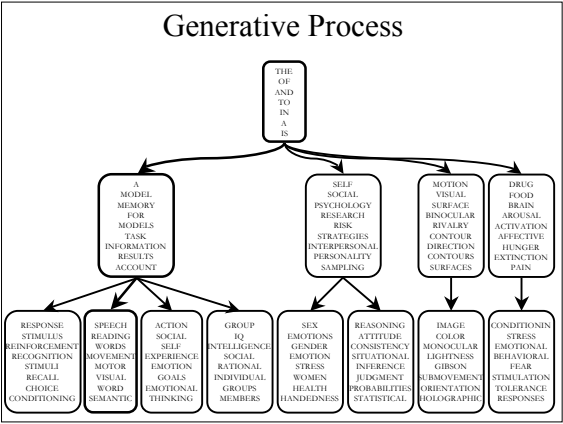
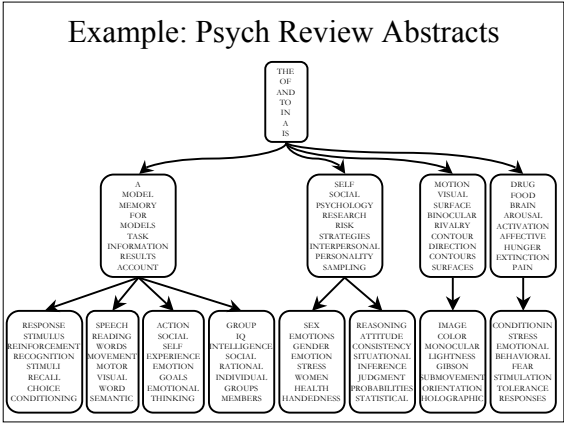
LANGUAGE:

[S] RESEARCHERS GIVE THE SPEECH
 [S] THE SOUND FEEL NO LISTENERS
 [S] WHICH WAS TO BE MEANING
 [S] HER VOCABULARIES STOPPED WORDS
 [S] HE EXPRESSLY WANTED THAT BETTER VOWEL

Nested Chinese Restaurant Process

Topic Hierarchies

- In regular topic model, no relations between topics
- Nested Chinese Restaurant Process
 - Blei, Griffiths, Jordan, Tenenbaum (2004)
 - Learn hierarchical structure, as well as topics within structure



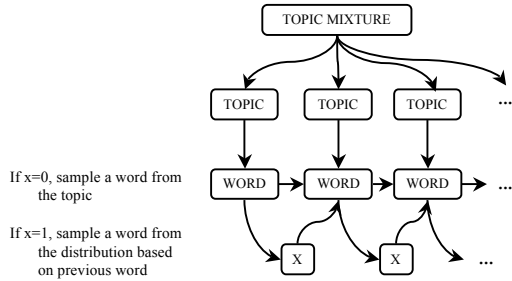
Collocation Topic Model

What about collocations?

- Why are these words related?
 - PLAY - GROUND
 - DOW - JONES
 - BUMBLE - BEE
- Suggests at least two routes for association:
 - Semantic
 - Collocation

→ Integrate collocations into topic model

Collocation Topic Model




Collocation Topic Model

Example:
"DOW JONES RISES"

JONES is more likely explained as a word following DOW than as word sampled from topic

Result: DOW_JONES recognized as collocation



Examples Topics from New York Times

Terrorism

SEPT_11
WAR
SECURITY
IRAQ
TERRORISM
NATION
KILLED
AFGHANISTAN
ATTACKS
OSAMA_BIN_LADEN
AMERICAN
ATTACK
NEW_YORK_REGION
NEW
MILITARY
NEW_YORK
WORLD
NATIONAL
QAEDA
TERRORIST_ATTACKS

Wall Street Firms

WALL_STREET
ANALYSTS
INVESTORS
FIRM
GOLDMAN_SACHS
FIRMS
INVESTMENT
MERRILL_LYNCH
COMPANIES
SECURITIES
RESEARCH
STOCK
BUSINESS
ANALYST
WALL_STREET_FIRMS
SALOMON_SMITH_BARNNEY
CLIENTS
INVESTMENT_BANKING
INVESTMENT_BANKERS
INVESTMENT_BANKS

Stock Market

WEEK
DOW_JONES
POINTS
10_YR_TREASURY_YIELD
PERCENT
CLOSE
NASDAQ_COMPOSITE
STANDARD_POOR
CHANGE
FRIDAY
DOW_INDUSTRIALS
GRAPH_TRACKS
EXPECTED
BILLION
NASDAQ_COMPOSITE_INDEX
EST_02
PHOTO_YESTERDAY
YIN
10
500_STOCK_INDEX

Bankruptcy

BANKRUPTCY
CREDITORS
BANKRUPTCY_PROTECTION
ASSETS
COMPANY
FILED
BANKRUPTCY_FILING
ENRON
BANKRUPTCY_COURT
KIMART
CHAPTER_11
FILING
COOPER
BILLIONS
COMPANIES
BANKRUPTCY_PROCEEDINGS
DEBTS
RESTRUCTURING
CASE
GROUP