Algorithms for NLP

Sentiment Analysis

Yulia Tsvetkov – CMU

Slides: Dan Jurafsky – Stanford

What is Sentiment Analysis?

Positive or negative movie review?

- unbelievably disappointing
- Full of zany characters and richly applied satire, and some great plot twists
- this is the greatest screwball comedy ever filmed
- It was pathetic. The worst part about it was the boxing scenes.

Google Product Search

HP Officejet 6500A Plus e-All-in-One Color Ink-jet - Fax / copier / printer / scanner

$89 online, **$100** nearby ★★★★☆ 377 reviews

September 2010 - Printer - HP - Inkjet - Office - Copier - Color - Scanner - Fax - 250 sh

Reviews

Summary - Based on 377 reviews

1 star | 2 | 3 | 4 stars | 5 stars

What people are saying

ease of use	"This was very easy to setup to four computers."
value	"Appreciate good quality at a fair price."
setup	"Overall pretty easy setup."
customer service	"I DO like honest tech support people."
size	"Pretty Paper weight."
mode	"Photos were fair on the high quality mode."
colors	"Full color prints came out with great quality."

Bing Shopping

HP Officejet 6500A E710N Multifunction Printer

Product summary   Find best price   **Customer reviews**   Specifications   Related items

$121.53 - $242.39 (14 stores)

☐ Compare

Average rating ★★★½☆ (144)

★★★★★	(55)
★★★★☆	(54)
★★★☆☆	(10)
★★☆☆☆	(6)
★☆☆☆☆	(23)
☆☆☆☆☆	(0)

Most mentioned

Performance	(57)
Ease of Use	(43)
Print Speed	(39)
Connectivity	(31)

More ▾

Show reviews by source

- Best Buy (140)
- CNET (5)
- Amazon.com (3)

Twitter sentiment versus Gallup Poll of Consumer Confidence

Brendan O'Connor, Ramnath Balasubramanyan, Bryan R. Routledge, and Noah A. Smith. 2010.
From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In ICWSM-2010

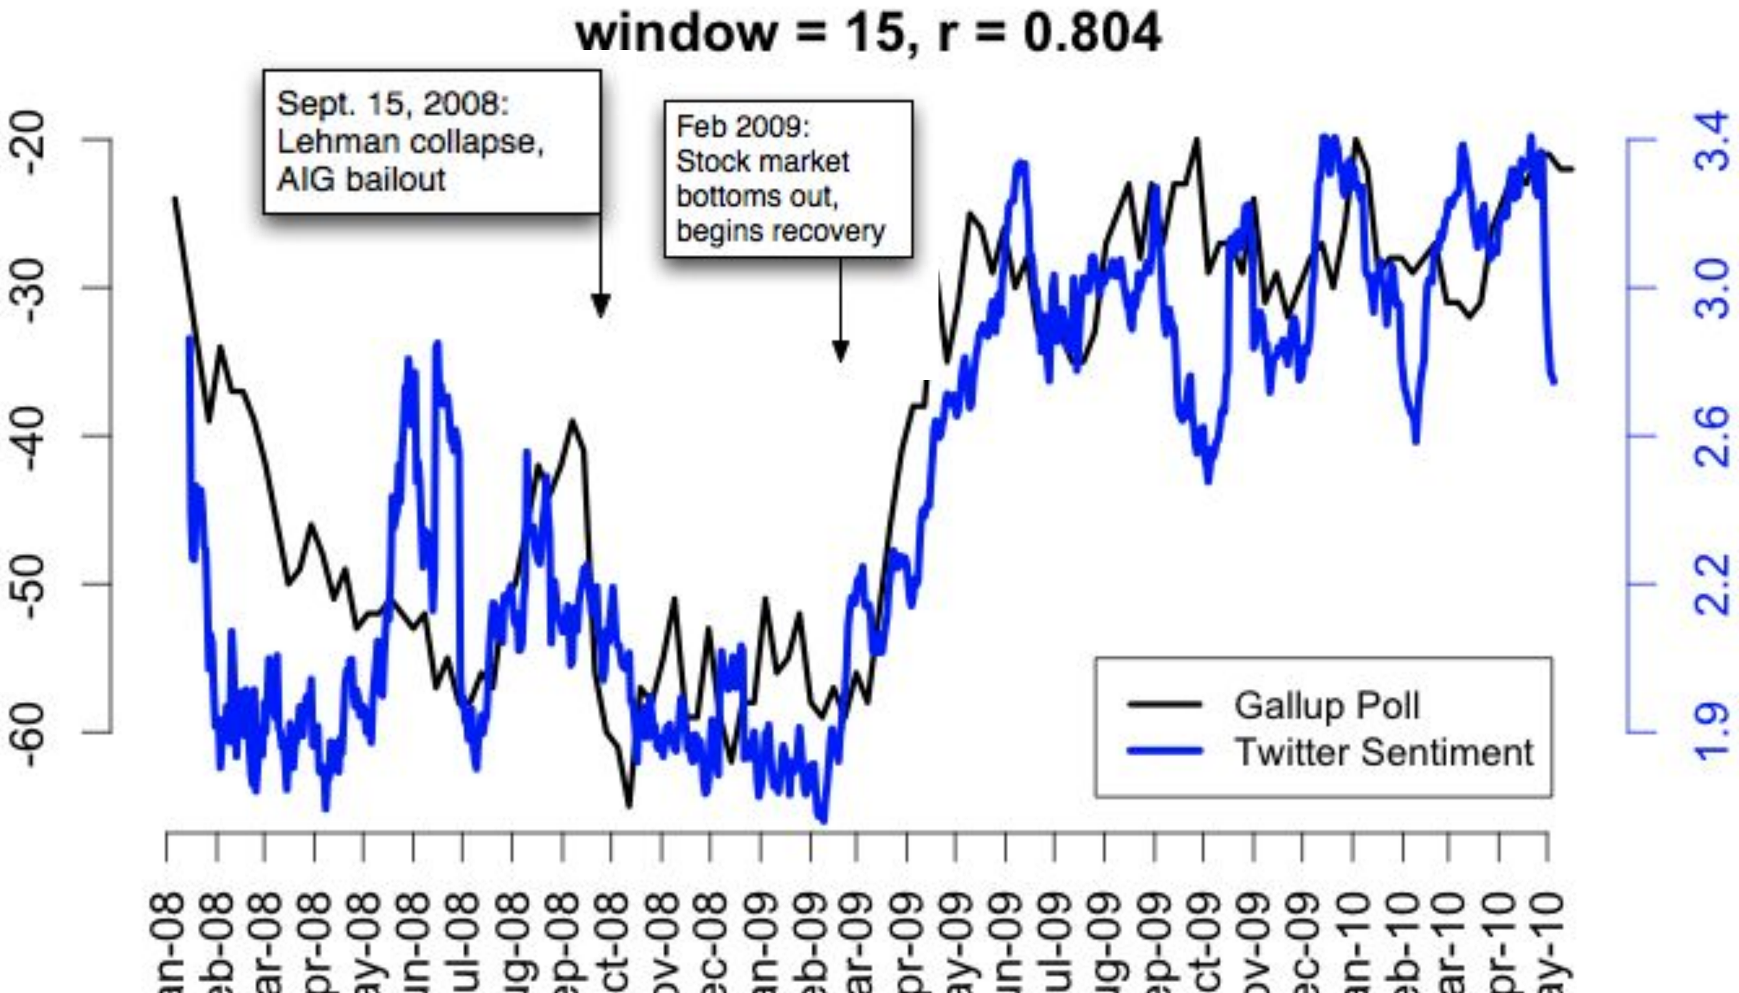

Twitter sentiment

Johan Bollen, Huina Mao, Xiaojun Zeng. 2011. [Twitter mood predicts the stock market,](#) Journal of Computational Science 2:1, 1-8. 10.1016/j.jocs.2010.12.007.

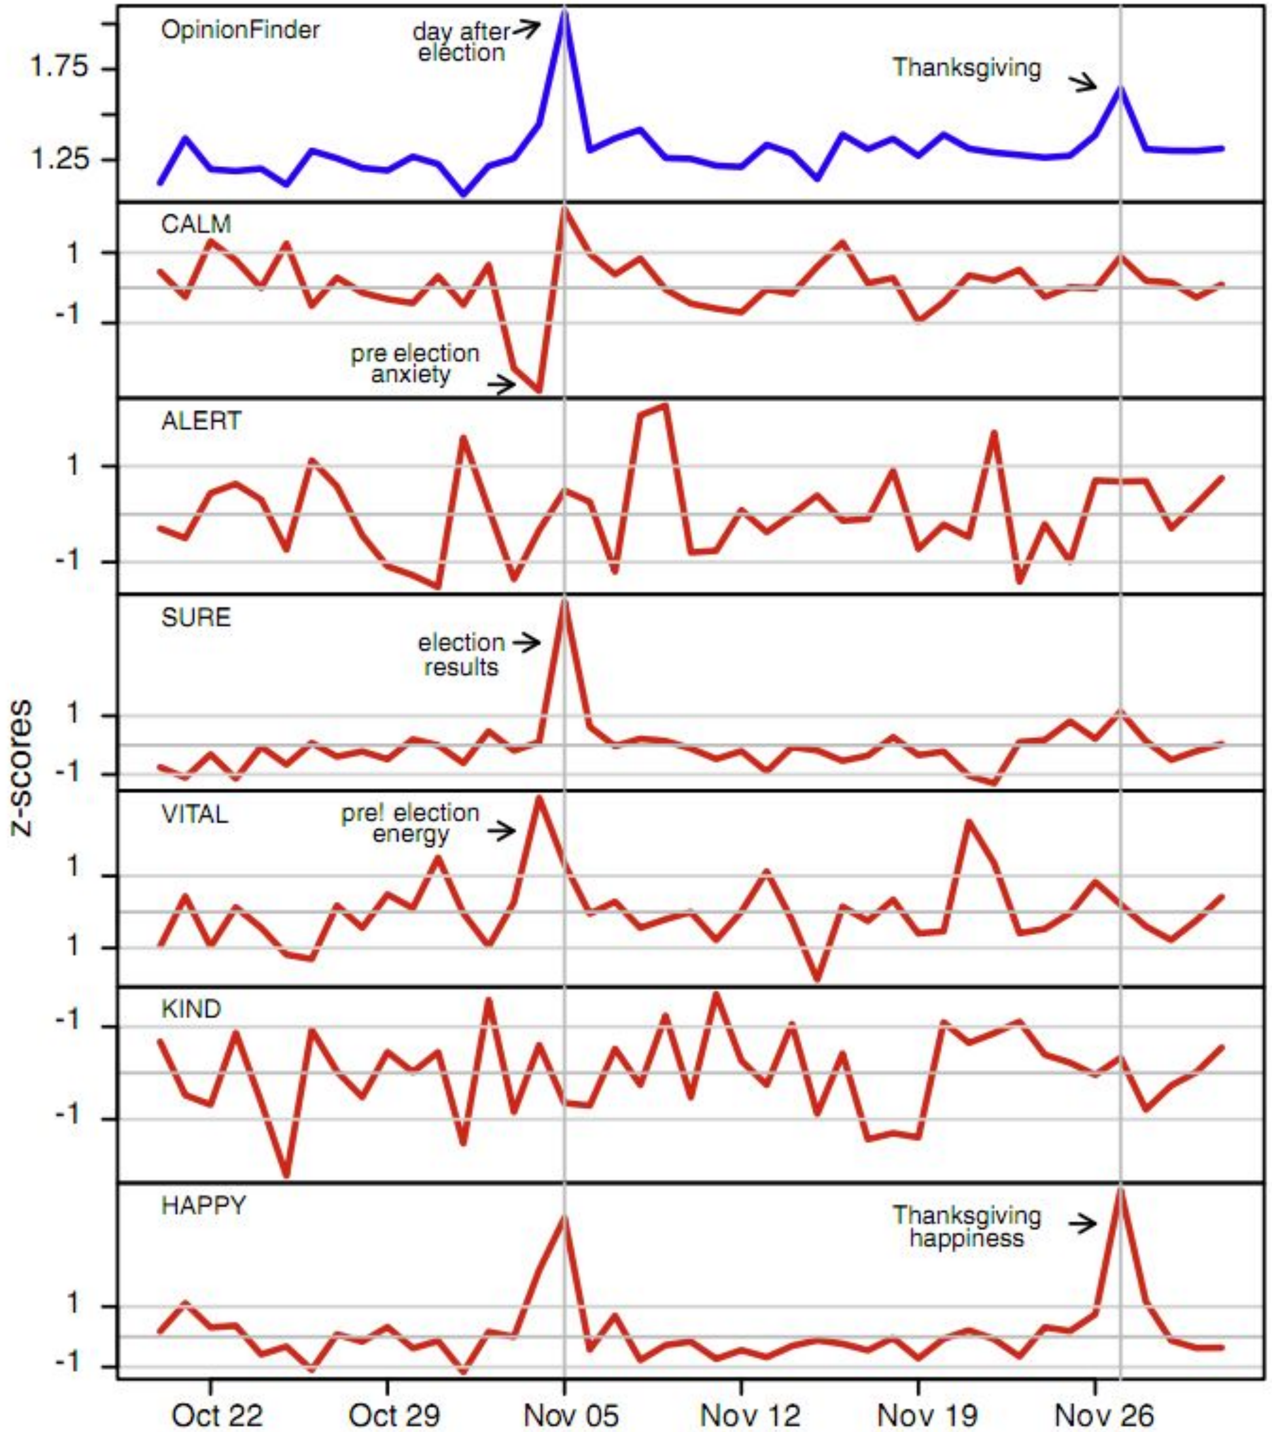

Bollen et al. (2011)

- CALM predicts DJIA 3 days later
- At least one current hedge fund uses this algorithm

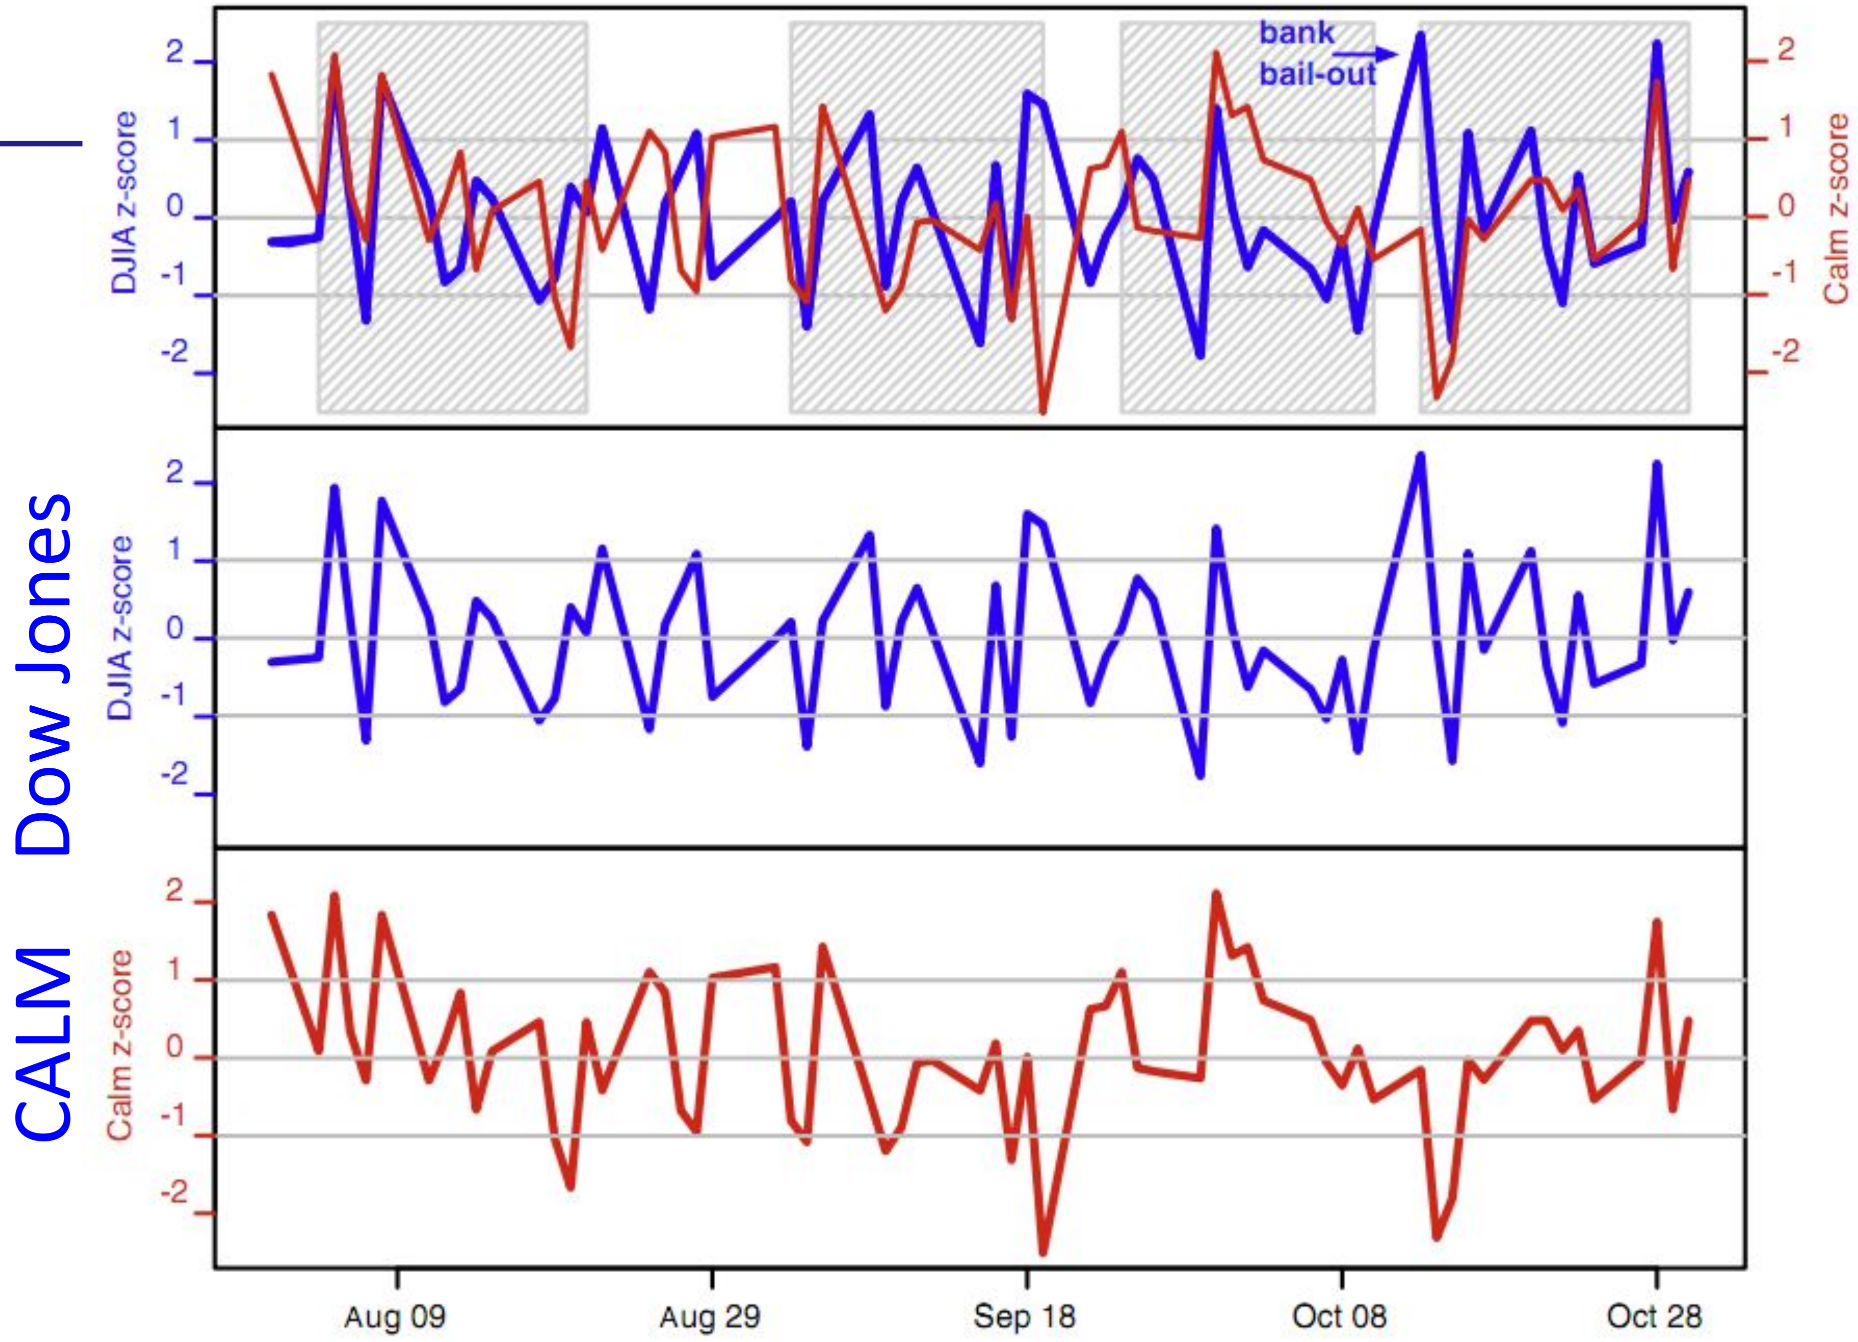

Target Sentiment on Twitter

- [Twitter Sentiment App]
- Alec Go, Richa Bhayani, Lei Huang. 2009. Twitter Sentiment Classification using Distant Supervision

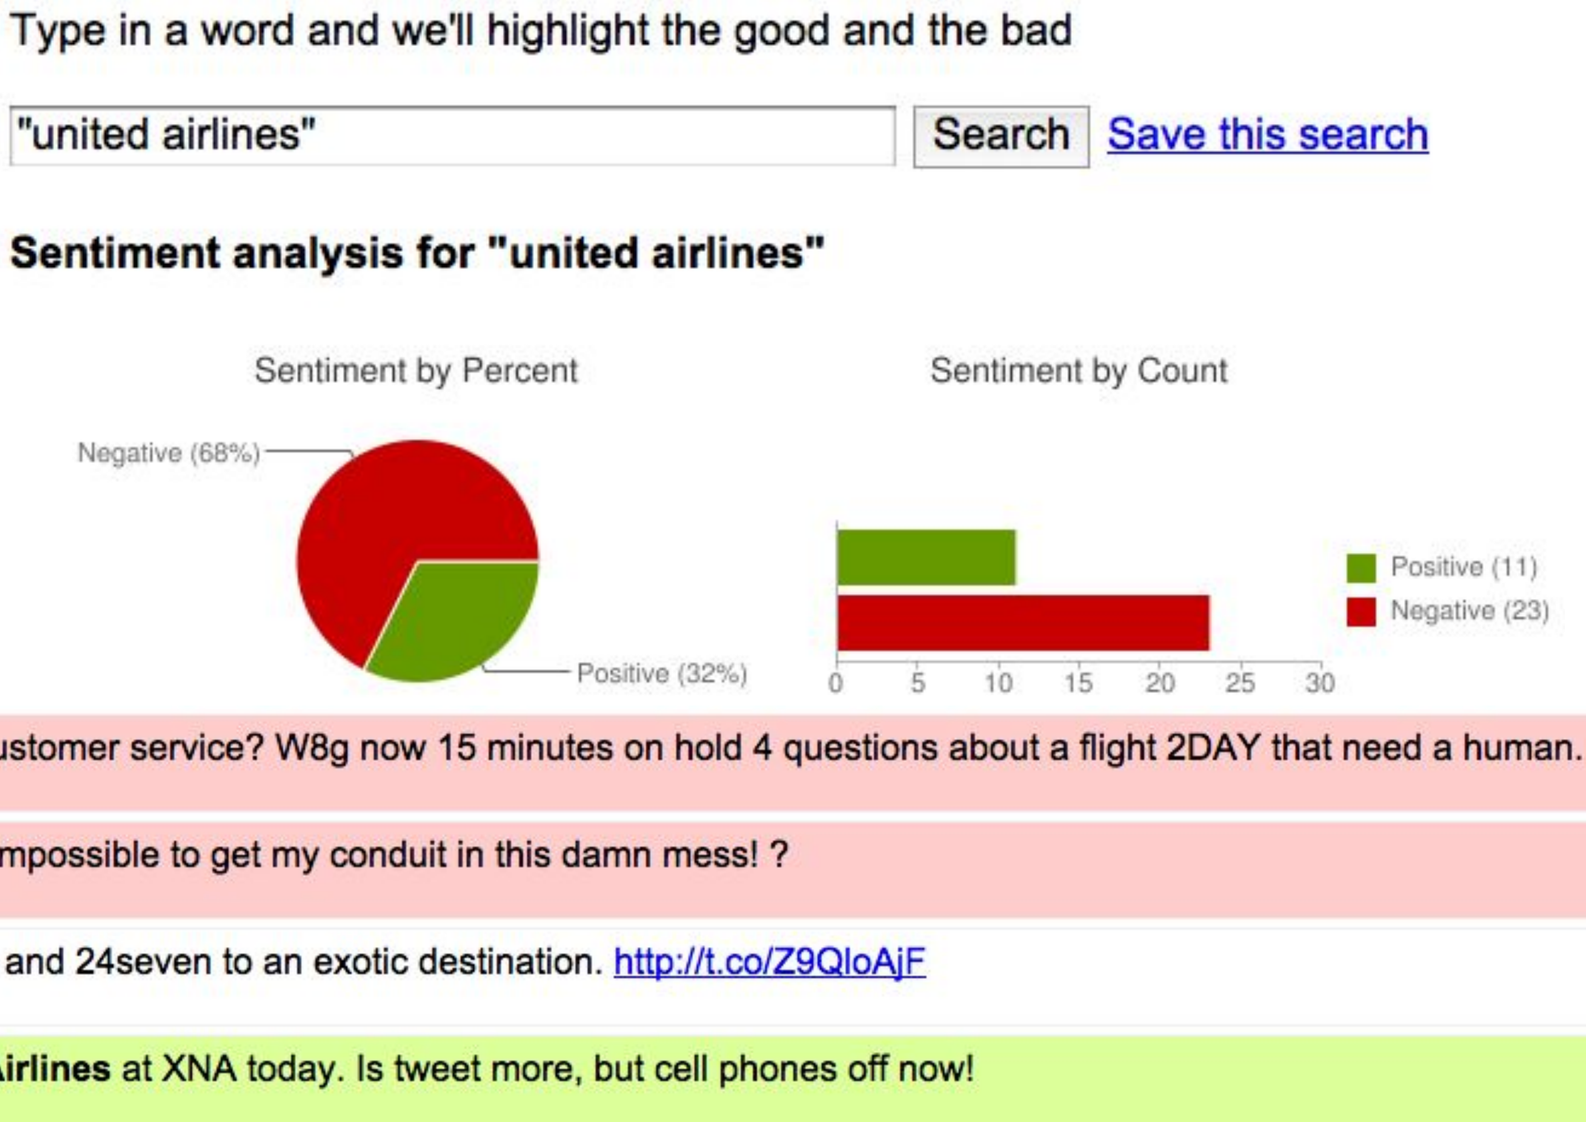

Sentiment analysis has many other names

- Opinion extraction
- Opinion mining
- Sentiment mining
- Subjectivity analysis

Why sentiment analysis?

- *Movie*: is this review positive or negative?
- *Products*: what do people think about the new iPhone?
- *Public sentiment*: how is consumer confidence? Is despair increasing?
- *Politics*: what do people think about this candidate or issue?
- *Prediction*: predict election outcomes or market trends from sentiment

Why sentiment analysis?

+ an interesting use-case for modeling natural language understanding

- sentiment
- emotion, mood, attitude, personality
- negation
- metaphor, non-literal language
- sarcasm

Scherer Typology of Affective States

- **Emotion**: brief organically synchronized … evaluation of a major event
  - *angry, sad, joyful, fearful, ashamed, proud, elated*
- **Mood**: diffuse non-caused low-intensity long-duration change in subjective feeling
  - *cheerful, gloomy, irritable, listless, depressed, buoyant*
- **Interpersonal stances**: affective stance toward another person in a specific interaction
  - *friendly, flirtatious, distant, cold, warm, supportive, contemptuous*
- **Attitudes**: enduring, affectively colored beliefs, dispositions towards objects or persons
  - *liking, loving, hating, valuing, desiring*
- **Personality traits**: stable personality dispositions and typical behavior tendencies
  - *nervous, anxious, reckless, morose, hostile, jealous*

Scherer Typology of Affective States

- **Emotion**: brief organically synchronized … evaluation of a major event
  - *angry, sad, joyful, fearful, ashamed, proud, elated*
- **Mood**: **diffuse** non-caused low-intensity long-duration change in subjective feeling
  - *cheerful, gloomy, irritable, listless, depressed, buoyant*
- **Interpersonal stances**: affective stance toward another person in a specific interaction
  - *friendly, flirtatious, distant, cold, warm, supportive, contemptuous*
- **Attitudes**: enduring, affectively colored beliefs, dispositions towards objects or persons
  - *liking, loving, hating, valuing, desiring*
- **Personality traits:** stable personality dispositions and typical behavior tendencies
  - *nervous, anxious, reckless, morose, hostile, jealous*

Sentiment Analysis

- Sentiment analysis is the detection of **attitudes**

  "enduring, affectively colored beliefs, dispositions towards objects or persons"

1. **Holder (source)** of attitude
2. **Target (aspect)** of attitude
3. **Type** of attitude
   - From a set of types
     - *Like, love, hate, value, desire,* etc.
   - Or (more commonly) simple weighted **polarity**:
     - *positive, negative, neutral,* together with *strength*
4. **Text** containing the attitude
   - Sentence or entire document

Sentiment Analysis

- Simplest task:
  - Is the attitude of this text positive or negative?
- More complex:
  - Rank the attitude of this text from 1 to 5
- Advanced:
  - Detect the target, source, or complex attitude types

Sentiment Analysis:
A Baseline Algorithm

Sentiment Classification in Movie Reviews

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.
Bo Pang and Lillian Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. ACL, 271-278

- Polarity detection:
  - Is an IMDB movie review positive or negative?
- Data: *Polarity Data 2.0:*
  - http://www.cs.cornell.edu/people/pabo/movie-review-data

IMDB data in the Pang and Lee database

✓

when _star wars_ came out some twenty years ago , the image of traveling throughout the stars has become a commonplace image . [. . .]

when han solo goes light speed , the stars change to bright lines , going towards the viewer in lines that converge at an invisible point .

cool .

october sky offers a much simpler image–that of a single white dot , traveling horizontally across the night sky . [. . .]

✗

" snake eyes " is the most aggravating kind of movie : the kind that shows so much potential then becomes unbelievably disappointing .

it's not just because this is a brian depalma film , and since he's a great director and one who's films are always greeted with at least some fanfare .

and it's not even because this was a film starring nicolas cage and since he gives a brauvara performance , this film is hardly worth his talents .

Baseline Algorithm (adapted from Pang and Lee)

- Tokenization
- Feature Extraction
- Classification using different classifiers
  - Naïve Bayes
  - MaxEnt
  - SVM

Sentiment Tokenization Issues

- Deal with HTML and XML markup
- Twitter mark-up (names, hash tags)
- Capitalization (preserve for
  words in all caps)
- Phone numbers, dates
- Emoticons
- Useful code:
  - [Christopher Potts sentiment tokenizer]()
  - [Brendan O'Connor twitter tokenizer]()

Potts emoticons

```
[<>]?                            # optional hat/brow
[:;=8]                           # eyes
[\-o\*\']?                       # optional nose
[\)\]\(\[dDpP/\:\}\{@\|\\]       # mouth
|                                #### reverse orientation
[\)\]\(\[dDpP/\:\}\{@\|\\]       # mouth
[\-o\*\']?                       # optional nose
[:;=8]                           # eyes
[<>]?                            # optional hat/brow
```

Extracting Features for Sentiment Classification

- How to handle negation
  - `I` **`didn't`** `like this movie`

  vs
  - `I really like this movie`
- Which words to use?
  - Only adjectives
  - All words
    - All words turns out to work better, at least on this data

Negation

Das, Sanjiv and Mike Chen. 2001. Yahoo! for Amazon: Extracting market sentiment from stock message boards. In Proceedings of the Asia Pacific Finance Association Annual Conference (APFA).
Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

Add NOT_ to every word between negation and following punctuation:

didn't like this movie , but I

didn't NOT_like NOT_this NOT_movie but I

Cross-Validation

- Break up data into 10 folds
  - (Equal positive and negative inside each fold?)
- For each fold
  - Choose the fold as a temporary test set
  - Train on 9 folds, compute performance on the test fold
- Report average performance of the 10 runs

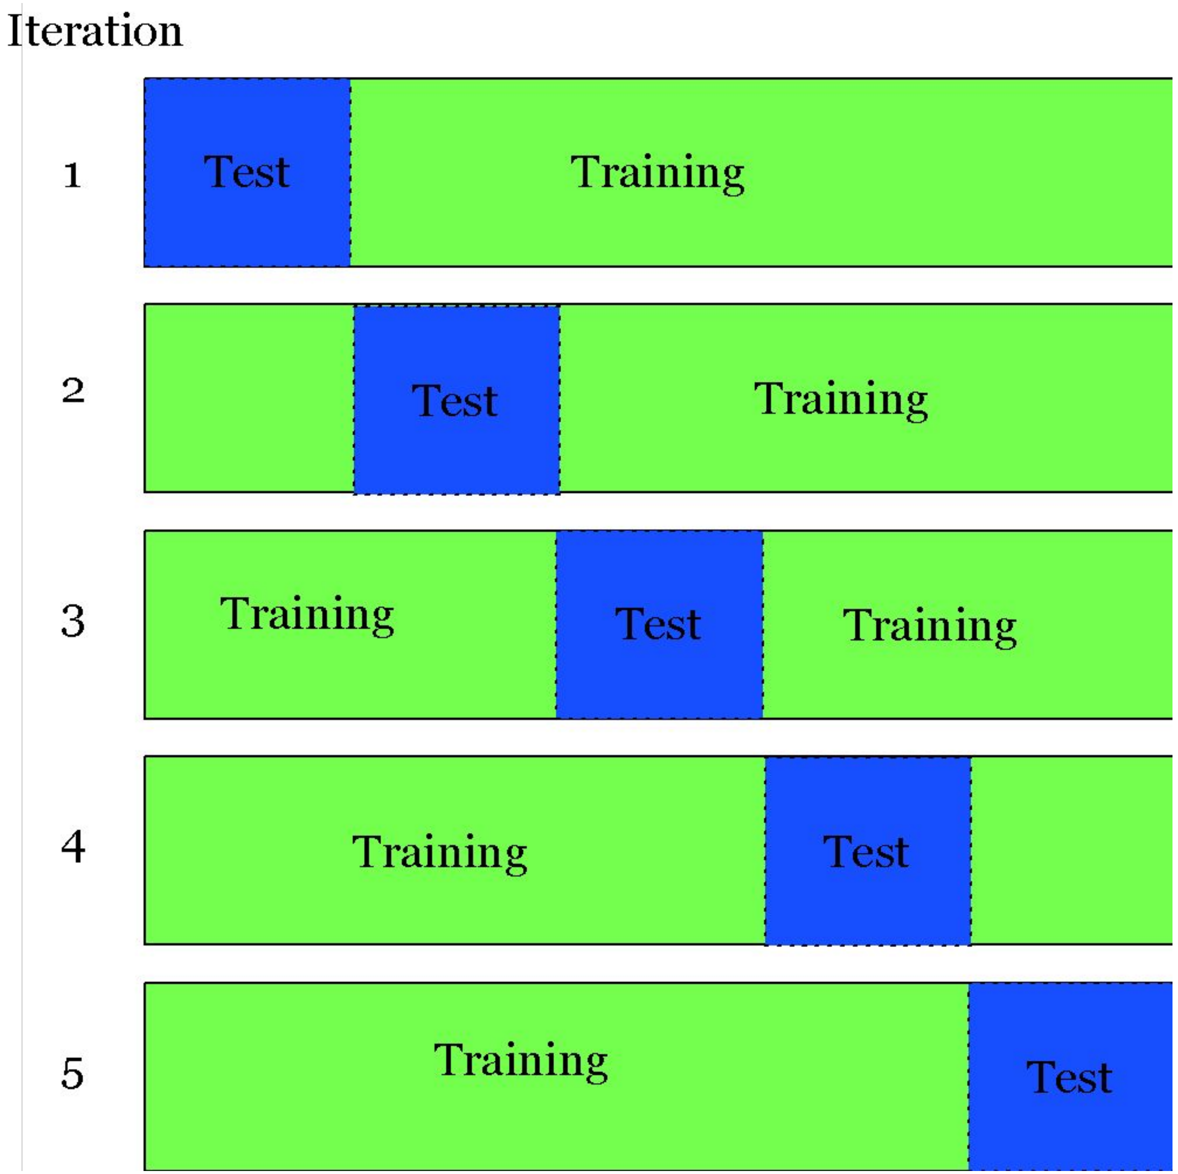

	Features	# of features	frequency or presence?	NB	ME	SVM
(1)	unigrams	16165	freq.	**78.7**	N/A	72.8
(2)	unigrams	"	pres.	81.0	80.4	**82.9**
(3)	unigrams+bigrams	32330	pres.	80.6	80.8	**82.7**
(4)	bigrams	16165	pres.	77.3	**77.4**	77.1
(5)	unigrams+POS	16695	pres.	81.5	80.4	**81.9**
(6)	adjectives	2633	pres.	77.0	**77.7**	75.1
(7)	top 2633 unigrams	2633	pres.	80.3	81.0	**81.4**
(8)	unigrams+position	22430	pres.	81.0	80.1	**81.6**

Figure 3: Average three-fold cross-validation accuracies, in percent. Boldface: best performance for a given setting (row). Recall that our baseline results ranged from 50% to 69%.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

Problems: What makes reviews hard to classify?

- Subtlety:
  - Perfume review in *Perfumes: the Guide*:
    - “If you are reading this because it is your darling fragrance, please wear it at home exclusively, and tape the windows shut.”
  - Dorothy Parker on Katherine Hepburn
    - “She runs the gamut of emotions from A to B”

- “This film should be brilliant. It sounds like a great plot, the actors are first grade, and the supporting cast is good as well, and Stallone is attempting to deliver a good performance. However, it **can’t hold up**.”
- Well as usual Keanu Reeves is nothing special, but surprisingly, the very talented Laurence Fishbourne is **not so good** either, I was surprised.

Computing with Affective Lexicons

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

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

Connotation in the lexicon

- Words have connotation as well as sense
- Can we build lexical resources that represent these connotations?
- And use them in these computational tasks?

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

Affective Lexicons

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**

The General Inquirer

Philip J. Stone, Dexter C Dunphy, Marshall S. Smith, Daniel M. Ogilvie. 1966. The General Inquirer: A Computer Approach to Content Analysis. MIT Press

- Home page: http://www.wjh.harvard.edu/~inquirer
- List of Categories: http://www.wjh.harvard.edu/~inquirer/homecat.htm
- Spreadsheet: http://www.wjh.harvard.edu/~inquirer/inquirerbasic.xls
- Categories:
  - Positiv (1915 words) and Negativ (2291 words)
  - Strong vs Weak, Active vs Passive, Overstated versus Understated
  - Pleasure, Pain, Virtue, Vice, Motivation, Cognitive Orientation, etc
- Free for Research Use

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

- 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*)
- $30 or $90 fee

Sample LIWC Features

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

Feature	**Type**	**Example**
Anger words	LIWC	hate, kill, pissed
Metaphysical issues	LIWC	God, heaven, coffin
Physical state/function	LIWC	ache, breast, sleep
Inclusive words	LIWC	with, and, include
Social processes	LIWC	talk, us, friend
Family members	LIWC	mom, brother, cousin
Past tense verbs	LIWC	walked, were, had
References to friends	LIWC	pal, buddy, coworker

MPQA Subjectivity Cues Lexicon

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann (2005). Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. Proc. of HLT-EMNLP-2005.

Riloff and Wiebe (2003). Learning extraction patterns for subjective expressions. EMNLP-2003.

- Home page: http://www.cs.pitt.edu/mpqa/subj_lexicon.html
- 6885 words from 8221 lemmas
  - 2718 positive
  - 4912 negative
- Each word annotated for intensity (strong, weak)
- GNU GPL

Bing Liu Opinion Lexicon

Minqing Hu and Bing Liu. Mining and Summarizing Customer Reviews. ACM SIGKDD-2004.

- Bing Liu's Page on Opinion Mining
- http://www.cs.uic.edu/~liub/FBS/opinion-lexicon-English.rar

- 6786 words
  - 2006 positive
  - 4783 negative

SentiWordNet

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010 SENTIWORDNET 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. LREC-2010

- Home page: http://sentiwordnet.isti.cnr.it/
- All WordNet synsets automatically annotated for degrees of positivity, negativity, and neutrality/objectiveness
- [estimable(J,3)] "may be computed or estimated"

  `Pos 0 Neg 0 Obj 1`

- [estimable(J,1)] "deserving of respect or high regard"

  `Pos .75 Neg 0 Obj .25`

Other Affective Lexicons

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

Two families of theories of emotion

- Atomic basic emotions
  - A finite list of 6 or 8, from which others are generated
- Dimensions of emotion
  - Valence (positive negative)
  - Arousal (strong, weak)
  - Control

Ekman's 6 basic emotions:

Surprise, happiness, anger, fear, disgust, sadness

Valence/Arousal Dimensions

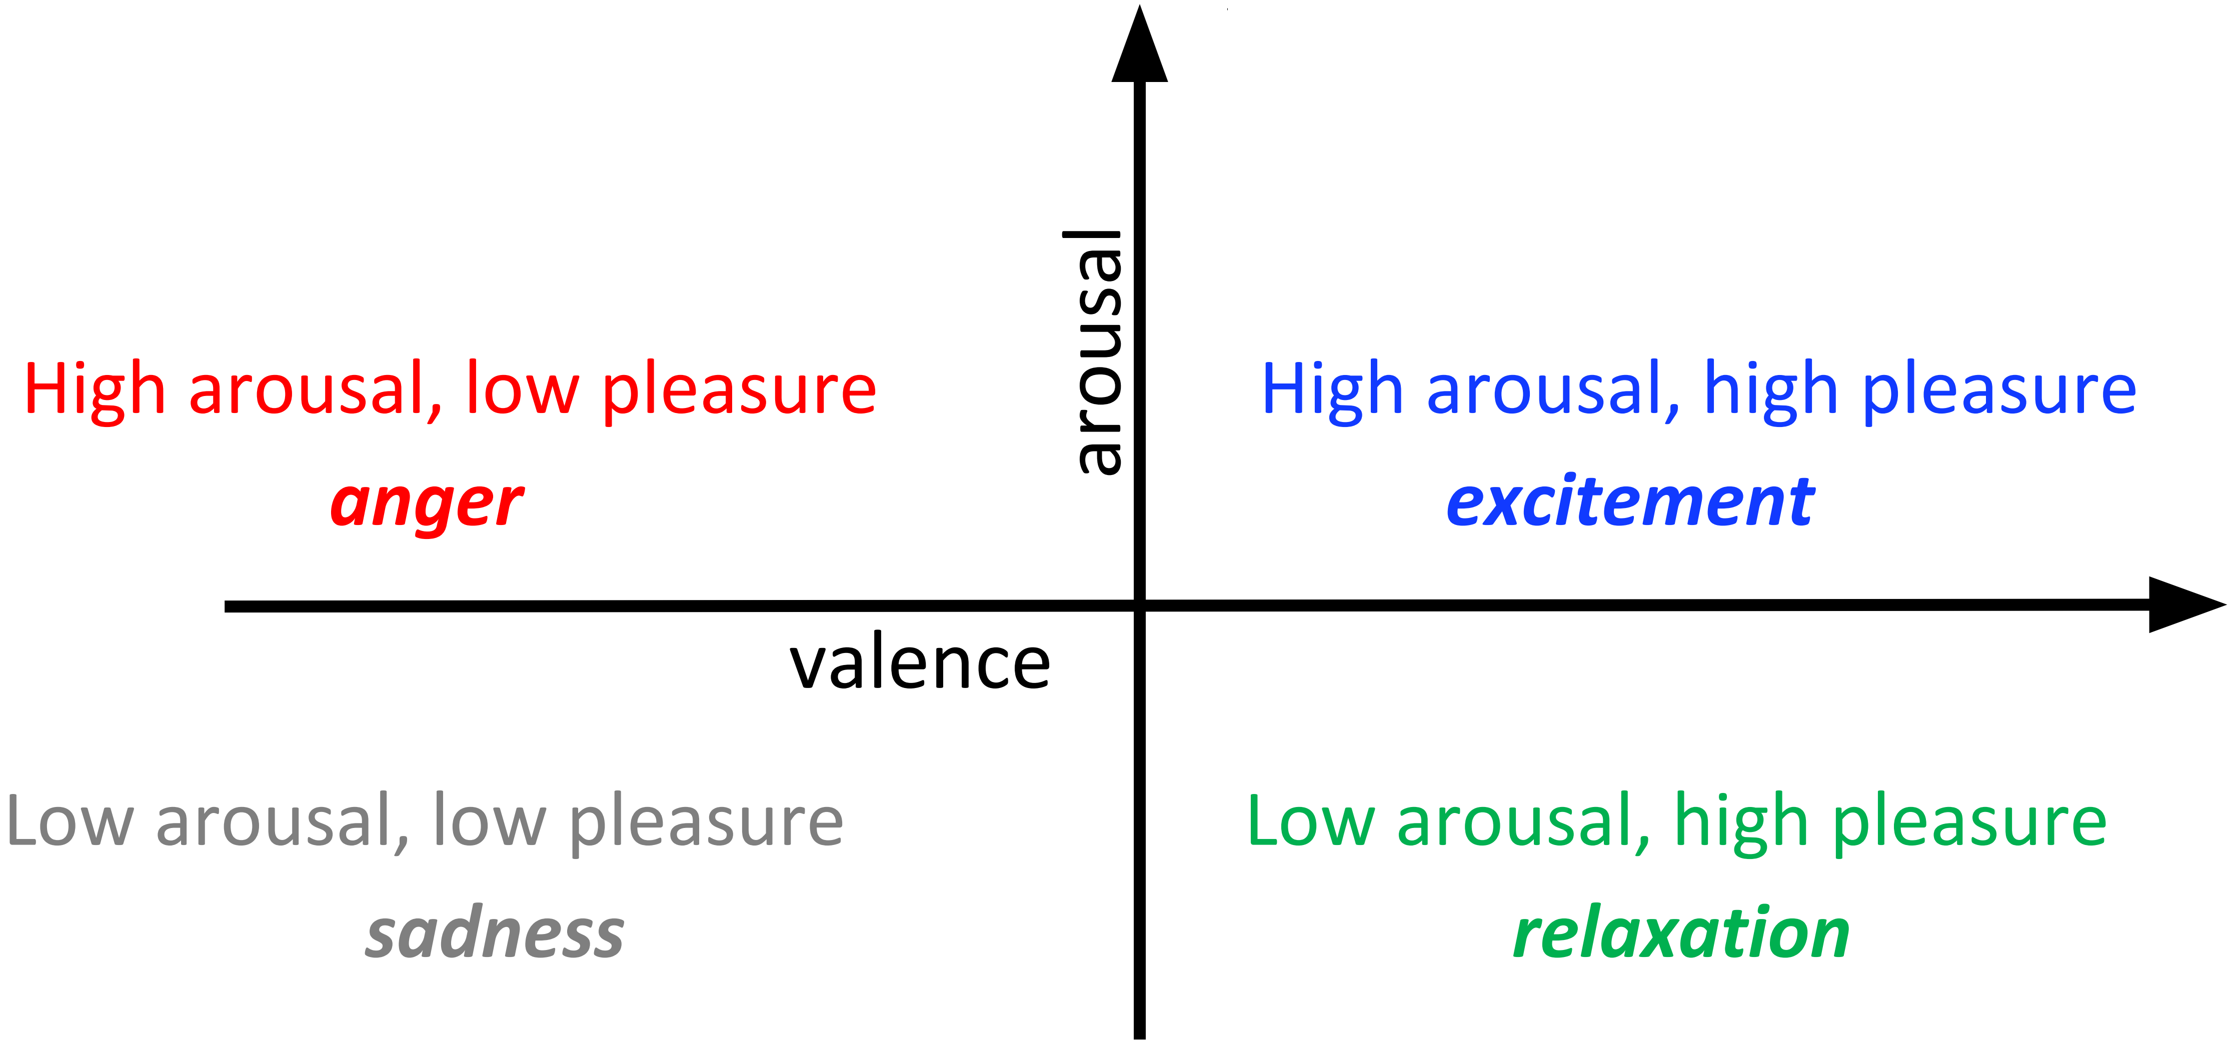

Atomic units vs. Dimensions

Distinctive

- Emotions are units.
- Limited number of basic emotions.
- Basic emotions are innate and universal

Dimensional

- Emotions are dimensions.
- Limited # of labels but unlimited number of emotions.
- Emotions are culturally learned.

Adapted from Julia Braverman

One emotion lexicon from each paradigm!

1. 8 basic emotions:
   - NRC Word-Emotion Association Lexicon (Mohammad and Turney 2011)
2. Dimensions of valence/arousal/dominance
   - Warriner, A. B., **Kuperman**, V., and Brysbaert, M. (2013)

- Both built using Amazon Mechanical Turk

Plutchick's wheel of emotion

- 8 basic emotions
- in four opposing pairs:
  - joy–sadness
  - anger–fear
  - trust–disgust
  - anticipation–surprise

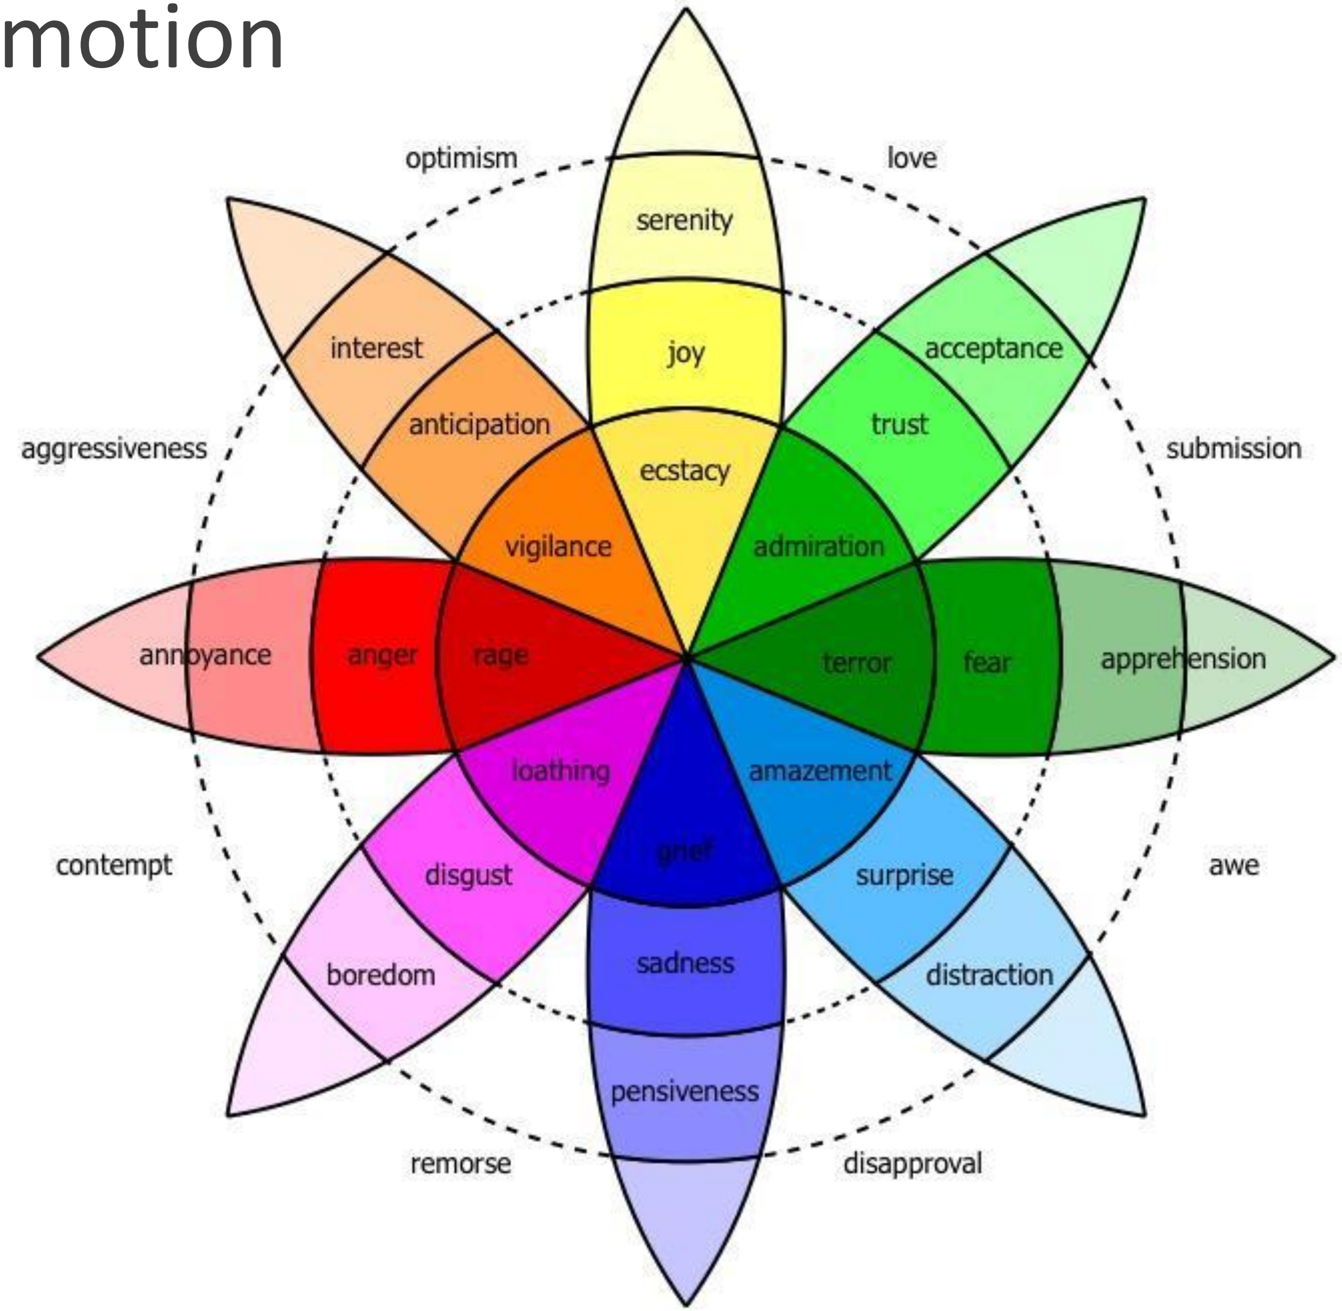

NRC Word-Emotion Association Lexicon

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
```

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**

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

...

Lexicon of valence, arousal, and dominance

- Warriner, A. B., **Kuperman**, V., and Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods 45, 1191-1207.*
- Supplementary data: This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License.

- **Ratings for 14,000 words for emotional dimensions:**
  - **valence** (the pleasantness of the stimulus)
  - **arousal** (the intensity of emotion provoked by the stimulus)
  - **dominance** (the degree of control exerted by the stimulus)

Lexicon of valence, arousal, and dominance

- **valence** (the pleasantness of the stimulus)
  9: happy, pleased, satisfied, contented, hopeful
  1: unhappy, annoyed, unsatisfied, melancholic, despaired, or bored
- **arousal** (the intensity of emotion provoked by the stimulus)
  9: stimulated, excited, frenzied, jittery, wide-awake, or aroused
  1: relaxed, calm, sluggish, dull, sleepy, or unaroused;
- **dominance** (the degree of control exerted by the stimulus)
  9: in control, influential, important, dominant, autonomous, or controlling
  1: controlled, influenced, cared-for, awed, submissive, or guided
- Again produced by AMT

Lexicon of valence, arousal, and dominance: Examples

Valence		Arousal		Dominance	
vacation	8.53	rampage	7.56	self	7.74
happy	8.47	tornado	7.45	incredible	7.74
whistle	5.7	zucchini	4.18	skillet	5.33
conscious	5.53	dressy	4.15	concur	5.29
torture	1.4	dull	1.67	earthquake	2.14

Algorithms for Learning Affective Lexicons

Semi-supervised learning of lexicons

- Use a small amount of information
  - A few labeled examples
  - A few hand-built patterns
- To bootstrap a lexicon

Hatzivassiloglou and McKeown 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

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…

Hatzivassiloglou & McKeown 1997
Step 2

- Expand seed set to conjoined adjectives

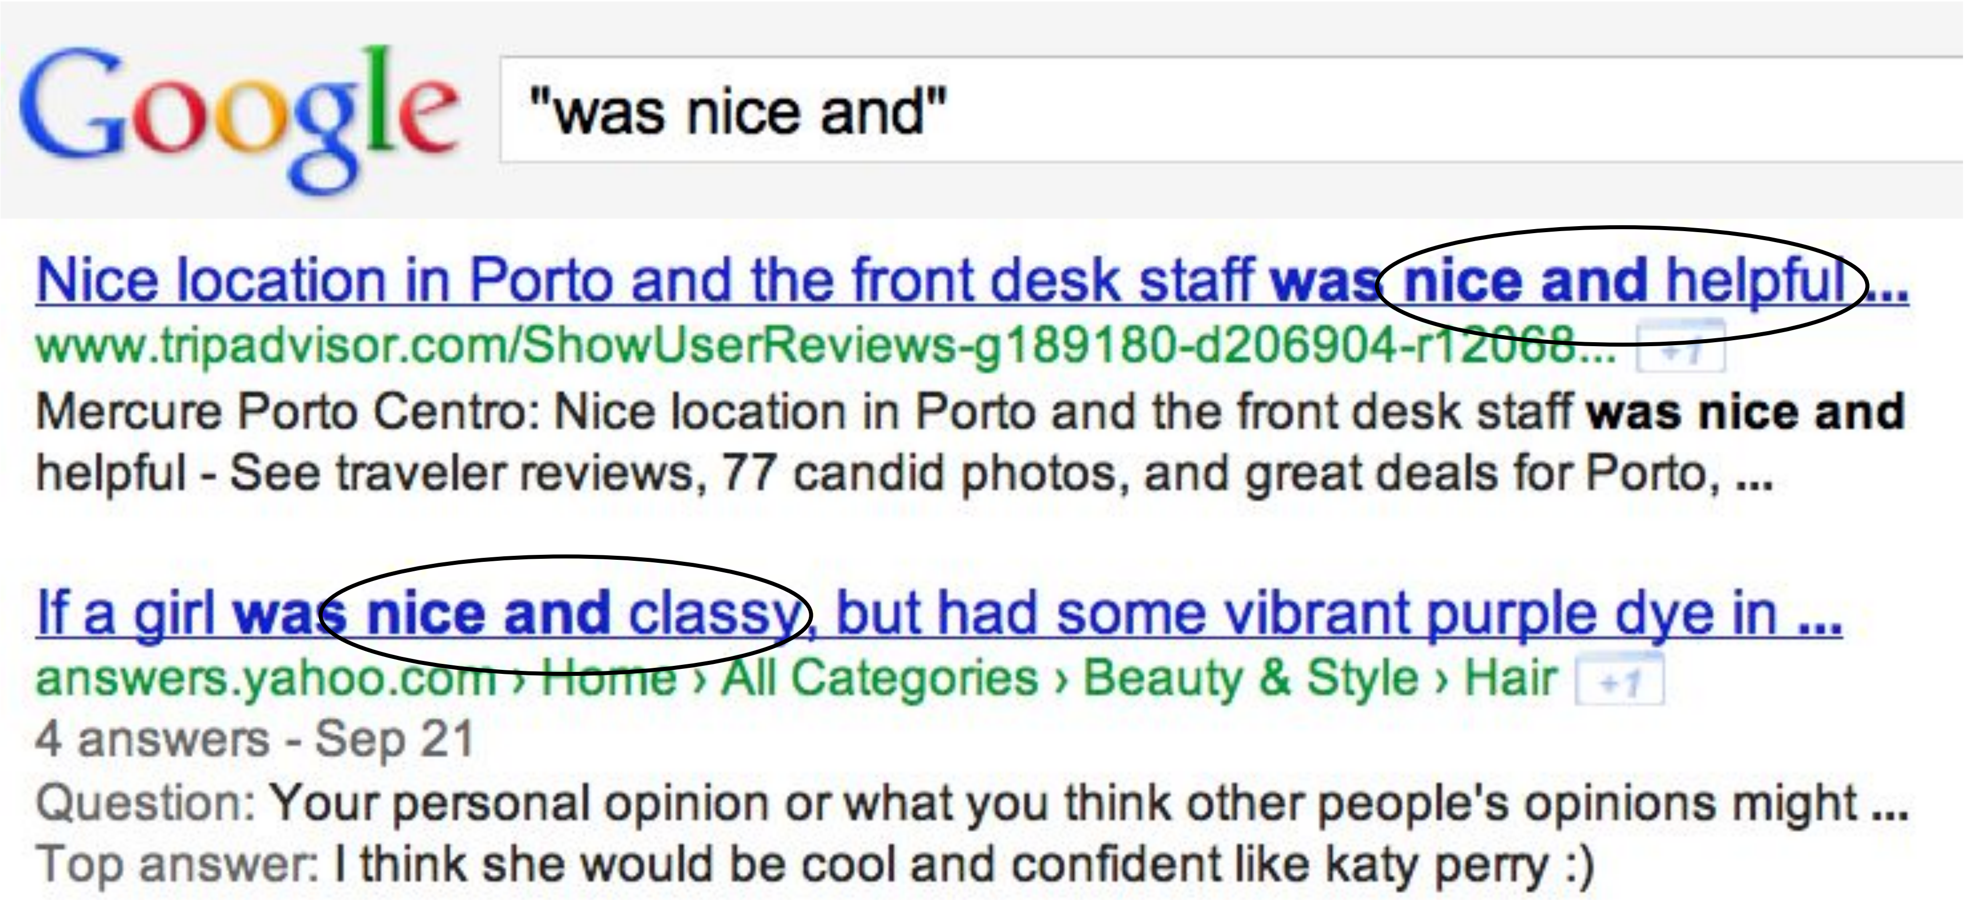
Google
"was nice and"

Nice location in Porto and the front desk staff **was nice and** helpful ...
www.tripadvisor.com/ShowUserReviews-g189180-d206904-r12068...
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
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 :)

nice, helpful

nice, classy

Hatzivassiloglou & McKeown 1997
Step 3

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

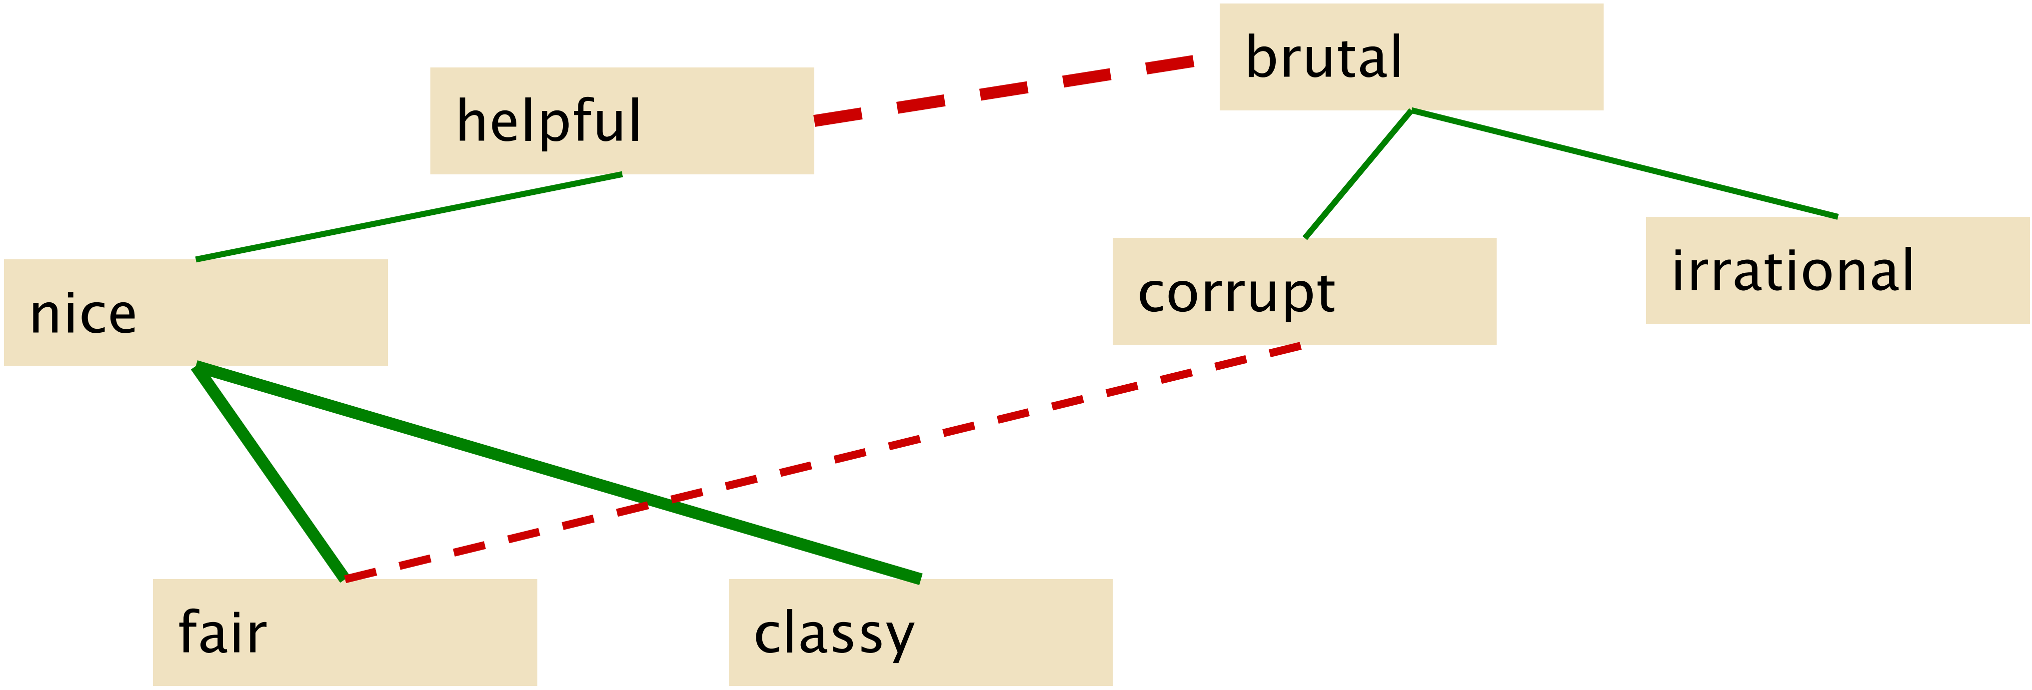

Hatzivassiloglou & McKeown 1997
Step 4

- Clustering for partitioning the graph into two

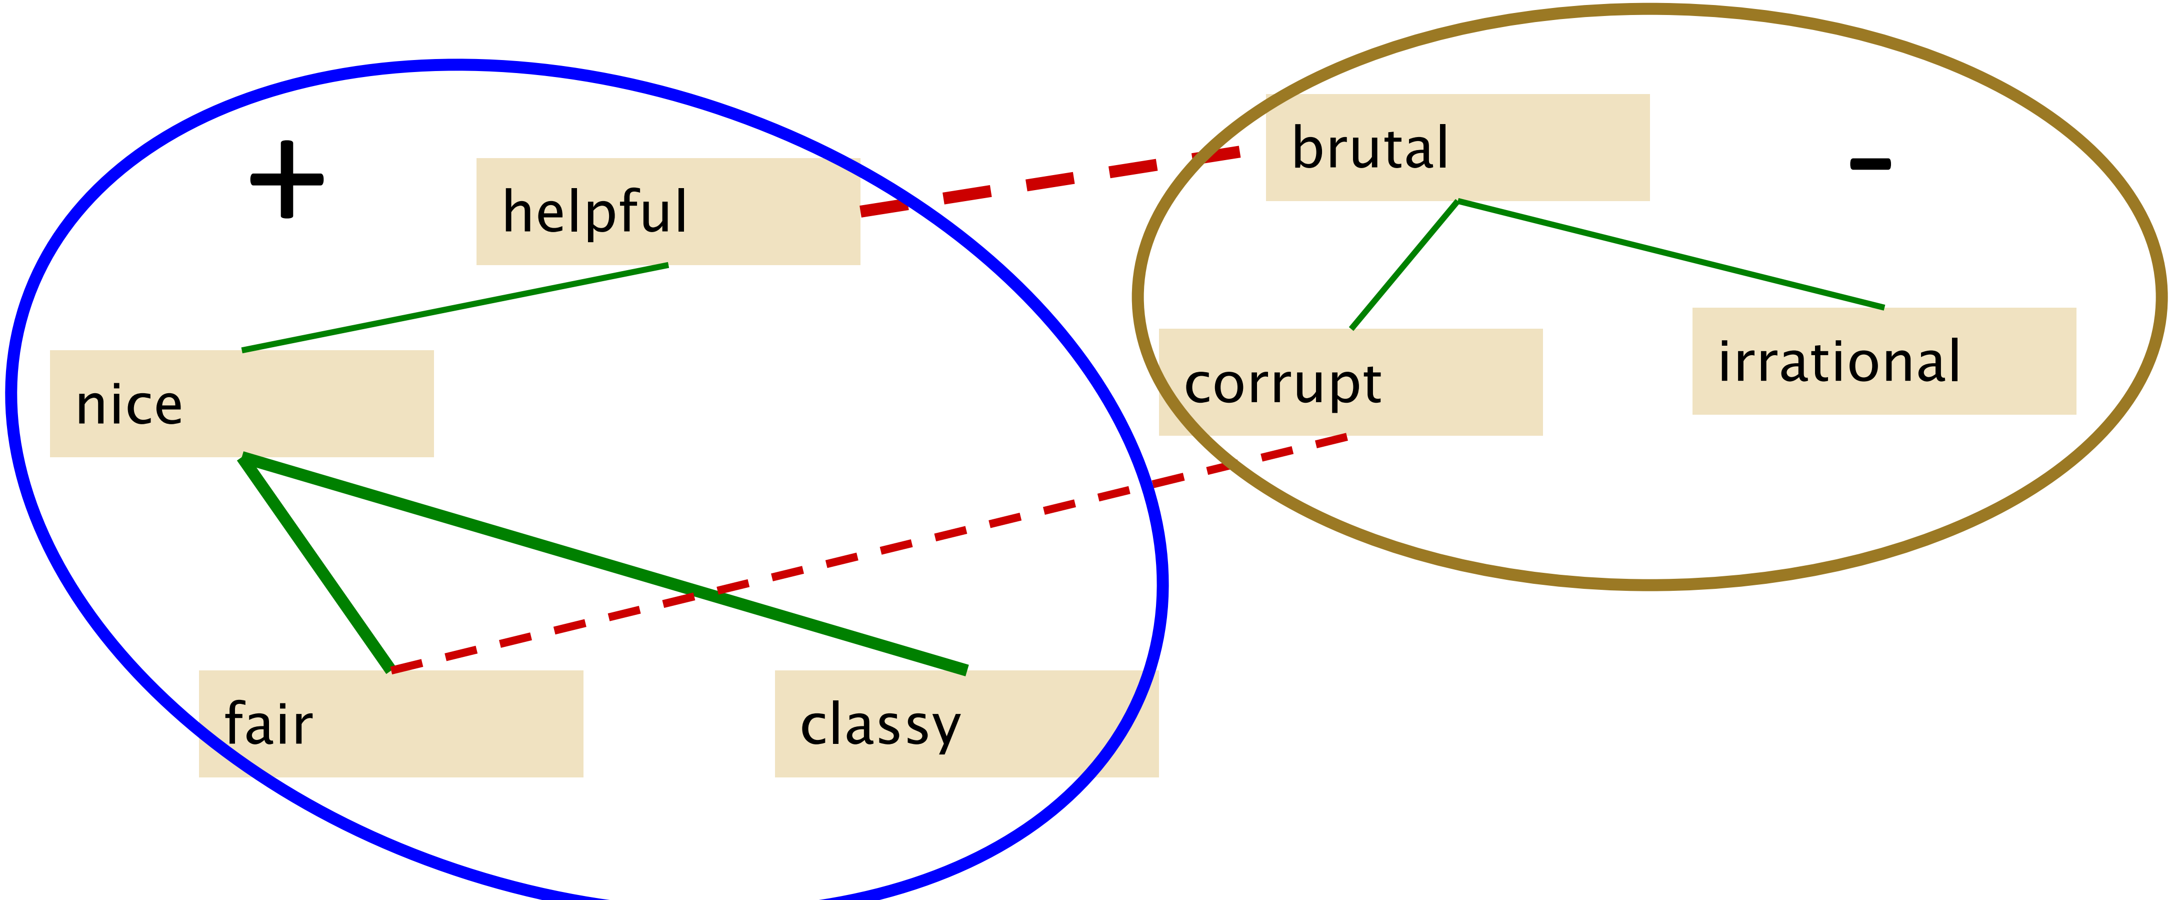

Output polarity lexicon

- 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…

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

Extract two-word phrases with adjectives

First Word	Second Word	Third Word (not extracted)
JJ	NN or NNS	anything
RB, RBR, RBS	JJ	Not NN nor NNS
JJ	JJ	Not NN or NNS
NN or NNS	JJ	Nor NN nor NNS
RB, RBR, or RBS	VB, VBD, VBN, VBG	anything

How to measure polarity of a phrase?

- Positive phrases co-occur more with *"excellent"*
- Negative phrases co-occur more with *"poor"*
- But how to measure co-occurrence?

Pointwise Mutual Information

- **PMI between two words**:
  - How much more do two words co-occur than if they were independent?

$$\mathrm{PMI}(word_1, word_2) = \log_2 \frac{P(word_1, word_2)}{P(word_1)P(word_2)}$$

Does phrase appear more with "poor" or "excellent"?

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

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

Phrases from a thumbs-down review

Phrase	POS tags	Polarity
direct deposits	JJ NNS	5.8
online web	JJ NN	1.9
very handy	RB JJ	1.4
…		
virtual monopoly	JJ NN	-2.0
lesser evil	RBR JJ	-2.3
other problems	JJ NNS	-2.8
low funds	JJ NNS	-6.8
unethical practices	JJ NNS	-8.5
Average		-1.2

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

Using WordNet to learn polarity

S.M. Kim and E. Hovy. 2004. Determining the sentiment of opinions. COLING 2004
M. Hu and B. Liu. Mining and summarizing customer reviews. In Proceedings of KDD, 2004

- WordNet: online thesuarus
- Create positive (“good”) and negative seed-words (“terrible”)
- Find Synonyms and Antonyms
  - Positive Set: Add synonyms of positive words (“well”) and antonyms of negative words
  - Negative Set: Add synonyms of negative words (“awful”) and antonyms of positive words (”evil”)
- Repeat, following chains of synonyms
- Filter

Summary on semi-supervised lexicon learning

- Advantages:
  - Can be domain-specific
  - Can be more robust (more words)
- Intuition
  - Start with a seed set of words (‘good’, ‘poor’)
  - Find other words that have similar polarity:
    - Using “and” and “but”
    - Using words that occur nearby in the same document
    - Using WordNet synonyms and antonyms

    - Use seeds and semi-supervised learning to induce lexicons

Supervised Learning of Sentiment Lexicons

Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.
Potts 2011 NSF Workshop talk.

- Review datasets
  - IMDB, Goodreads, Open Table, Amazon, Trip Advisor
- Each review has a score (1-5, 1-10, etc)
- Just count how many times each word occurs with each score
  - (and normalize)

Analyzing the polarity of each word in IMDB

Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.

- How likely is each word to appear in each sentiment class?
- Count("bad") in 1-star, 2-star, 3-star, etc.
- But can't use raw counts:
- Instead, **likelihood:**

$$P(w \mid c) = \frac{f(w,c)}{\sum_{w \in c} f(w,c)}$$

- Make them comparable between words
  - **Scaled likelihood:**

$$\frac{P(w \mid c)}{P(w)}$$

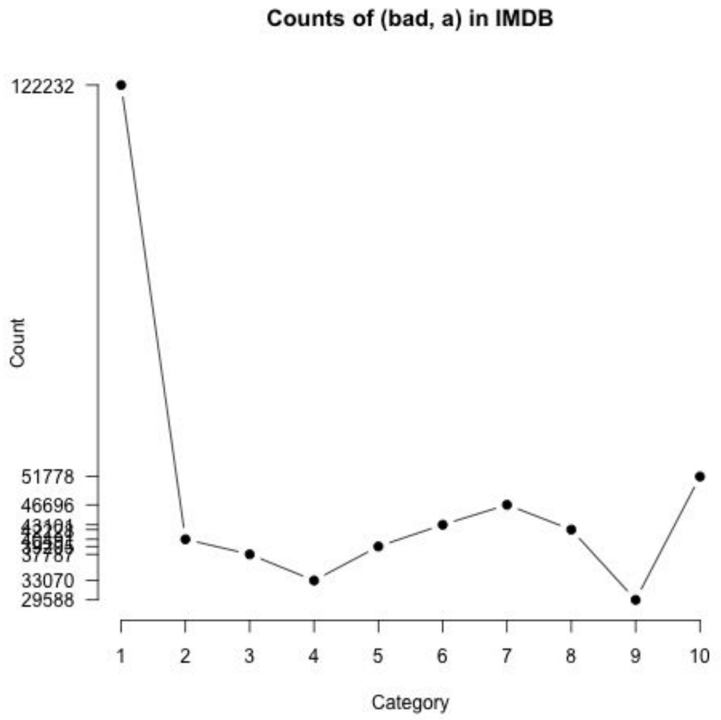

“Potts diagrams”

Potts, Christopher. 2011. NSF workshop on restructuring adjectives.

Positive scalars | **Negative scalars** | **Emphatics** | **Attenuators**

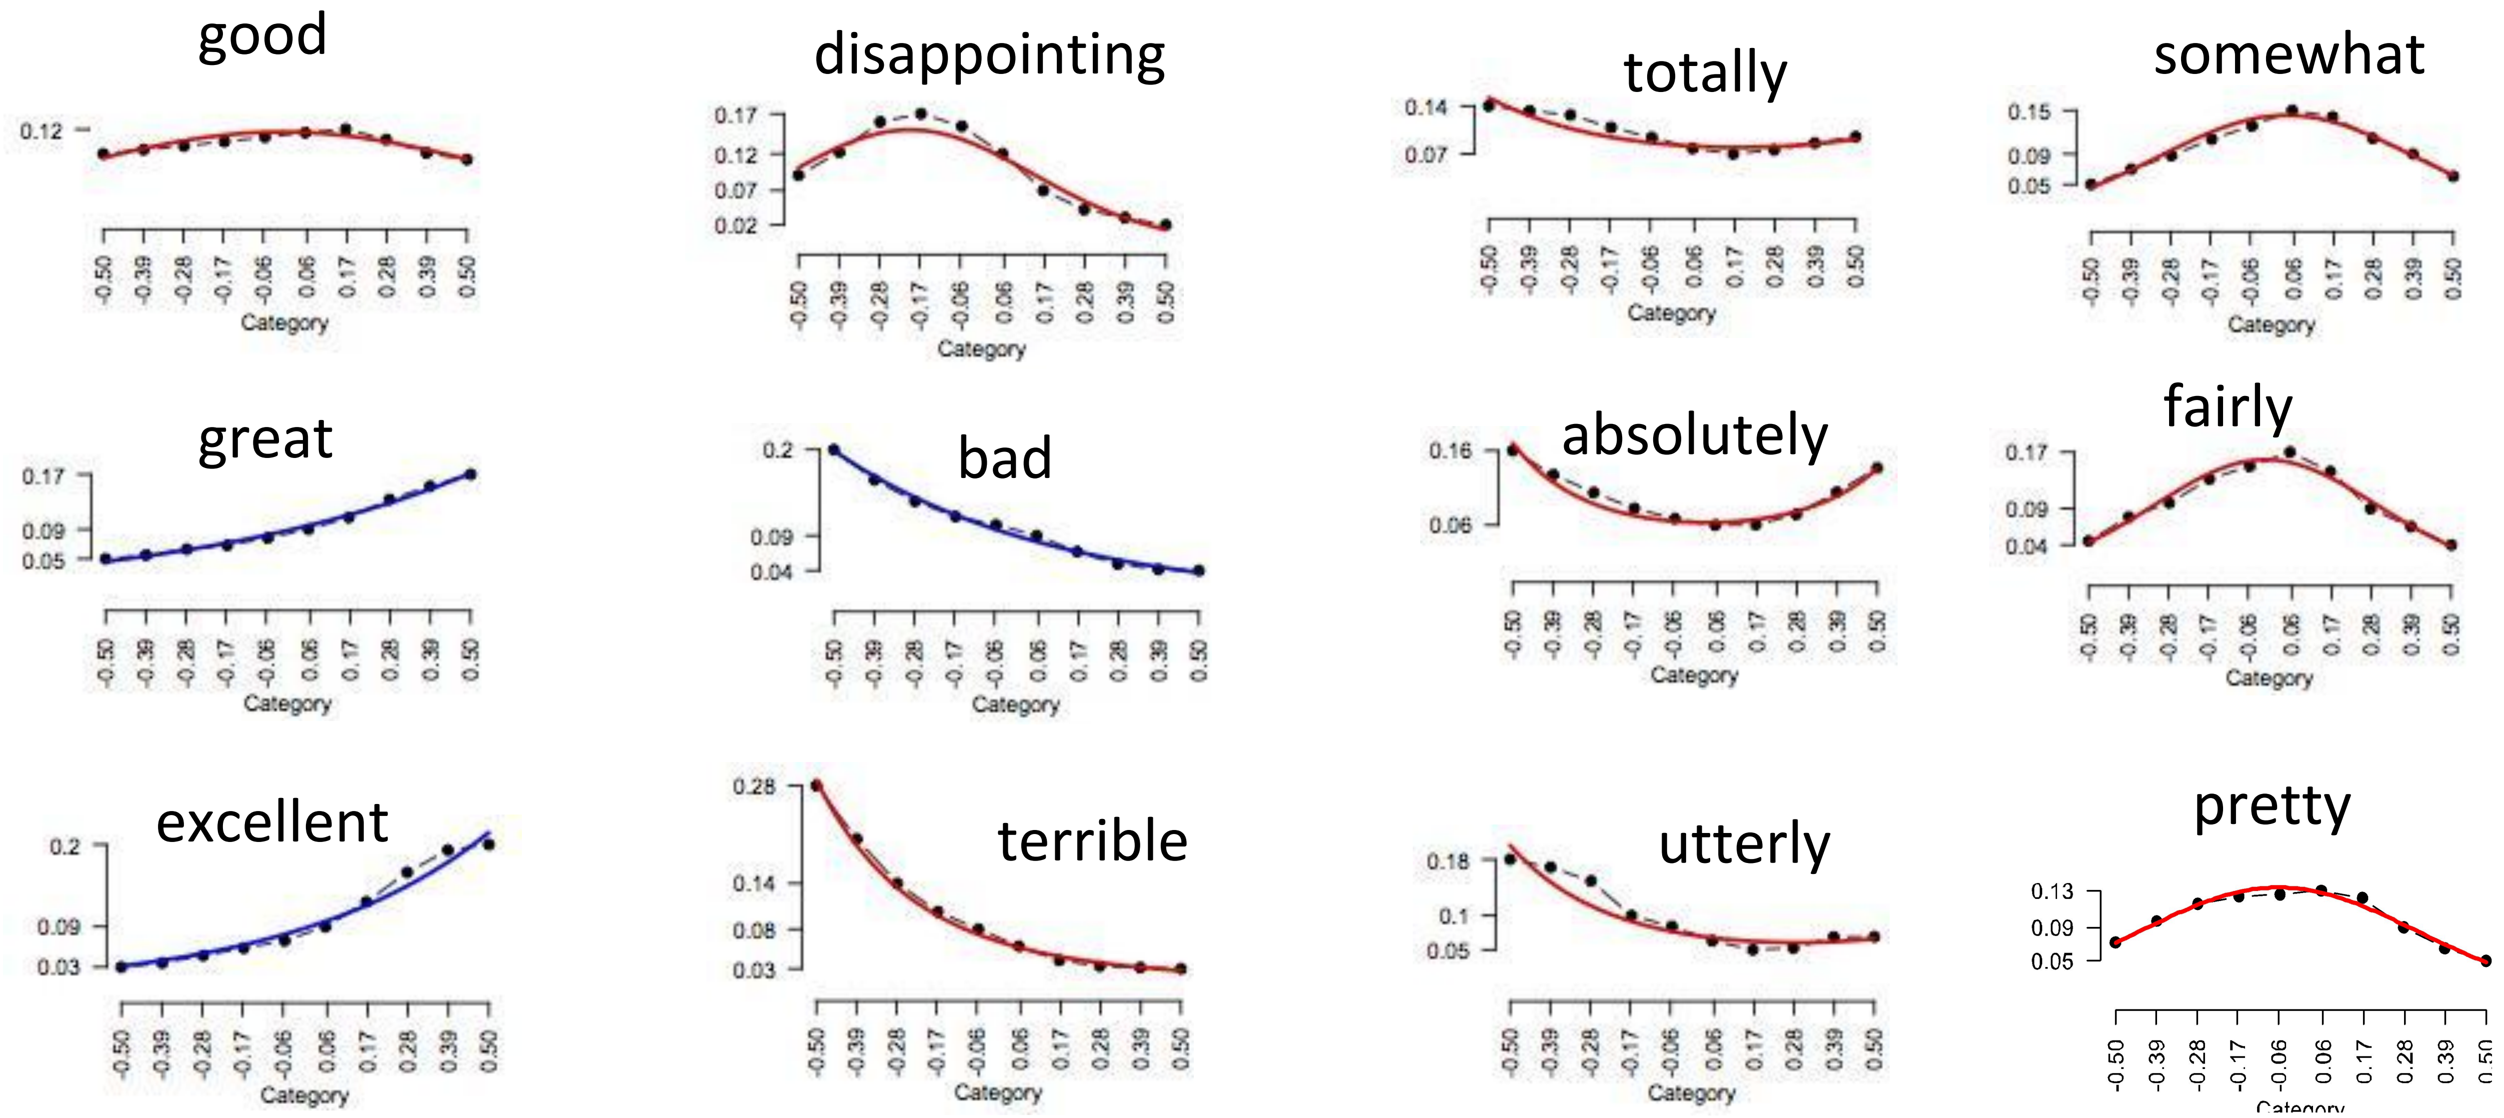

Supervised Sentiment Analysis

Baseline algorithm: Sentiment Classification in Movie Reviews

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.
Bo Pang and Lillian Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. ACL, 271-278

- Build a classifier
  - Predict sentiment given features
  - Use "counts of lexicon categories" as a features
  - Handle negation
  - Use counts of **all** the words and bigrams in the training set

  - This is hard to beat
  - But only works if the training and test sets are very similar

Sentiment Classification with Recursive Neural Networks

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher Manning, Andrew Ng and Christopher Potts
Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank
Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)

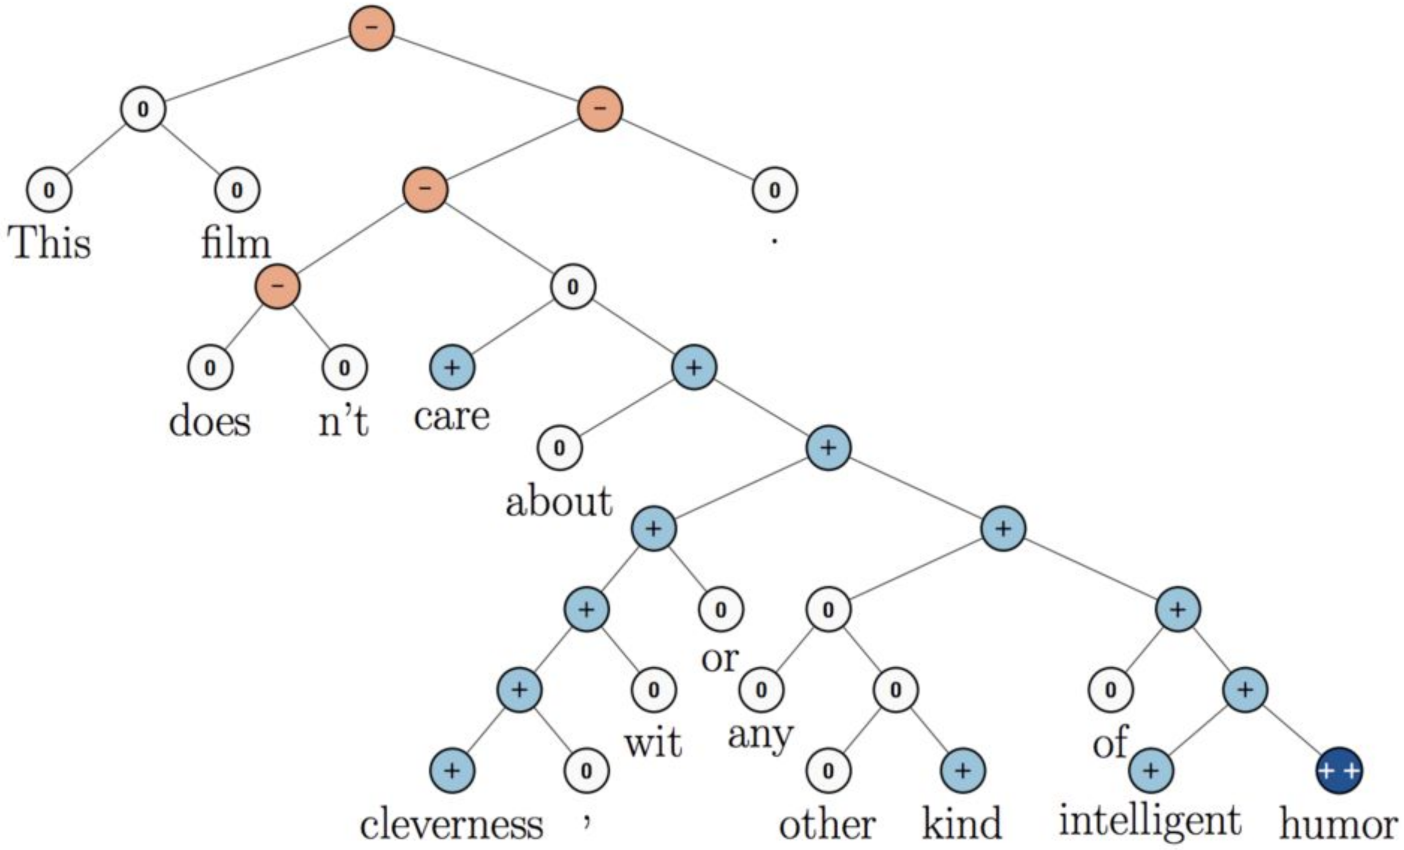

Stanford Sentiment Treebank

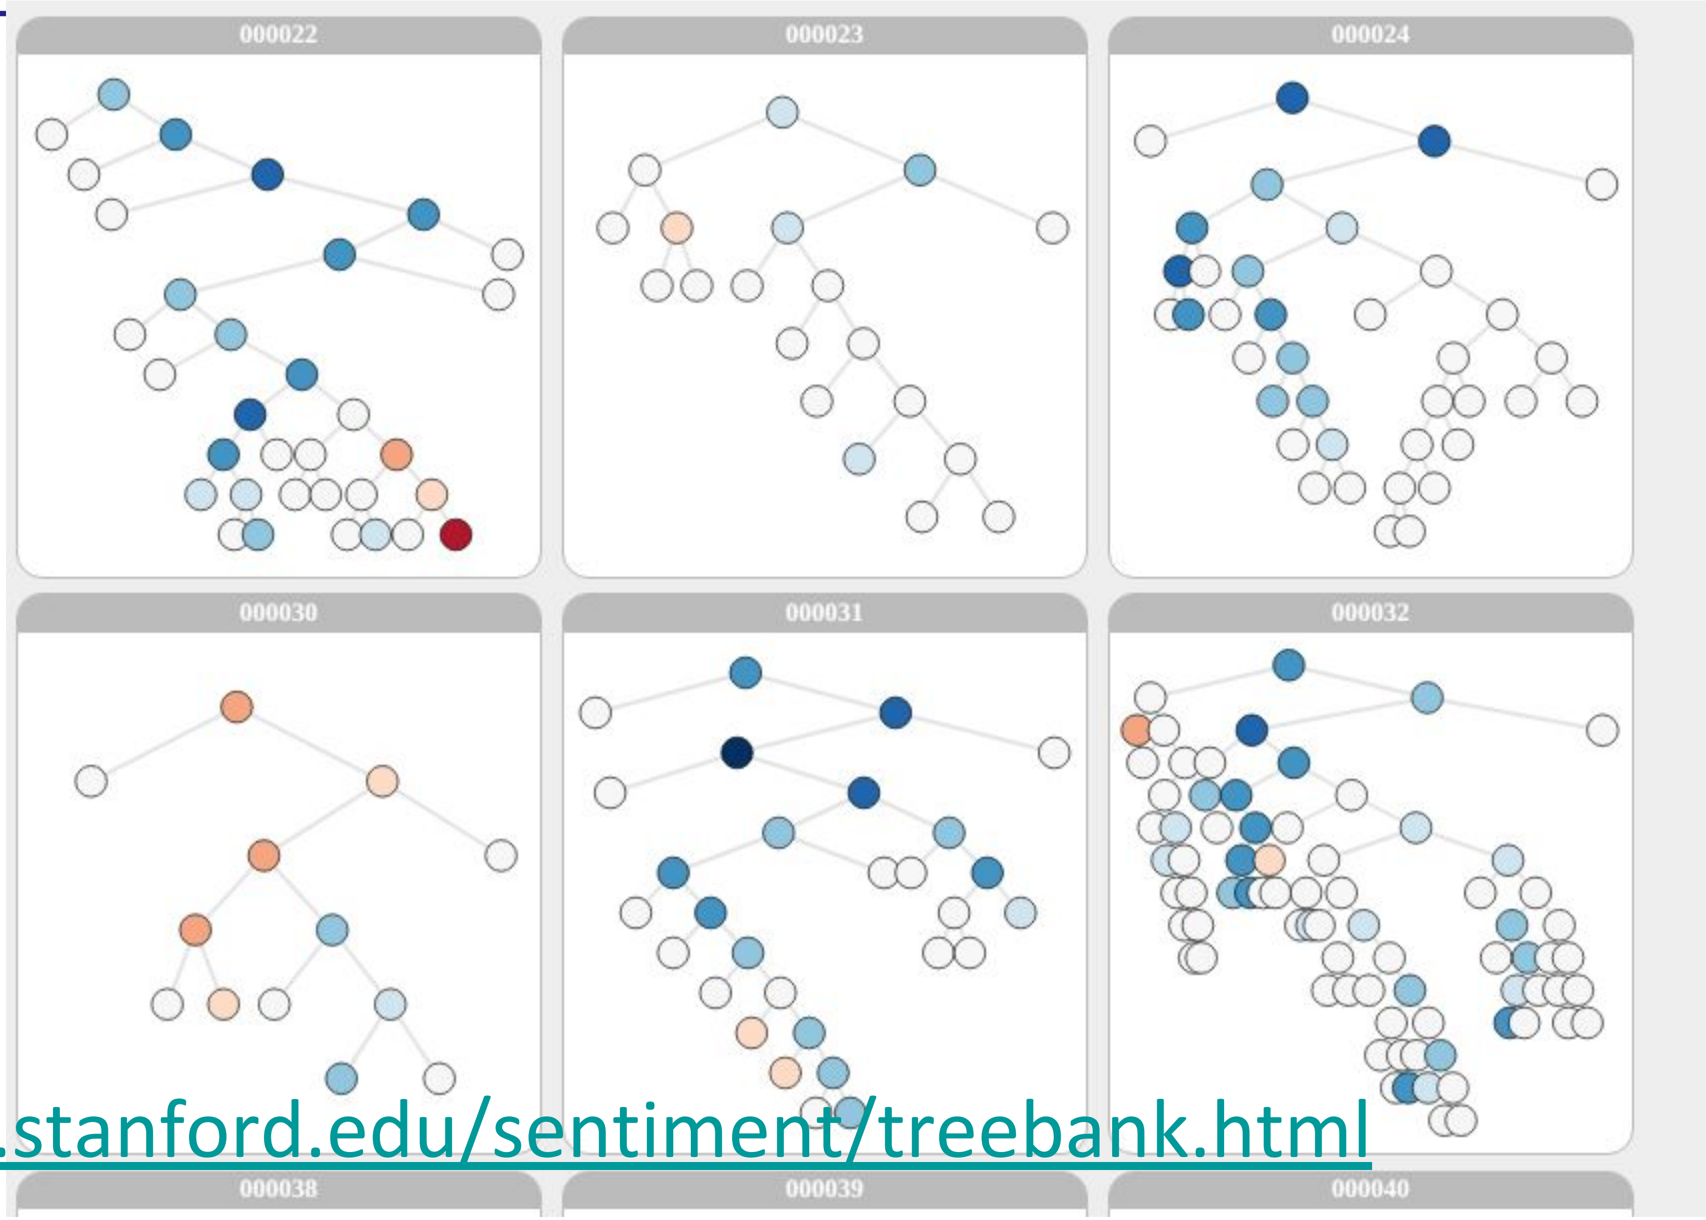

https://nlp.stanford.edu/sentiment/treebank.html

Recursive Neural Networks

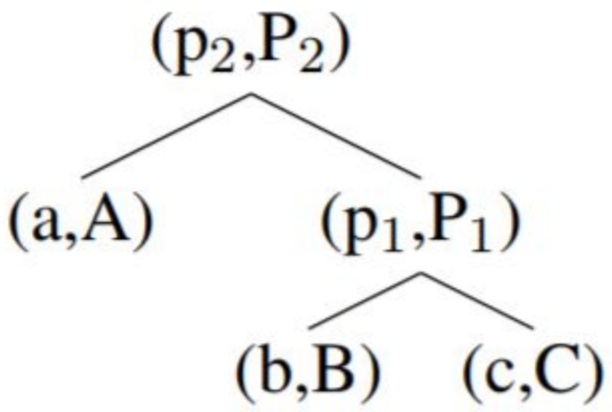

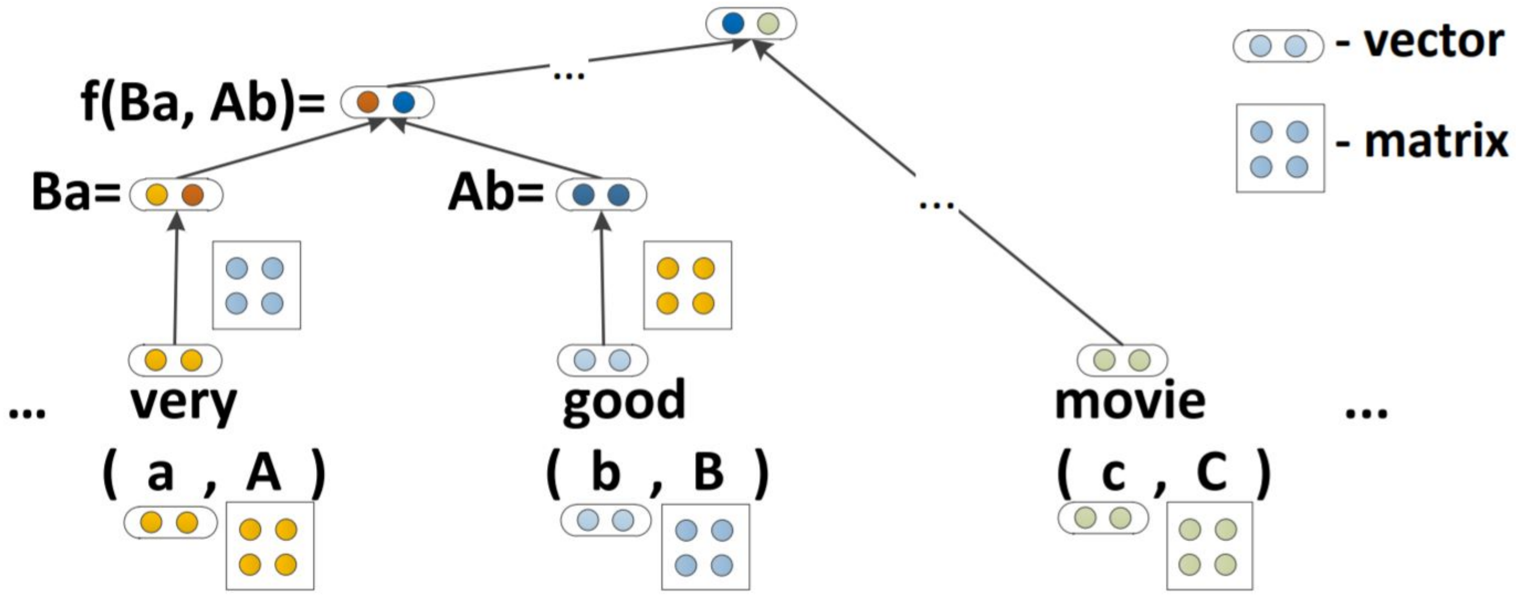

Sentiment Classification with Recursive Neural Networks

Model	Fine-grained		Positive/Negative	
	All	Root	All	Root
NB	67.2	41.0	82.6	81.8
SVM	64.3	40.7	84.6	79.4
BiNB	71.0	41.9	82.7	83.1
VecAvg	73.3	32.7	85.1	80.1
RNN	79.0	43.2	86.1	82.4
MV-RNN	78.7	44.4	86.8	82.9
RNTN	**80.7**	**45.7**	**87.6**	**85.4**

Table 1: Accuracy for fine grained (5-class) and binary predictions at the sentence level (root) and for all nodes.

SOTA Methods

SemEval Competitions: 2007, 2010, 2014, 2015, 2016, 2017

http://alt.qcri.org/semeval2016

SemEval-2016 : Semantic Evaluation Exercises

International Workshop on Semantic Evaluation (SemEval-2016)

Track II. Sentiment Analysis Track

- Task 4: Sentiment Analysis in Twitter
- Task 5: Aspect-Based Sentiment Analysis
- Task 6: Detecting Stance in Tweets
- Task 7: Determining Sentiment Intensity of English and Arabic Phrases

SemEval 2016

SwissCheese at SemEval-2016 Task 4: Sentiment Classification Using an Ensemble of Convolutional Neural Networks with Distant Supervision

Jan Deriu*
ETH Zurich
Switzerland
jderiu@student.ethz.ch

Maurice Gonzenbach*
ETH Zurich
Switzerland
mauriceg@student.ethz.ch

Fatih Uzdilli
Zurich University of Applied Sciences
Switzerland
uzdi@zhaw.ch

Aurelien Lucchi
ETH Zurich
Switzerland
alucchi@inf.ethz.ch

Valeria De Luca
ETH Zurich
Switzerland
vdeluca@vision.ee.ethz.ch

Martin Jaggi
ETH Zurich
Switzerland
jaggi@inf.ethz.ch

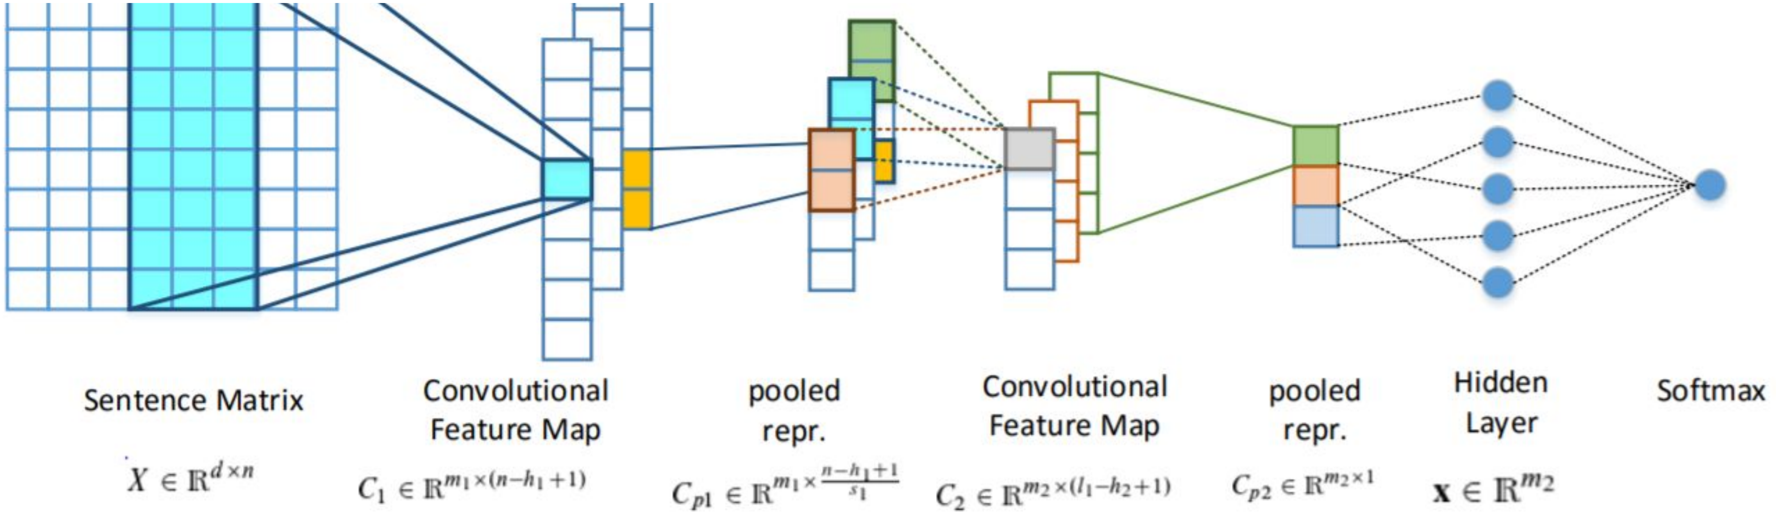

Figure 1: The architecture of the CNNs used in our approach.

SemEval 2017

http://alt.qcri.org/semeval2017/task4/index.php?id=results

SemEval-2017 Task 4

Sentiment Analysis in Twitter

Results

1. All training data can be found here.
2. The test data can be found here.
3. The gold labels, submissions and scores for all teams can be found here.
4. The task paper can be found here.

```
@InProceedings{SemEval:2017:task4,
  author    = {Sara Rosenthal and Noura Farra and Preslav Nakov},
  title     = {{SemEval}-2017 Task 4: Sentiment Analysis in {T}witter},
  booktitle = {Proceedings of the 11th International Workshop on Semantic Evaluation},
  series    = {SemEval '17},
  month     = {August},
  year      = {2017},
  address   = {Vancouver, Canada},
  publisher = {Association for Computational Linguistics},
}
```

Connotation Frames

[Hannah Rashkin, Sameer Singh, Yejin Choi (2016) Connotation Frames: A Data-Driven Investigation. ACL]

Attacked
vs
Fought

Frame Semantic Parsing + Sentiment

English Verb: survive

Example Tweets

"US teenager … also survived Boston Marathon bombing"

Connotation Frame for *surviving* verbs:

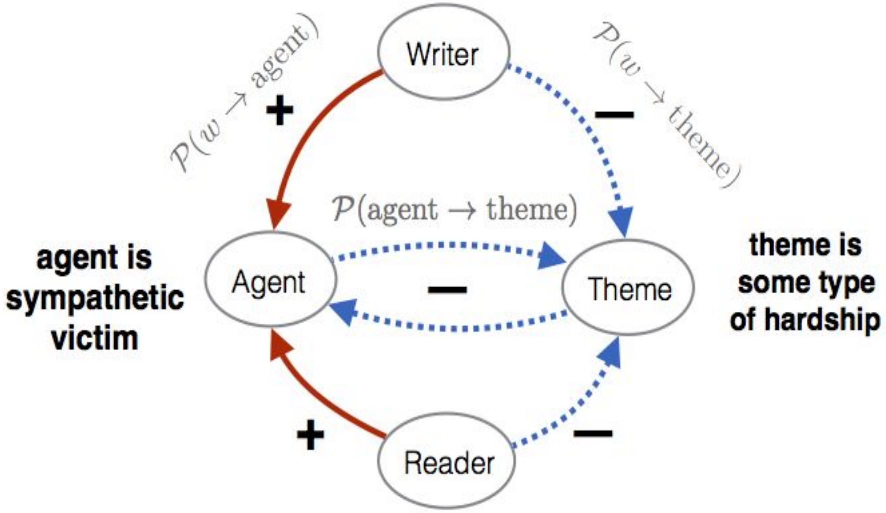

Writer's Perspective: how the writer feels about the entities interacting through the predicate

Reader's Perspective: how the reader will likely feel about the entities

Entity's Perspective: how the entities feel about one another

Cross-Domain Sentiment Classification

Domain-Adversarial Training of Neural Networks
Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, Victor Lempitsky
Journal of Machine Learning Research 2016, vol. 17, p. 1-35

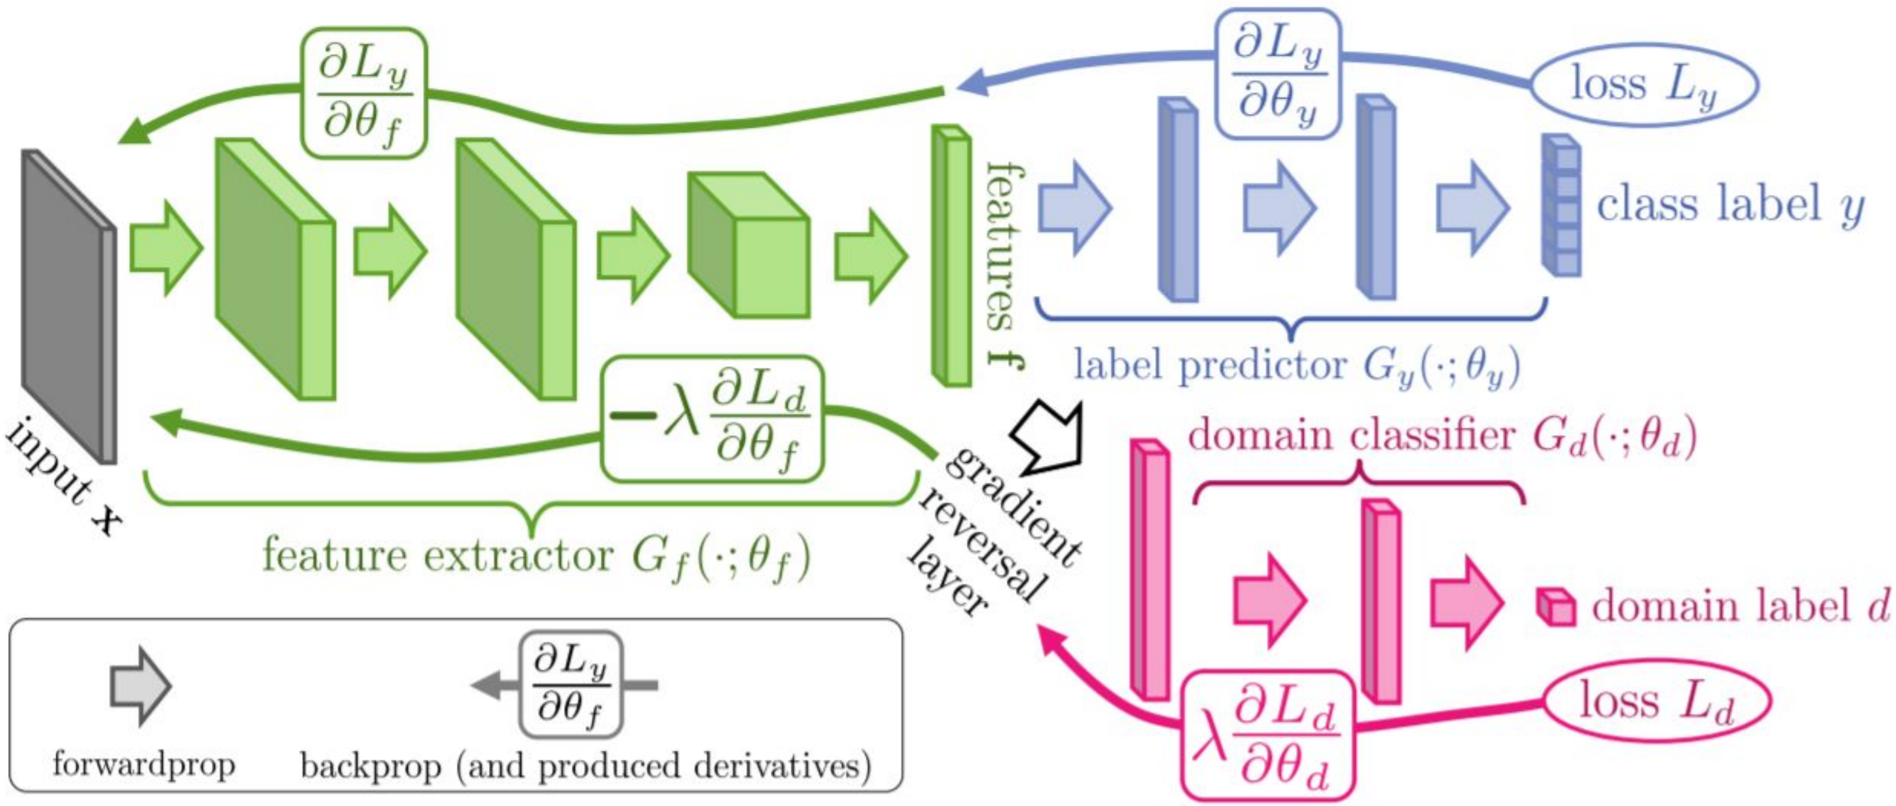

Multilingual Sentiment Lexicons

https://polyglot.readthedocs.io/en/latest/Sentiment.html

Building Sentiment Lexicons for All Major Languages
Yanqing Chen and Steven Skiena
ACL 2014

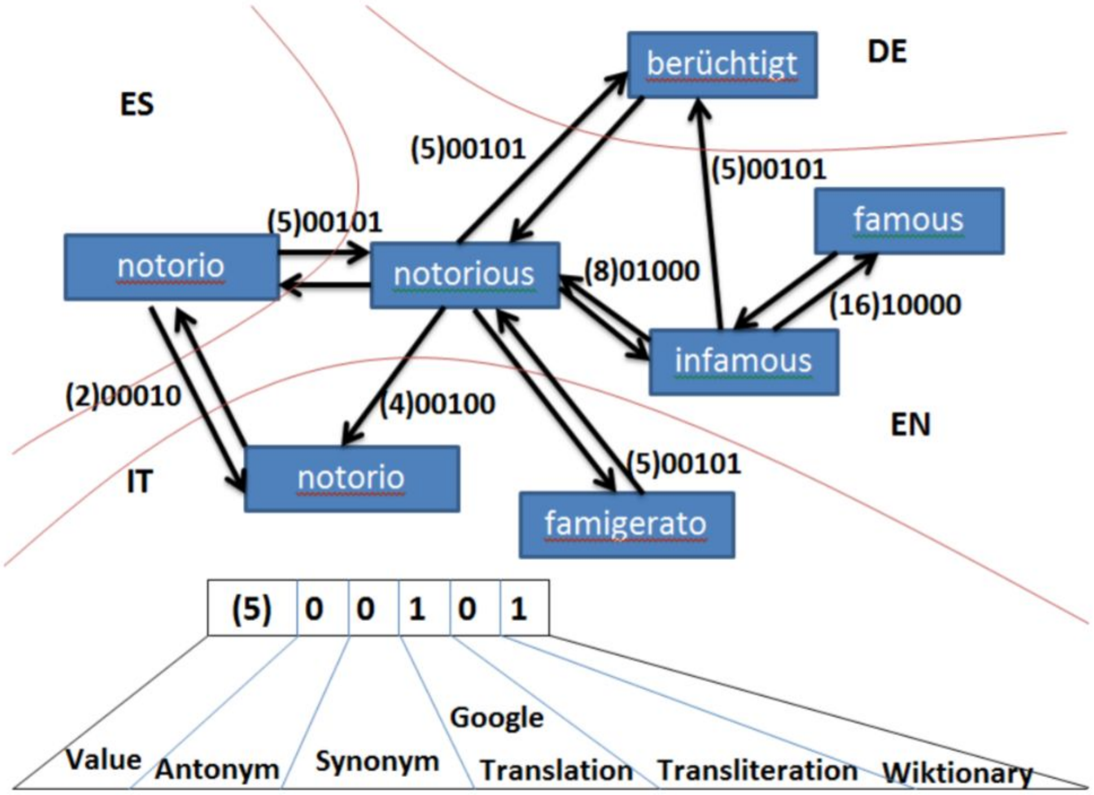

Languages Coverage

```
from polyglot.downloader import downloader
print(downloader.supported_languages_table("sentiment2", 3))
```

```
  1. Turkmen                    2. Thai                       3. Latvian
  4. Zazaki                     5. Tagalog                    6. Tamil
  7. Tajik                      8. Telugu                     9. Luxembourgish, Letzeb...
 10. Alemannic                 11. Latin                     12. Turkish
 13. Limburgish, Limburgan...  14. Egyptian Arabic           15. Tatar
 16. Lithuanian                17. Spanish; Castilian        18. Basque
 19. Estonian                  20. Asturian                  21. Greek, Modern
 22. Esperanto                 23. English                   24. Ukrainian
 25. Marathi (Marāṭhī)         26. Maltese                   27. Burmese
 28. Kapampangan               29. Uighur, Uyghur            30. Uzbek
 31. Malagasy                  32. Yiddish                   33. Macedonian
 34. Urdu                      35. Malayalam                 36. Mongolian
 37. Breton                    38. Bosnian                   39. Bengali
 40. Tibetan Standard, Tib...  41. Belarusian                42. Bulgarian
 43. Bashkir                   44. Vietnamese                45. Volapük
 46. Gan Chinese               47. Manx                      48. Gujarati
 49. Yoruba                    50. Occitan                   51. Scottish Gaelic; Gaelic
 52. Irish                     53. Galician                  54. Ossetian, Ossetic
 55. Oriya                     56. Walloon                   57. Swedish
 58. Silesian                  59. Lombard language          60. Divehi; Dhivehi; Mald...
 61. Danish                    62. German                    63. Armenian
 64. Haitian; Haitian Creole   65. Hungarian                 66. Croatian
 67. Bishnupriya Manipuri      68. Hindi                     69. Hebrew (modern)
 70. Portuguese                71. Afrikaans                 72. Pashto, Pushto
 73. Amharic                   74. Aragonese                 75. Bavarian
 76. Assamese                  77. Panjabi, Punjabi          78. Polish
 79. Azerbaijani               80. Italian                   81. Arabic
 82. Icelandic                 83. Ido                       84. Scots
 85. Sicilian                  86. Indonesian                87. Chinese Word
 88. Interlingua               89. Waray-Waray               90. Piedmontese language
 91. Quechua                   92. French                    93. Dutch
 94. Norwegian Nynorsk         95. Norwegian                 96. Western Frisian
 97. Upper Sorbian             98. Nepali                    99. Persian
100. Ilokano                  101. Finnish                  102. Faroese
103. Romansh                  104. Javanese                 105. Romanian, Moldavian, ...
106. Malay                    107. Japanese                 108. Russian
109. Catalan; Valencian       110. Fiji Hindi               111. Chinese
112. Cebuano                  113. Czech                    114. Chuvash
115. Welsh                    116. West Flemish             117. Kirghiz, Kyrgyz
118. Kurdish                  119. Kazakh                   120. Korean
121. Kannada                  122. Khmer                    123. Georgian
124. Sakha                    125. Serbian                  126. Albanian
127. Swahili                  128. Chechen                  129. Sundanese
130. Sanskrit (Saṁskṛta)      131. Venetian                 132. Northern Sami
```

Summary

- Lexicons
- Lexicon learning
- Sentiment classification
- Sentiment analysis ++
  - More labels
  - Aspect-based sentiment
  - Stance detection
  - Emotion classification
  - Sentiment + syntactic parsing
  - Sentiment + semantic frames
  - Multilingual sentiment analysis