Affect Associations

The Building Blocks of Sentiment Analysis

Saif Mohammad
National Research Council Canada

Word Associations

Beyond literal meaning, words have other associations that add to their meanings.

- Associations with sentiment
- Associations with emotions
- Associations with social overtones
- Associations with cultural implications
- Associations with colours

Connotations.

Word-Sentiment Associations

- Adjectives
  - reliable and stunning are typically associated with positive sentiment
  - rude, and broken are typically associated with negative sentiment
- Nouns and verbs
  - holiday and smiling are typically associated positive sentiment
  - death and crying are typically associated with negative sentiment

Goal: Capture word-sentiment associations.

Word-Emotion Associations

Words have associations with emotions:

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

Goal: Capture word-emotion associations.

Sentiment Analysis

- Is a given piece of text positive, negative, or neutral?
  - The text may be a sentence, a tweet, an SMS message, a customer review, a document, and so on.

Sentiment Analysis

- Is a given piece of text positive, negative, or neutral?
  - The text may be a sentence, a tweet, an SMS message, a customer review, a document, and so on.

Emotion Analysis

- What emotion is being expressed in a given piece of text?
  - Example emotions: joy, sadness, fear, anger, 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 happiness and well-being; early detection of cyber bullying
- Measuring the impact of activist movements through text generated in social media.
- Improving automatic dialogue systems
- Improving automatic tutoring systems
- Detecting how people use emotion-bearing-words and metaphors to persuade and coerce others
- Developing affect-sensitive characters in computer games

Uses of Word-Affect Associations

- Sentence-, tweet-, message-level affect classification (aka traditional sentiment analysis)
- Tracking the distribution of affect words in text
  - tracking sentiment towards products, people, issues, etc.
  - literary analysis
- Linguistic studies
  - how words are used to convey affect
  - understanding sentiment composition
- Information visualization
  - digital humanities

This talk

- methods for capturing word-affect associations
- applications of word-affect associations

Which Emotions?

Charles Darwin

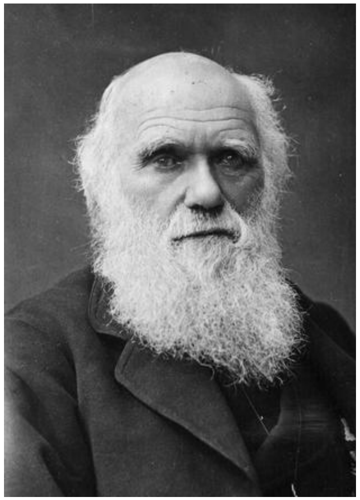

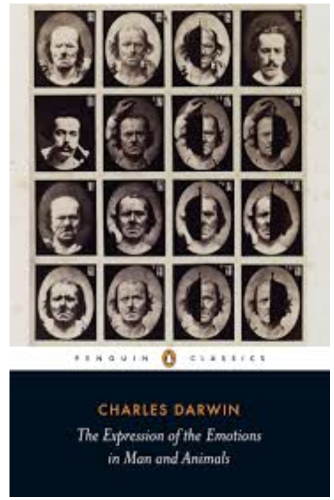

FIG. 20.—Terror,
from a photograph by Dr. Duchenne.

- published *The Expression of the Emotions in Man and Animals* in 1872
- seeks to trace the animal origins of human characteristics
  - pursing of the lips in concentration
  - tightening of the muscles around the eyes in anger
- claimed that certain facial expressions are universal
  - these facial expressions are associated with emotions

Debate: Universality of Perception of Emotions

Margaret Mead
Cultural anthropologist

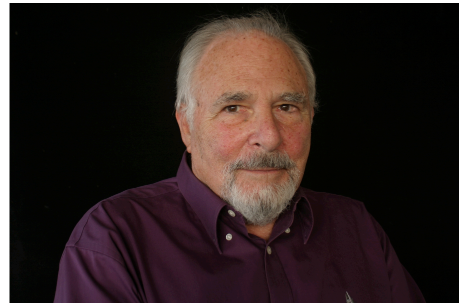

Paul Ekman
Psychologist and discoverer of micro expressions.

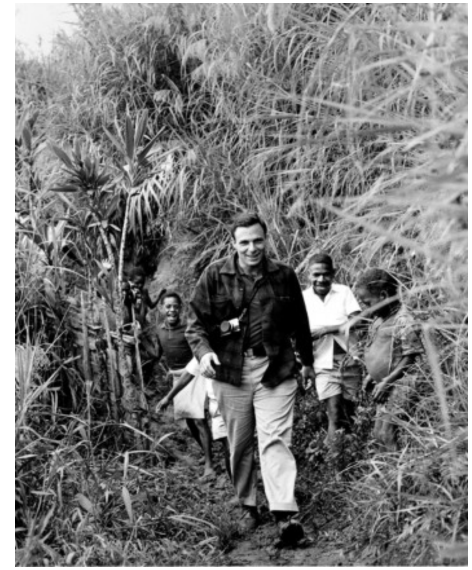

- Circa 1950's, Margaret Mead and others believed facial expressions and their meanings were culturally determined
  - behavioural learning processes
- Paul Ekman provided the strongest evidence to date that Darwin, not Margaret Mead, was correct in claiming facial expressions are universal
- Found universality of six emotions

Paul Ekman, 1971: Six Basic Emotions

- Anger
- Disgust
- Fear
- Joy
- Sadness
- Surprise

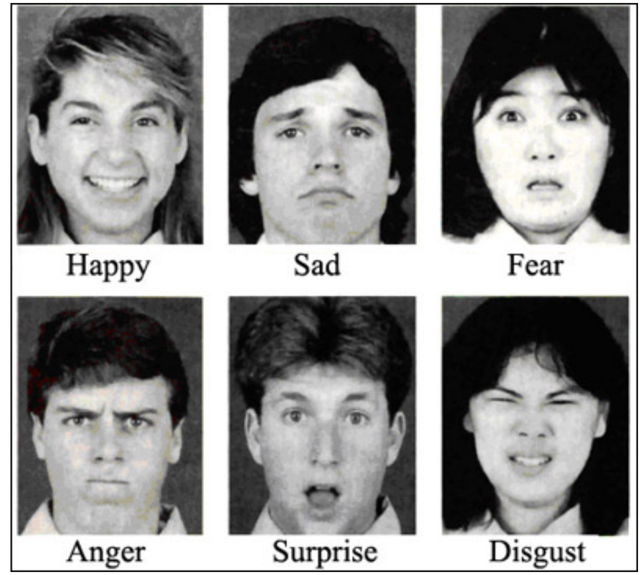

Plutchik, 1980: Eight Basic Emotions

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

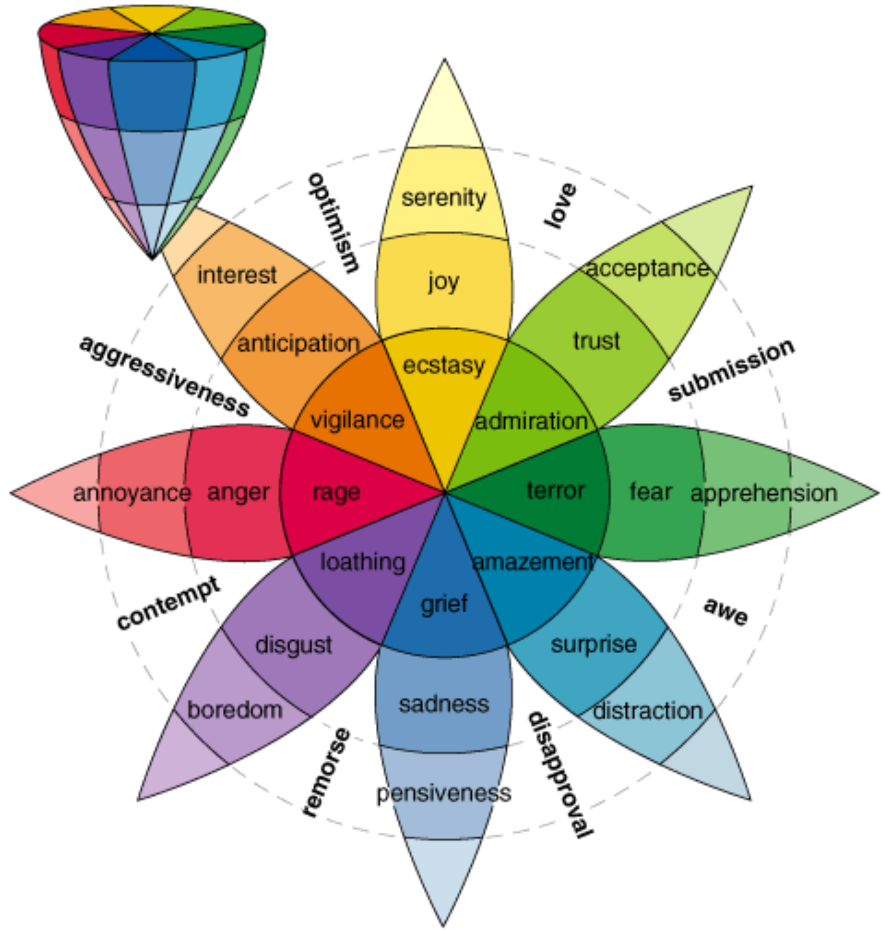

Goal: We chose to capture word-emotion associations for the 8 Plutchik emotions.

Annotations by 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.

Word-Choice Question

Peter Turney, AI2

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

- *car*
- *tree*
- *tears*
- *olive*

- Generated automatically
  - Near-synonym taken from thesaurus
  - Distractors are randomly chosen
- Guides Turkers to desired sense
- Aids 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 when asked ‘associated with’ rather than ‘evoke’.

Emotion Lexicon

- NRC Emotion Lexicon
  - Associations of ~14,000 words with 8 emotions and 2 polarities

Available at: www.saifmohammad.com

Paper:

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

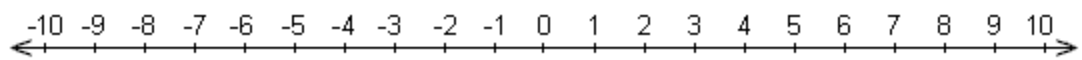

What if you want to capture fine-grained intensity of emotion/sentiment?

How to manually create sentiment lexicons with intensity values?

- Humans are not good at giving real-valued scores?
  - hard to be consistent across multiple annotations
  - difficult to maintain consistency across annotators
    - 0.8 for annotator may be 0.7 for another

Comparative Annotations

Paired Comparisons (Thurstone, 1927; David, 1963):

If X is the property of interest (positive, useful, etc.), give two terms and ask which is more X

- less cognitive load
- helps with consistency issues
- requires a large number of annotations
  - order N^2, where N is number of terms to be annotated

Comparative Annotations

Paired Comparisons (Thurstone, 1927; David, 1963)**:**

If X is the property of interest (positive, useful, etc.), give two terms and ask which is more X

Need a method that preserves the comparison aspect, without greatly increasing the number of annotations needed.

Possible solution:

Best–Worst Scaling (Louviere & Woodworth, 1990)**:**
(a.k.a. Maximum Difference Scaling or MaxDiff)

Best–Worst Scaling (BWS)

aka Maximum Difference Scaling (MaxDiff)

- The annotator is presented with four words (say, A, B, C, and D) and asked:
  - which word is the most positive (least negative)
  - which is the least positive (most negative)

- By answering just these two questions, five out of the six inequalities are known
  - For e.g.:
    - If A is most positive
    - and D is least positive, then we know:

$$A > B, A > C, A > D, B > D, C > D$$

Example BWS Annotation Instance

Focus words:

1. worse 2. was not sufficient 3. more afraid 4. banish

Q1. Identify the word that is associated with the MOST amount of POSITIVE sentiment (or, least amount of negative sentiment) -- the most positive term.

- ◯ worse
- ◯ was not sufficient
- ◯ more afraid
- ◯ banish

Q2. Identify the word that is associated with the MOST amount of NEGATIVE sentiment (or, least amount of positive sentiment) -- the most negative term.

- ◯ worse
- ◯ was not sufficient
- ◯ more afraid
- ◯ banish

Best–Worst Scaling

- Each of these BWS questions can be presented to multiple annotators.
- The responses to the BWS questions can then be easily translated into:
  - a real-valued score for all the terms (Orme, 2009)

    score(w) = (#mostPositive(w) - #mostNegative(w)) / #annotations(w)

    the scores range from:
    -1 (least association with positive sentiment)
    to 1 (most association with positive sentiment)

  - the scores can then be used to rank of all the terms

Comparative Annotations

Paired Comparisons (Thurstone, 1927; David, 1963)**:**

If X is the property of interest (positive, useful, etc.), give two terms and ask which is more X

Best–Worst Scaling (Louviere & Woodworth, 1990)**:**
(a.k.a. Maximum Difference Scaling or MaxDiff)

Give k terms and ask which is most X, and which is least X
(*k is usually 4 or 5*)

- preserves the comparative nature
- keeps the number of annotations down to about 2N
- leads to more reliable annotations
  - less biased and more discriminating (Cohen, 2003)

Best-Worst Scaling
Sentiment Lexicons

Svetlana Kiritchenko
NRC

Lexicon	Language	Domain
1. SemEval-2015 English Twitter Sentiment Lexicon	English	Twitter
2. SemEval-2016 Arabic Twitter Sentiment Lexicon	Arabic	Twitter
3. Sentiment Composition Lexicon for Negators, Modals, and Adverbs (SCL-NMA)	English	General
4. Sentiment Composition Lexicon for Opposing Polarity Phrases (SCL-OPP)	English	General

Lexicons and papers available at:
http://saifmohammad.com/WebPages/SCL.html

Robustness of the Annotations

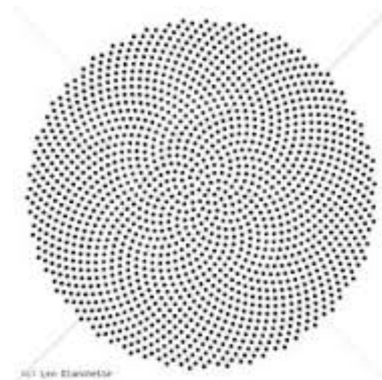

BWS annotations

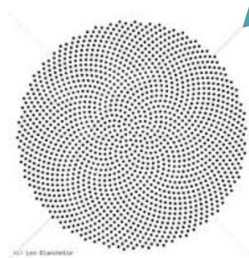

half of the BWS annotations

↓

ranking of words by sentiment

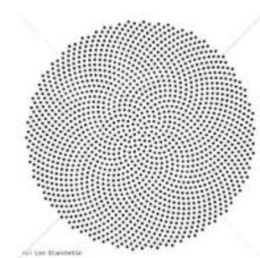

half of the BWS annotations

↓

ranking of words by sentiment

Robustness of the Annotations

The two rankings were very similar:

- average difference in scores was 0.04
- Spearman Rank Correlation coefficient between the two rankings was about 0.98 for all four lexicons

Least Perceptible Difference

- Least perceptible difference aka just-noticeable difference
  - a concept from psychophysics
  - the amount by which something that can be measured (e.g., weight or sound intensity) needs to be changed in order for the difference to be noticeable by a human (Fechner, 1966)

- With our BWS annotations, which are comparative in nature, we can measure the least perceptible difference in sentiment
  - 0.08 (~4% in range -1..1)
  - useful in studying sentiment composition
    - e.g., to determine whether a modifier significantly impacts the sentiment of the word it modifies

Two of the lexicons we created were
for Twitter terms.

SemEval-2015 English Twitter Sentiment Lexicon

- Includes fine-grained sentiment associations for 1,515 terms from tweets:
  - regular English words: peace, jumpy
  - tweet-specific terms
    - hashtags and conjoined words: #inspiring, #needsleep
    - misspellings: appriciate
    - creative spellings:раssword gooooood, cant w8
    - abbreviations: smfh, lol
    - emoticons: :'(, <33
  - negated terms: not nice, nothing better, can't wait

SemEval-2016 Arabic Twitter Sentiment Lexicon: 1367 terms

Examples of sentiment scores

Term	Sentiment Score -1 (most negative) to 1 (most positive)
awesomeness	0.827
#happygirl	0.625
cant waitttt	0.601
don't worry	0.152
not true	-0.226
cold	-0.450
#getagrip	-0.587
#sickening	-0.722

Two of the lexicons we created were **sentiment composition lexicons.**

Sentiment Composition Lexicon

Sentiment composition lexicon (SCL): a list of phrases and their constituent words with association to positive (negative) sentiment

would not be happy	-0.6
happy	0.9

These lexicons are useful for studying sentiment composition.

Sentiment Composition Lexicon
for Negators, Modals, and Adverbs (SCL-NMA)

- SCL-NMA provides fine-grained sentiment associations for 3207 terms:
  - phrases involving negators (e.g., did not harm)
  - phrases involving modal verbs (e.g., should be better)
  - phrases involving degree adverbs (e.g., certainly agree)
  - phrases involving combinations (e.g., would be very easy)
  - their constituent content words (e.g., harm, better, agree, easy)

- Use SCL-NMA to help understand how modifiers (negators, modal verbs, degree adverbs) affect sentiment in phrases

Sentiment Composition Lexicon
for Negators, Modals, and Adverbs (SCL-NMA)

On combination with a negator:

- positive words become negative but sentiment intensity is reduced
- some negative words become positive, but many just become less negative

On combination with a modal:

- intensity of sentiment is reduced

LiveSlides web content

To view

Download the add-in.

liveslides.com/download

Start the presentation.

Sentiment Composition Lexicon for Opposing Polarity Phrases

a.k.a. SemEval-2016 English Twitter Mixed Polarity Lexicon

- Sentiment composition is easy when both terms have same polarity, or if one or both terms are neutral
  - we wanted to create a dataset particularly challenging for determining sentiment composition
- Opposing Polarity Phrase (OPP): includes at least one positive word and at least one negative word
- Lexicon includes 1,661 English terms:
  - 851 OPP bigrams and trigrams: happy accident, guilty pleasures, best winter break
  - 810 unigrams that are part of the selected ngrams: happy, accident, winter

LiveSlides web content

To view

Download the add-in.

liveslides.com/download

Start the presentation.

Papers:

- Capturing Reliable Fine-Grained Sentiment Associations by Crowdsourcing and Best-Worst Scaling. Svetlana Kiritchenko and Saif M. Mohammad. In Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. June 2016. San Diego, CA.
- Sentiment Composition of Words with Opposing Polarities. Svetlana Kiritchenko and Saif M. Mohammad. In Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. June 2016. San Diego, CA.
- The Effect of Negators, Modals, and Degree Adverbs on Sentiment Composition. Svetlana Kiritchenko and Saif M. Mohammad, In Proceedings of the NAACL 2016 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA), June 2014, San Diego, California.
- Semeval-2016 Task 7: Determining Sentiment Intensity of English and Arabic Phrases. Svetlana Kiritchenko, Saif M. Mohammad, and Mohammad Salameh. In Proceedings of the International Workshop on Semantic Evaluation (SemEval '16). June 2016. San Diego, California.

Visualizing Emotions in Text

SHAPES OF STORIES

Tracking Emotions in Stories

- Can we automatically track the emotions of characters?
- Can we track the change in distribution of emotion words?
- Are there some canonical shapes common to most stories?

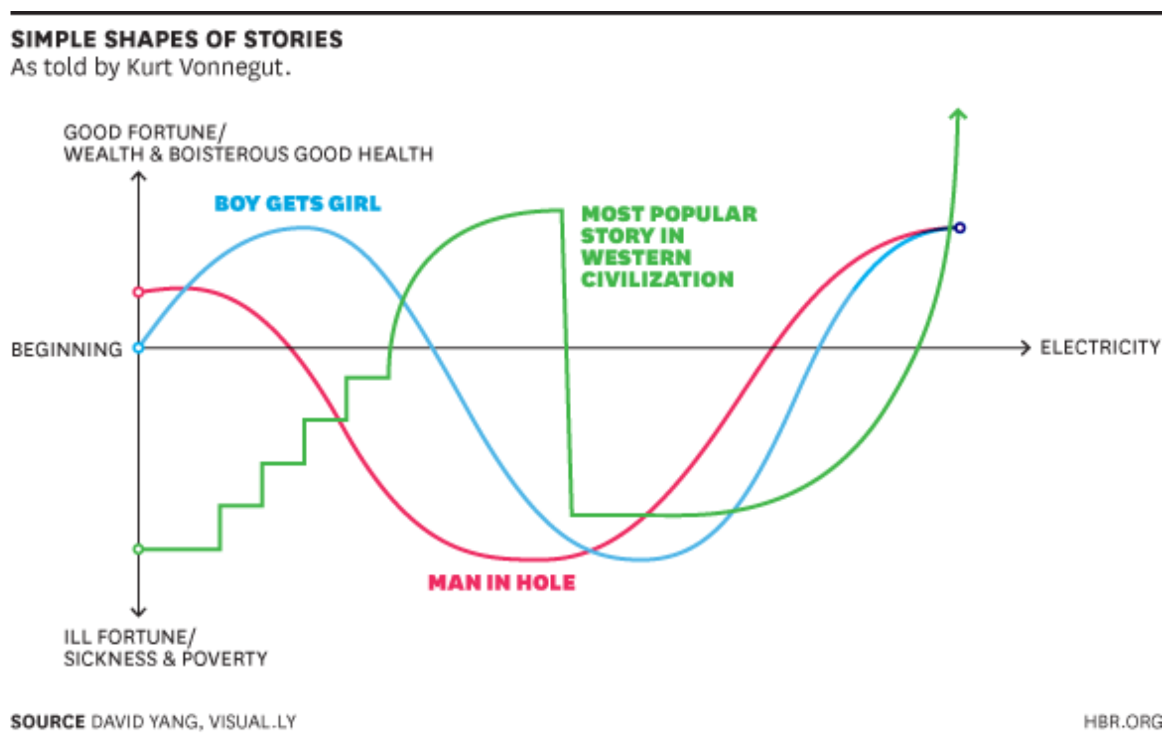

LiveSlides web content

To view

Download the add-in.

liveslides.com/download

Start the presentation.

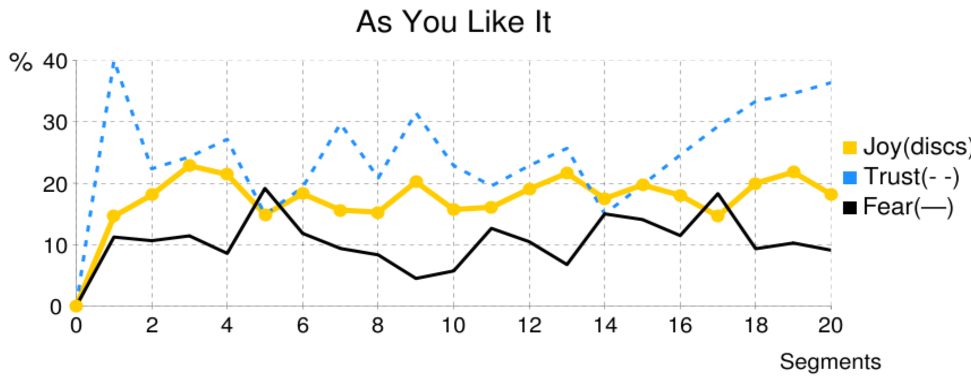

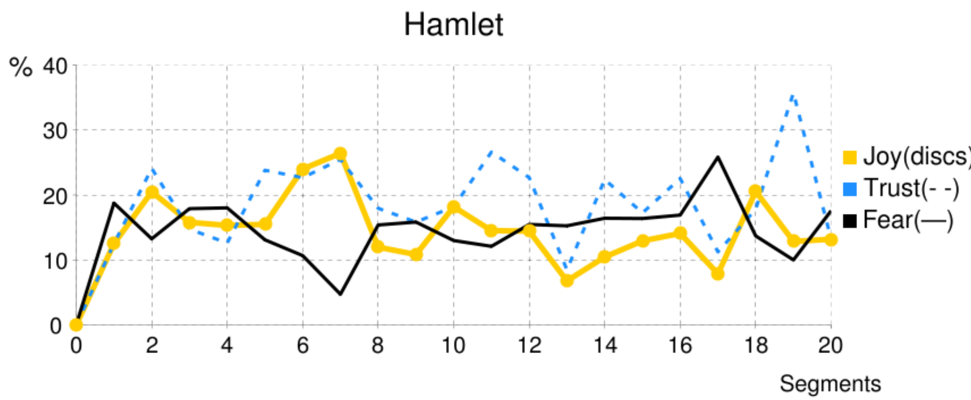

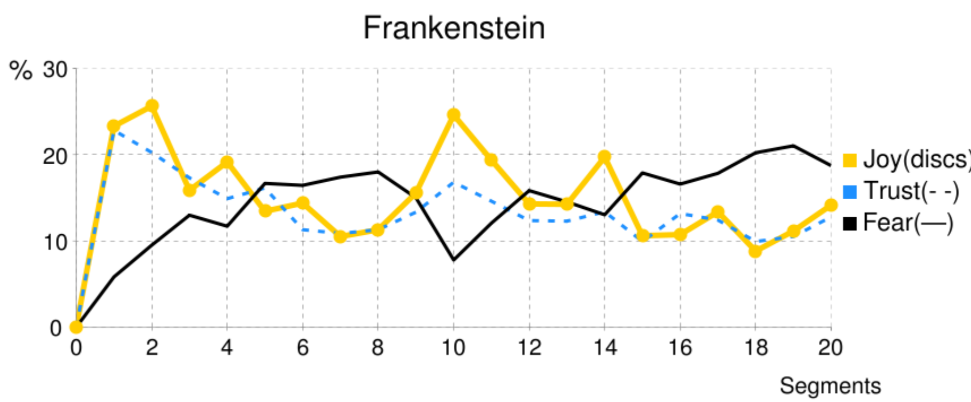

Tony Yang,
Simon Fraser University

Work on shapes of stories

- 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.
- Character-based kernels for novelistic plot structure. Elsner, M., 2012, April. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 634-644). Association for Computational Linguistics.
- A novel method for detecting plot. M. Jockers http://www.matthewjockers.net/2014/06/05/a-novel-method-for-detecting-plot/, June 2014.
- The emotional arcs of stories are dominated by six basic shapes. Reagan, A.J., Mitchell, L., Kiley, D., Danforth, C.M. and Dodds, P.S., 2016. EPJ Data Science, 5(1), p.31.

Generating music from text

Paper:

- Generating Music from Literature. Hannah Davis and Saif M. Mohammad, In Proceedings of the EACL Workshop on Computational Linguistics for Literature, April 2014, Gothenburg, Sweden.

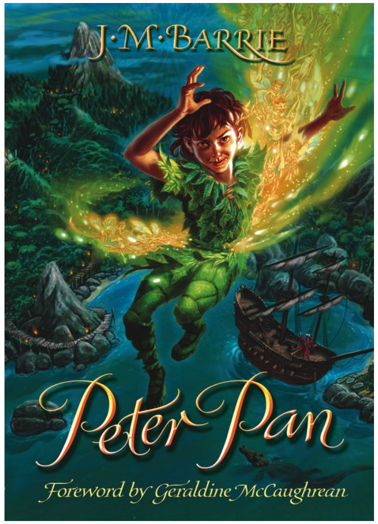

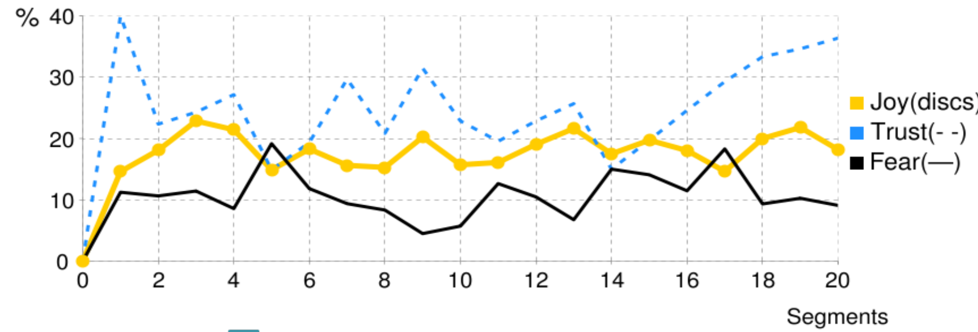

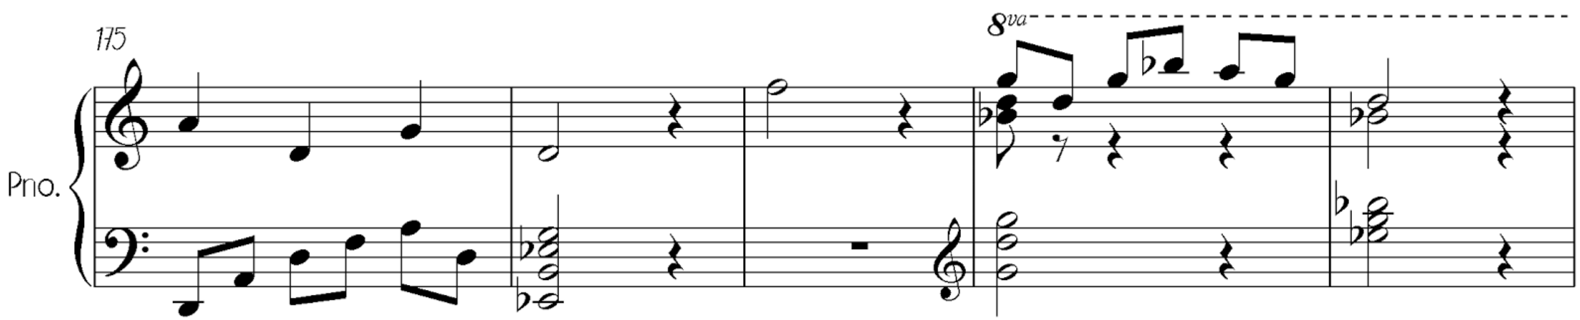

A method to generate music from literature.

- music that captures the change in the distribution of emotion words.

Challenges

- Not changing existing music -- generating novel pieces
- Paralysis of choice
- Has to sound good
- No one way is the right way -- evaluation is tricky

Music-Emotion Associations

Hannah Davis
Artist/Programmer

- Major and Minor Keys
  - major keys: happiness
  - minor keys: sadness
- Tempo
  - fast tempo: happiness or excitement
- Melody
  - a sequence of consonant notes: joy and calm
  - a sequence of dissonant notes: excitement, anger, or unpleasantness

Hunter et al., 2010, Hunter et al., 2008, Ali and Peynirciolu, 2010, Gabrielsson and Lindstrom, 2001, Webster and Weir, 2005

TransProse

Automatically generates three simultaneous piano melodies pertaining to the dominant emotions in the text, using the NRC Emotion Lexicon.

Examples

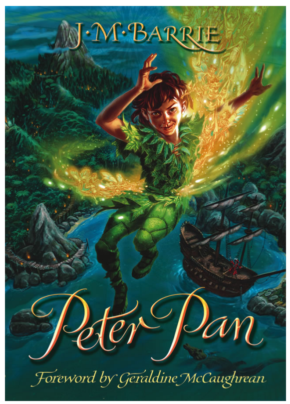

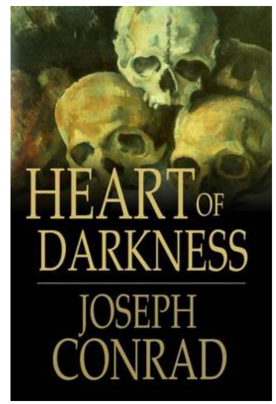

TransProse: www.musicfromtext.com
Music played 300,000 times since website launched in April 2014.

TransProse Music Played by an Orchestra, at the Louvre Museum, Paris

A symphony orchestra performs under the glass of the Louvre museum in Paris on Sept. 20. Accenture Strategy has created a symphonic experience enabled by human insight and artificial intelligence technology. (Michel Euler/AP)

Debate: Universality of Perception of Emotions

Margaret Mead
Cultural anthropologist

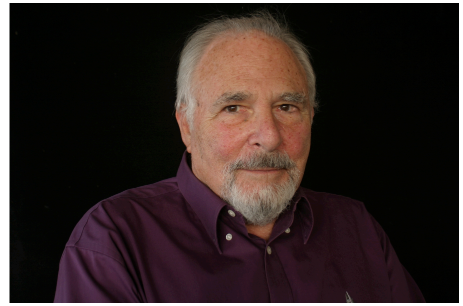

Paul Ekman
Psychologist and discoverer of micro expressions.

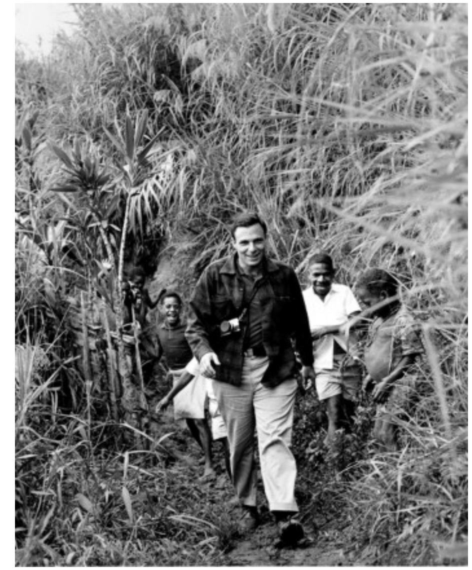

Lisa Barrett
University Distinguished Professor of Psychology, Northeastern University

- Grad school experiment on people's ability to distinguish photos of depression from anxiety
  - one is based on sadness, and the other on fear
  - found agreement to be poor
- Agreement also drops for Ekman emotions when participants are given:
  - Just the pictures (no emotion word options)
  - Or say, two scowling faces and asked if the two are feeling the same emotion

Some Emotions more basic than others?

may be not…

Hashtagged Tweets

- Hashtagged words are good labels of sentiments and emotions

  Some jerk just stole my photo on #tumblr #grrr #anger

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

    The reviewers want me to re-annotate the data. #joy

Paper:

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

Generating lexicon for 500 emotions

NRC Hashtag Emotion Lexicon: About 20,000 words associated with about 500 emotions

Papers:

- 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.
- Using Hashtags to Capture Fine Emotion Categories from Tweets. Saif M. Mohammad, Svetlana Kiritchenko, Computational Intelligence, in press.

Sentiment Analysis of Tweets

and the role of word-sentiment associations

Papers:

- Sentiment Analysis of Short Informal Texts. Svetlana Kiritchenko, Xiaodan Zhu and Saif Mohammad. Journal of Artificial Intelligence Research, 50, August 2014.
- 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.

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

Sentiment Lexicons

Created a sentiment lexicon using a Turney (2003) inspired method that uses PMI of a word with co-occurring positive and negative seed hashtags.

Positive
spectacular 0.91
okay 0.3

Negative
lousy -0.84
unpredictable -0.17

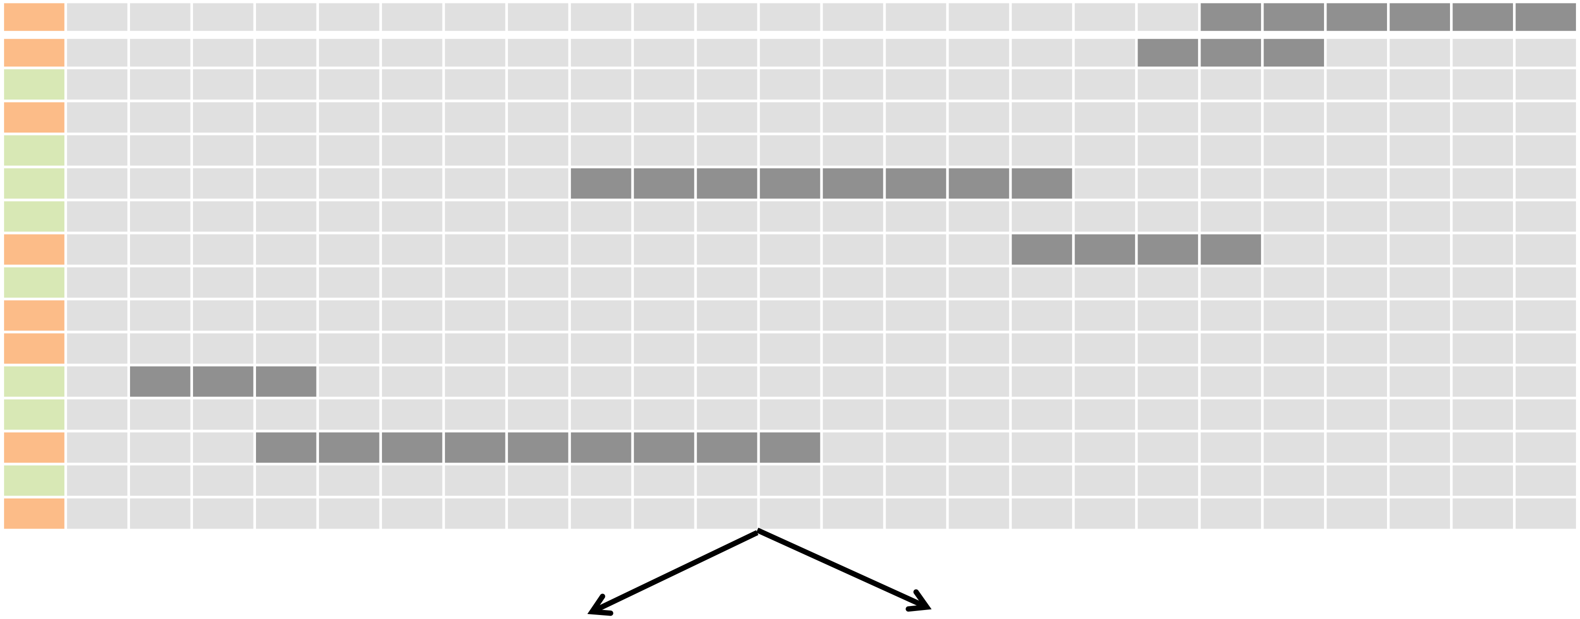

All the affirmative contexts

Generate sentiment lexicon for words in affirmative context

All the negated contexts

Generate sentiment lexicon for words in negated context

Term	**Sentiment140 Lexicons**		
	Base	**AffLex**	**NegLex**
Positive terms			
great	1.177	1.273	-0.367
nice	0.974	1.149	-0.912
honest	0.391	0.431	-0.123
Negative terms			
terrible	-1.766	-1.850	-0.890
bad	-1.297	-1.674	0.021
negative	-0.090	-0.261	0.389

Table 3: Example sentiment scores from the Sentiment140 Base, Affirmative Context (AffLex) and Negated Context (NegLex) Lexicons.

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
all-caps	YES, COOL
punctuation	#!+: 1, #?+: 0, #!?+: 0
word clusters	probably, definitely, probly
emoticons	:D, >:(
elongated words	soooo, yaayyy

Sentiment Analysis Competition

SemEval-2013: Classify Tweets, 44 teams

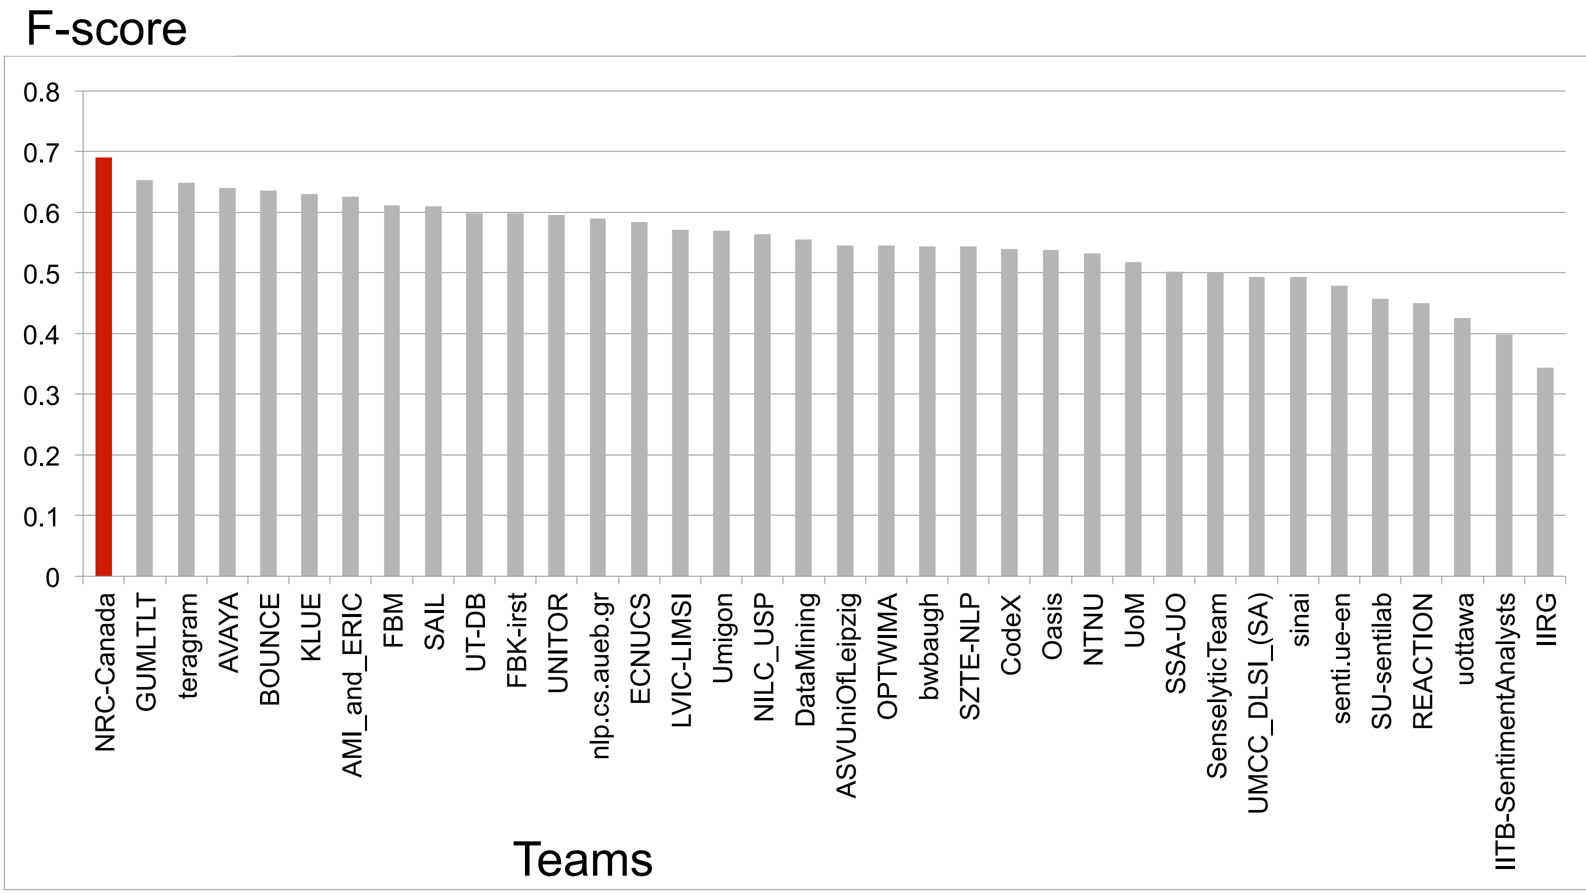

Sentiment Analysis Competition

SemEval-2013: Classify SMS messages, 30 teams

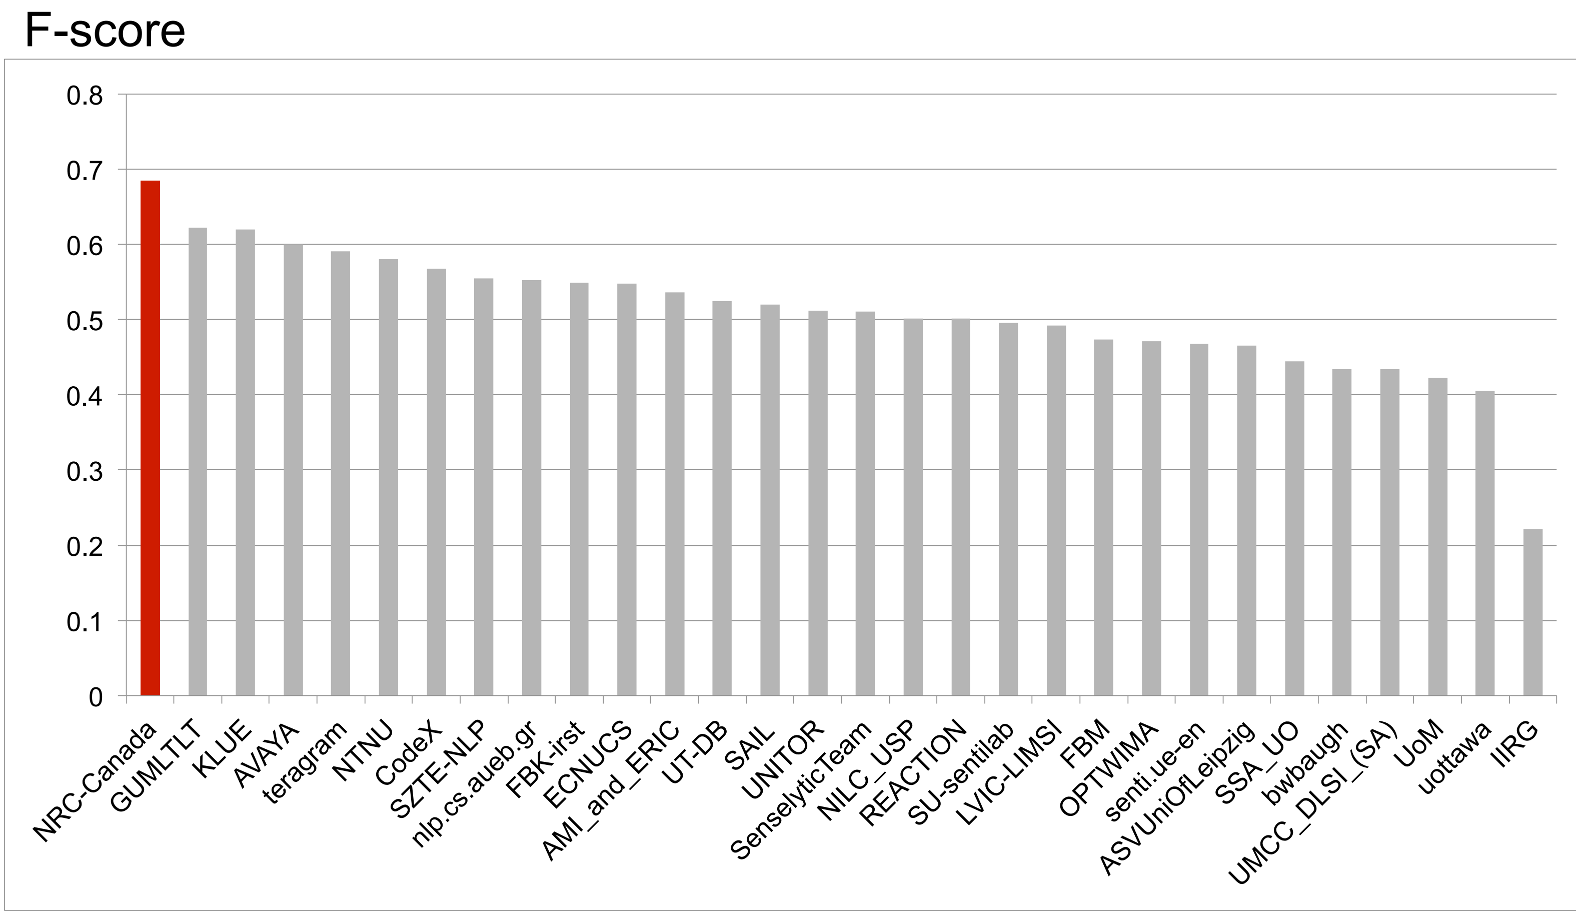

NRC-Canada in SemEval-2013, Task 2

Released description of features.
Released resources created (tweet-specific sentiment lexicons).
www.purl.com/net/sentimentoftweets

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.

Sentiment Analysis of Social Media Texts (SemEval-2014 Task 9)

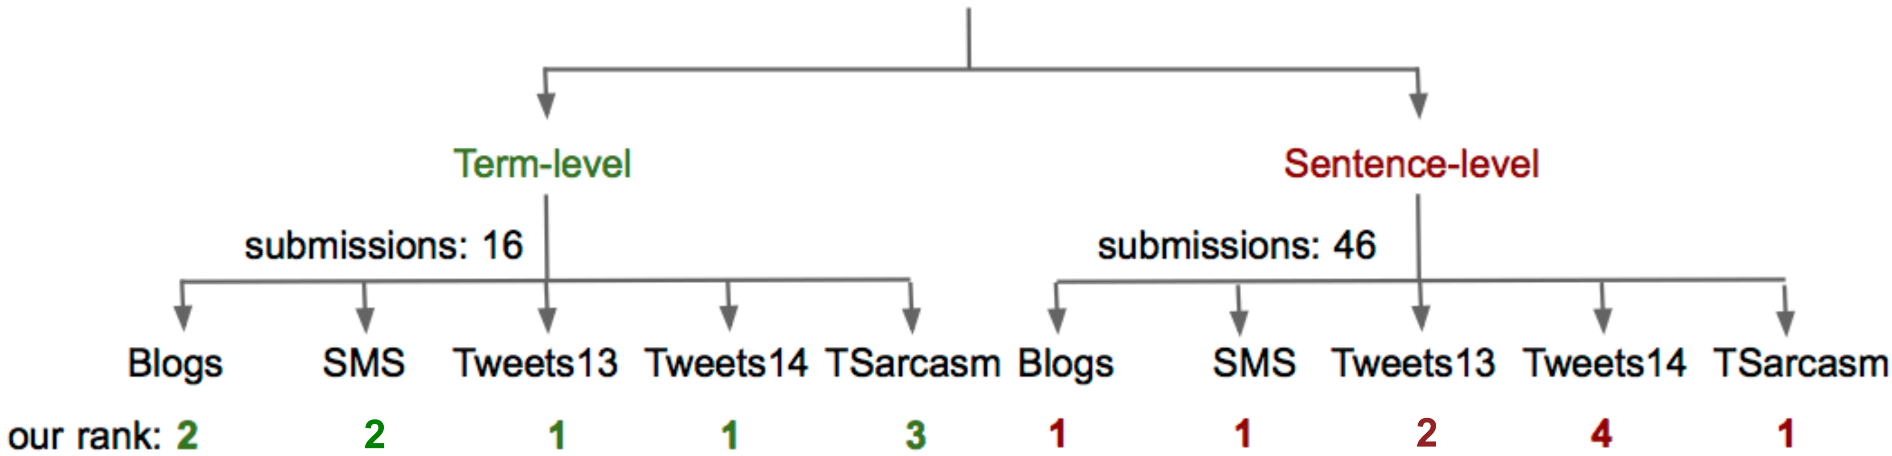

Sentiment Analysis of Social Media Texts (SemEval-2014 Task 9)

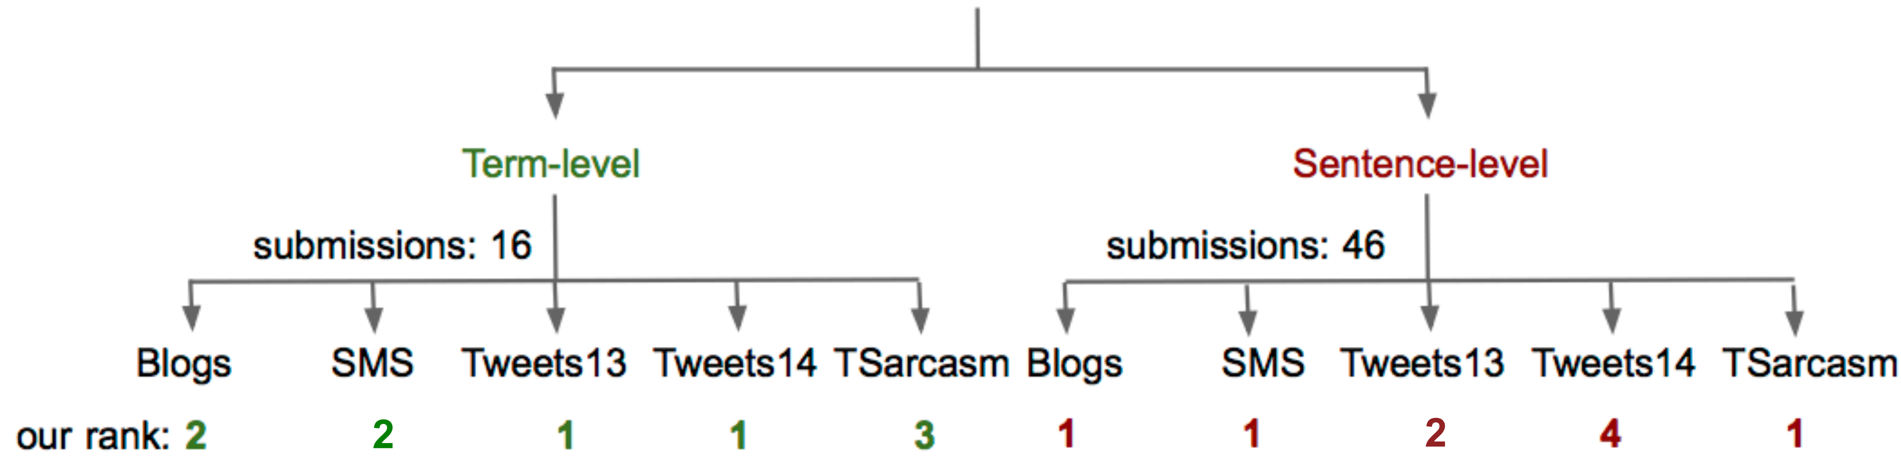

Aspect-Based Sentiment Analysis (SemEval-2014 Task 4)

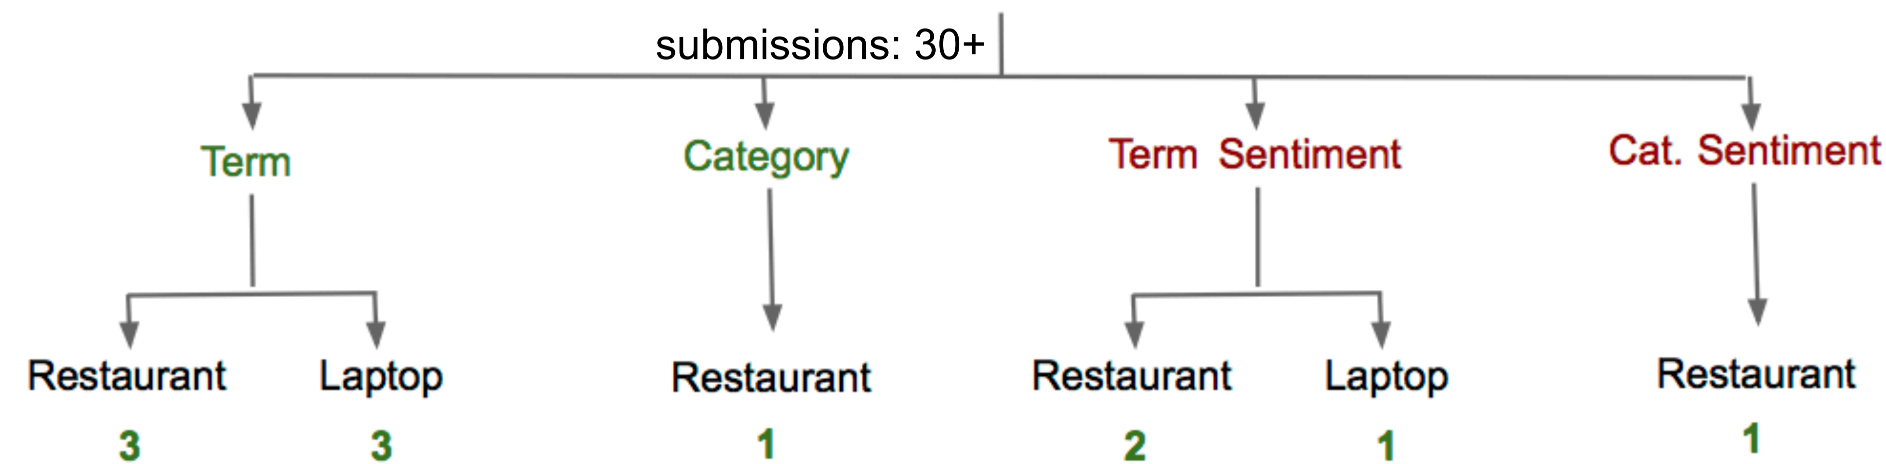

Feature Contributions (on Tweets)

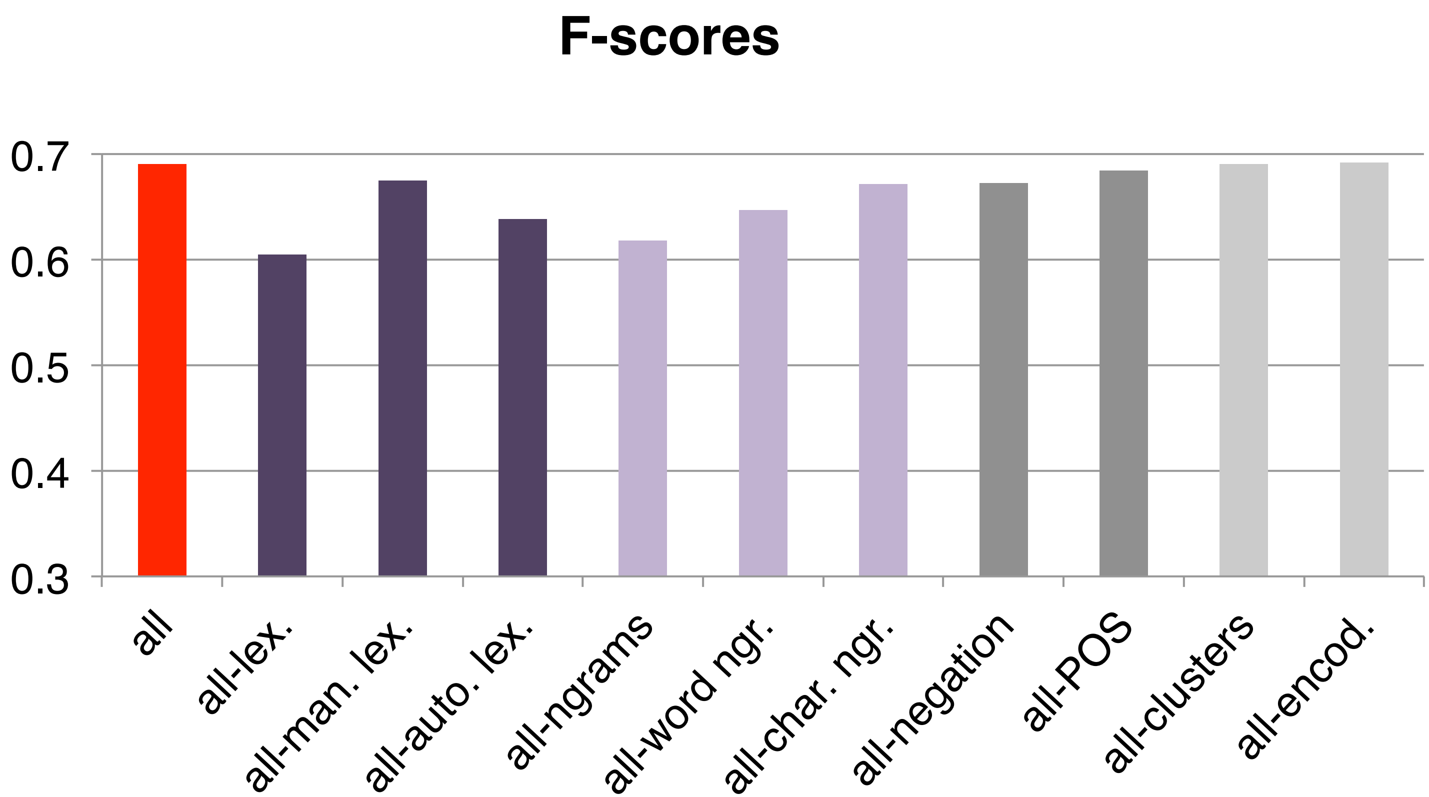

Movie Reviews

- Data from rottentomatoes.com (Pang and Lee, 2005)
- Socher et al. (2013) training and test set up
- Message-level task
  - Two-way classification: positive or negative

	System	Accuracy
(a)	Majority baseline	50.1
(b)	SVM-unigrams	71.9
(c)	Previous best result (Socher et al., 2013)	85.4
(d)	Our system	**85.5**

Other Recent Approaches to Creating Sentiment Lexicons

- Using neural networks and deep learning techniques
  - Duyu Tang, Furu Wei, Bing Qin, Ming Zhou and Ting Liu (2014)
- Constructing domain-specific sentiment
  - Sheng Huanga, Zhendong Niua, and Chongyang Shi (2014)
  - Ilia Chetviorkin and Natalia Loukachevitch (2014)
- Others:
  - Hassan Saif, Miriam Fernandez, Yulan He, and Harith Alani (2014): SentiCircles for Contextual and Conceptual Semantic Sentiment Analysis of Twitter.
  - Shi Feng, Kaisong Song, Daling Wang, Ge Yu (2014): A word-emoticon mutual reinforcement ranking model for building sentiment lexicon from massive collection of microblogs.
  - Raheleh Makki, Stephen Brooks and Evangelos E. Milios (2014): Context-Specific Sentiment Lexicon Expansion via Minimal User Interaction.
  - Yanqing Chen and Steven Skiena (2014): Building Sentiment Lexicons for All Major Languages
  - Bandhakavi et al. (2014): Generating a Word-Emotion Lexicon from #Emotional Tweets -- EM with Mixture of Classes Model.

Organizing Shared Tasks

- Detecting sentiment intensity of words and phrases: SemEval 2015, 2016
- Detecting sentiment of Tweets: SemEval 2015
- Detecting Stance in Tweets: SemEval 2016

Detecting Stance in Tweets

Given a tweet text and a target determine whether:

- the tweeter is in favor of the given target
- the tweeter is against the given target
- neither inference is likely

Example 1:

Target: Jeb Bush

Tweet: Jeb Bush is the only sane candidate in this republican lineup.

Systems have to deduce that the tweeter is likely in favor of the target.

Detecting Stance in Tweets

favor against neither

Given a tweet text and a target determine whether:

- the tweeter is in favor of the given target
- the tweeter is against the given target
- neither inference is likely

Example 2:

Target: Donald Trump
Tweet: Jeb Bush is the only sane candidate in this republican lineup.

Systems have to deduce that the tweeter is likely against the target.

Example 3:

Target: pro-life movement
Tweet: The pregnant are more than walking incubators, and have rights!

Systems have to deduce that the tweeter is likely against the target.

Stance vs. Sentiment

- positive language $\nRightarrow$ favor; negative language $\nRightarrow$ against
- the target can be expressed in different ways
  - impacts whether the instance is labeled favor or against
- the target of interest may not be mentioned in the text
  - especially for issue targets: legalization of abortion
- the target of interest may not be the target of opinion in the text

Example 3:

Target: Donald Trump

Tweet: Jeb Bush is the only sane candidate in this republican lineup.

Applications of automatic stance detection:
information retrieval, text summarization, textual entailment, social media analytics.

SemEval-2016 Task#6: Detecting Stance in Tweets

Task A: Supervised Framework

- provided training and test data for five targets
- atheism, climate change is a real concern, feminist movement, Hillary Clinton, legalization of abortion

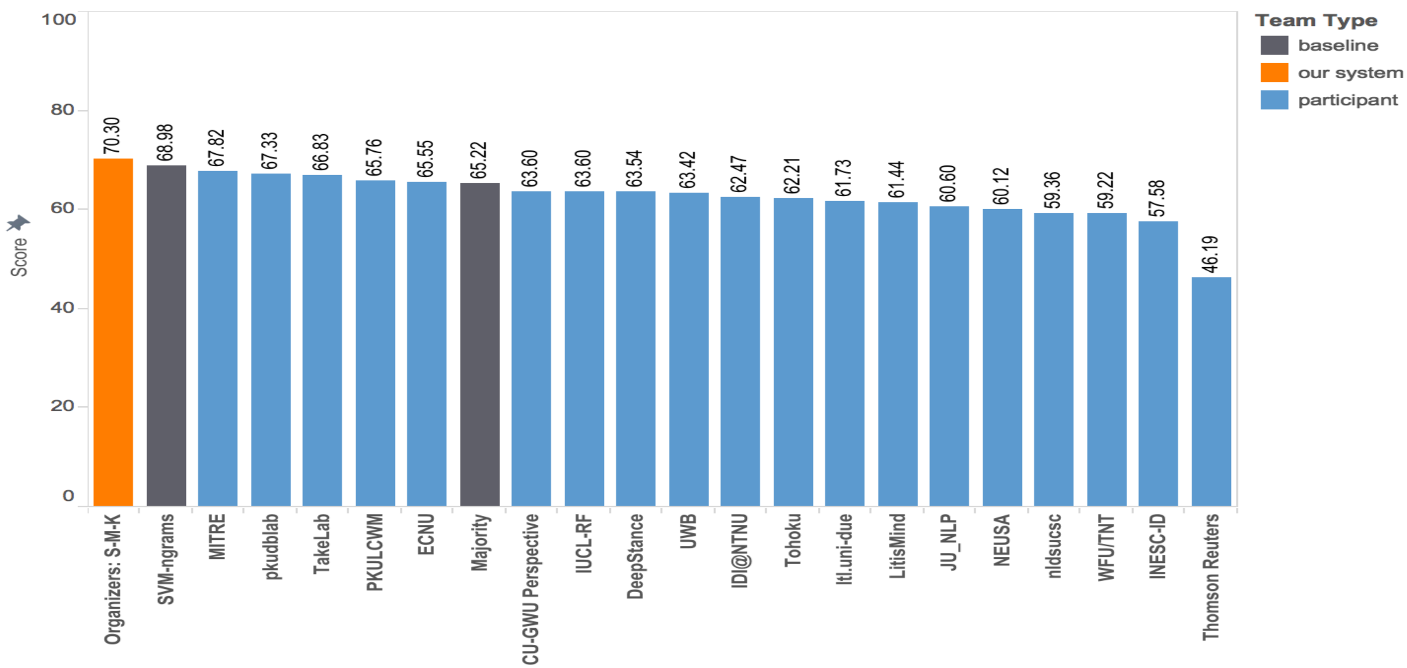

Detecting Stance in Tweets

Parinaz Sobhani

Svetlana Kiritchenko

- Best score on Task A: 70.3
  - SVM, word embeddings, word and character ngrams

Paper: Stance and Sentiment in Tweets. Saif M. Mohammad, Parinaz Sobhani, and Svetlana Kiritchenko. *Special Section of the ACM Transactions on Internet Technology on Argumentation in Social Media,* In Press.

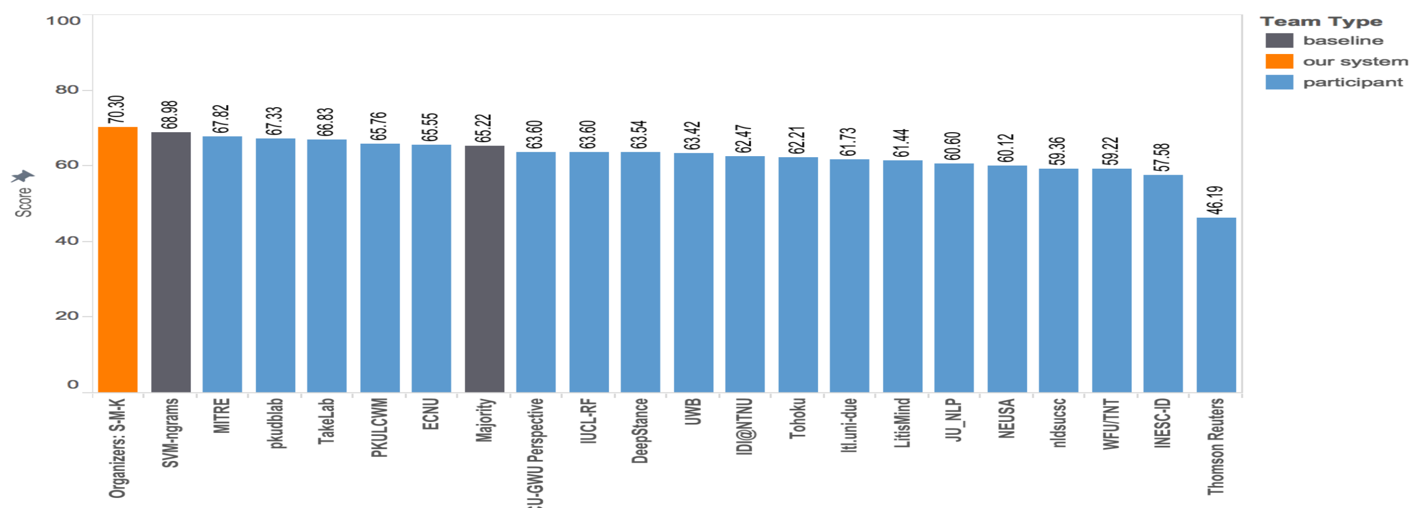

SemEval-2016 Task#6: Detecting Stance in Tweets

Task B: Weakly Supervised Framework

- test data: tweets for one target 'Donald Trump'
- training data: none
- unlabeled data: 78k tweets
  - associated with 'Donald Trump' to various degrees

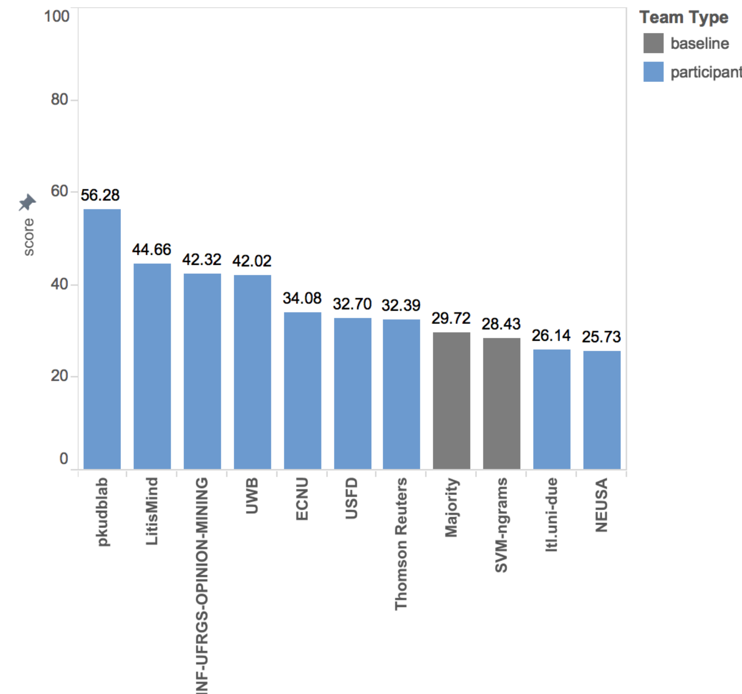

Stance Project Homepage

http://www.saifmohammad.com/WebPages/StanceDataset.htm

- Complete Stance Dataset with annotation for both stance and sentiment
- Interactive visualization
- Papers

Summary: Created Affect Association Lexicons

- Manually
  - Traditional ratings
    - NRC Emotion Lexicon: ~14,000 words, 8 emotions, 2 sentiments
  - Best Worst Scaling
    - Twitter Lexicons for English and Arabic
    - Sentiment composition lexicons
      - for phrases with negators, modals, and adverbs
      - for opposing polarity phrases
- Automatically
  - for hundreds of affect categories
  - using hashtag words and emoticons

Summary: Uses of Affect Associations

- Sentence-, tweet-, message-level affect classification
- Tracking the distribution of affect words in text
  - literary analysis
  - information visualization
- As gold datasets to evaluate automatic methods of generating affect lexicons
  - can also be used as seeds for generating large sentiment lexicons
- Linguistic studies
  - understanding sentiment composition

Resources Available at: www.saifmohammad.com/ResearchAreas.html

- word-emotion and word-sentiment association lexicons
  - manually created
    - best-worst scaling, sentiment composition
  - automatically generated
    - from tweets and hashtags
- word-colour association lexicon
- metaphor-emotion data
- interactive visualizations
- tutorials and book chapters on sentiment analysis

Saif M Mohammad

www.saifmohammad.com

saif.mohammad@nrc-cnrc.gc.ca

Measuring the Least Perceptible Difference

- Least perceptible difference in sentiment scores is a point *d* at which we can say with high confidence that the two terms do not have the same sentiment associations

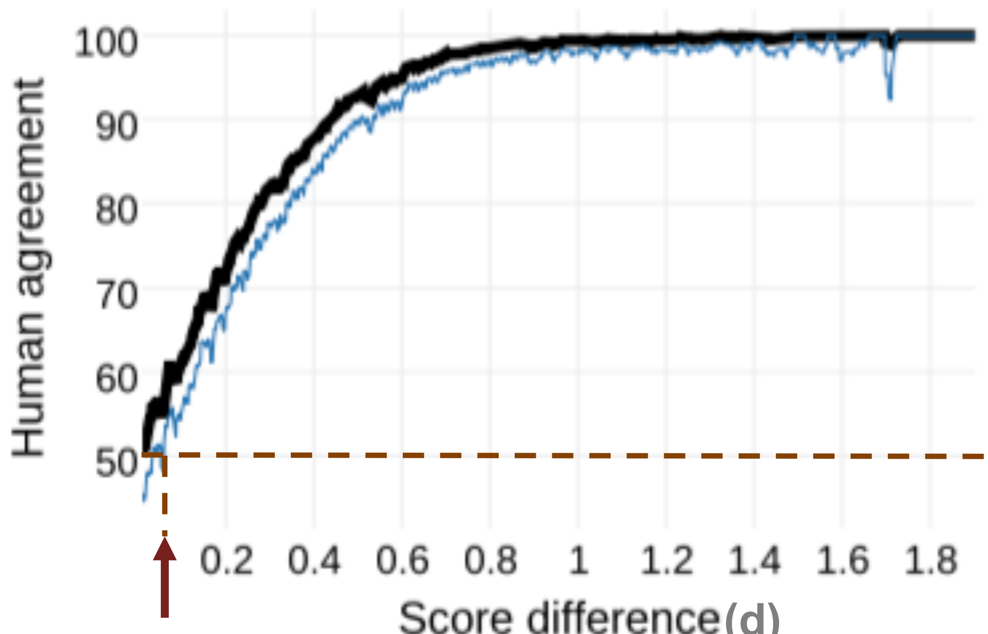

least perceptible difference

Least Perceptible Differences in lexicons:

General English	: 0.069
English Twitter	: 0.080
Arabic Twitter	: 0.087

(~4% in range -1..1)