Crowdsourcing Word-Emotion Associations

Saif Mohammad

National Research Council Canada

Includes joint work with Peter Turney, Tony Yang, Svetlana Kiritchenko, and Xiaodan Zhu.

Word-Emotion Associations

Words have associations with emotions:

- attack and public speaking typically associated with fear
- yummy and vacation typically associated with joy
- death and crying typically associated with sadness
- result and wait typically associated anticipation

Similarly, some words are associated with positive sentiment and some with negative sentiment.

Goal: Create large word-emotion and word-sentiment association lexicons.

Sentiment Analysis

- Is a given sentence positive, negative, or neutral?
- Is a word within a sentence positive, negative, or neutral?

Sentiment Analysis

- Is a given sentence positive, negative, or neutral?
- Is a word within a sentence positive, negative, or neutral?

Emotion Analysis

- What emotion is being expressed in a given sentence?
  - Basic emotions: joy, sadness, fear, anger,…
  - Other emotions: guilt, pride, optimism, frustration,…

Applications of Sentiment Analysis and Emotion Detection

- Tracking sentiment towards politicians, movies, products
- Improving customer relation models
- Identifying what evokes strong emotions in people
- Detecting personality
- Detecting happiness and well-being
- Measuring the impact of activist movements through text generated in social media.
- Improving automatic dialogue systems
- Detecting how people use emotion-bearing-words and metaphors to persuade and coerce others

Applications of Sentiment Analysis and Emotion Detection

- Tracking sentiment towards politicians, movies, products
- Improving customer relation models
- Identifying what evokes strong emotions in people
- Detecting personality
- Detecting happiness and well-being
- Measuring the impact of activist movements through text generated in social media.
- Improving automatic dialogue systems
- Detecting how people use emotion-bearing-words and metaphors to persuade and coerce others

This Talk

- Creating large word-emotion and word-sentiment association lexicons by crowdsourcing to Mechanical Turk
  - Tracking and visualizing sentiment and emotions in text
  - Creating a large word-colour association lexicon

- Creating large word-emotion and word-sentiment association lexicons from hashtagged tweets
  - Detecting sentiment and emotions in phrases, sentences, and tweets
  - Using hundreds of fine emotion categories to detect personality traits

Challenges in detecting emotions

- Not explicitly stated
  - Need world knowledge and context
- No tone, pitch, or other prosodic information
- Text may have sarcasm, exaggeration, etc

Which Emotions?

Plutchik, 1980: Eight Basic Emotions

- Joy
- Trust
- Fear
- Surprise
- Sadness
- Disgust
- Anger
- Anticipation

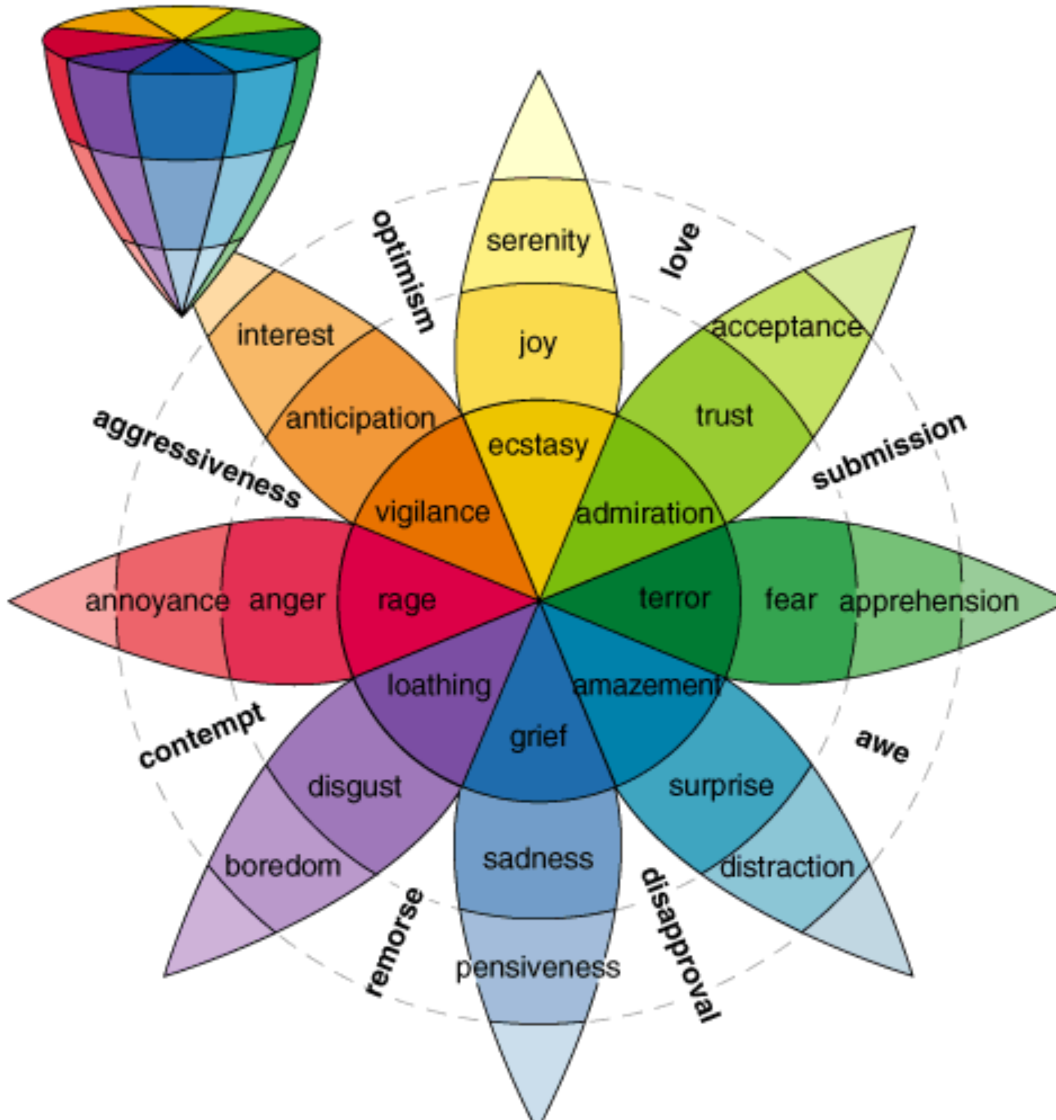

Amazon's Mechanical Turk

- Requester
  - breaks task into small independent units – HITs
  - specifies:
    - compensation for solving each HIT
    - # of independent annotations required for each HIT
- Turkers
  - attempt as many HITs as they wish

Crowdsourcing

- Benefits
  - Inexpensive
  - Convenient and time-saving
    - Especially for large-scale annotation
- Challenges
  - Quality control
    - Malicious annotations
    - Inadvertent errors
  - Words used in different senses are associated with different emotions.
  - Working with subjective questions

Target Words

- Must be:
  - in *Roget's Thesaurus*
  - high-frequency term in the Google n-gram corpus

Word-Choice Question

Q1. Which word is closest in meaning to *cry*?

• *car* • *tree* • *tears* • *olive*

Peter Turney, NRC

- Generated automatically
  - Near-synonym taken from thesaurus
  - Distractors are randomly chosen
- Guides Turkers to desired sense
- Aides quality control
  - If Q1 is answered incorrectly:
    - Responses to the remaining questions for the word are discarded

Association Questions

Q2. How much is *cry* associated with the emotion sadness?

(for example, *death* and *gloomy* are strongly associated with sadness)

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

- Eight such questions for the eight basic emotions.
- Two such questions for positive or negative sentiment.
- Better agreement with *associated* rather than *evoked*.

Emotion Lexicon

- Each word-sense pair is annotated by 5 Turkers
- About 10% of the assignments were discarded due to incorrect response to Q1 (gold question)
- Targets with less than 3 valid assignments removed
- NRC Emotion Lexicon
  - sense-level lexicon
    - word sense pairs: 24,200
  - word-level lexicon
    - union of emotions associated with different senses
    - word types: 14,200

Crowdsourcing a Word-Emotion Association Lexicon, Saif Mohammad and Peter Turney, *Computational Intelligence*, 29 (3), pages 436-465, 2013.

Tony (Wenda) Yang
Simon Fraser University

Visualizing Emotions in Text

- Tracking Sentiment in Mail: How Genders Differ on Emotional Axes, Saif Mohammad and Tony Yang, In Proceedings of the ACL 2011 Workshop on ACL 2011 Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA), June 2011, Portland, OR.

- From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales, Saif Mohammad, In Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), June 2011, Portland, OR.

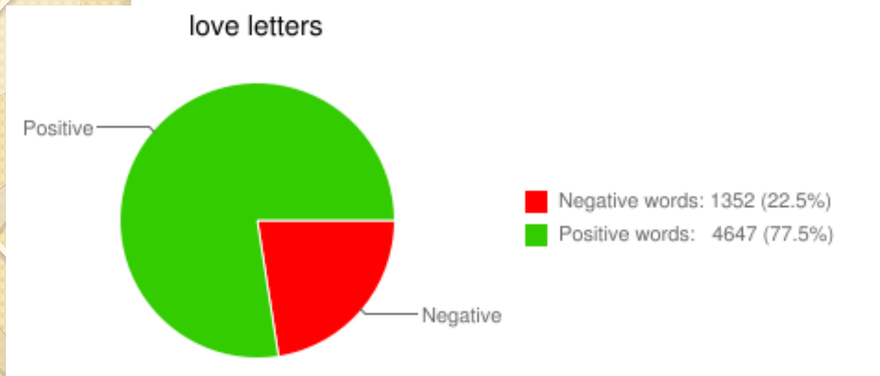
love letters
Positive
Negative
Negative words: 1352 (22.5%)
Positive words: 4647 (77.5%)

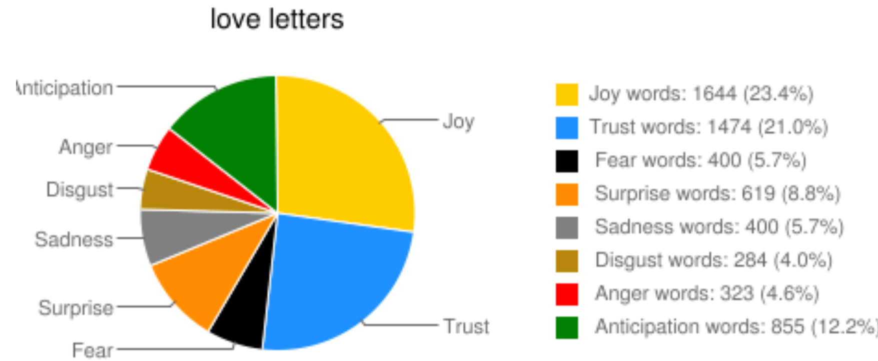
love letters
nticipation
Anger
Disgust
Sadness
Surprise
Fear
Joy
Trust
Joy words: 1644 (23.4%)
Trust words: 1474 (21.0%)
Fear words: 400 (5.7%)
Surprise words: 619 (8.8%)
Sadness words: 400 (5.7%)
Disgust words: 284 (4.0%)
Anger words: 323 (4.6%)
Anticipation words: 855 (12.2%)

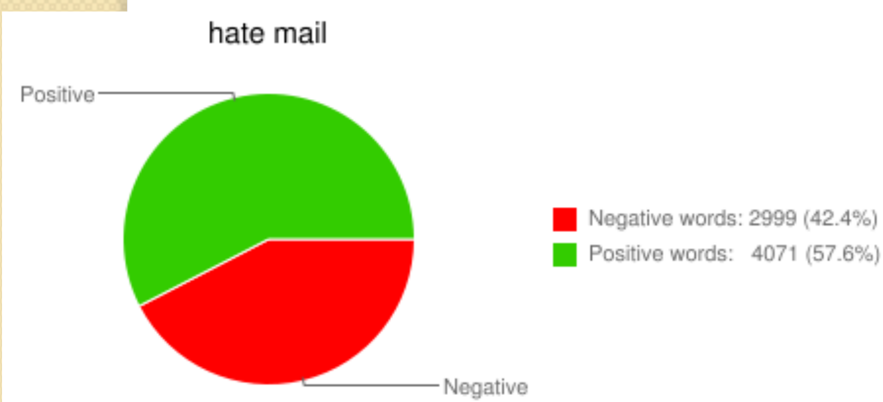
hate mail
Positive
Negative
Negative words: 2999 (42.4%)
Positive words: 4071 (57.6%)

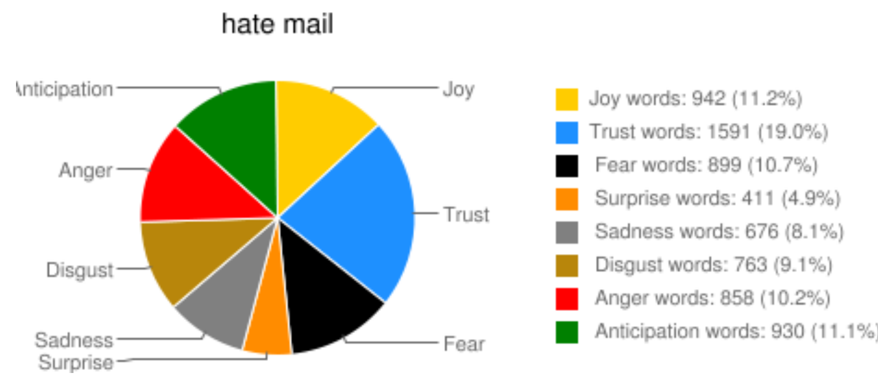
hate mail
nticipation
Anger
Disgust
Sadness
Surprise
Joy
Trust
Fear
Joy words: 942 (11.2%)
Trust words: 1591 (19.0%)
Fear words: 899 (10.7%)
Surprise words: 411 (4.9%)
Sadness words: 676 (8.1%)
Disgust words: 763 (9.1%)
Anger words: 858 (10.2%)
Anticipation words: 930 (11.1%)

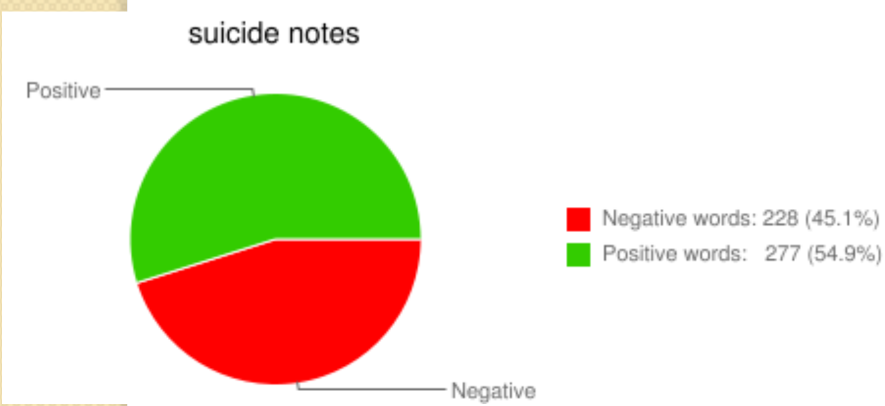
suicide notes
Positive
Negative
Negative words: 228 (45.1%)
Positive words: 277 (54.9%)

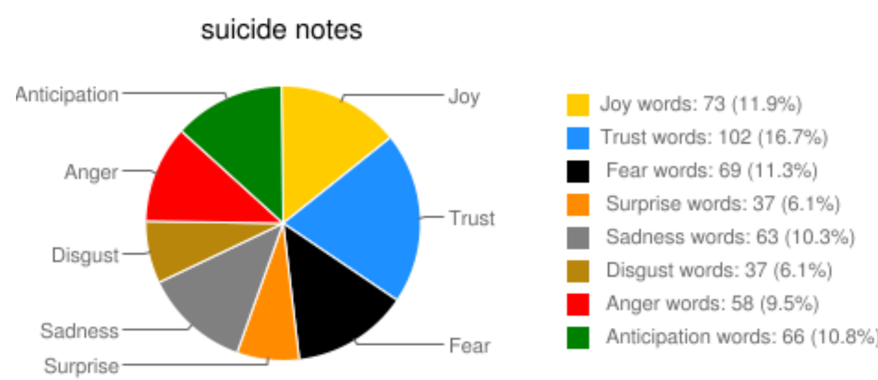
suicide notes
Anticipation
Anger
Disgust
Sadness
Surprise
Joy
Trust
Fear
Joy words: 73 (11.9%)
Trust words: 102 (16.7%)
Fear words: 69 (11.3%)
Surprise words: 37 (6.1%)
Sadness words: 63 (10.3%)
Disgust words: 37 (6.1%)
Anger words: 58 (9.5%)
Anticipation words: 66 (10.8%)

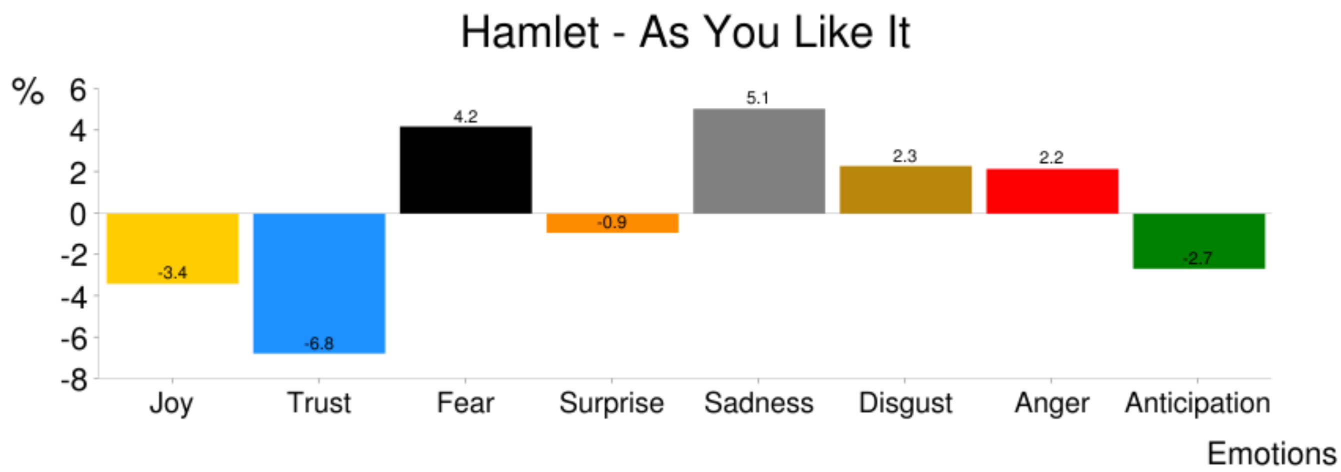

relative salience of trust words

relative salience of sadness words

Flow of Emotions

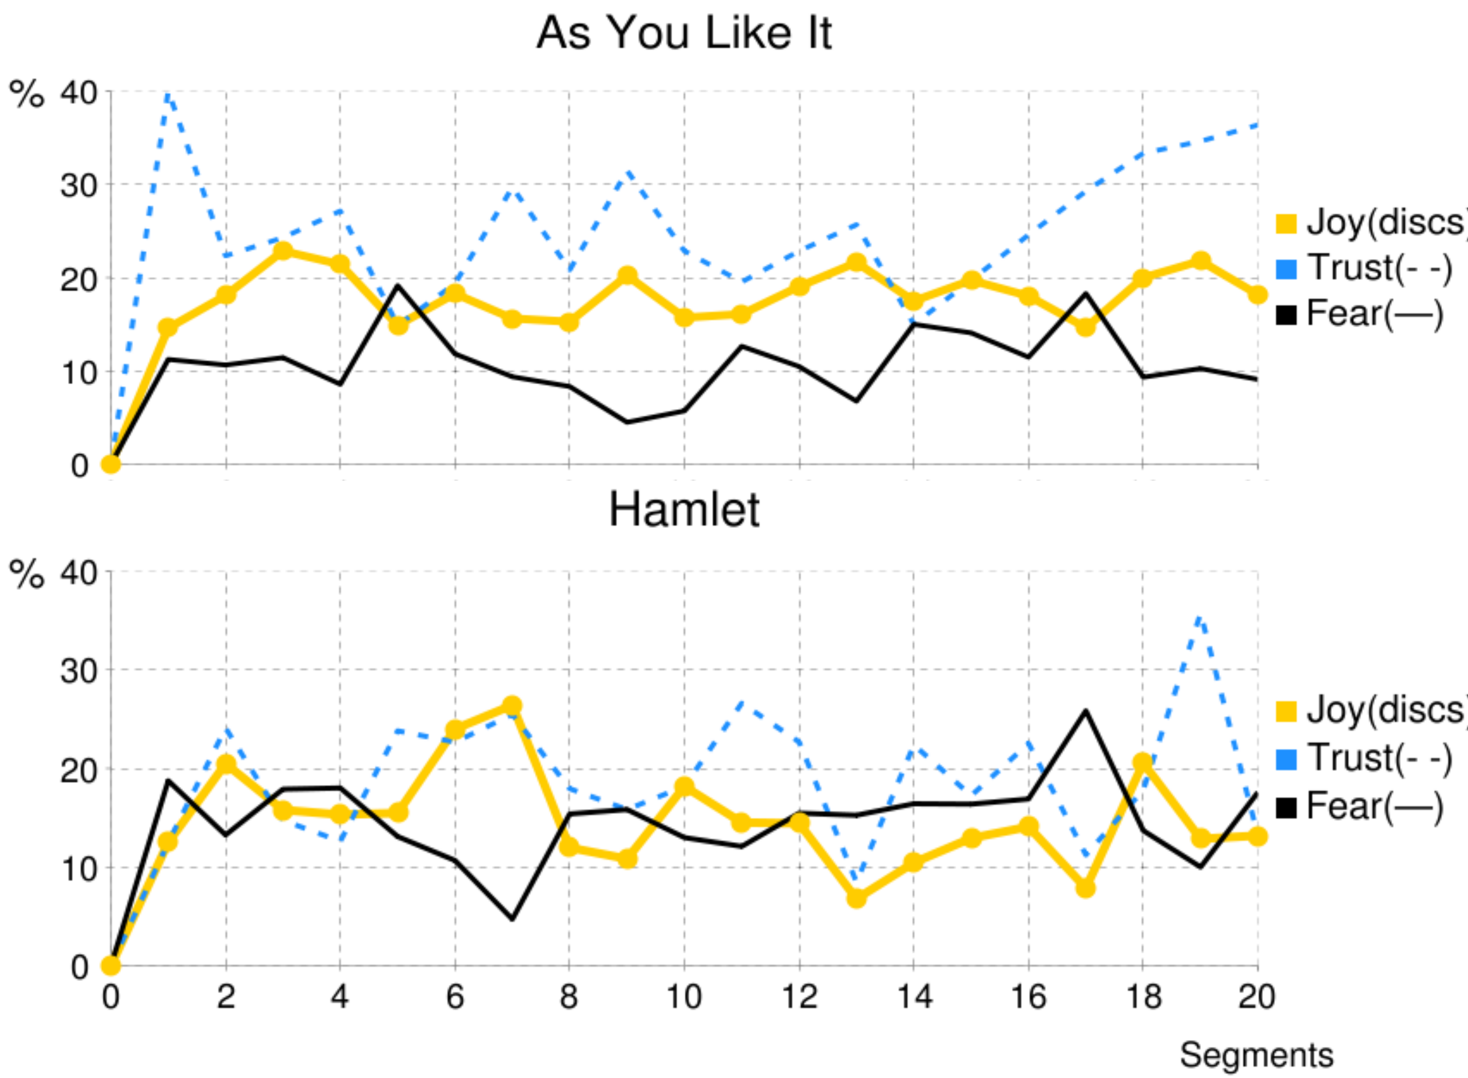

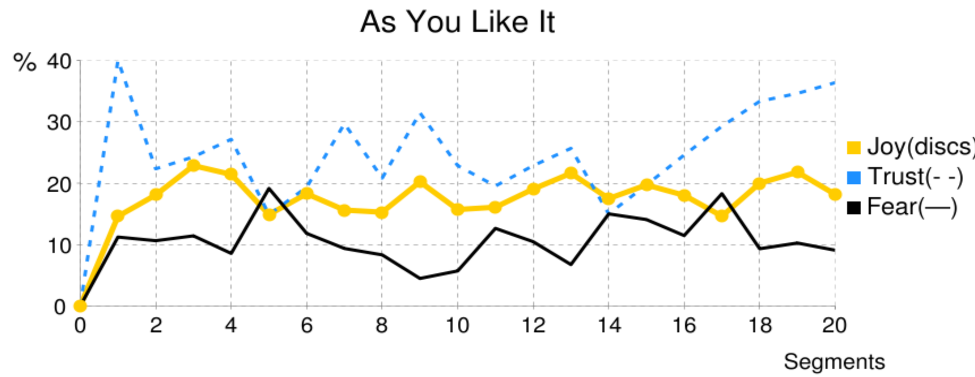
As You Like It
%
0
10
20
30
40
0
2
4
6
8
10
12
14
16
18
20
Segments
Joy(discs)
Trust(- -)
Fear(—)

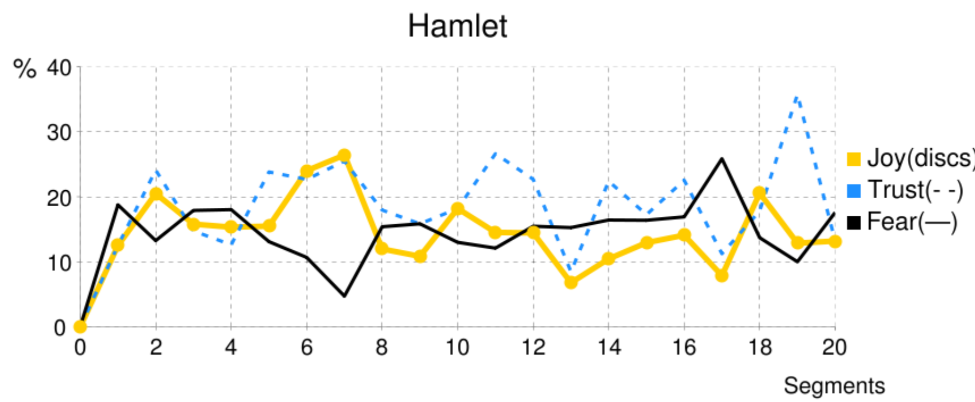
Hamlet
%
0
10
20
30
40
0
2
4
6
8
10
12
14
16
18
20
Segments
Joy(discs)
Trust(- -)
Fear(—)

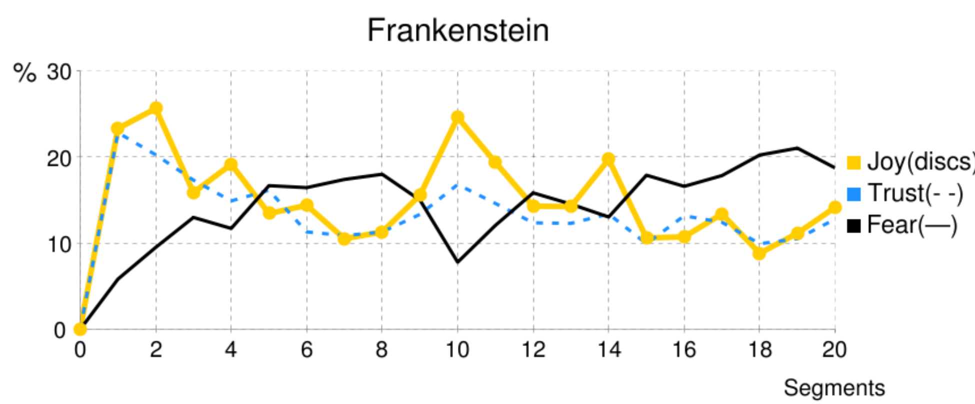
Frankenstein
%
0
10
20
30
0
2
4
6
8
10
12
14
16
18
20
Segments
Joy(discs)
Trust(- -)
Fear(—)

Emotion Word Density

average number of emotion words in every X words

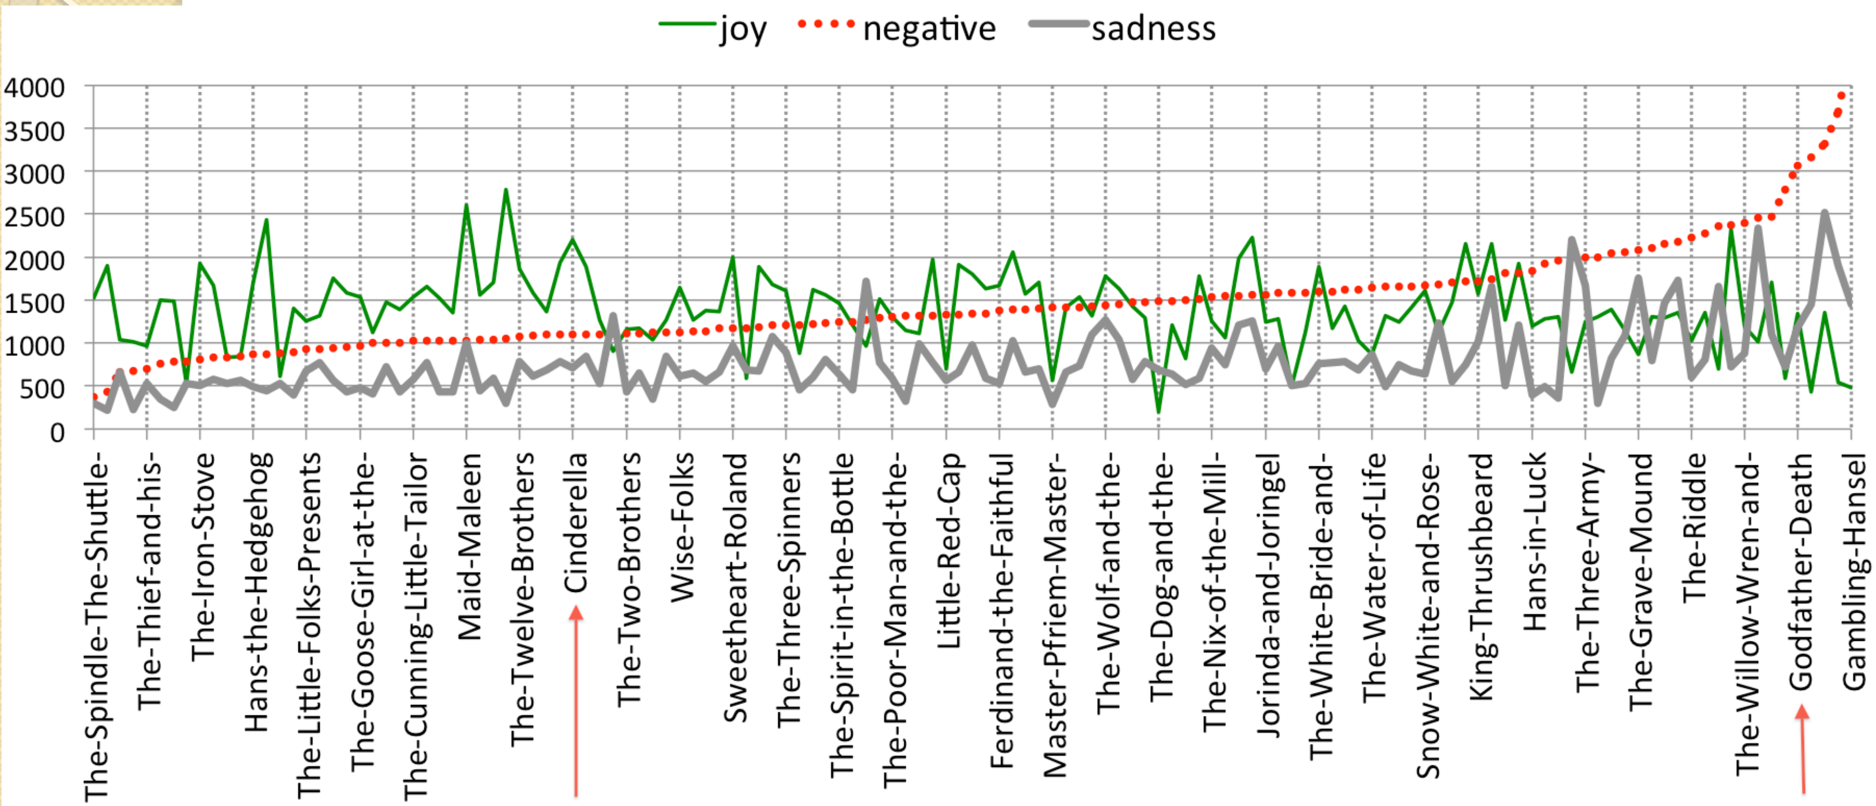

Brothers Grimm fairy tales ordered as per increasing negative word density. X = 10,000.

Analysis of Emotion Words in Books

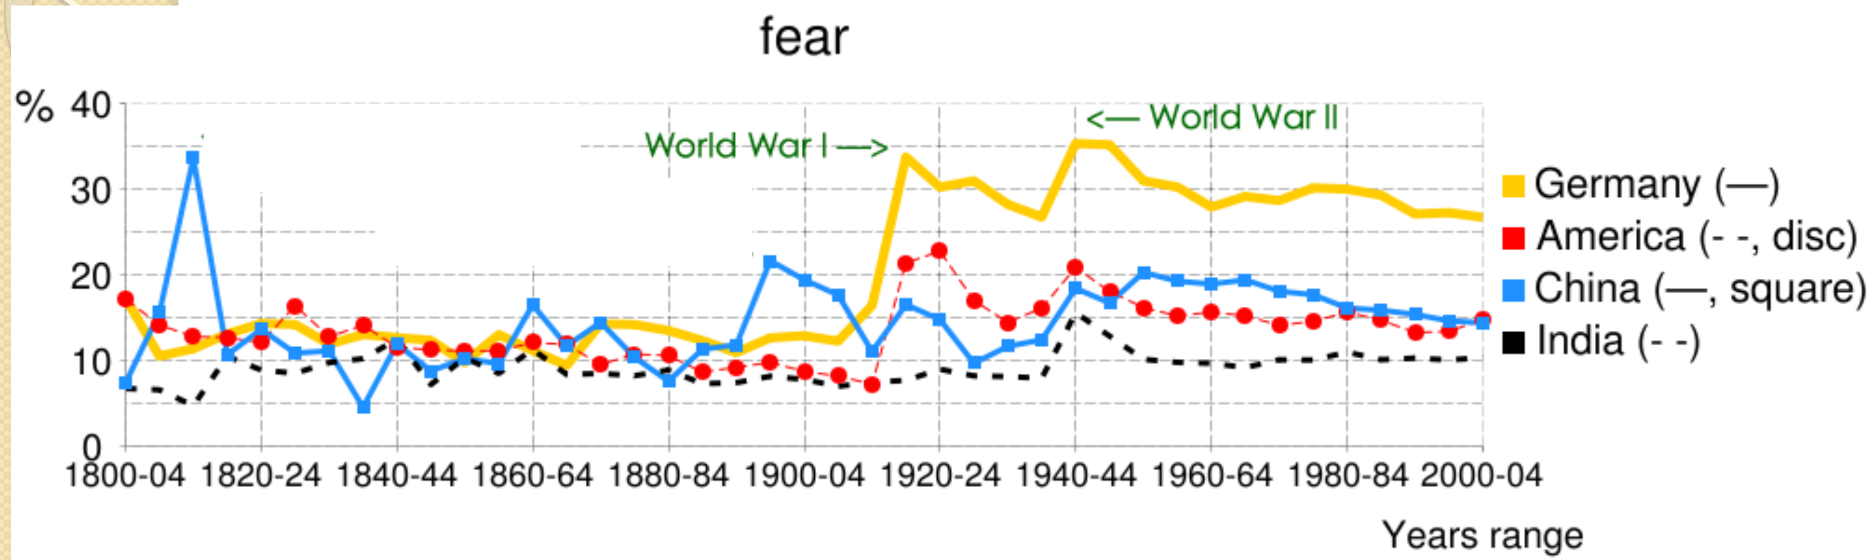

Percentage of fear words in close proximity to occurrences of America, China, Germany, and India in books.

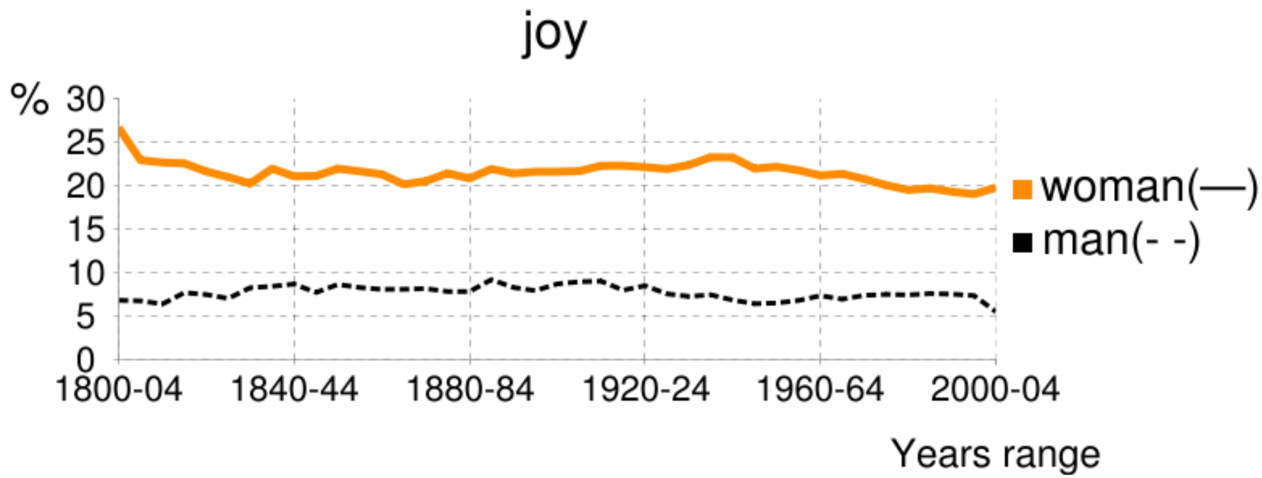

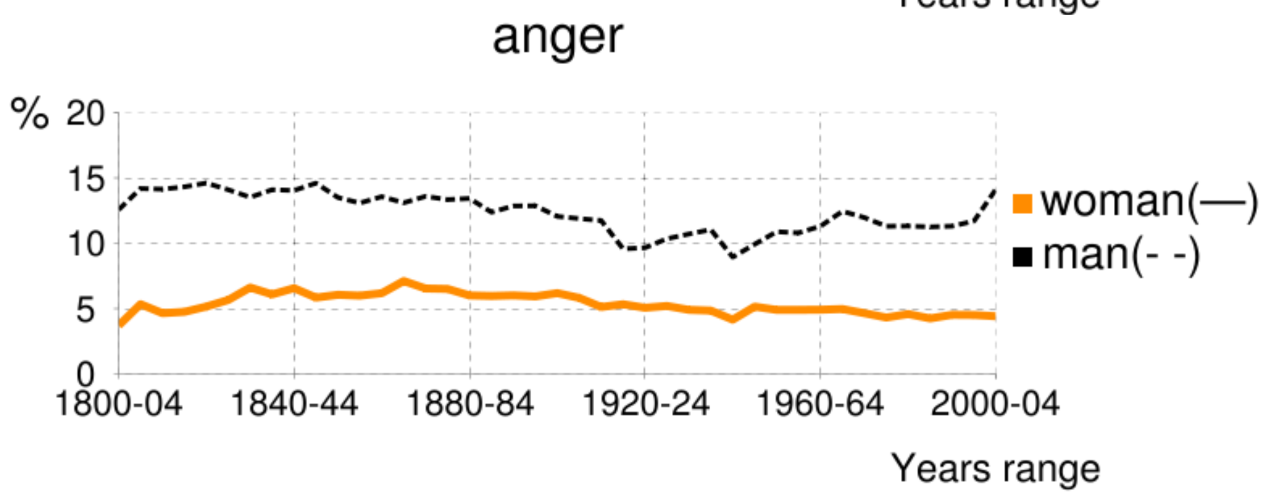

Percentage of joy and anger words in close proximity to occurrences of man and woman in books.

Words Have Associations with Colours Too

- Colourful Language: Measuring Word-Colour Associations, Saif Mohammad, In Proceedings of the ACL 2011 Workshop on Cognitive Modeling and Computational Linguistics (CMCL), June 2011, Portland, OR.
- Even the Abstract have Colour: Consensus in WordColour Associations, Saif Mohammad, In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, June 2011, Portland, OR.

Concrete Concepts

iceberg → white

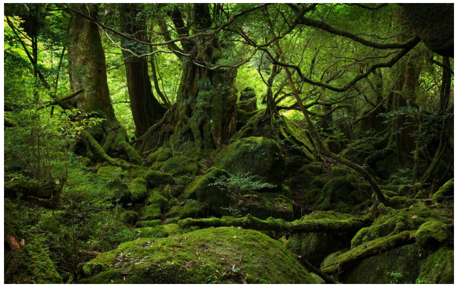

vegetation → green

Abstract Concepts

danger

→ red

honesty

→ white

Colours Add to Linguistic Information

- Strengthens the message (improves semantic coherence)
- Eases cognitive load on the receiver
- Conveys the message quickly
- Evokes the desired emotional response

Related Work

- Berlin and Kay, 1969, and later Kay and Maffi (1999)
  - If a language has only two colours: white and black.
  - If a language has three: white, black, red.
  - And so on till eleven colours.
- Berlin and Kay order:

  1. white, 2. black, 3. red, 4. green, 5. yellow, 6. blue,
  7. brown, 8. pink, 9. purple, 10. orange, 11. grey
- We used these eleven colours in our annotations.
  - Hundreds more: http://en.wikipedia.org/wiki/List_of_colors

Just the A's

Color names

Name	Hex triplet	Red	Green	Blue	Hue	Satur	Light	Satur	Value
Air Force blue	#5D8AA8	36%	54%	66%	204°	30%	51%	45%	66%
Alice blue	#F0F8FF	94%	97%	100%	208°	100%	97%	6%	100%
Alizarin crimson	#E32636	82%	10%	26%	231°	78%	46%	187%	110%
Almond	#EFDECD	94%	87%	80%	30°	52%	87%	14%	94%
Amaranth	#E52B50	90%	17%	31%	348°	78%	53%	81%	90%
Amber	#FFBF00	100%	75%	0%	45°	100%	50%	100%	100%
Amber (SAE/ECE)	#FF7E00	100%	49%	0%	30°	100%	50%	100%	100%
American rose	#FF033E	100%	1%	24%	345°	100%	51%	99%	87%
Amethyst	#9966CC	60%	40%	80%	270°	50%	60%	50%	80%
Android Green	#A4C639	64%	78%	22%	74°	55%	50%	7%	78%
Anti-flash white	#F2F3F4	95%	95%	96%	210°	8%	95%	1%	96%
Antique brass	#CD9575	80%	58%	46%	22°	47%	63%	43%	80%
Antique fuchsia	#915C83	57%	36%	51%	316°	22%	47%	37%	57%
Antique white	#FAEBD7	98%	92%	84%	34°	78%	91%	14%	98%
Ao (English)	#008000	0%	50%	0%	120°	100%	25%	100%	50%
Apple green	#8DB600	55%	71%	0%	74°	100%	36%	100%	50%
Apricot	#FBCEB1	98%	81%	69%	24°	90%	84%	29%	98%
Aqua	#00FFFF	0%	100%	100%	180°	100%	50%	100%	100%
Aquamarine	#7FFFD0	50%	100%	83%	160°	100%	75%	50%	100%
Army green	#4B5320	29%	33%	13%	69°	44%	23%	61%	33%
Arsenic	#3B444B	23%	27%	29%	206°	12%	26%	21%	29%
Arylide yellow	#E9D66B	91%	84%	42%	51°	74%	67%	54%	91%
Ash grey	#B2BEB5	70%	75%	71%	135°	9%	72%	6%	75%
Asparagus	#87A96B	53%	66%	42%	93°	27%	54%	37%	66%
Atomic tangerine	#FF9966	100%	60%	40%	20°	100%	70%	60%	100%
Auburn	#6D351A	43%	21%	10%	20°	62%	27%	76%	43%
Aureolin	#FDEE00	99%	93%	0%	56°	100%	50%	100%	99%

Crowdsourcing

- Questionnaire:

  Q. Which colour is associated with *sleep*?

  • black • green • purple…
  … (11 colour options in random order)

- No "not associated with any colour" option.

Associations with Colours

% of annotations

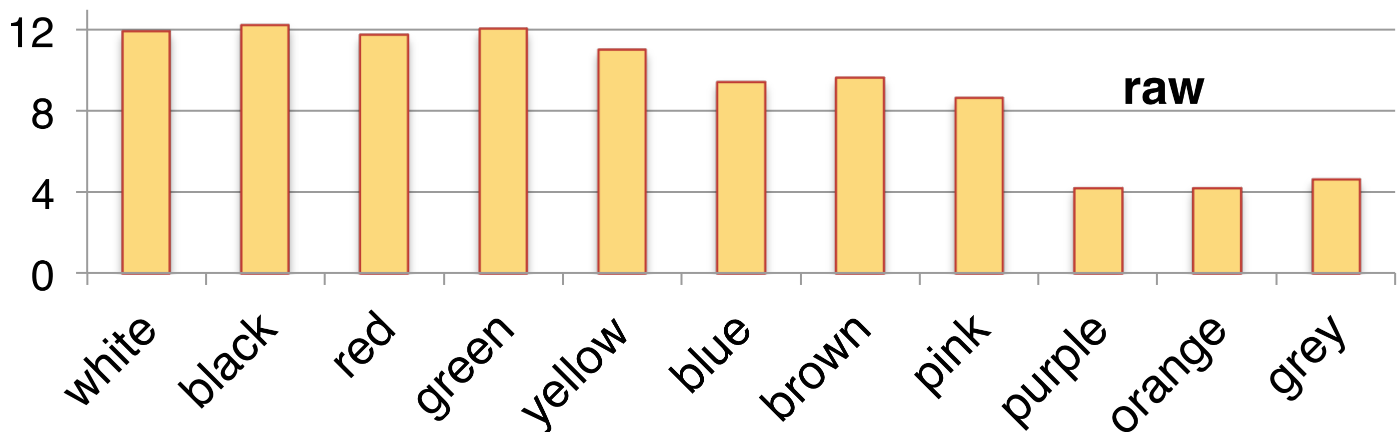

% of terms

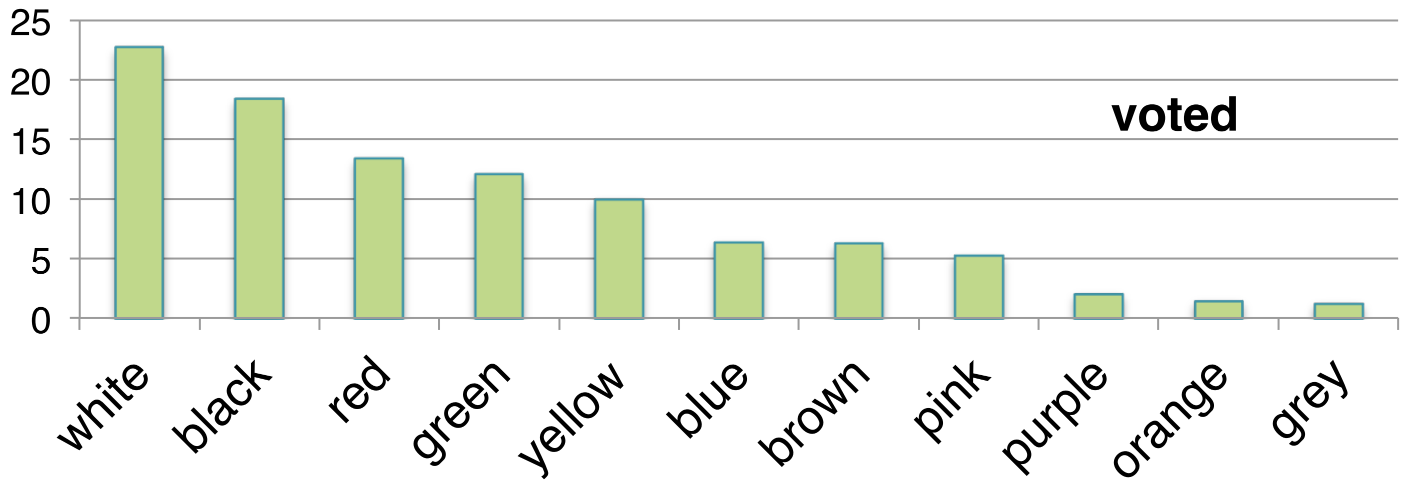

Berlin and Kay order

Obtaining Word-Sentiment and Word-Emotion Associations from Tweets

Hashtagged Tweets

- Hashtagged words are good labels of sentiments and emotions

  Can't wait to have my own Google glasses #awesome

  Some jerk just stole my photo on #tumblr. #grr **#anger**

- Hashtags are not always good labels:
  - hashtag used sarcastically

    The reviewers want me to re-annotate the data. **#joy**
  - hashtagged emotion not in the rest of the message

    Mika used my photo on tumblr. **#anger**

#Emotional Tweets, Saif Mohammad, In Proceedings of the First Joint Conference on Lexical and Computational Semantics (*Sem), June 2012, Montreal, Canada.

Automatically Generated Sentiment Lexicons

- Lists of word--sentiment pairs, with scores indicating the degree of association

spectacular positive 0.91	spectacular 0.91
okay positive 0.3	okay 0.3
lousy negative 0.84	lousy -0.84
unpredictable negative 0.17	unpredictable -0.17

Automatically Generated New Sentiment Lexicons

- Created a list of seed sentiment words by looking up synonyms of *excellent, good, bad*, and *terrible*:
  - 30 positive words
  - 46 negative words
- Polled the Twitter API for tweets with seed-word hashtags
  - A set of 775,000 tweets was compiled from April to December 2012

Automatically Generated New Lexicons

- A tweet is considered:
  - positive if it has a positive hashtag
  - negative if it has a negative hashtag

- For every word w in the set of 775,000 tweets, an association score is generated:

$$score(w) = PMI(w,positive) - PMI(w,negative)$$

PMI = pointwise mutual information

If $score(w) > 0$, then w is positive

If $score(w) < 0$, then w word is negative

NRC Hashtag Sentiment Lexicon

- w can be:
  - any unigram in the tweets: 54,129 entries
  - any bigram in the tweets: 316,531 entries
  - non-contiguous pairs (any two words) from the same tweet: 308,808 entries

- Multi-word entries incorporate context:

  unpredictable story 0.4

  unpredictable steering -0.7

Generating lexicon for 500 emotions

- Collected hundreds of thousands of tweets with #joy, #jealousy, #frustrated, etc
- Determined words that tend to co-occur with each of these emotion-word hashtags.
  - NRC Hashtag Emotion Lexicon: About 5000 words associated with about 500 emotions

SemEval-2013, Task 2

Svetlana Kiritchenko, NRC

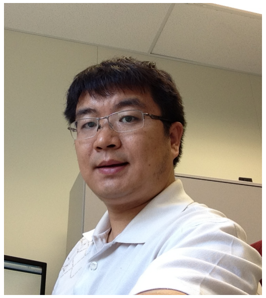

Xiaodan Zhu, NRC

- Is a given message positive, negative, or neutral?
  - tweet or SMS
- Is a given term within a message positive, negative, or neutral?

International competition on sentiment analysis of tweets:

- SemEval-2013 (co-located with NAACL-2013)
- 44 teams

NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets, Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu, In Proceedings of the seventh international workshop on Semantic Evaluation Exercises (SemEval-2013), June 2013, Atlanta, USA.

Setup

- **Pre-processing:**
  - URL -> http://someurl
  - UserID -> @someuser
  - Tokenization and part-of-speech (POS) tagging (CMU Twitter NLP tool)
- **Classifier:**
  - SVM with linear kernel
- **Evaluation:**
  - Macro-averaged F-pos and F-neg

Features

Features	Examples
sentiment lexicon	#positive: 3, scorePositive: 2.2; maxPositive: 1.3; last: 0.6, scoreNegative: 0.8, scorePositive_neg: 0.4
word n-grams	spectacular, like documentary
char n-grams	spect, docu, visua
part of speech	#N: 5, #V: 2, #A:1
negation	#Neg: 1; ngram:perfect → ngram:perfect_neg, polarity:positive → polarity:positive_neg
word clusters	probably, definitely, def
all-caps	YES, COOL
punctuation	#!+: 1, #?+: 0, #!?+: 0
word clusters	probably, definitely, probly
emoticons	:D, >:(
elongated words	soooo, yaayyy

Sentiment Analysis Competition

Classify Tweets

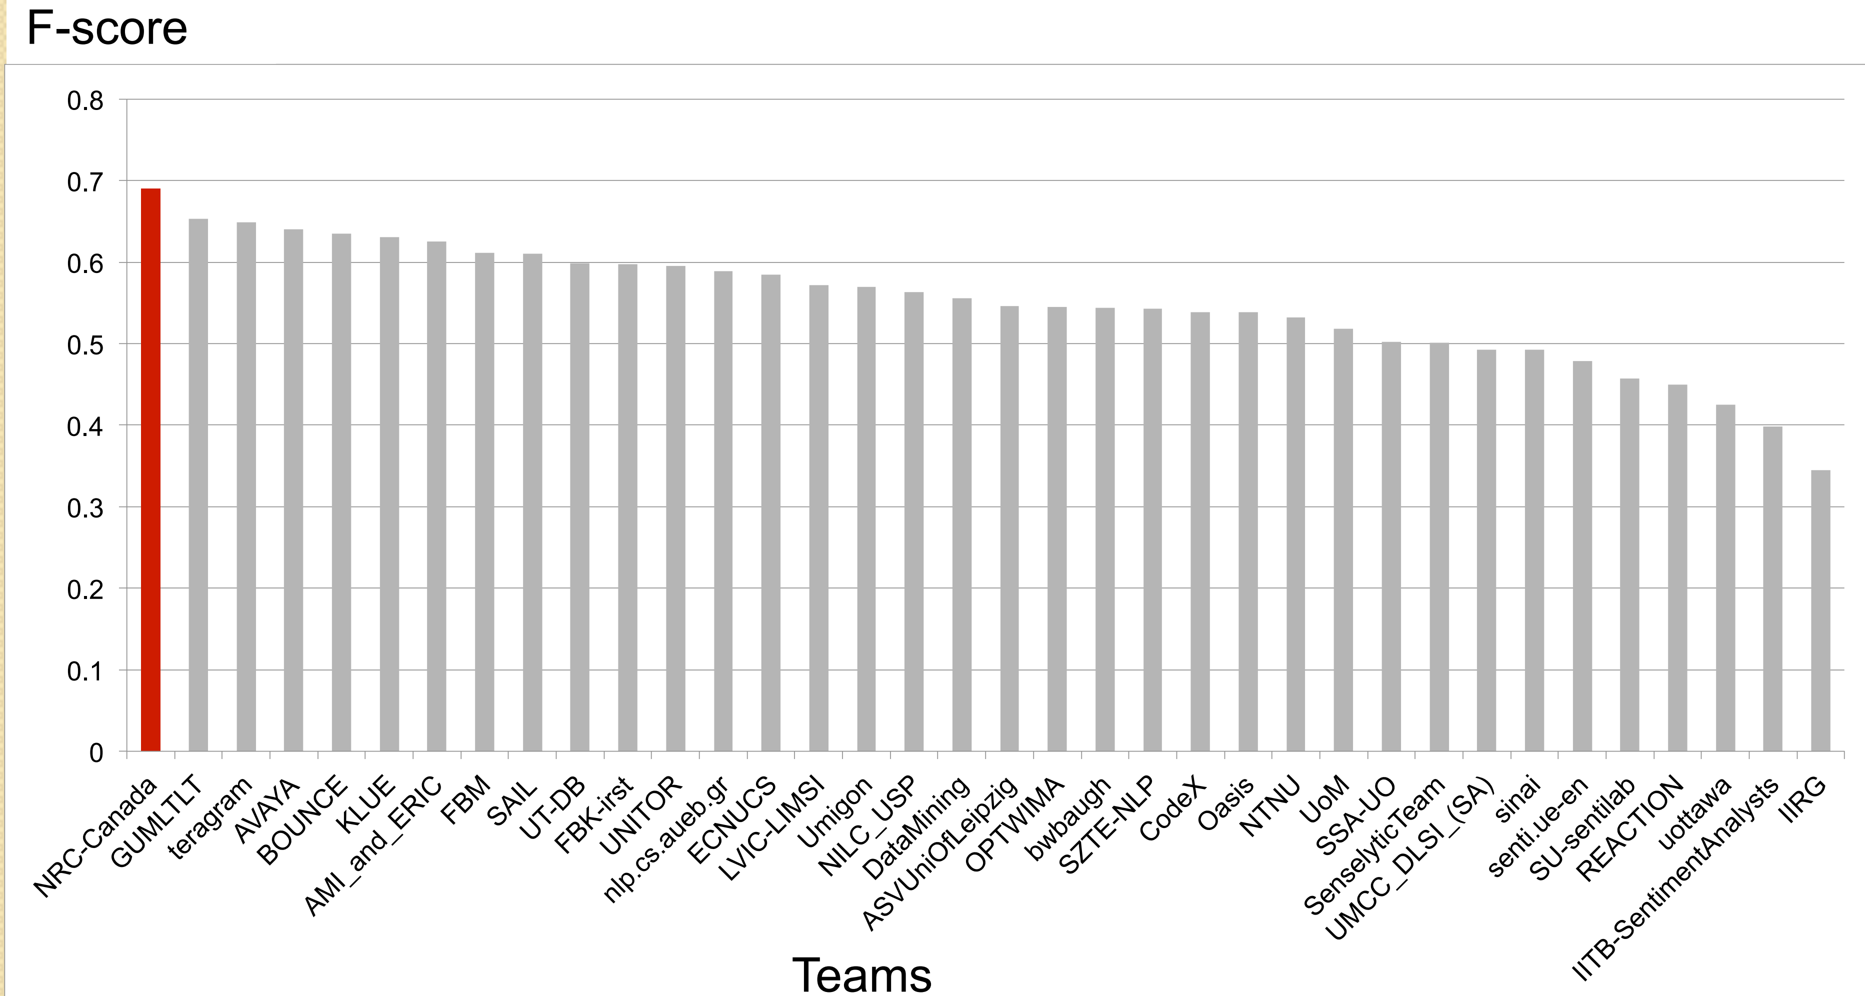

Sentiment Analysis Competition
Classify SMS

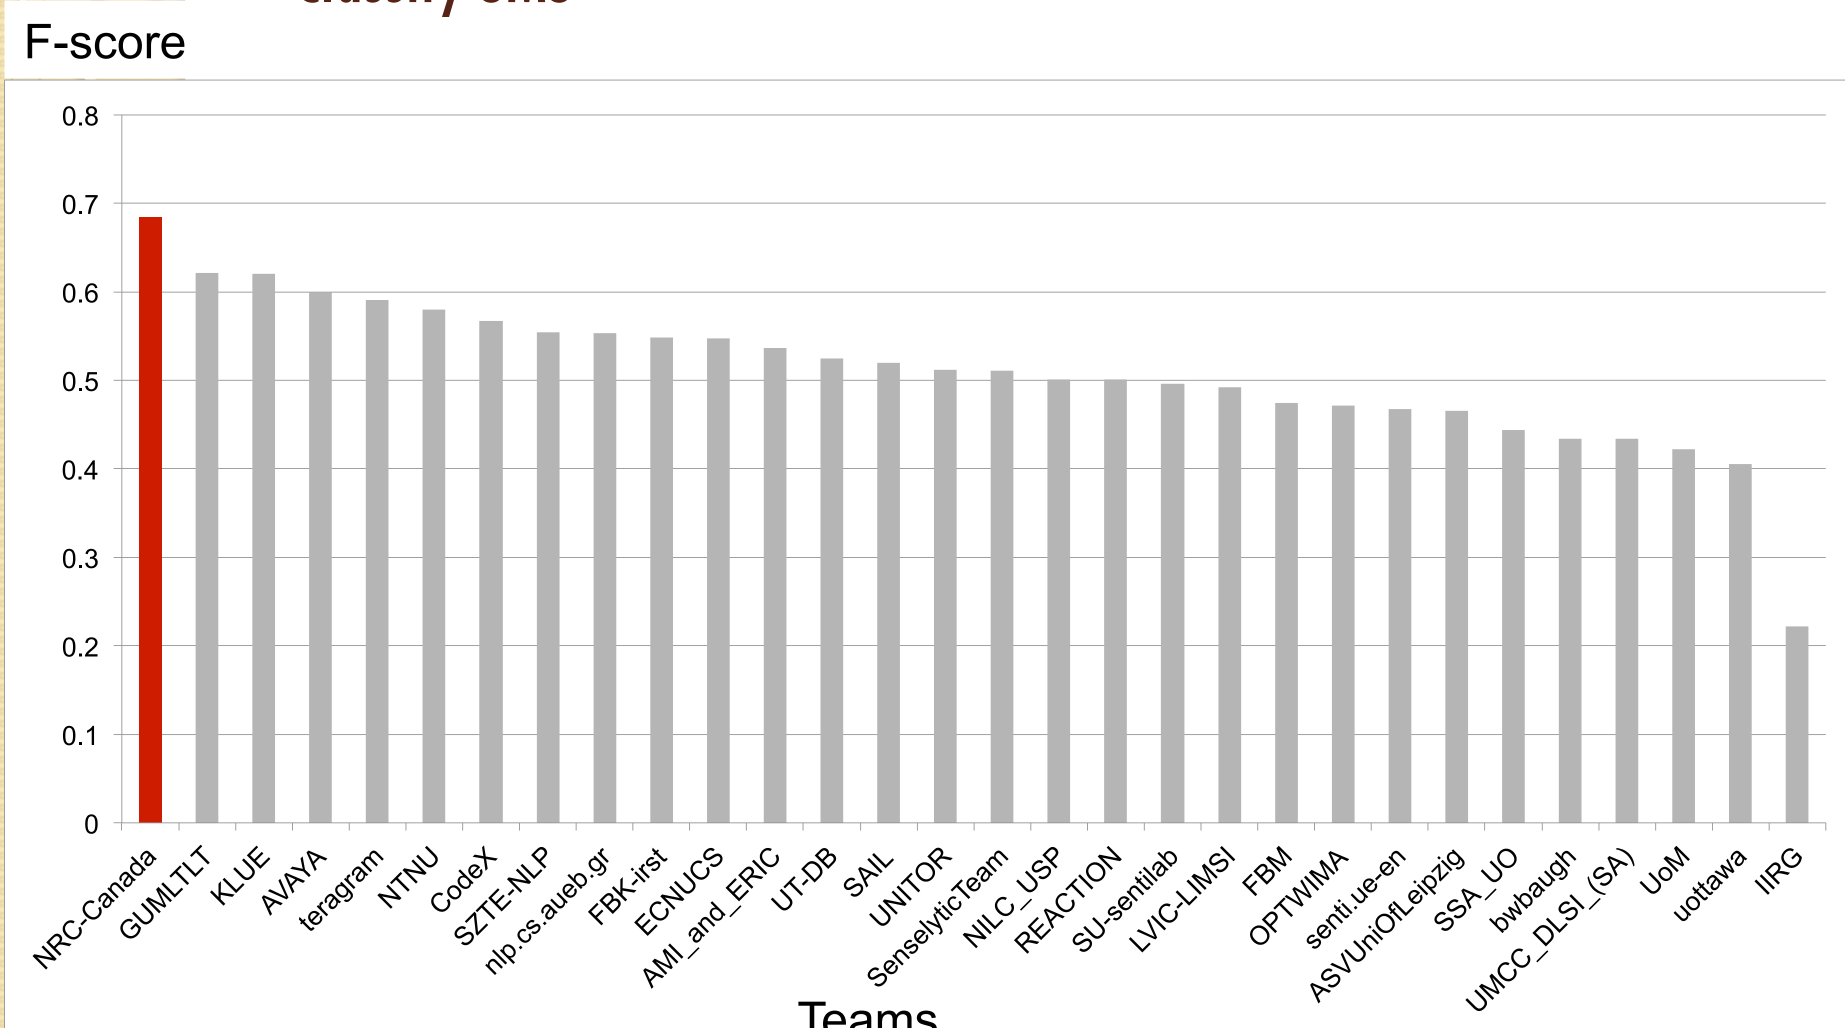

Sentiment Analysis Competition
Classify expression in Tweet

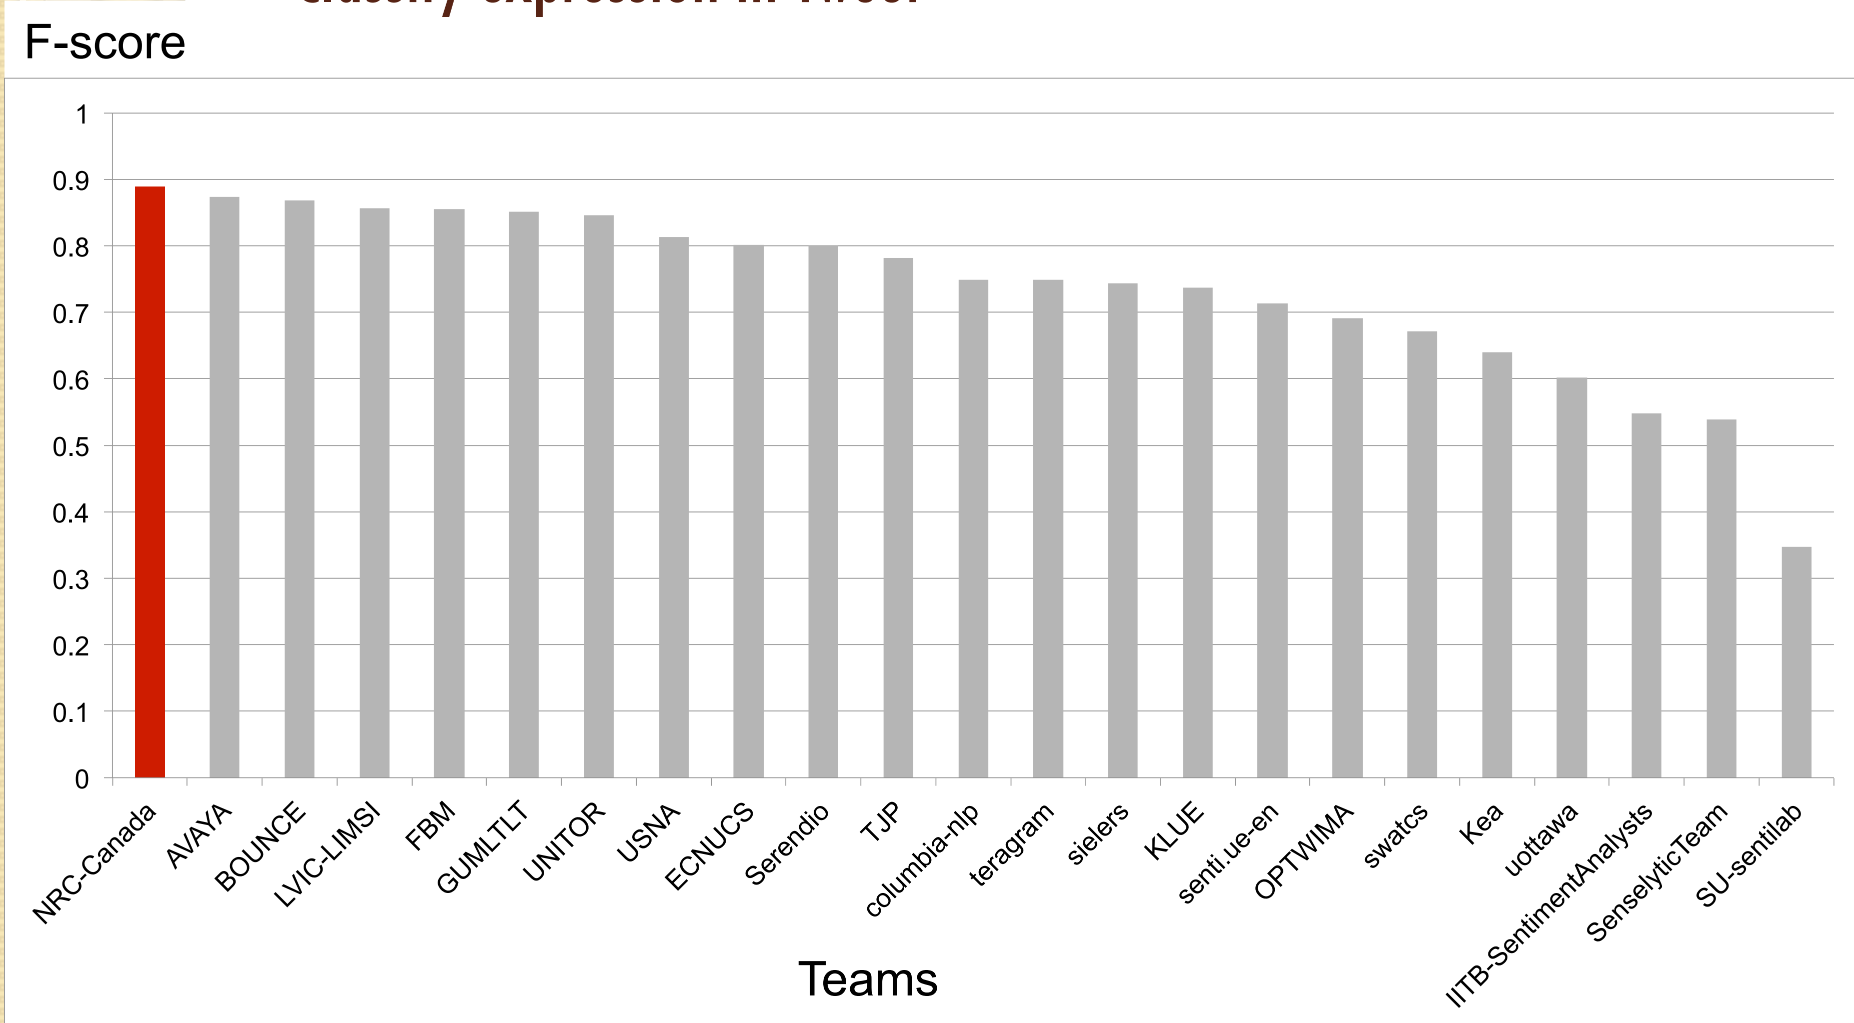

Sentiment Analysis Competition
Classify expression in SMS

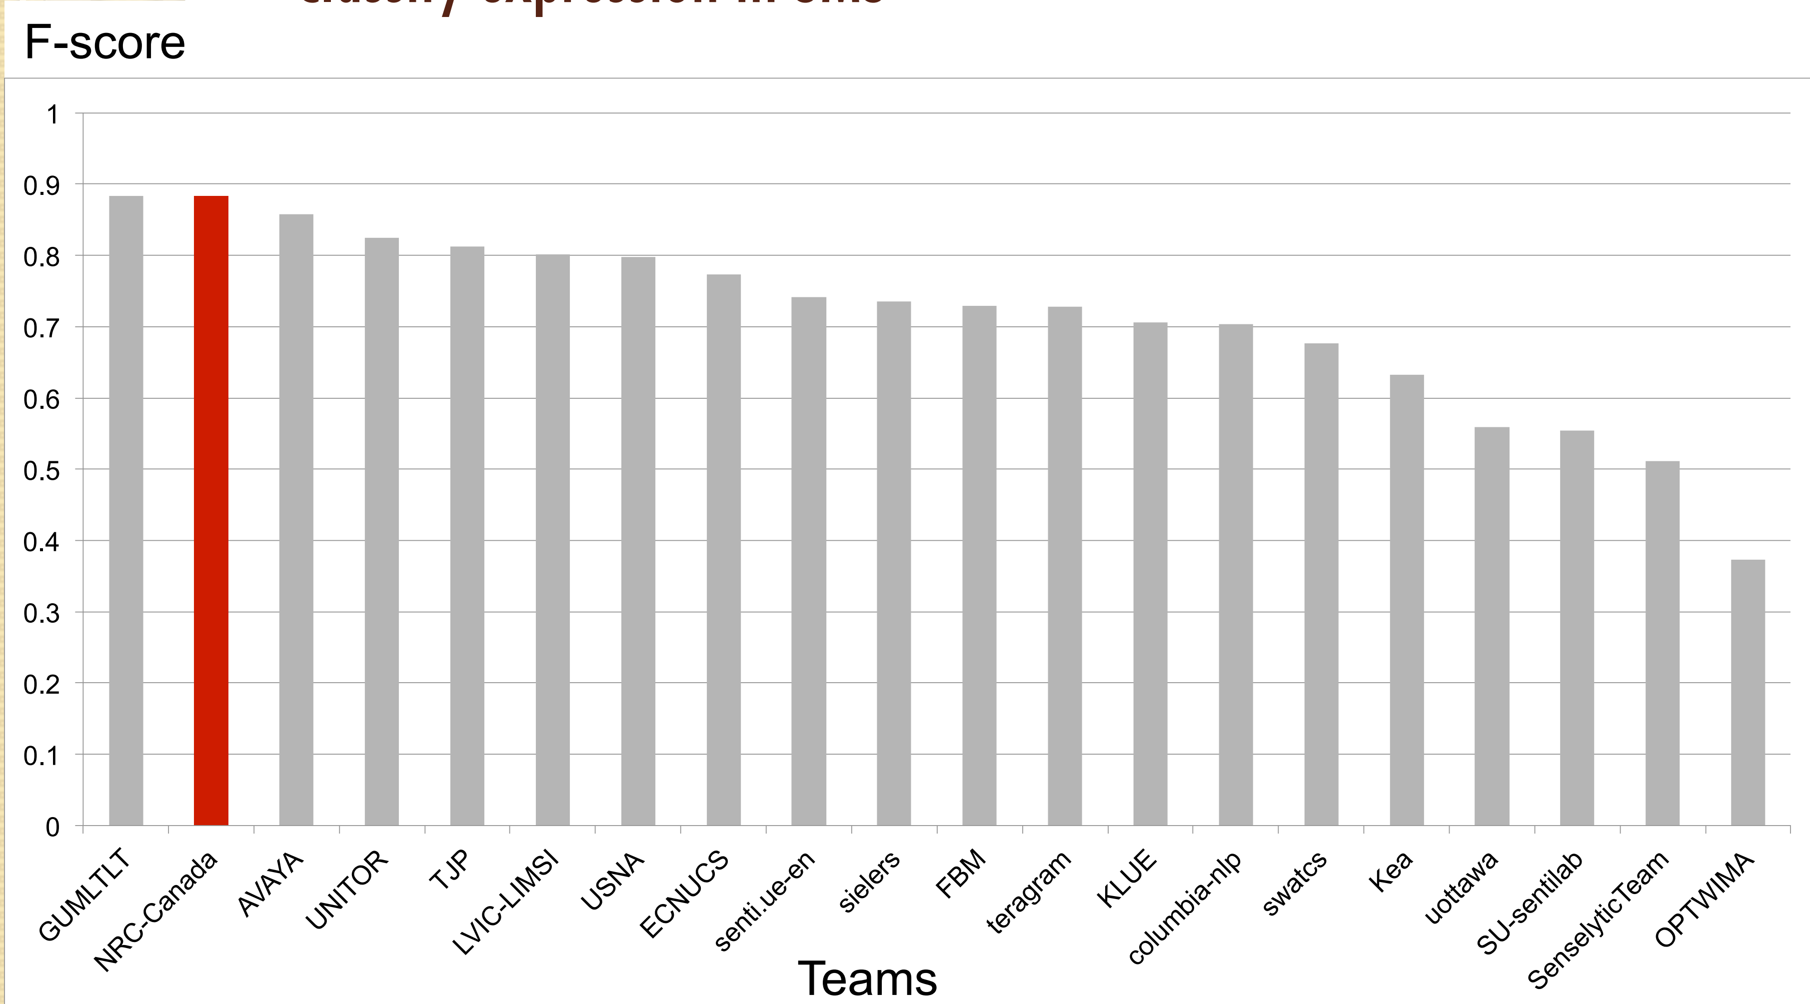

Feature Contributions (on Tweets)

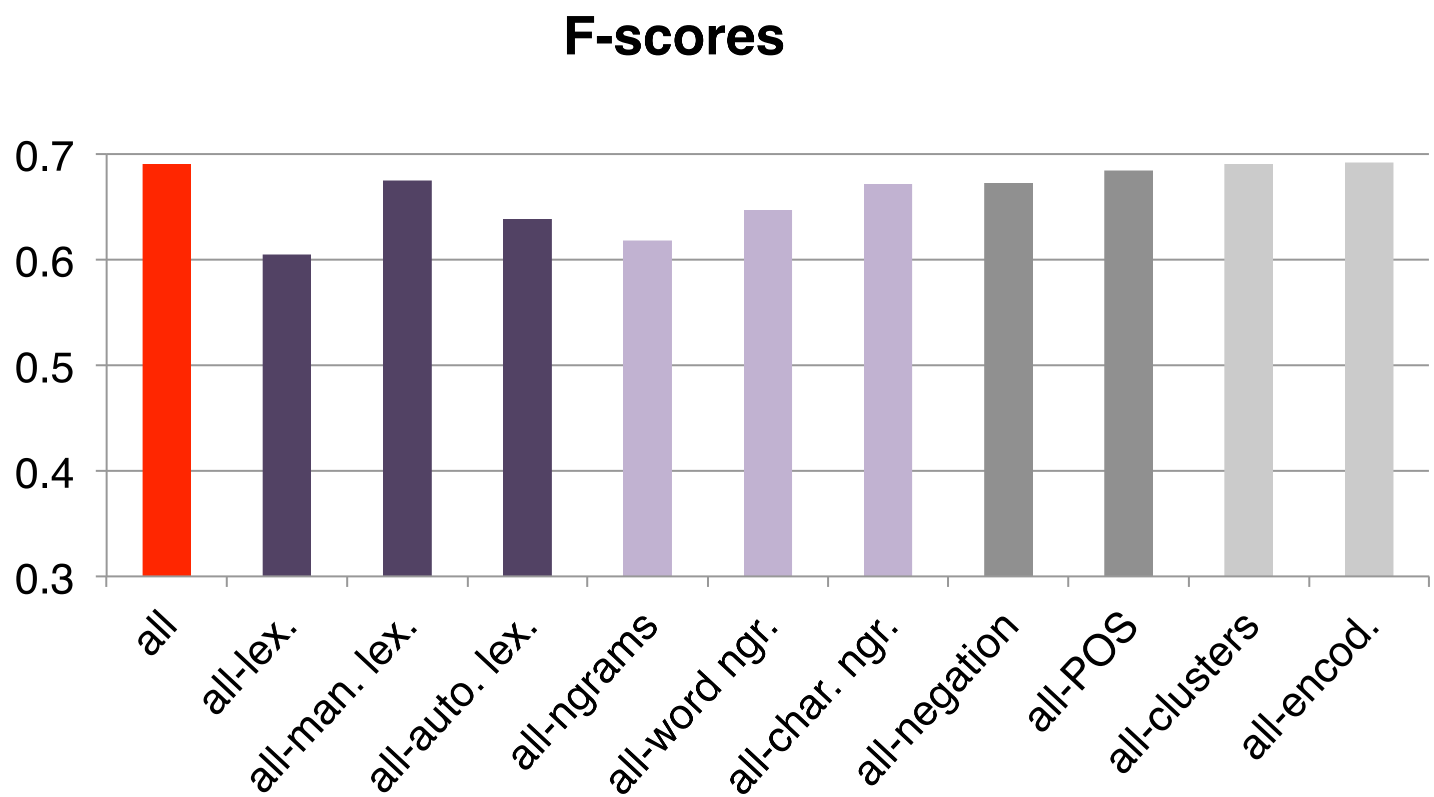

Do we really need to model hundreds of emotions?

- Is modeling sentiment not enough?
- Is modeling basic emotions not enough?

Personality Detection

- Detecting personality from written text
  - Essays, email, blogs, tweets
- Applications
  - Health, marketing, intelligence
- Normally personality determined through questionnaires
  - Various personality models proposed

Wanted to explore the relationship between emotions and personality.

Big 5 Personality Model

- extroversion vs. introversion
  - *sociable, assertive vs. aloof, shy*
- neuroticism vs. emotional stability
  - *insecure, anxious vs. calm, unemotional*
- agreeability vs. disagreeability
  - *friendly, co-operative vs. antagonistic, fault-finding*
- conscientiousness vs. unconscientiousness
  - *self-disciplined, organized vs. inefficient, careless*
- openness to experience vs. conventionality
  - *intellectual, insightful vs. shallow, unimaginative*

Personality Detection

- Data:
  - Essays
  - Facebook posts
- Features:
  - baseline: surface form and lexical categories (Mairesse et al. 2007)
  - added ngrams
    - no improvement
  - added fine emotion categories from the NRC Hashtag Emotion Lexicon
    - significant improvement

Using Nuances of Emotion to Identify Personality, Saif M. Mohammad and Svetlana Kiritchenko, In Proceedings of the ICWSM Workshop on Computational Personality Recognition, July 2013, Boston, USA.

Summary

- Created large word-emotion and word-sentiment association lexicons for sentiment analysis
  - Crowdsourcing to Mechanical Turk
  - Automatically from hashtagged tweets
- Crowdsourcing method
  - Involved direct manual annotation
  - Used word-choice question for quality control and to guide annotators to desired sense
- Hashtag method
  - Showed scalability to hundreds of affect categories
  - Incorporated context through multi-word entries
  - Provided competitive edge in the SemEval-2013 sentiment analysis competition

Summary (continued)

- Tracked emotions in email, novels, and Google Books corpus.
- Showed that modeling hundreds of fine emotion categories significantly improves results in personality detection.
- Found that when colors are listed by frequency of associations with words, they obey the Berlin-Kay order.

Questions!