Lexical semantics, sentiment lexicons, lexicon expansion

CS 585, Fall 2017

Introduction to Natural Language Processing

http://people.cs.umass.edu/~brenocon/inlp2017

Brendan O'Connor
College of Information and Computer Sciences
University of Massachusetts Amherst

Word sense disambiguation

- Task: Choose a word's sense in context
- Given KB and text:
  Want to tag spans in text with concept IDs
- Disambiguation problem
  - "I saw the bank" => bank#1 or bank#2?
  - "Michael Jordan was here" => ?

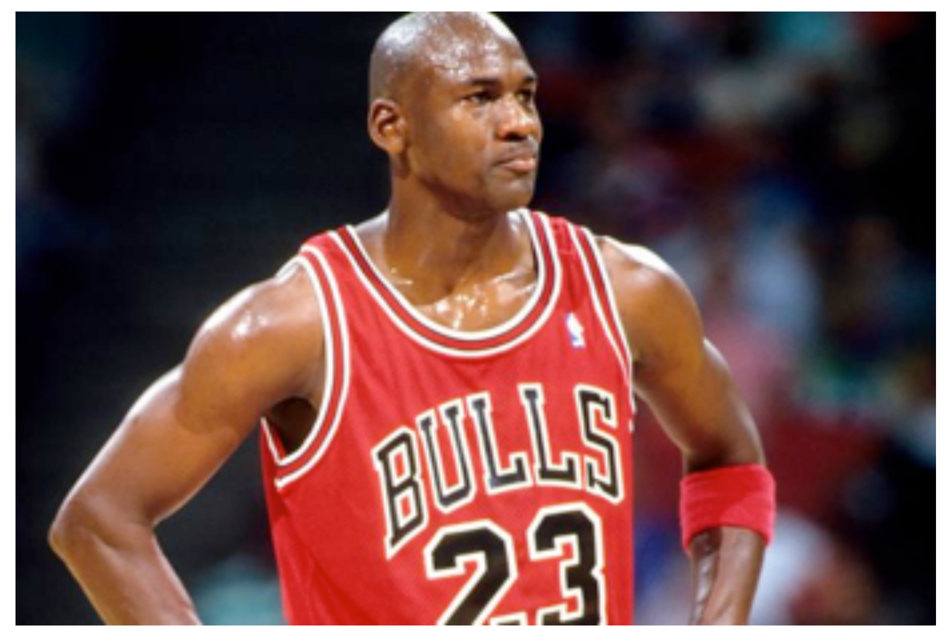

Word sense disambiguation

- Task: Choose a word's sense in context
- Given KB and text:
  Want to tag spans in text with concept IDs
- Disambiguation problem
  - "I saw the <u>bank</u>" => bank#1 or bank#2?
  - "<u>Michael Jordan</u> was here" => ?

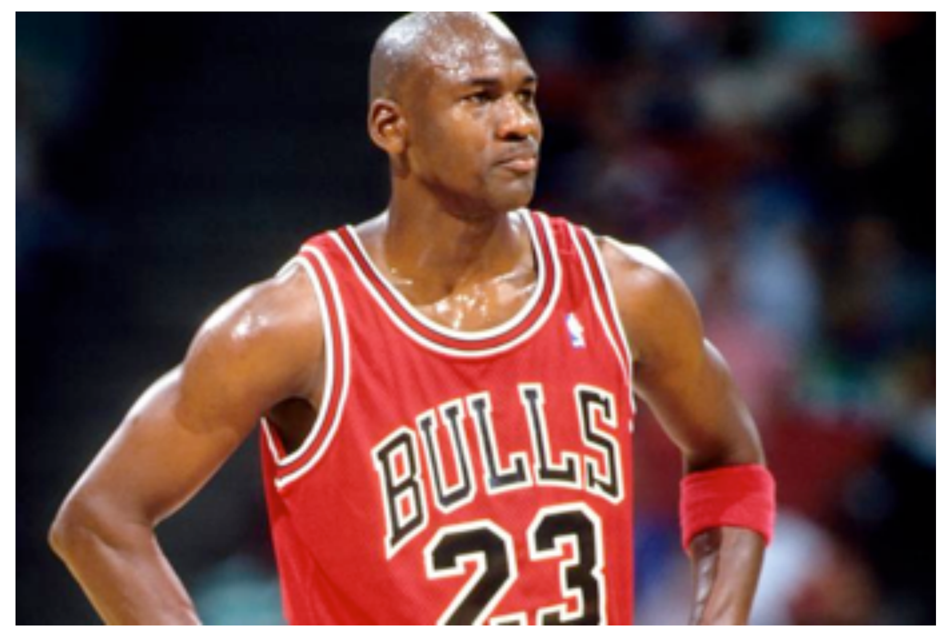

- Many terms for this: concept tagging, entity linking, "wikification", WSD

Word sense disambiguation

- Supervised setting: need ground-truth concept IDs for words in text
- Main approach: use *contextual information* to disambiguate.

Intuition from Warren Weaver (1955):

"If one examines the words in a book, one at a time as through an opaque mask with a hole in it one word wide, then it is obviously impossible to determine, one at a time, the meaning of the words…

But if one lengthens the slit in the opaque mask, until one can see not only the central word in question but also say N words on either side, then if N is large enough one can unambiguously decide the meaning of the central word…

The practical question is : ``What minimum value of N will, at least in a tolerable fraction of cases, lead to the correct choice of meaning for the central word?"

[slide: SLP3]

Two kinds of features in the vectors

- **Collocational** features and **bag-of-words** features
  - **Collocational**
    - Features about words at **specific** positions near target word
      - Often limited to just word identity and POS
  - **Bag-of-words**
    - Features about words that occur anywhere in the window (regardless of position)
      - Typically limited to frequency counts

[slide: SLP3]

Examples

- Example text (WSJ):

  An electric guitar and **bass** player stand off to one side not really part of the scene

- Assume a window of +/- 2 from the target

[slide: SLP3]

Examples

- Example text (WSJ)

  An electric guitar and **bass** player stand off to one side not really part of the scene,

- Assume a window of +/- 2 from the target

[slide: SLP3]

Collocational features

- Position-specific information about the words and collocations in window
- guitar and bass player stand

$$\left[w_{i-2}, \mathrm{POS}_{i-2}, w_{i-1}, \mathrm{POS}_{i-1}, w_{i+1}, \mathrm{POS}_{i+1}, w_{i+2}, \mathrm{POS}_{i+2}, w_{i-2}^{i-1}, w_{i}^{i+1}\right]$$

```
[guitar, NN, and, CC, player, NN, stand, VB, and guitar, player stand]
```

- word 1,2,3 grams in window of ± 3 is common

[slide: SLP3]

Bag-of-words features

- “an unordered set of words” – position ignored
- Counts of words occur within the window.
- First choose a vocabulary
- Then count how often each of those terms occurs in a given window
  - sometimes just a binary “indicator” 1 or 0

[slide: SLP3]

Word sense disambiguation

- Supervised setting: need ground-truth concept IDs for words in text
- Contextual features
  - Word immediately to left ... to right ...
  - Word within 10 word window (20 word window? entire document?)
- Features from matching a concept description, if your KB has one
  - *Michael Jeffrey Jordan (born February 17, 1963), also known by his initials, MJ,[1] is an American former professional basketball player. He is also a businessman, and principal owner and chairman of the Charlotte Hornets. Jordan played 15 seasons in the National Basketball Association (NBA) for theChicago Bulls and Washington Wizards.*
- Overall (prior) sense frequency
  - For WN, hard to beat Most Frequent Sense baseline (?!)
- Contrast to distributional semantics:
  unsupervised learning of word meanings

Affect in text

Affective meaning

- Drawing on literatures in
  - affective computing (Picard 95)
  - linguistic subjectivity (Wiebe and colleagues)
  - social psychology (Pennebaker and colleagues)
- Can we model the lexical semantics relevant to:
  - sentiment
  - emotion
  - personality
  - mood
  - attitudes

[slide: SLP3]

Why compute affective meaning?

- Detecting:
  - sentiment towards politicians, products, countries, ideas
  - frustration of callers to a help line
  - stress in drivers or pilots
  - depression and other medical conditions
  - confusion in students talking to e-tutors
  - emotions in novels (e.g., for studying groups that are feared over time)
- Could we generate:
  - emotions or moods for literacy tutors in the children's storybook domain
  - emotions or moods for computer games
  - personalities for dialogue systems to match the user

[slide: SLP3]

Scherer's typology of affective states

Emotion: relatively brief episode of synchronized response of all or most organismic subsystems in response to the evaluation of an event as being of major significance

angry, sad, joyful, fearful, ashamed, proud, desperate

Mood: diffuse affect state …change in subjective feeling, of low intensity but relatively long duration, often without apparent cause

cheerful, gloomy, irritable, listless, depressed, buoyant

Interpersonal stance: affective stance taken toward another person in a specific interaction, coloring the interpersonal exchange

distant, cold, warm, supportive, contemptuous

Attitudes: relatively enduring, affectively colored beliefs, preferences predispositions towards objects or persons

liking, loving, hating, valuing, desiring

Personality traits: emotionally laden, stable personality dispositions and behavior tendencies, typical for a person

nervous, anxious, reckless, morose, hostile, envious, jealous

[slide: SLP3]

Sentiment/affect lexicons

- Overall text affect analyzers/generators/etc. tend to be domain-specific
- Sentiment/affect lexicons: attempt to be give useful word-level information across many situations

- Long list of polarity lexicons in Reagan et al., 2016
  - https://arxiv.org/pdf/1512.00531.pdf

Dictionary	# Fixed	# Stems	Total	Range	# Pos	# Neg	Construction	License	Ref.
labMT	10222	0	10222	$1.3 \to 8.5$	7152	2977	Survey: MT, 50 ratings	CC	[5]
ANEW	1034	0	1034	$1.2 \to 8.8$	584	449	Survey: FSU Psych 101	Free for research	[7]
LIWC07	2145	2338	4483	[-1,0,1]	406	500	Manual	Paid, commercial	[8]
MPQA	5587	1605	7192	[-1,0,1]	2393	4342	Manual + ML	GNU GPL	[9]
OL	6782	0	6782	[-1,1]	2003	4779	Dictionary propagation	Free	[10]
WK	13915	0	13915	$1.3 \to 8.5$	7761	5945	Survey: MT, at least 14 ratings	CC	[11]
LIWC01	1232	1090	2322	[-1,0,1]	266	344	Manual	Paid, commercial	[8]
LIWC15	4071	2478	6549	[-1,0,1]	642	746	Manual	Paid, commercial	[8]
PANAS-X	20	0	20	[-1,1]	10	10	Manual	Copyrighted paper	[21]
Pattern	1528	0	1528	$-1.0 \to 1.0$	575	679	Unspecified	BSD	[22]
SentiWordNet	147700	0	147700	$-1.0 \to 1.0$	17677	20410	Synset synonyms	CC BY-SA 3.0	[23]
AFINN	2477	0	2477	[-5,-4, ...,4,5]	878	1598	Manual	ODbL v1.0	[24]
GI	3629	0	3629	[-1,1]	1631	1998	Harvard-IV-4	Unspecified	[25]
WDAL	8743	0	8743	$0.0 \to 3.0$	6517	1778	Survey: Columbia students	Unspecified	[26]
EmoLex	14182	0	14182	[-1,0,1]	2231	3243	Survey: MT	Free for research	[27]
MaxDiff	1515	0	1515	$-1.0 \to 1.0$	775	726	Survey: MT, MaxDiff	Free for research	[28]
HashtagSent	54129	0	54129	$-6.9 \to 7.5$	32048	22081	PMI with hashtags	Free for research	[29]
Sent140Lex	62468	0	62468	$-5.0 \to 5.0$	38312	24156	PMI with emoticons	Free for research	[30]
SOCAL	7494	0	7494	$-30.2 \to 30.7$	3325	4169	Manual	GNU GPL	[31]
SenticNet	30000	0	30000	$-1.0 \to 1.0$	16715	13285	Label propogation	Citation requested	[32]
Emoticons	132	0	132	[-1,0,1]	58	48	Manual	Open source code	[33]
SentiStrength	1270	1345	2615	[-5,-4, ...,4,5]	601	2002	LIWC+GI	Unknown	[34]
VADER	7502	0	7502	$-3.9 \to 3.4$	3333	4169	MT survey, 10 ratings	Freely available	[35]
Umigon	927	0	927	[-1,1]	334	593	Manual	Public Domain	[36]
USent	592	0	592	[-1,1]	63	529	Manual	CC	[37]
EmoSenticNet	13188	0	13188	[-10,-2,-1,0,1,10]	9332	1480	Bootstrapped extension	Non-commercial	[38]

LIWC (Linguistic Inquiry and Word Count)

Pennebaker, J.W., Booth, R.J., & Francis, M.E. (2007). Linguistic Inquiry and Word Count: LIWC 2007. Austin, TX

- Very commonly used, very commonly criticized.
  Created by psychologists (not linguists...)
- Home page: http://www.liwc.net/
- 2300 words, >70 classes
- **Affective Processes**
  - negative emotion (*bad, weird, hate, problem, tough*)
  - positive emotion (*love, nice, sweet*)
- **Cognitive Processes**
  - Tentative (*maybe, perhaps, guess*), Inhibition (*block, constraint*)
- **Pronouns, Negation** (*no, never*), **Quantifiers** (*few, many*)
- Costs money

[slide: SLP3]

NRC Word-Emotion Association Lexicon

Mohammad and Turney 2011

- 10,000 words chosen mainly from earlier lexicons
- Labeled by Amazon Mechanical Turk
- 5 Turkers per hit
- Give Turkers an idea of the relevant sense of the word
- Result:

```
amazingly   anger   0
amazingly   anticipation    0
amazingly   disgust 0
amazingly   fear    0
amazingly   joy 1
amazingly   sadness 0
amazingly   surprise    1
amazingly   trust   0
amazingly   negative    0
amazingly   positive    1
```

EmoLex	# of terms
EmoLex-Uni:	
Unigrams from Macquarie Thesaurus	
adjectives	200
adverbs	200
nouns	200
verbs	200
EmoLex-Bi:	
Bigrams from Macquarie Thesaurus	
adjectives	200
adverbs	187
nouns	200
verbs	200
EmoLex-GI:	
Terms from General Inquirer	
negative terms	2119
neutral terms	4226
positive terms	1787
EmoLex-WAL:	
Terms from WordNet Affect Lexicon	
anger terms	165
disgust terms	37
fear terms	100
joy terms	165
sadness terms	120
surprise terms	53
Union	**10170**

[slide: SLP3]

The AMT Hit

Prompt word: *startle*

Q1. Which word is closest in meaning (most related) to *startle*?

- *automobile*
- *shake*
- *honesty*
- *entertain*

Q2. How positive (good, praising) is the word *startle*?

- *startle* is not positive
- *startle* is weakly positive
- *startle* is moderately positive
- *startle* is strongly positive

Q3. How negative (bad, criticizing) is the word *startle*?

- *startle* is not negative
- *startle* is weakly negative
- *startle* is moderately negative
- *startle* is strongly negative

Q4. How much is *startle* associated with the emotion joy? (For example, *happy* and *fun* are strongly associated with joy.)

- *startle* is not associated with joy
- *startle* is weakly associated with joy
- *startle* is moderately associated with joy
- *startle* is strongly associated with joy

Q5. How much is *startle* associated with the emotion sadness? (For example, *failure* and *heart-break* are strongly associated with sadness.)

- *startle* is not associated with sadness
- *startle* is weakly associated with sadness
- *startle* is moderately associated with sadness
- *startle* is strongly associated with sadness

Q6. How much is *startle* associated with the emotion fear? (For example, *horror* and *scary* are strongly associated with fear.)

- Similar choices as in 4 and 5 above

Q7. How much is *startle* associated with the emotion anger? (For example, *rage* and *shouting* are strongly associated with anger.)

- Similar choices as in 4 and 5 above

Q8. How much is *startle* associated with the emotion trust? (For example, *faith* and *integrity* are strongly associated with trust.)

- Similar choices as in 4 and 5 above

Q9. How much is *startle* associated with the emotion disgust? (For example, *gross* and *cruelty* are strongly associated with disgust.)

- Similar choices as in 4 and 5 above

...

[slide: SLP3]

Sidenote

- Same author (Saif Mohammad) also has nice papers/webpages on logistic regression-based Twitter sentiment classifiers and other sentiment lexicons
  - http://saifmohammad.com/WebPages/lexicons.html

VADER

- Hutto and Gilbert (2014), freely available lexicon+software, esp for social media

- Crowdsourced lexicon

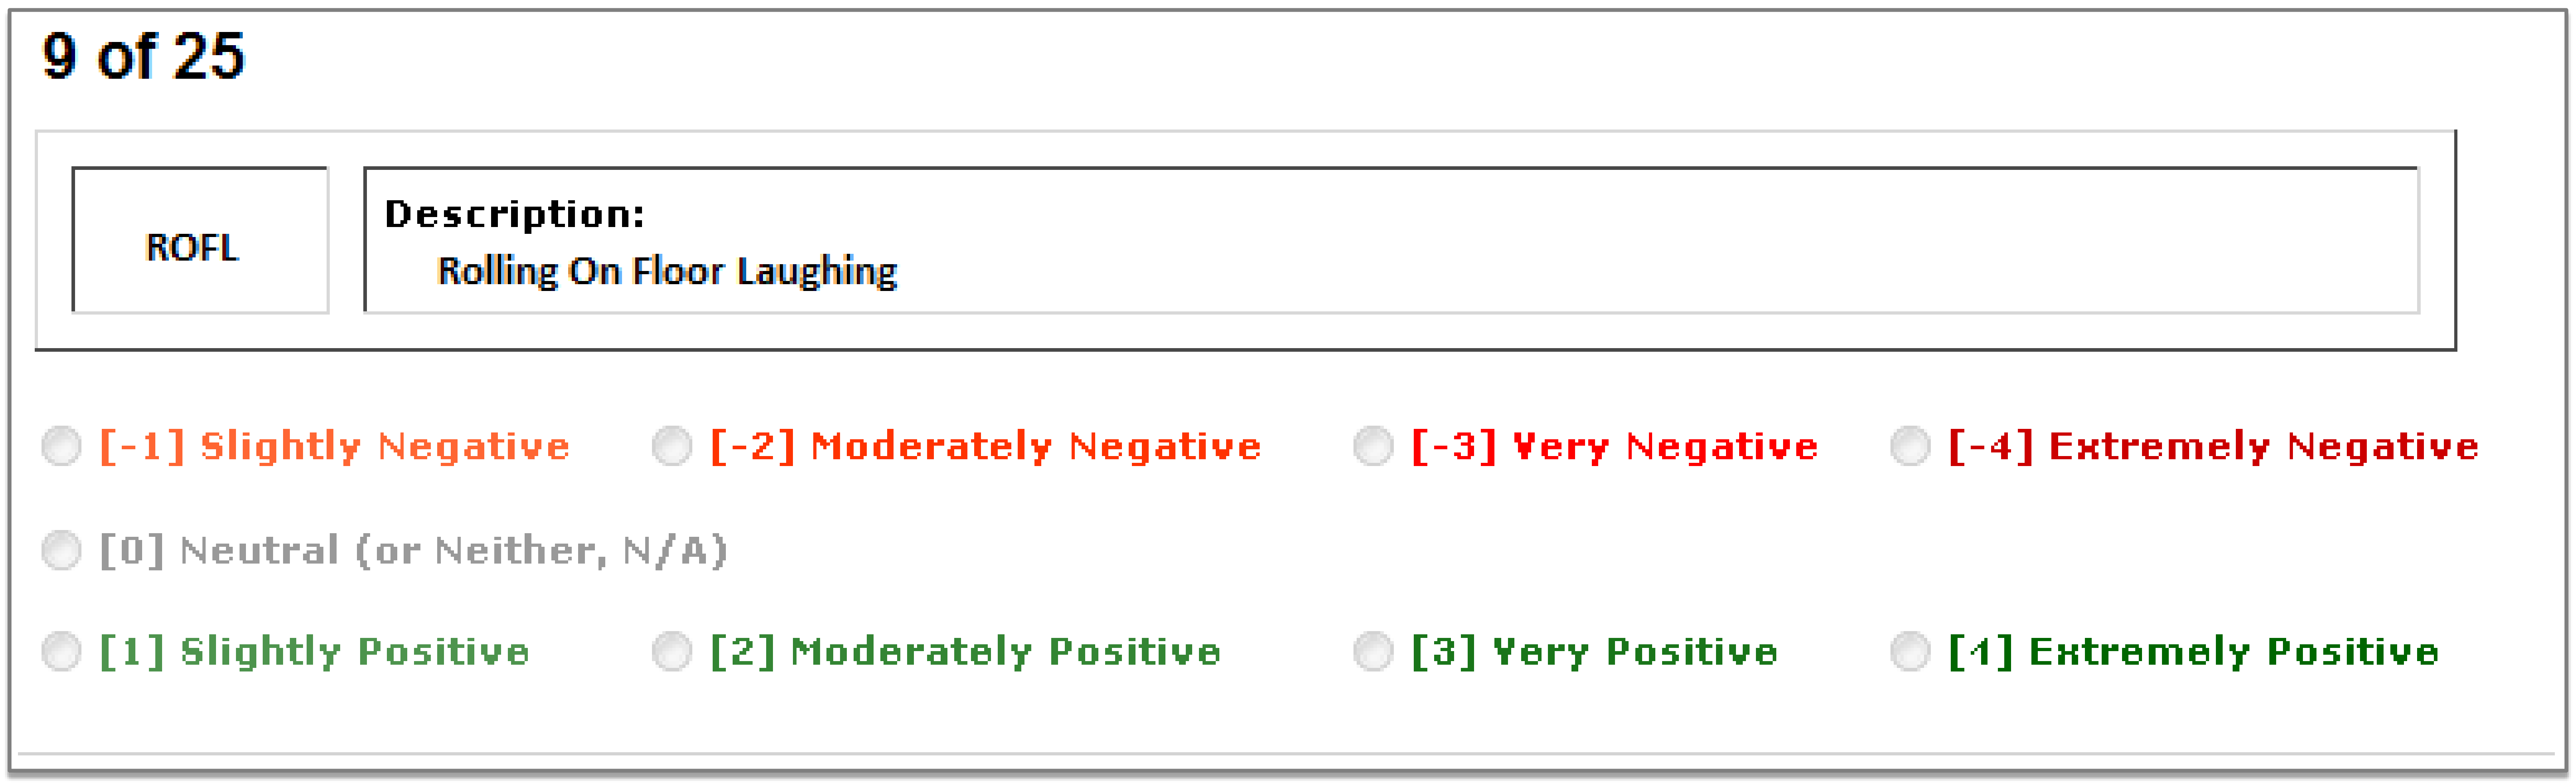

VADER

- Rule-based text classifier (not sup learning) on top of their sentiment lexicon
  - Punctuation, capitalization, degree modifiers / intensifiers, "but" as contrastive, negations
- Can exceed supervised learning performance
  - I'd expect sup learning wins if there's lots of in-domain training data... but that's not always feasible

	3-Class Classification Accuracy (F1 scores) Test Sets			
	Tweets	Movie	Amazon	NYT
VADER	**0.96**	0.61	**0.63**	**0.55**
NB (tweets)	0.84	0.53	0.53	0.42
ME (tweets)	0.83	0.56	0.58	0.45
SVM-C (tweets)	0.83	0.56	0.55	0.46
SVM-R (tweets)	0.65	0.49	0.51	0.46
NB (movie)	0.56	**0.75**	0.49	0.44
ME (movie)	0.56	**0.75**	0.51	0.45
NB (amazon)	0.69	0.55	0.61	0.48
ME (amazon)	0.67	0.55	0.60	0.43
SVM-C (amazon)	0.64	0.55	0.58	0.42
SVM-R (amazon)	0.54	0.49	0.48	0.44
NB (nyt)	0.59	0.56	0.51	0.49
ME (nyt)	0.58	0.55	0.51	0.50

Semi-supervised lexicon learning

- You have
  - 1. Large unlabeled corpus
  - 2. Some seed terms (positive and/or negative)
- Goal: expand your set of terms
- Intuition: use *co-occurrence* or *pattern* frequencies in corpus

Hatzivassiloglou and McKeown intuition for identifying word polarity

Vasileios Hatzivassiloglou and Kathleen R. McKeown. 1997. Predicting the Semantic Orientation of Adjectives. ACL, 174–181

- Adjectives conjoined by "*and*" have same polarity
  - Fair **and** legitimate, corrupt **and** brutal
  - *fair **and** brutal, *corrupt **and** legitimate
- Adjectives conjoined by "*but*" do not
  - fair **but** brutal

[slide: SLP3]

Hatzivassiloglou & McKeown 1997
Step 1

- Label **seed set** of 1336 adjectives (all >20 in 21 million word WSJ corpus)
  - 657 positive
    - adequate central clever famous intelligent remarkable reputed sensitive slender thriving…
  - 679 negative
    - contagious drunken ignorant lanky listless primitive strident troublesome unresolved unsuspecting…

[slide: SLP3]

Hatzivassiloglou & McKeown 1997 Step 2

- Expand seed set to conjoined adjectives

[slide: SLP3]

Hatzivassiloglou & McKeown 1997
Step 2

- Expand seed set to conjoined adjectives

Google "was nice and"

[Nice location in Porto and the front desk staff **was nice and** helpful ...]
www.tripadvisor.com/ShowUserReviews-g189180-d206904-r12068... +1
Mercure Porto Centro: Nice location in Porto and the front desk staff **was nice and** helpful - See traveler reviews, 77 candid photos, and great deals for Porto, ...

[If a girl **was nice and** classy, but had some vibrant purple dye in ...]
answers.yahoo.com › Home › All Categories › Beauty & Style › Hair +1
4 answers - Sep 21
Question: Your personal opinion or what you think other people's opinions might ...
Top answer: I think she would be cool and confident like katy perry :)

[slide: SLP3]

Hatzivassiloglou & McKeown 1997
Step 2

- Expand seed set to conjoined adjectives

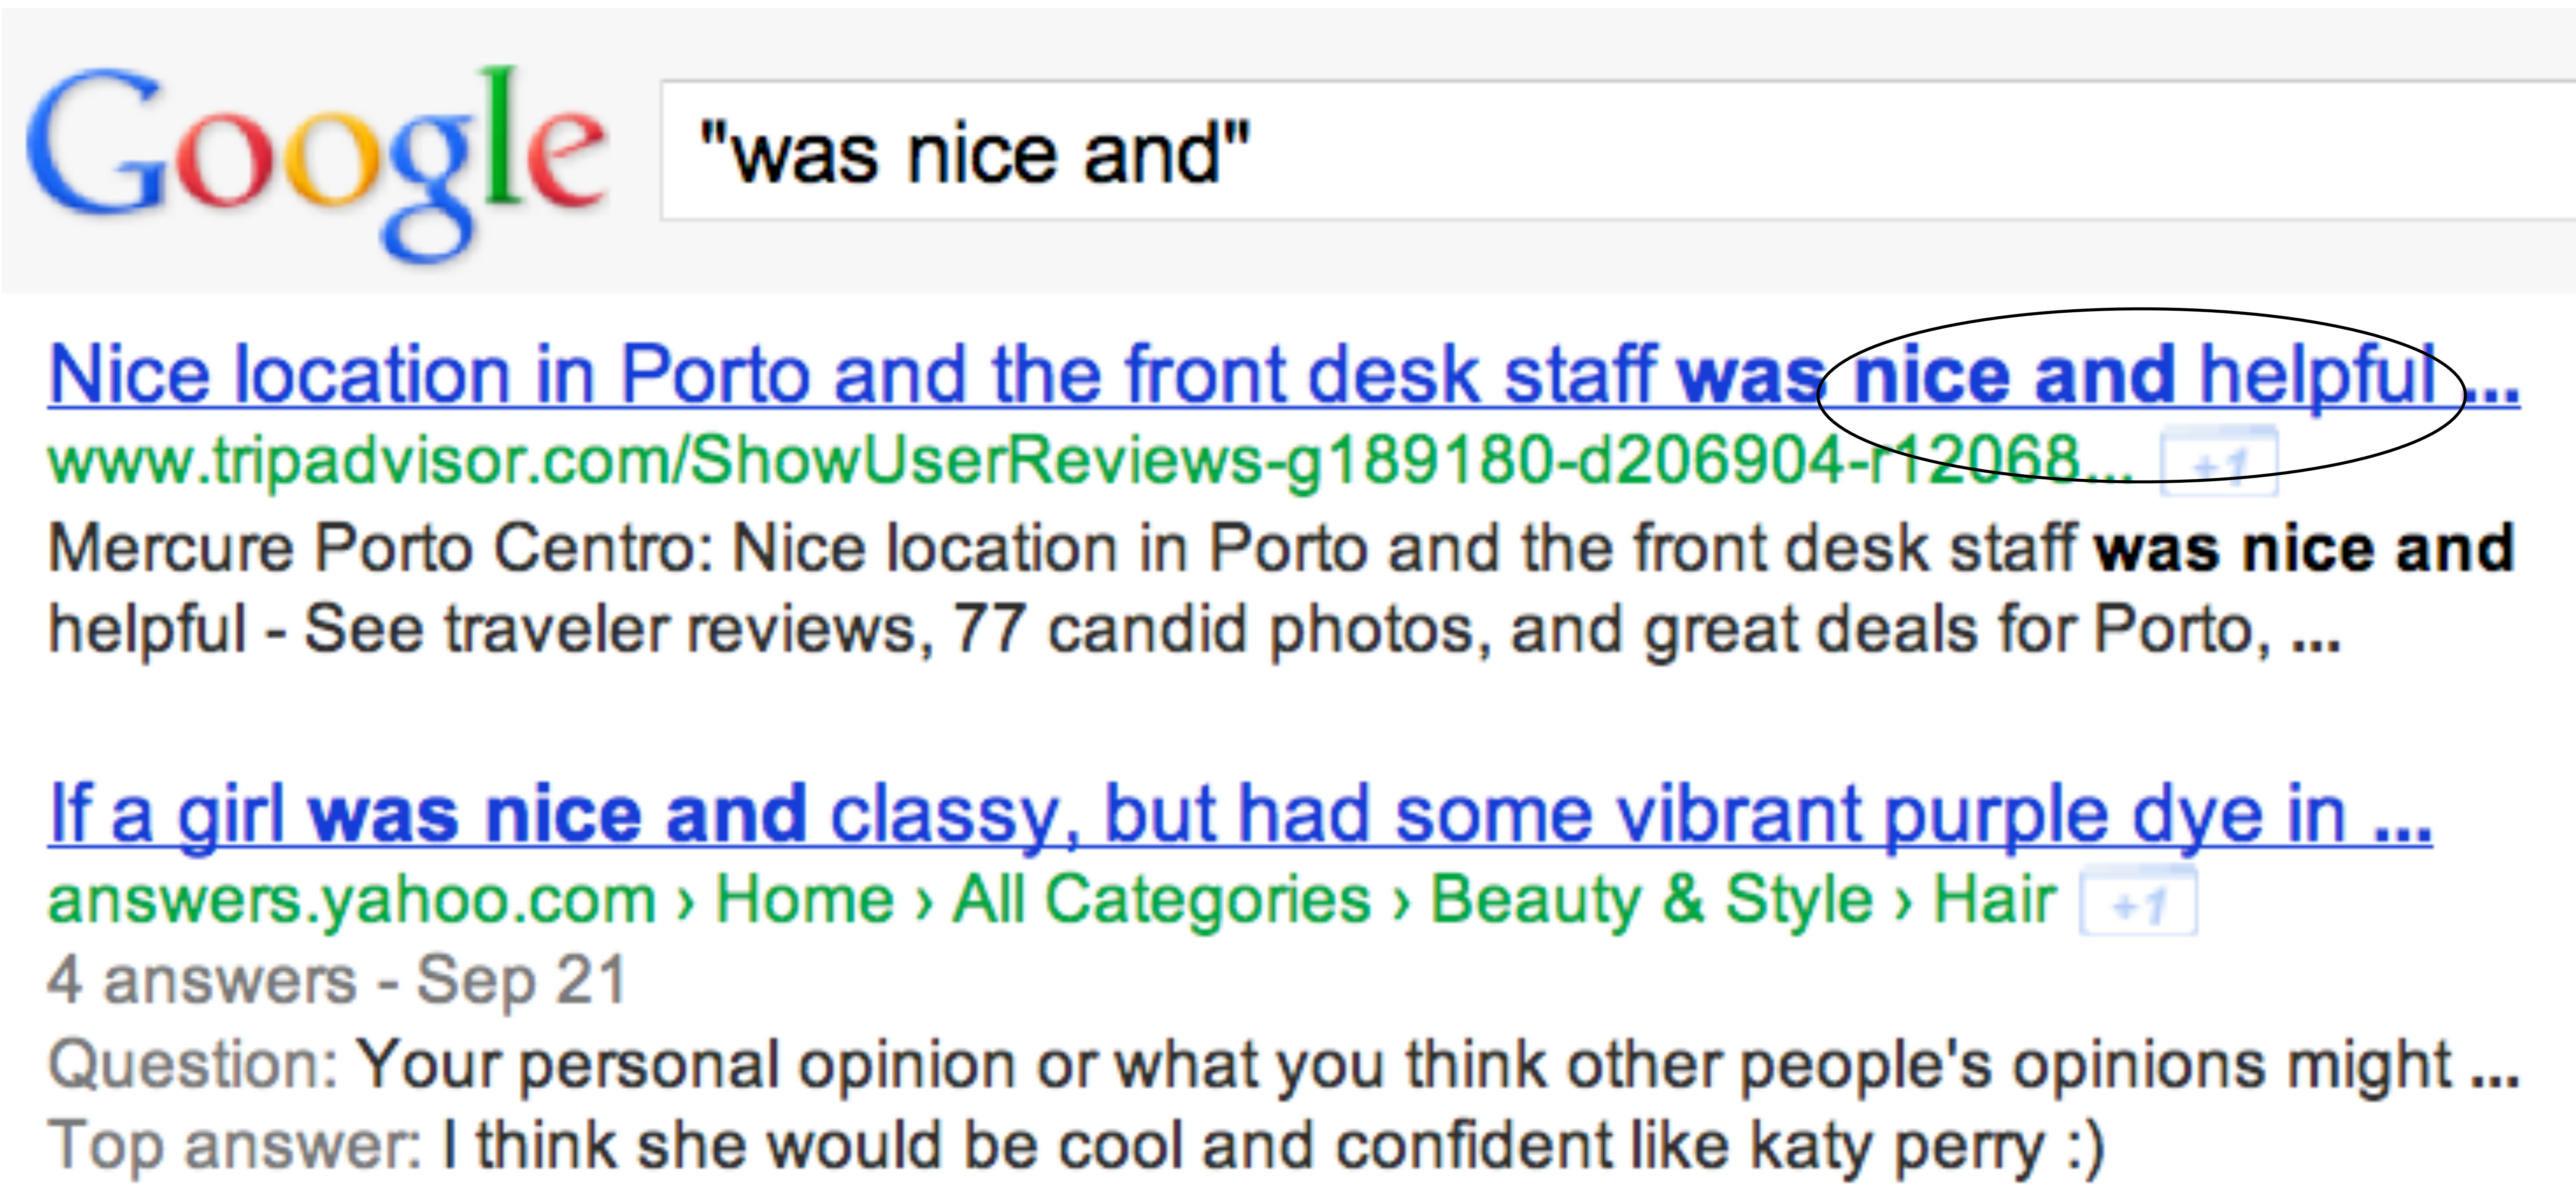

nice, helpful

[slide: SLP3]

Hatzivassiloglou & McKeown 1997
Step 2

- Expand seed set to conjoined adjectives

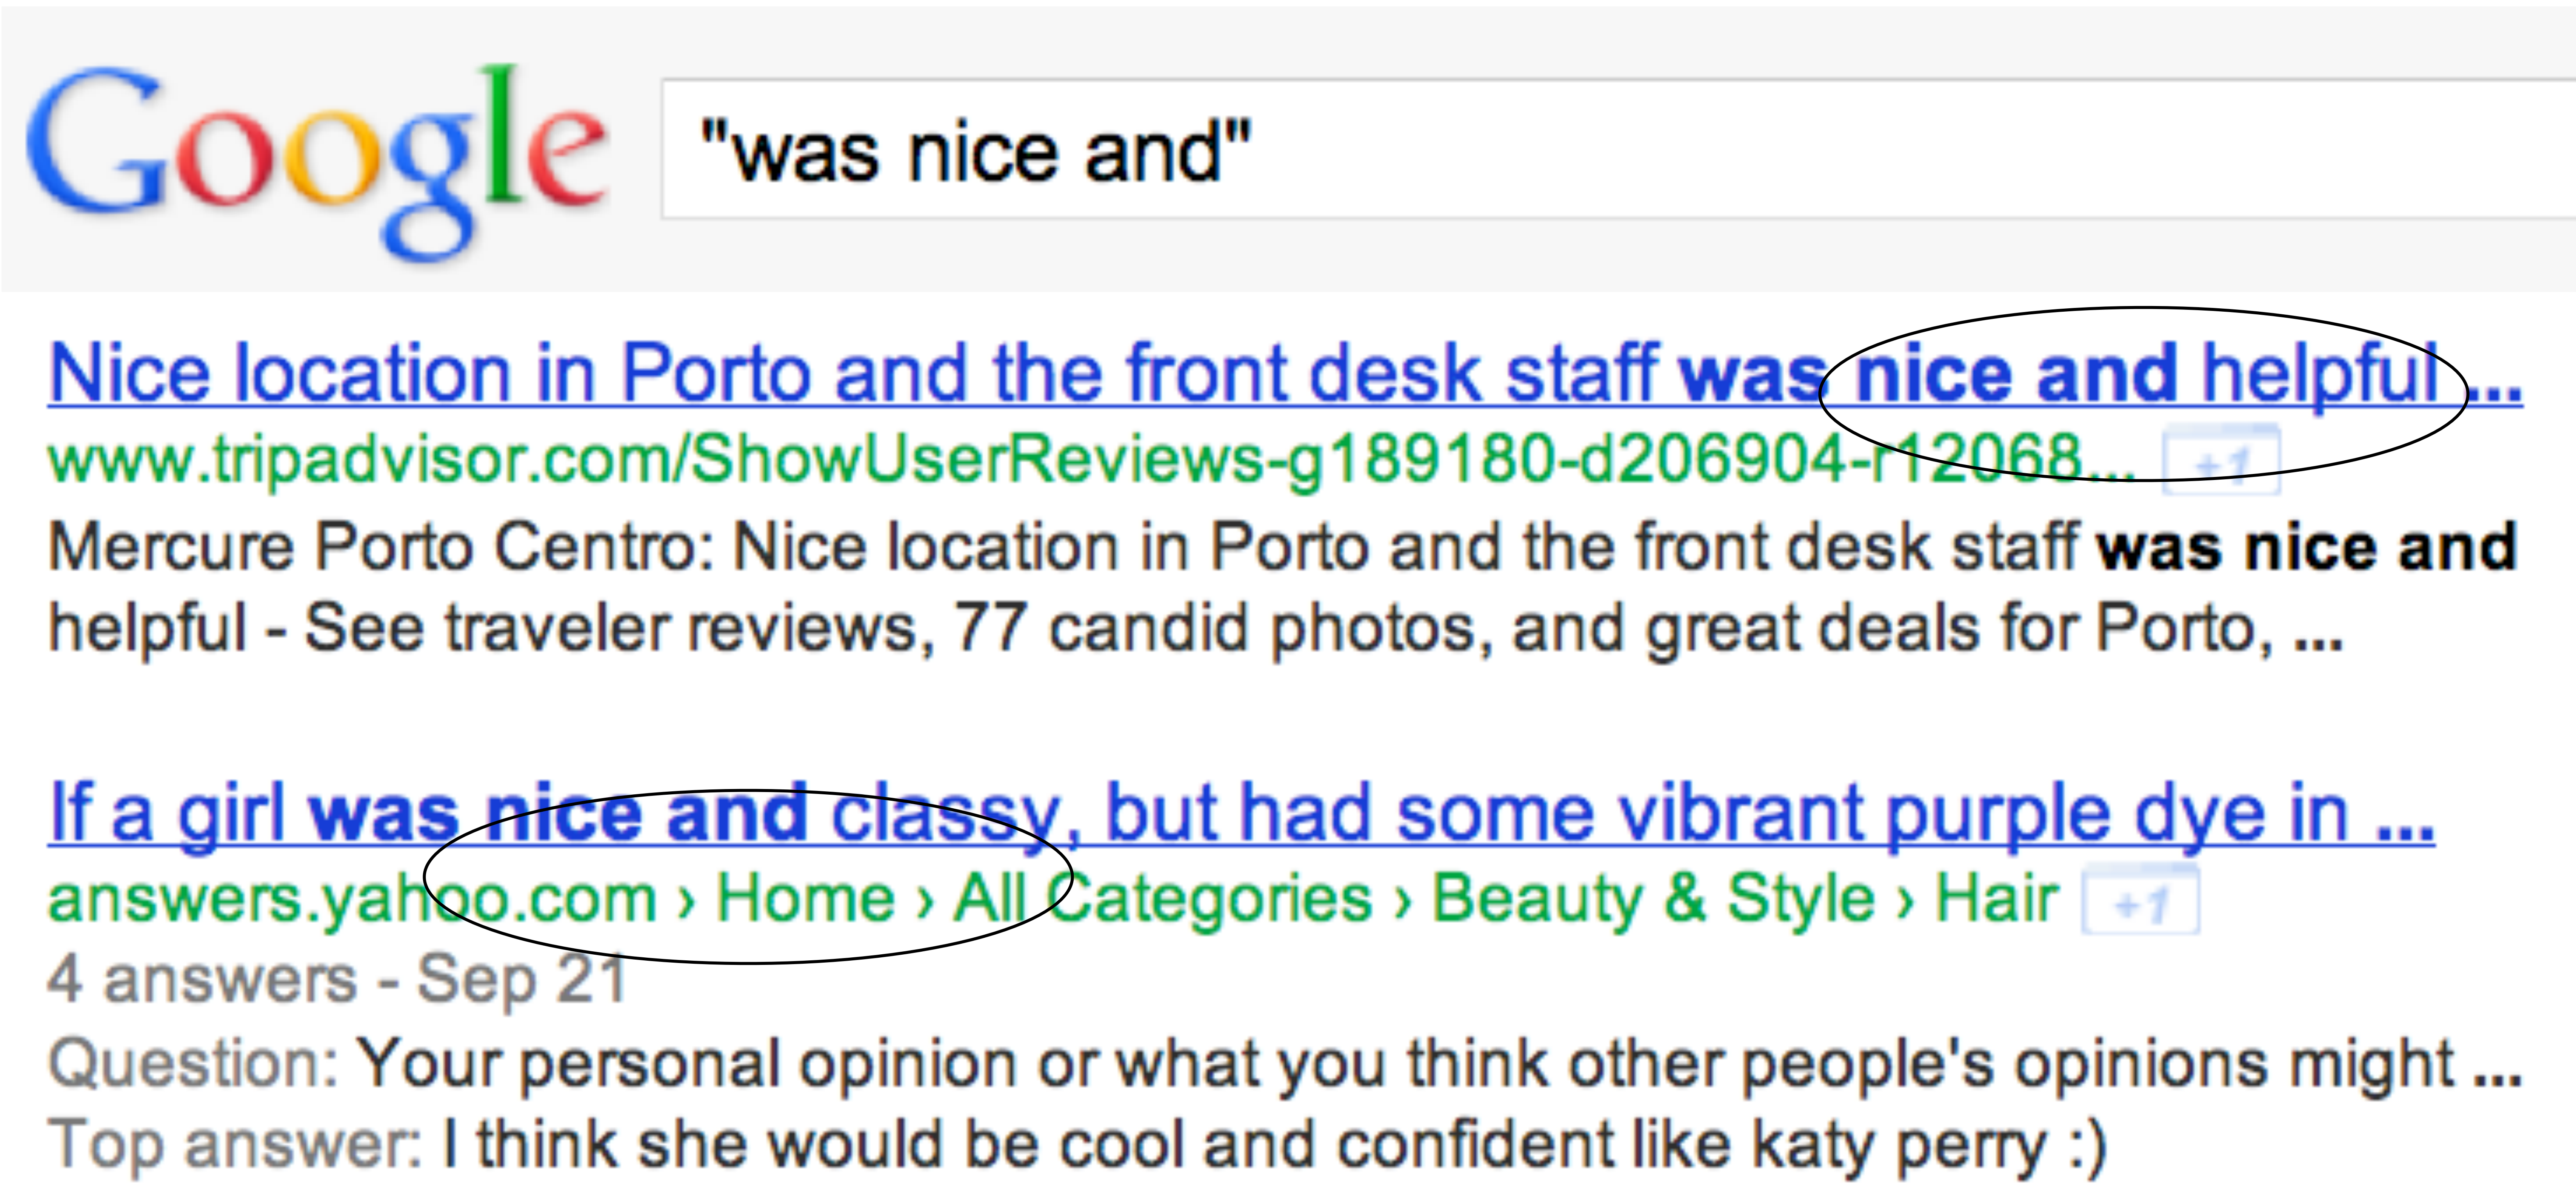

nice, helpful

nice, classy

[slide: SLP3]

Hatzivassiloglou & McKeown 1997 Step 3

- Supervised classifier assigns "polarity similarity" to each word pair, resulting in graph:

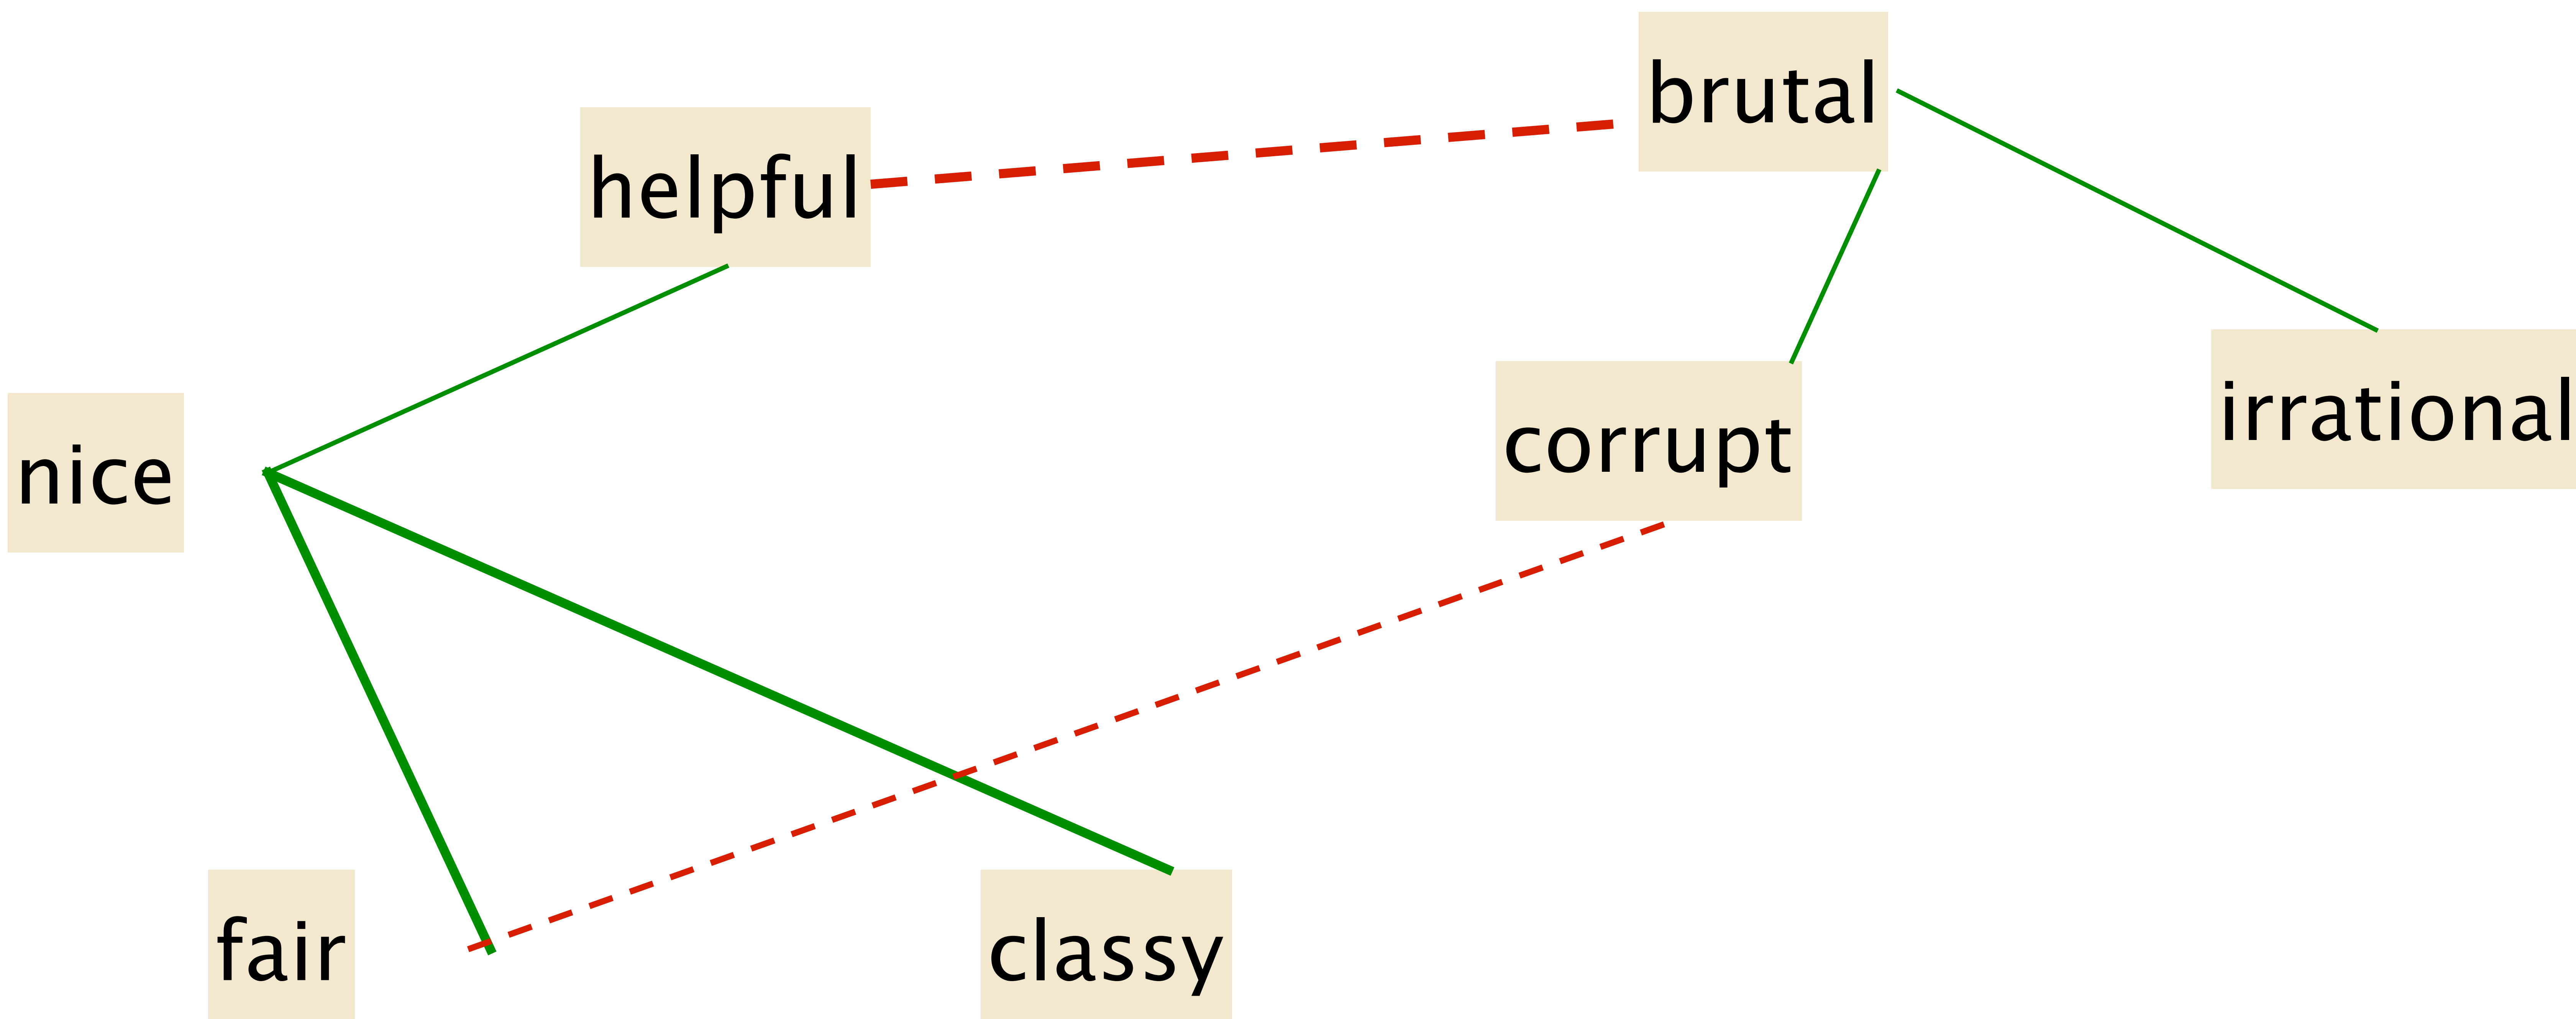

[slide: SLP3]

Hatzivassiloglou & McKeown 1997
Step 4

- Clustering for partitioning the graph into two

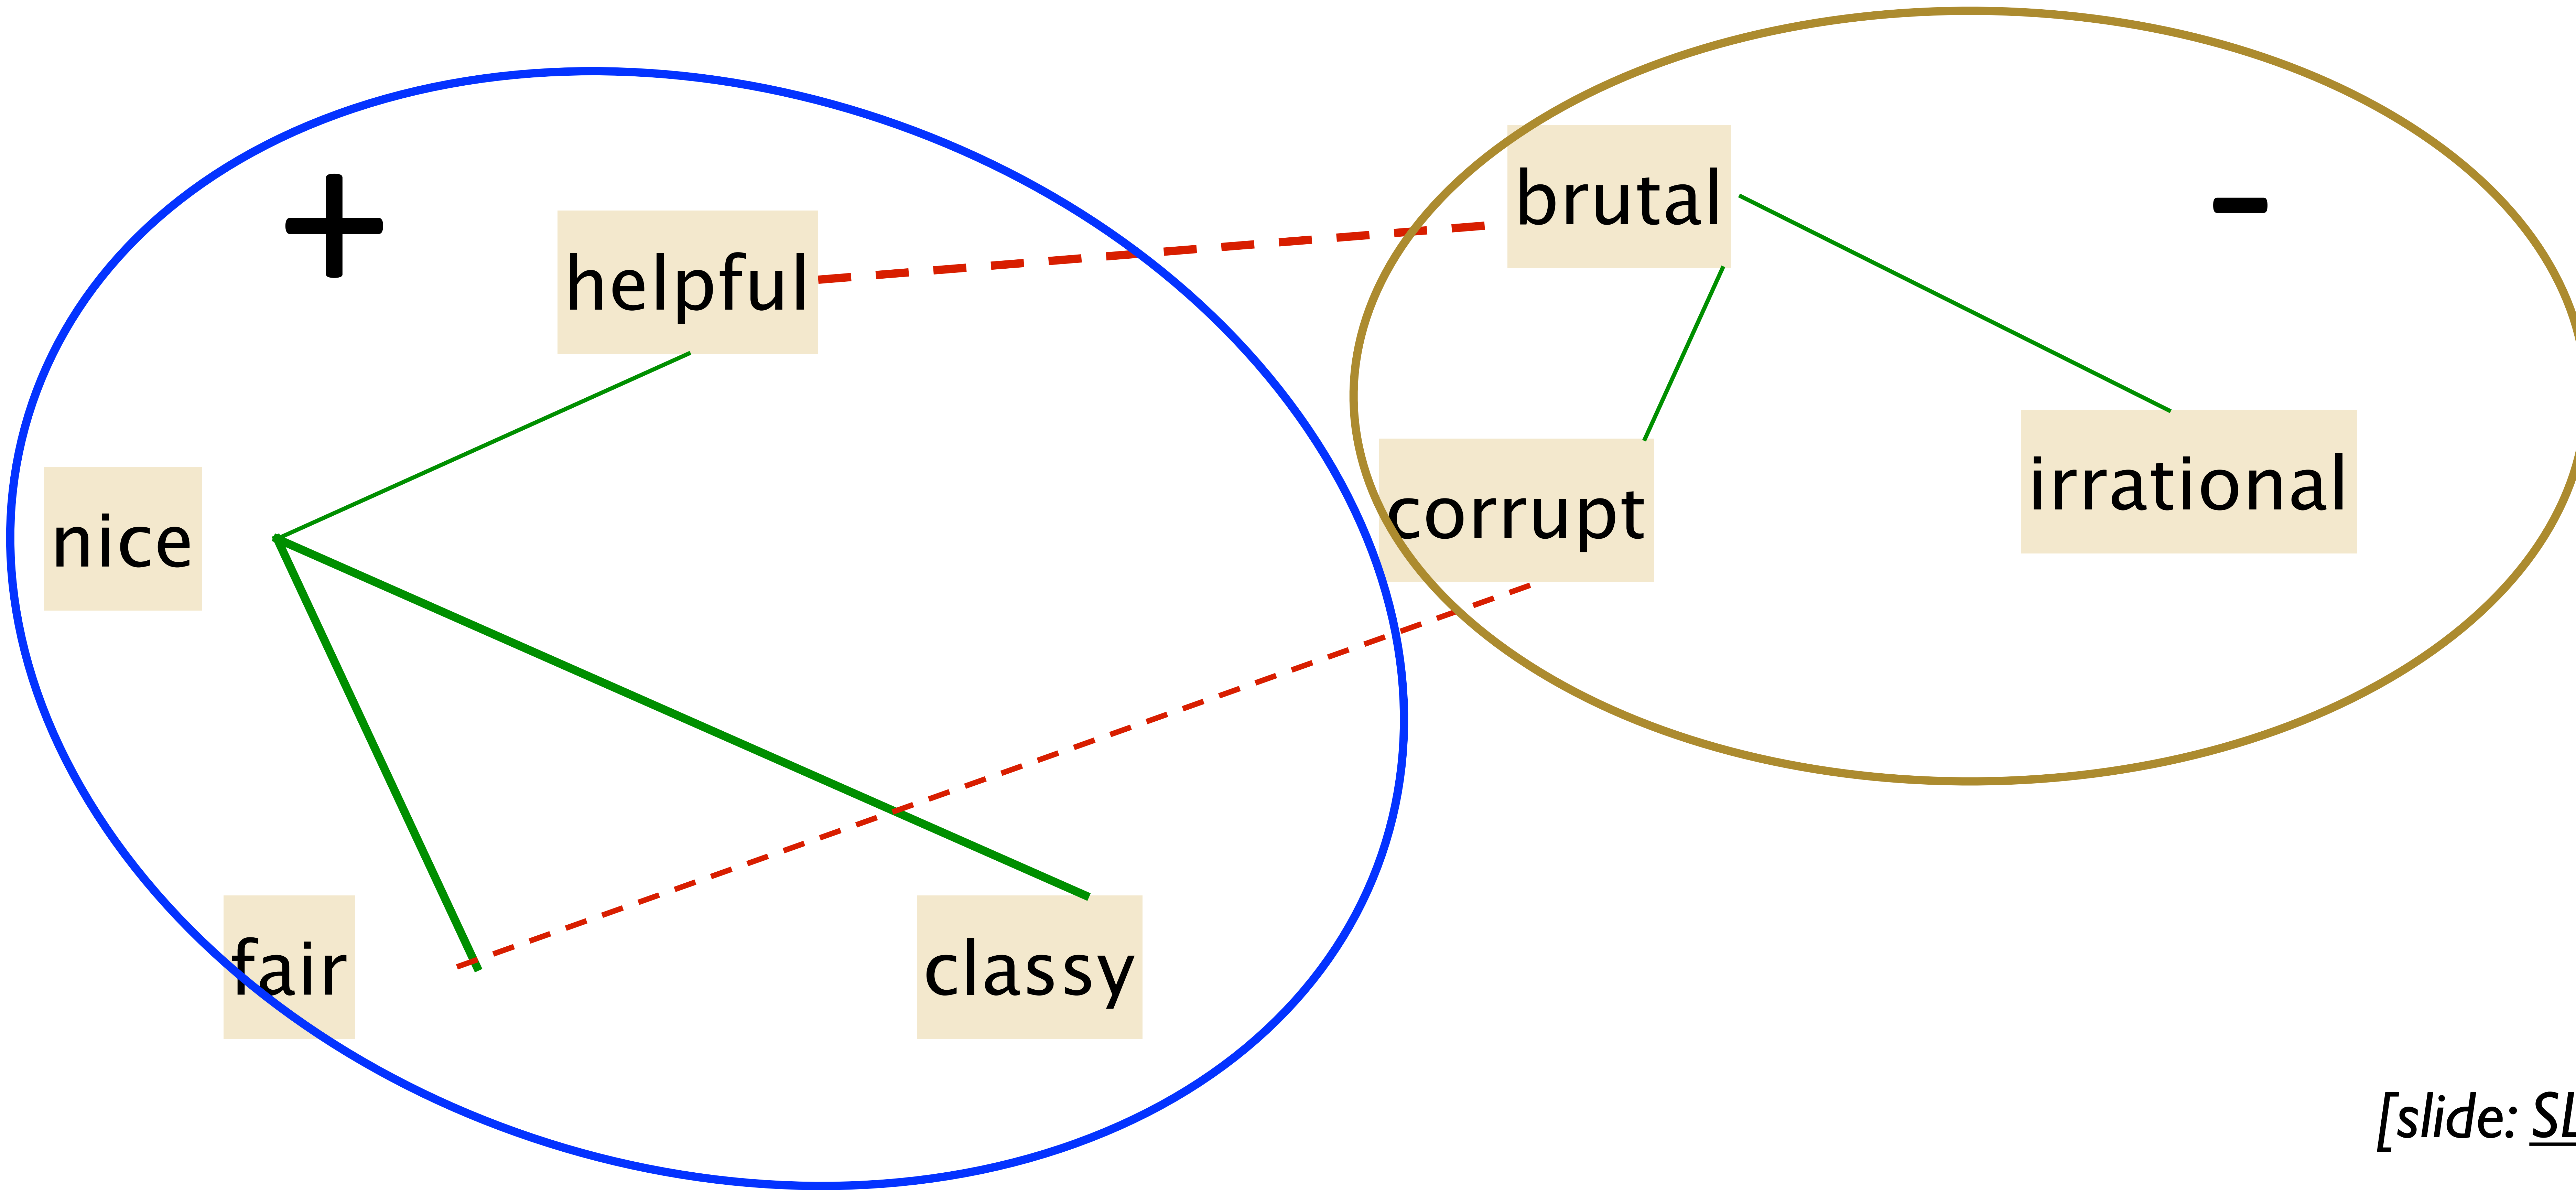

[slide: SLP3]

Output polarity lexicon

- Positive
  - bold decisive disturbing generous good honest important large mature patient peaceful positive proud sound stimulating straightforward strange talented vigorous witty…
- Negative
  - ambiguous cautious cynical evasive harmful hypocritical inefficient insecure irrational irresponsible minor outspoken pleasant reckless risky selfish tedious unsupported vulnerable wasteful…

[slide: SLP3]

Output polarity lexicon

- Positive
  - bold decisive **disturbing** generous good honest important large mature patient peaceful positive proud sound stimulating straightforward **strange** talented vigorous witty…
- Negative
  - ambiguous **cautious** cynical evasive harmful hypocritical inefficient insecure irrational irresponsible minor **outspoken pleasant** reckless risky selfish tedious unsupported vulnerable wasteful…

[slide: SLP3]

Turney Algorithm

Turney (2002): Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews

1. Extract a *phrasal lexicon* from reviews
2. Learn polarity of each phrase
3. Rate a review by the average polarity of its phrases

[slide: SLP3]

Pointwise Mutual Information

- Measure co-occurrence, but want to control for overall frequency (as opposed to raw count)
- How much more often do outcomes x and y co-occur, compared to chance?

[slide: SLP3]

Pointwise Mutual Information

- Measure co-occurrence, but want to control for overall frequency (as opposed to raw count)
- How much more often do outcomes x and y co-occur, compared to chance?

$$\mathrm{PMI}(x,y) = \log \frac{P(X=x, Y=y)}{P(X=x)P(Y=y)} = \log \frac{P(x,y)}{P(x)P(y)} = \log \frac{P(x \mid y)}{P(x)}$$

[slide: SLP3]

Pointwise Mutual Information

- Measure co-occurrence, but want to control for overall frequency (as opposed to raw count)
- How much more often do outcomes x and y co-occur, compared to chance?

$$\mathrm{PMI}(x, y) = \log \frac{P(X = x, Y = y)}{P(X = x)P(Y = y)} = \log \frac{P(x, y)}{P(x)P(y)} = \log \frac{P(x \mid y)}{P(x)}$$

- How much more often do words word1 and word2 co-occur (say, in same document), compared to chance?

[slide: SLP3]

Pointwise Mutual Information

- Measure co-occurrence, but want to control for overall frequency (as opposed to raw count)
- How much more often do outcomes x and y co-occur, compared to chance?

$$\mathrm{PMI}(x,y) = \log \frac{P(X=x, Y=y)}{P(X=x)P(Y=y)} = \log \frac{P(x,y)}{P(x)P(y)} = \log \frac{P(x \mid y)}{P(x)}$$

- How much more often do words word1 and word2 co-occur (say, in same document), compared to chance?

$$\mathrm{PMI}(word1, word2) = \log \frac{P(word1, word2)}{P(word1)P(word2)}$$

[slide: SLP3]

Pointwise Mutual Information

- Measure co-occurrence, but want to control for overall frequency (as opposed to raw count)
- How much more often do outcomes x and y co-occur, compared to chance?

$$\mathrm{PMI}(x,y) = \log \frac{P(X=x, Y=y)}{P(X=x)P(Y=y)} = \log \frac{P(x,y)}{P(x)P(y)} = \log \frac{P(x \mid y)}{P(x)}$$

- How much more often do words word1 and word2 co-occur (say, in same document), compared to chance?

$$\mathrm{PMI}(word1, word2) = \log \frac{P(word1, word2)}{P(word1)P(word2)}$$

- PMI is an easy, simple tool used a lot in NLP

[slide: SLP3]

Does phrase appear more with “poor” or “excellent”?

$$\text{Polarity}(\textit{phrase}) = \text{PMI}(\textit{phrase},\text{"excellent"}) - \text{PMI}(\textit{phrase},\text{"poor"})$$

$$= \log_2 \frac{\frac{1}{N} \textit{h}\text{its}(\textit{phrase} \text{ NEAR "excellent"})}{\frac{1}{N}\text{hits}(\textit{phrase})\frac{1}{N}\text{hits}(\text{"excellent"})} - \log_2 \frac{\frac{1}{N}\text{hits}(\textit{phrase} \text{ NEAR "poor"})}{\frac{1}{N}\text{hits}(\textit{phrase})\frac{1}{N}\text{hits}(\text{"poor"})}$$

[slide: SLP3]

Phrases from a thumbs-up review

Phrase	POS tags	Polarity
online service	JJ NN	2.8
online experience	JJ NN	2.3
direct deposit	JJ NN	1.3
local branch	JJ NN	0.42
…		
low fees	JJ NNS	0.33
true service	JJ NN	-0.73
other bank	JJ NN	-0.85
inconveniently located	JJ NN	-1.5
Average		0.32

[slide: SLP3]

Results of Turney algorithm

- 410 reviews from Epinions
  - 170 (41%) negative
  - 240 (59%) positive
- Majority class baseline: 59%
- Turney algorithm: 74%

- Phrases rather than words
- Learns domain-specific information

[slide: SLP3]

Summary

- Lexicons of connotations, not definitions: affect, polarity, etc.
  - Can be applied cross-domain
- Can be constructed by
  - Human judgments
  - Document-level supervised learning
  - Semi-supervised learning (co-occurrence)
    - Adapts a lexicon to a corpus
- Text analyzers
  - Simple: count/sum polarity scores of words in text
  - Better: also add rules/heuristics (e.g. VADER)
  - (Best?: supervised learning?)