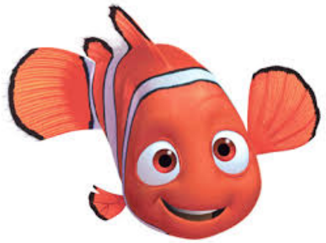

Finding Emotions in Language

Saif M. Mohammad

Senior Research Scientist
National Research Council Canada

Outline

- Introduction
  - emotion and language
  - the landscape of tasks and applications

Outline

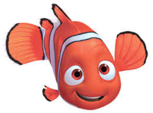

- Introduction
  - emotion and language
  - the landscape of tasks and applications
- Finding Emotions (humans)
  - annotating words, phrases, sentences, tweets
  - crowdsourcing
  - obtaining reliable fine-grained annotations

Outline

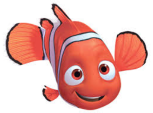

- Introduction
  - emotion and language
  - the landscape of tasks and applications
- Finding Emotions (humans)
  - annotating words, phrases, sentences, tweets
  - crowdsourcing
  - obtaining reliable fine-grained annotations
- Finding Emotions (machines)
  - automatic systems for emotion, sentiment, stance, personality, literary analysis, music generation, argumentation, pharmacovigilance
  - ethics in emotion analysis

Introduction

- emotion and language
- the landscape of tasks and applications

Emotion is any conscious experience characterized by intense mental activity and a high degree of pleasure or displeasure. Scientific discourse has drifted to other meanings and there is no consensus on a definition.

-- Wikipedia

Emotions and Language

- Emotions are central to human experience and behavior
- They condition our actions
- There is no cognition without emotion
- Emotions are central in organizing meaning
  - The three dimensions of meaning (Osgood, Suci, and Tannenbaum, 1957; Mehrabian and Russell,1974):
    - evaluativeness / valence
    - activation / arousal
    - potency / dominance

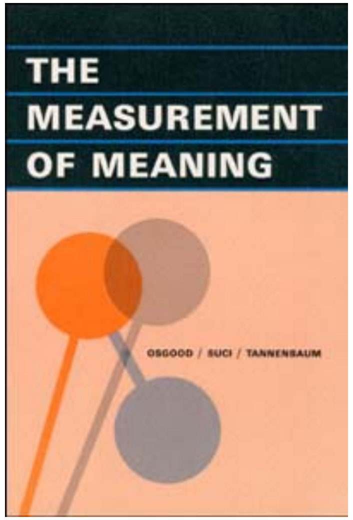

Which Emotions?

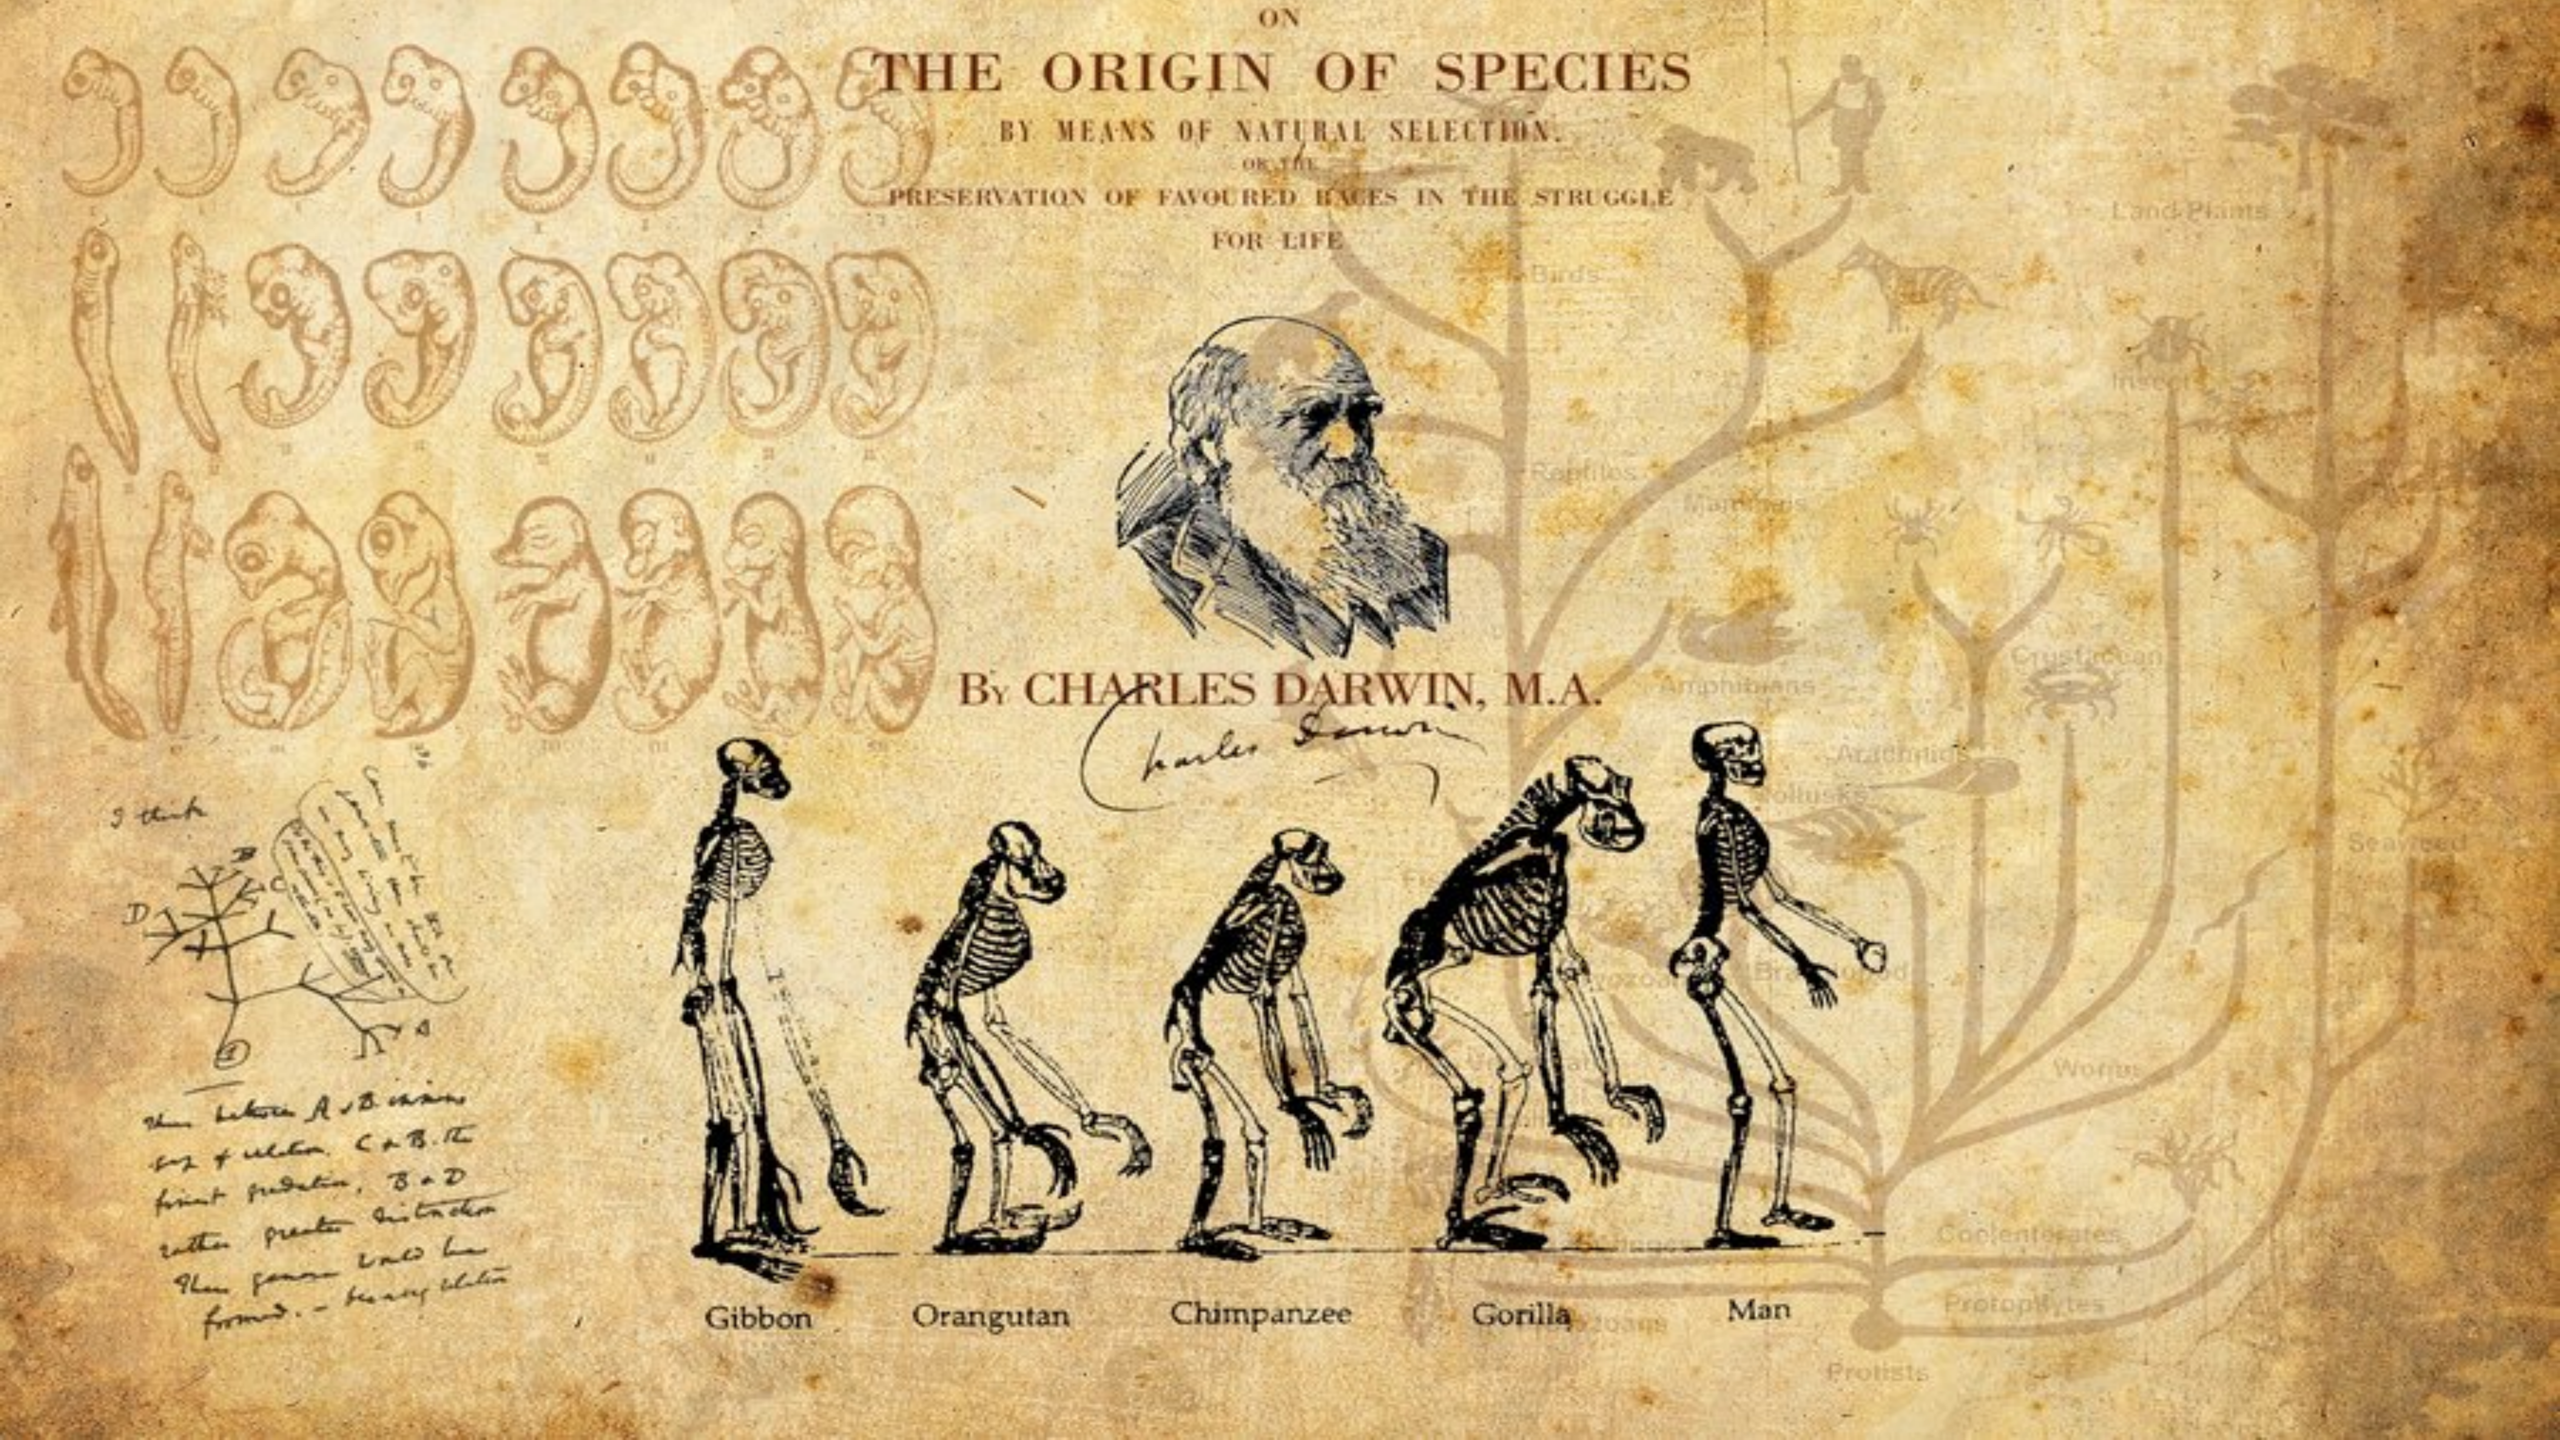
ON
THE ORIGIN OF SPECIES
BY MEANS OF NATURAL SELECTION,
OR THE
PRESERVATION OF FAVOURED RACES IN THE STRUGGLE
FOR LIFE
By CHARLES DARWIN, M.A.
Charles Darwin
I think
Gibbon
Orangutan
Chimpanzee
Gorilla
Man
Land Plants
Birds
Insects
Reptiles
Mammals
Crustacean
Amphibians
Arachnids
Mollusks
Seaweed
Worms
Coelenterates
Protophytes
Protists

Charles Darwin

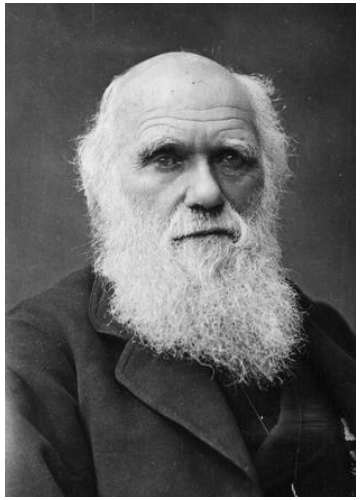

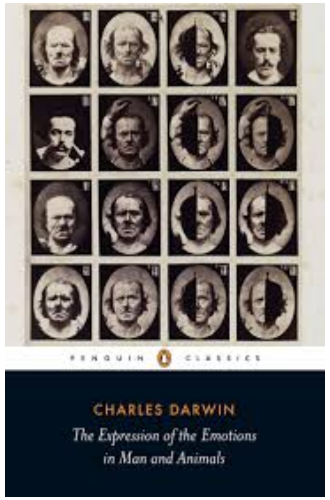

FIG. 20.—Terror,
from a photograph by Dr. Duchenne.

- 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

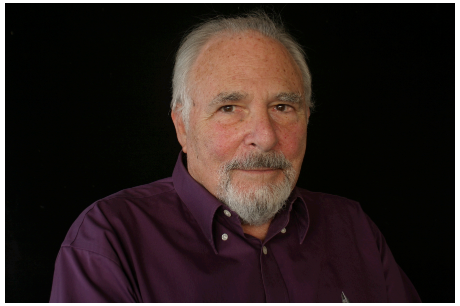

Paul Ekman
Psychologist and discoverer of micro expressions.

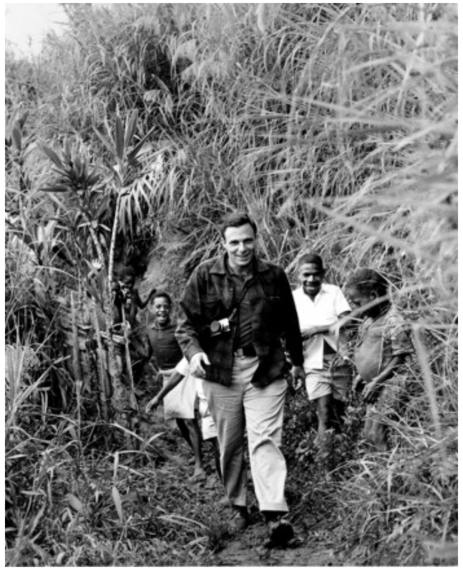

- Circa 1950's, Margaret Mead and others believed facial expressions and their meanings were culturally determined
  - behavioural learning processes
- Paul Ekman provided the strongest evidence to date that Darwin, not Margaret Mead, was correct in claiming facial expressions are universal
- Found universality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

Paul Ekman, 1971: Six Basic Emotions

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

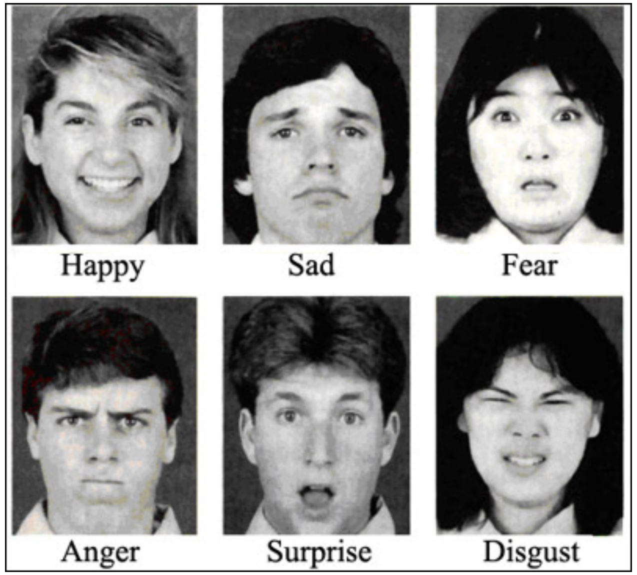

Plutchik, 1980: Eight Basic Emotions

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

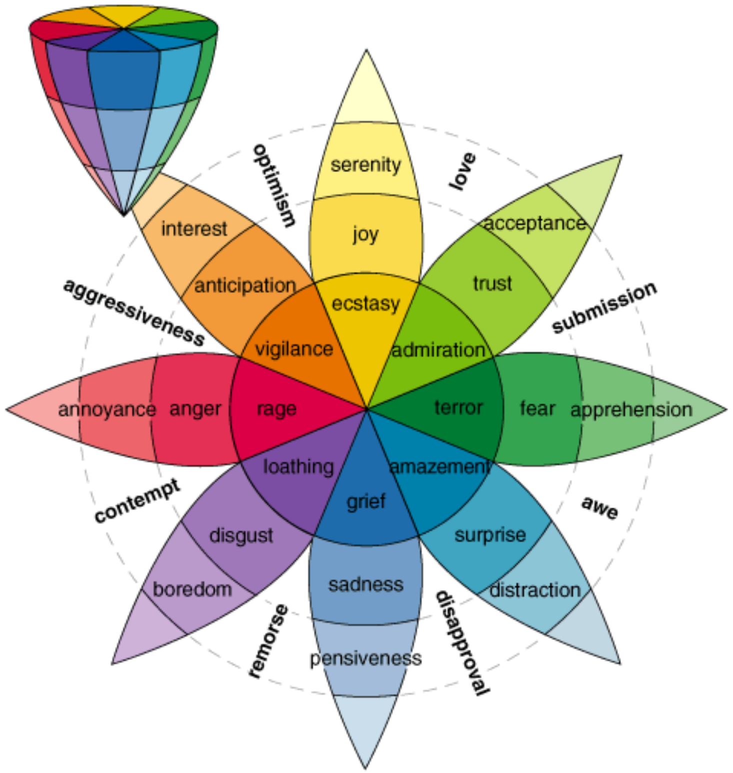

Plutchik, 1980: Eight Basic Emotions

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

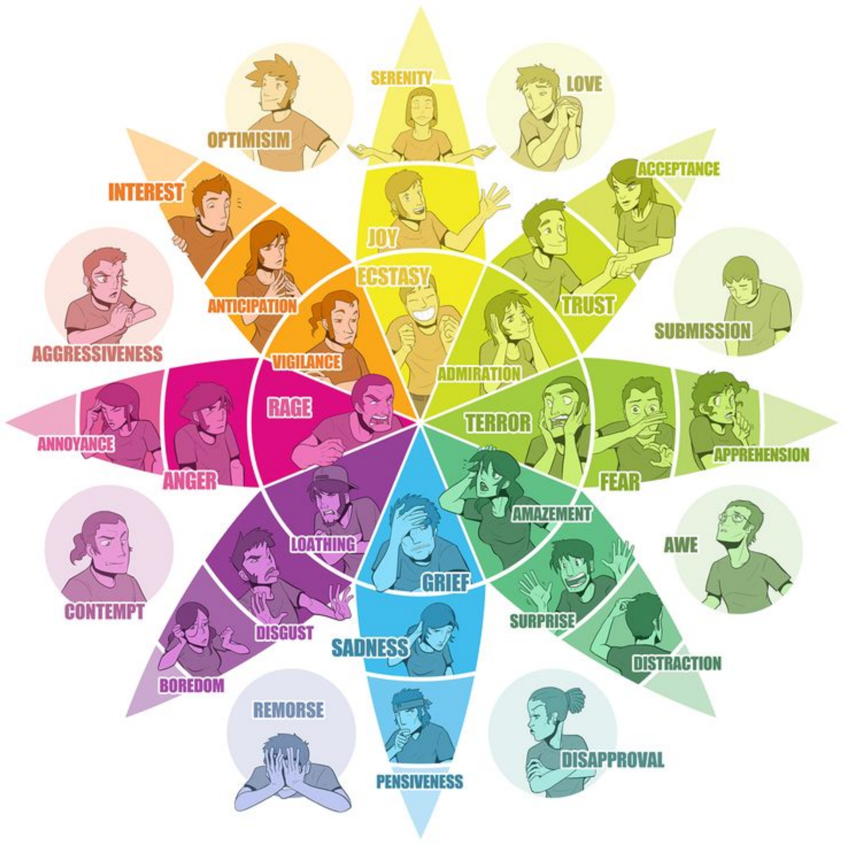

Image credit: Julia Belyanevych

Dimensional Model of Emotions

- small number of dimensions
- emotion is point in the multi-dimensional space

Circumplex Model (Russell, 1980)

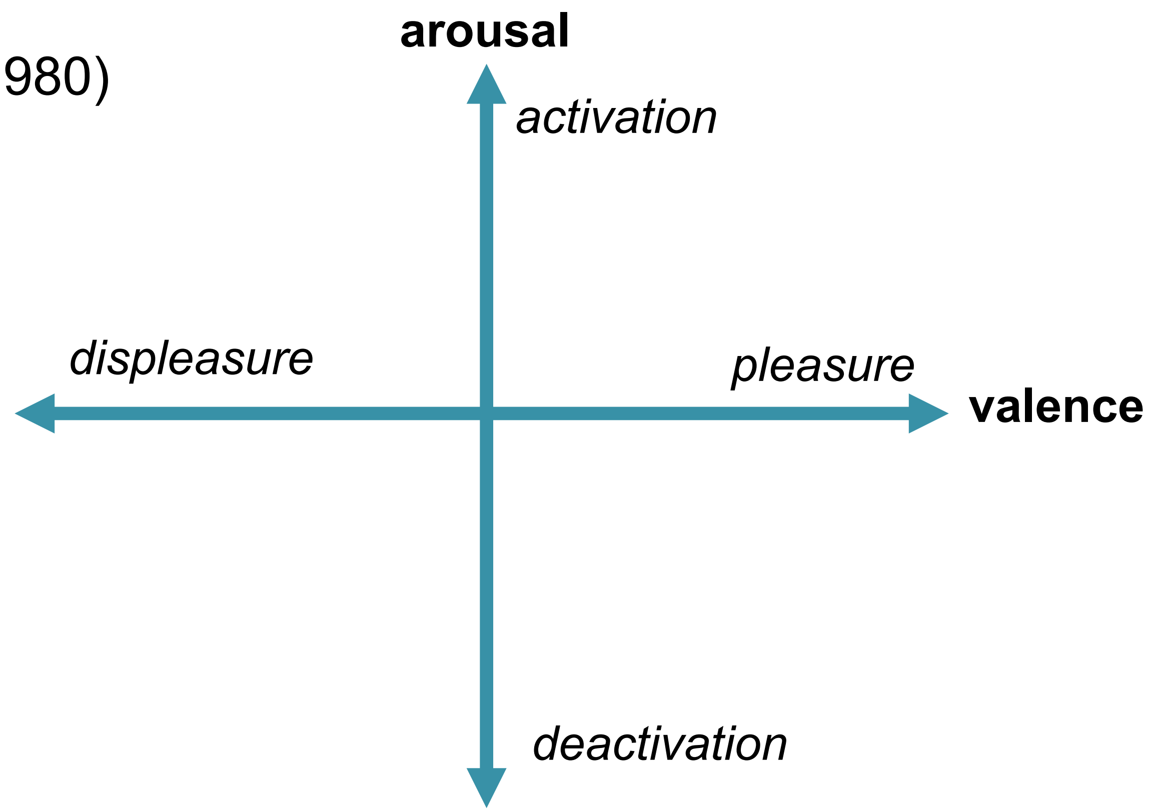

Dimensional Model of Emotions

- small number of dimensions
- emotion is point in the multi-dimensional space

Circumplex Model (Russell, 1980)

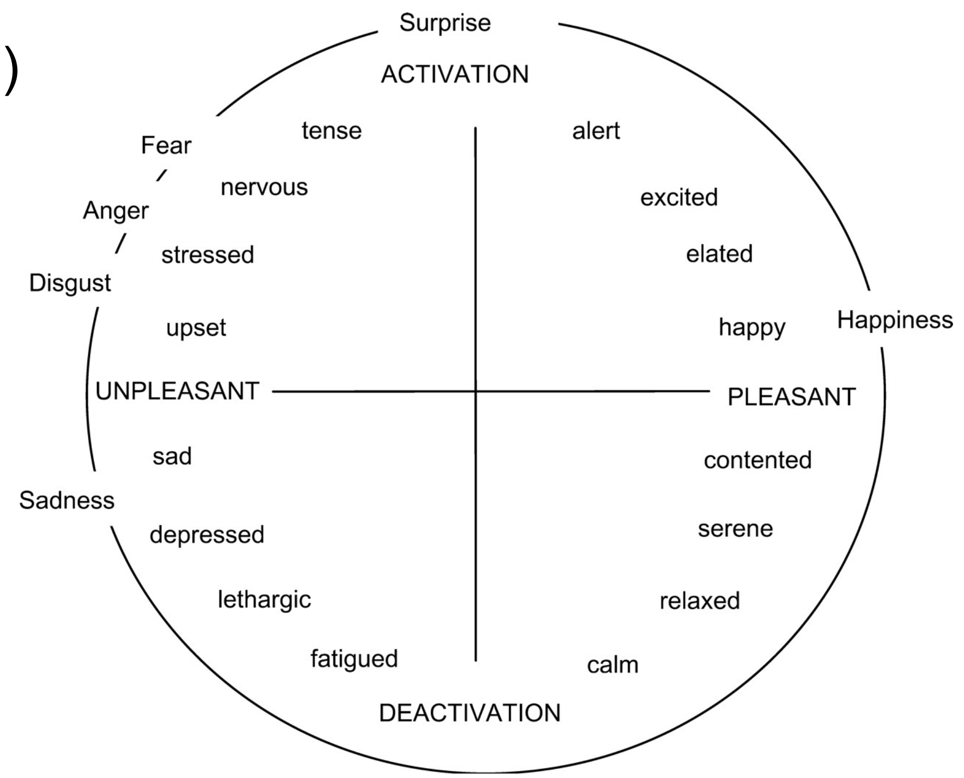

Sentiment Analysis

- 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 Automatically Detecting Sentiment and Emotions

- Computer games and animation

Applications of Automatically Detecting Sentiment and Emotions

- Computer games and animation
- Detecting cyber-bullying, depression

Applications of Automatically Detecting Sentiment and Emotions

- Computer games and animation
- Detecting depression, cyber-bullying
- Tutoring systems, writing assistants

Applications of Automatically Detecting Sentiment and Emotions

- Computer games and animation
- Detecting depression, cyber-bullying
- Tutoring systems, writing assistants
- Tracking sentiment towards products, events, people

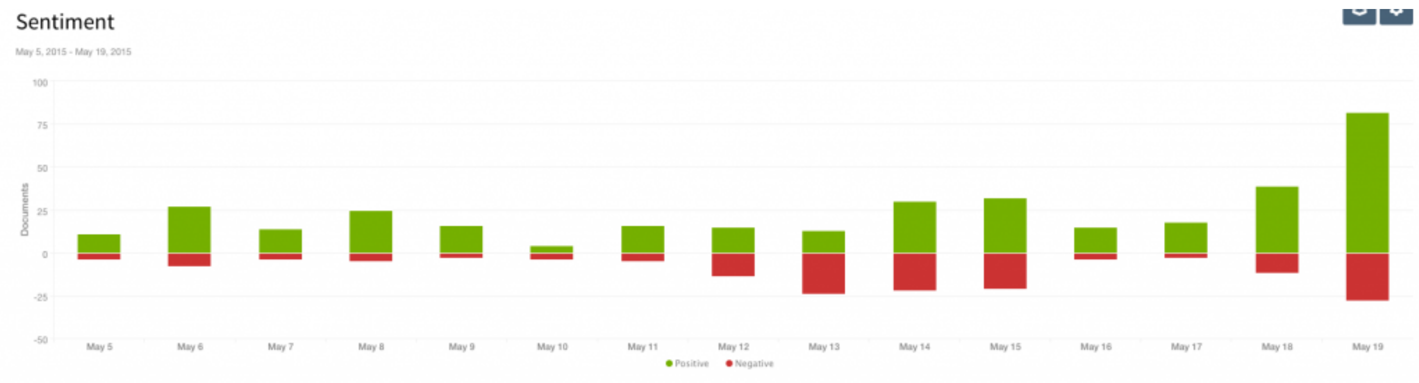

Applications of Automatically Detecting Sentiment and Emotions

- Computer games and animation
- Detecting depression, cyber-bullying
- Tutoring systems, writing assistants
- Tracking sentiment towards products, events, people
- Physiotherapy robots, virtual nurses

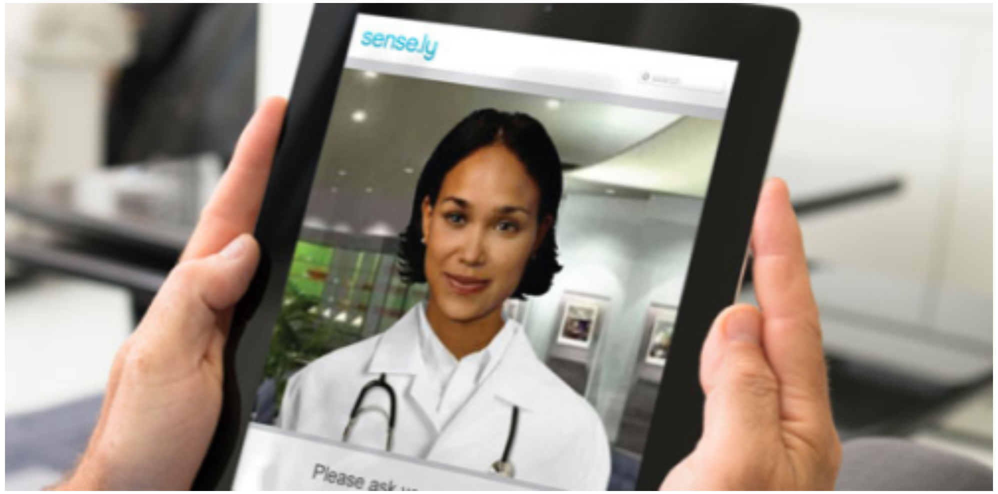

Sentiment and Emotion Tasks

- Is a given piece of text positive, negative, or neutral?
- Is a word within a sentence positive, negative, or neutral?
  - unpredictable movie plot vs. unpredictable steering
- What is the sentiment of the speaker/writer?
  - Is the objective to determine explicitly stated speaker sentiment or are we to infer likely speaker sentiment?
- What sentiment is evoked in the listener/reader?
- What is the sentiment of an entity mentioned in the text?

Sentiment and Emotion Tasks (continued)

- What is the sentiment towards specific aspects of a product?
  - sentiment towards the food and sentiment towards the service in a customer review of a restaurant
- What is the sentiment towards an entity such as a politician, government policy, company, issue, or product.
  - Stance detection: favorable or unfavorable
  - Framing: focusing on specific dimensions

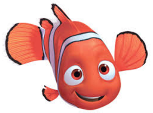

Finding Emotions (humans)

- annotating words, phrases, sentences, tweets
- crowdsourcing
- obtaining reliable fine-grained annotations

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

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.

Plutchik, 1980: Eight Basic Emotions

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

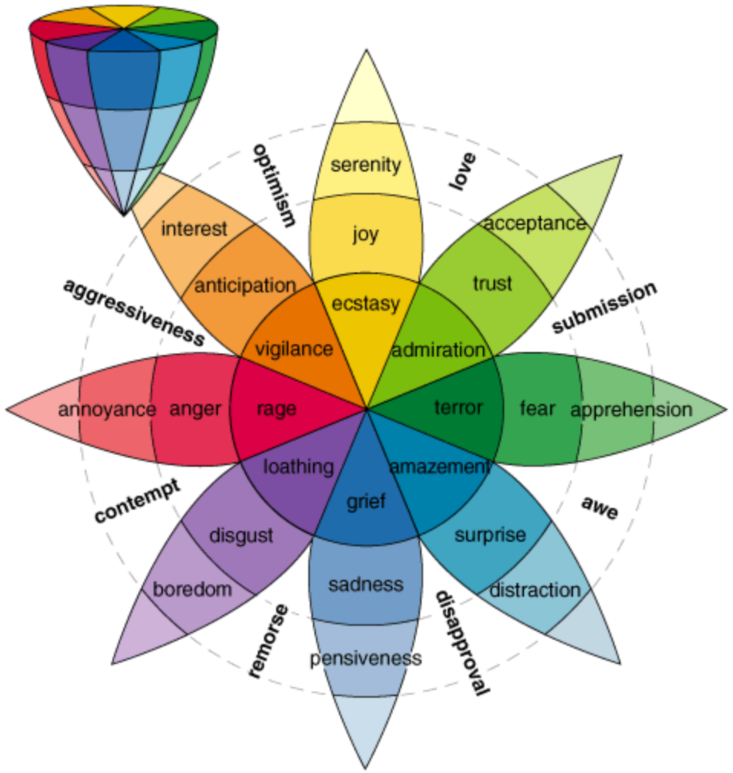

Goal: We chose to capture word-emotion associations for the 8 Pltchik emotions.

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

- Each word-sense pair is annotated by 5 Turkers
- 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

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.

How to capture fine-grained affect intensity associations reliably?

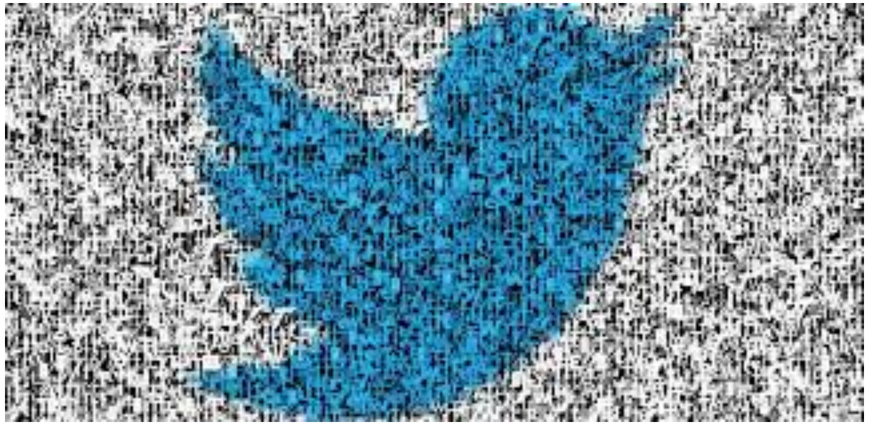

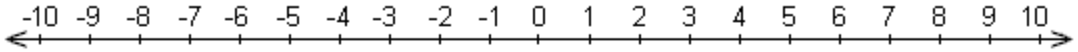

How to capture fine-grained affect intensity associations reliably?

Ranking Jelly Bean Flavours

- Black Pepper
- Booger
- Dirt
- Earthworm
- Earwax
- Rotten Egg
- Sausage
- Soap
- …

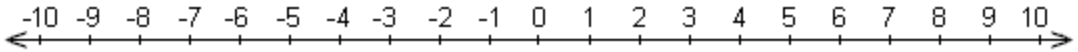

How to capture fine-grained affect intensity associations reliably?

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)

Comparative Annotations

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, less biased, more discriminating annotations
  (Kiritchenko and Mohammad, 2017, Cohen, 2003)

Best–Worst Scaling (BWS)

with example from Kiritchenko et al. 2014

- 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 all the terms

Best-Worst Scaling Lexicons

Svetlana Kiritchenko
NRC

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

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

English Twitter Lexicon:
Examples sentiment scores obtained using BWS

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

Measuring Quality of Annotations

- Less useful: standard inter-annotator agreement measures
  - when a tuple has two items that are close in emotion intensity
  - the disagreement is a useful signal for BWS

- More useful: a measure of reproducibility of the end result
  - repeat annotations
  - involve multiple respondents
  - if similar intensity rankings (and scores) are produced
    - one can be confident that the scores capture the true emotion intensities.

Reliability (Reproducibility) of Annotations

Average split-half reliability (SHR): a commonly used approach to determine consistency (Kuder and Richardson, 1937; Cronbach, 1946)

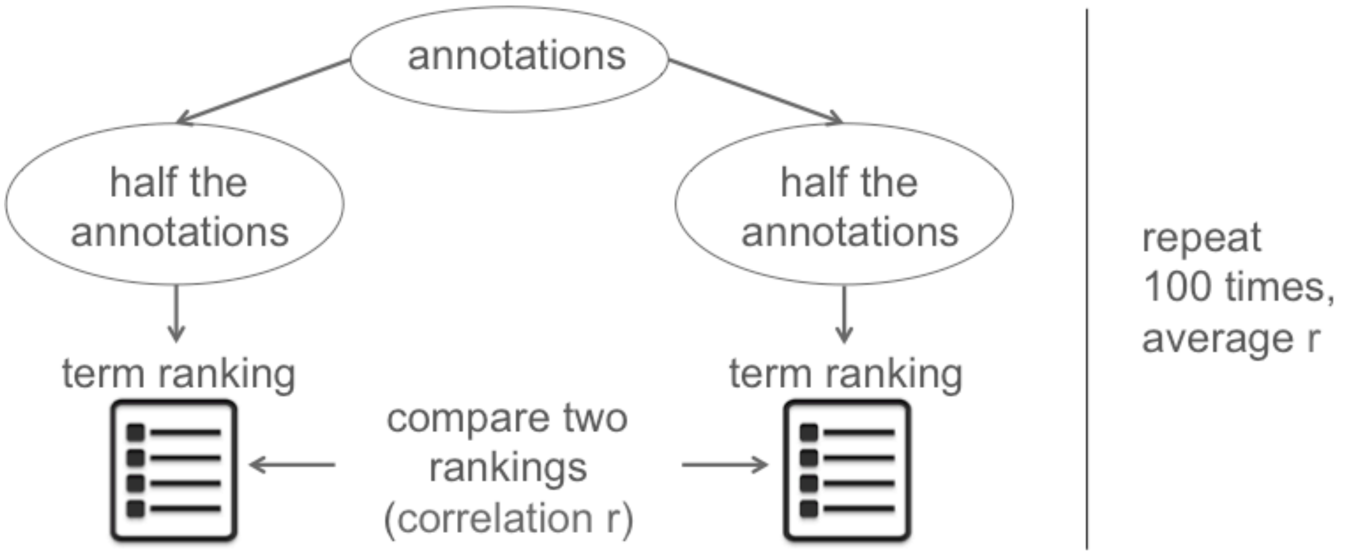

Robustness of the Annotations

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

Best-Worst Scaling vs. Rating Scales

Directly comparing the two methods:

- Which **manual** annotation method is better?

Measure reproducibility:

- Annotate the same dataset by both methods
- Show that for the number of responses, BWS results are more stable when experiment repeated

Best-Worst Scaling More Reliable than Rating Scales: A Case Study on Sentiment Intensity Annotation.
Svetlana Kiritchenko and Saif M. Mohammad. In Proceedings of the Association for Computational Linguistics. August 2017. Vancouver, Canada.

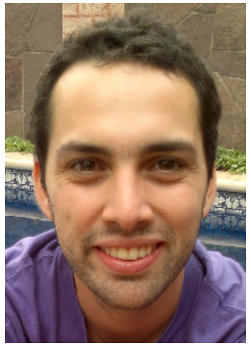

Felipe José Bravo Márquez

Emotion Intensity in Tweets

Paper:

WASSA-2017 Shared Task on Emotion Intensity. Saif M. Mohammad and Felipe Bravo-Marquez. In *Proceedings of the EMNLP 2017 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, September 2017, Copenhagen, Denmark.

WASSA-2017 Shared Task: Emotion Intensity in Tweets

Task:

Given a tweet and an emotion X, determine intensity of emotion X felt by the speaker,

- a real-valued score between 0 and 1
  - 1: the speaker is feeling the maximum amount of emotion X
  - 0: the speaker is feeling the least amount of emotion X

Data

- Annotated sentences using BWS

Task website:

http://saifmohammad.com/WebPages/EmotionIntensity-SharedTask.html

Our Example BWS Annotation Instance:
for tweet emotion intensity

Speaker 1: These days I see no light. Nothing is working out #depressed
Speaker 2: The refugees are the ones running from terror.
Speaker 3: Tim is sad that the business is not going to meet expectations.
Speaker 4: Too many people cannot make ends meet with their wages.

Q1. Which of the four speakers is likely to be the MOST SAD (or having a mental state most inclined towards sadness)

Q2. Which of the four speakers is likely to be the LEAST SAD (or having a mental state least inclined towards sadness)

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.

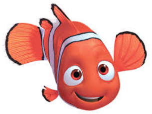

Finding Emotions (machines)

- automatic systems for emotion, sentiment, stance, personality, literary analysis, music generation, argumentation, pharmacovigilance
- ethics in emotion analysis

Tony Yang, Simon Fraser University

Visualizing Emotions in Text

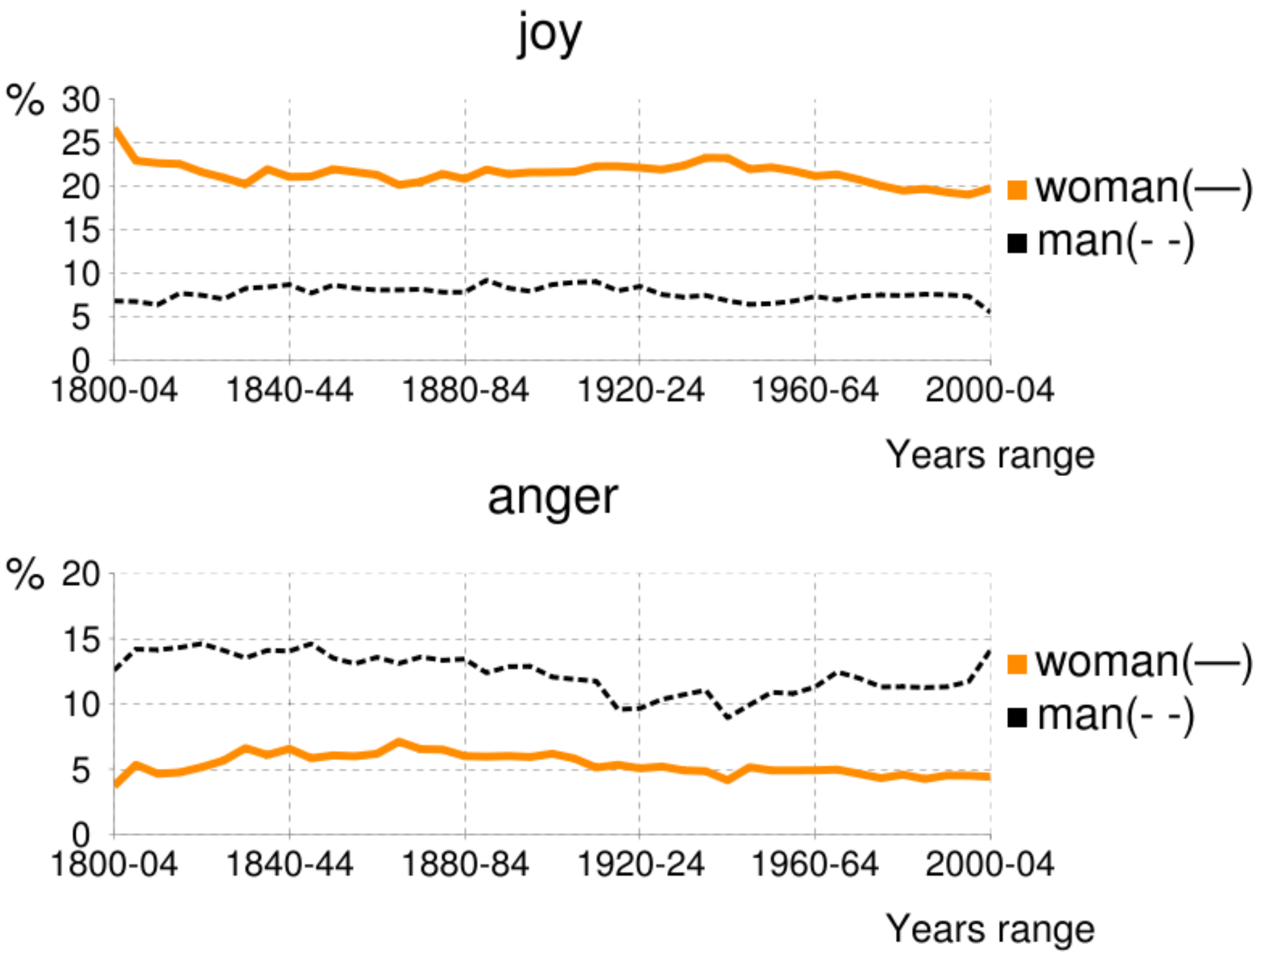

Percentage of joy and anger words in close proximity to occurrences of man and woman in books.

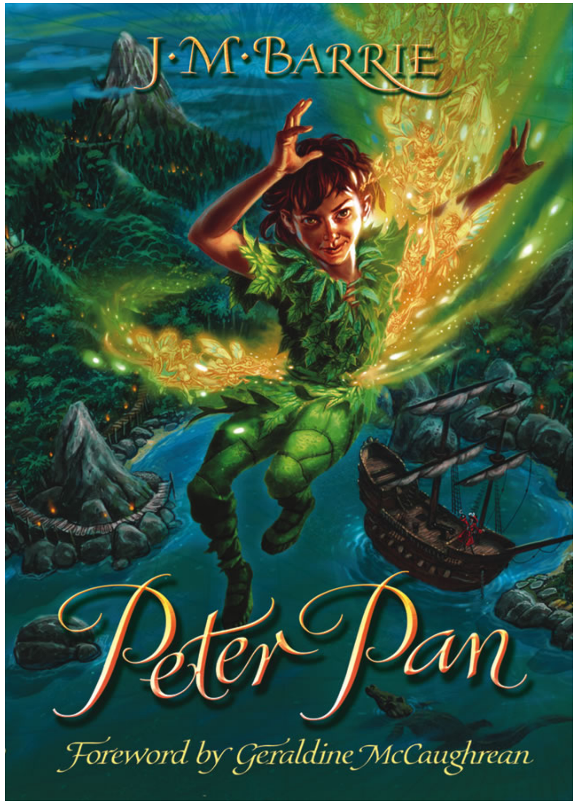

Stories

STORIES

Tracking Emotions in Stories

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

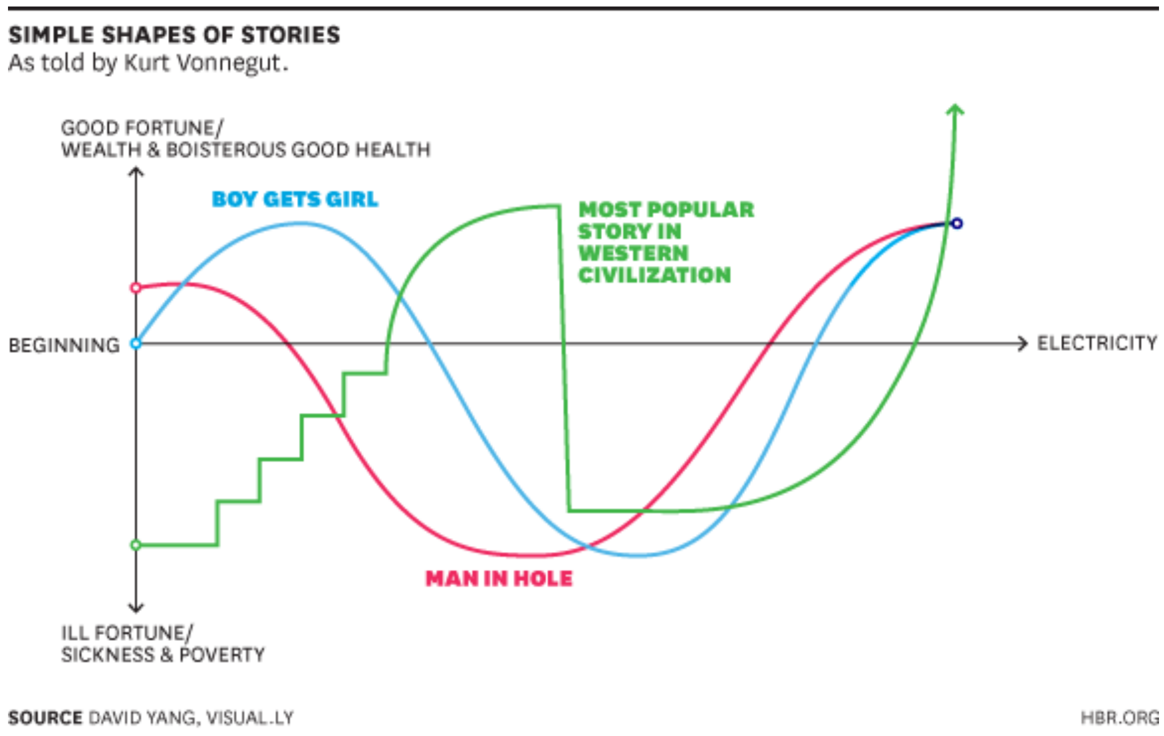

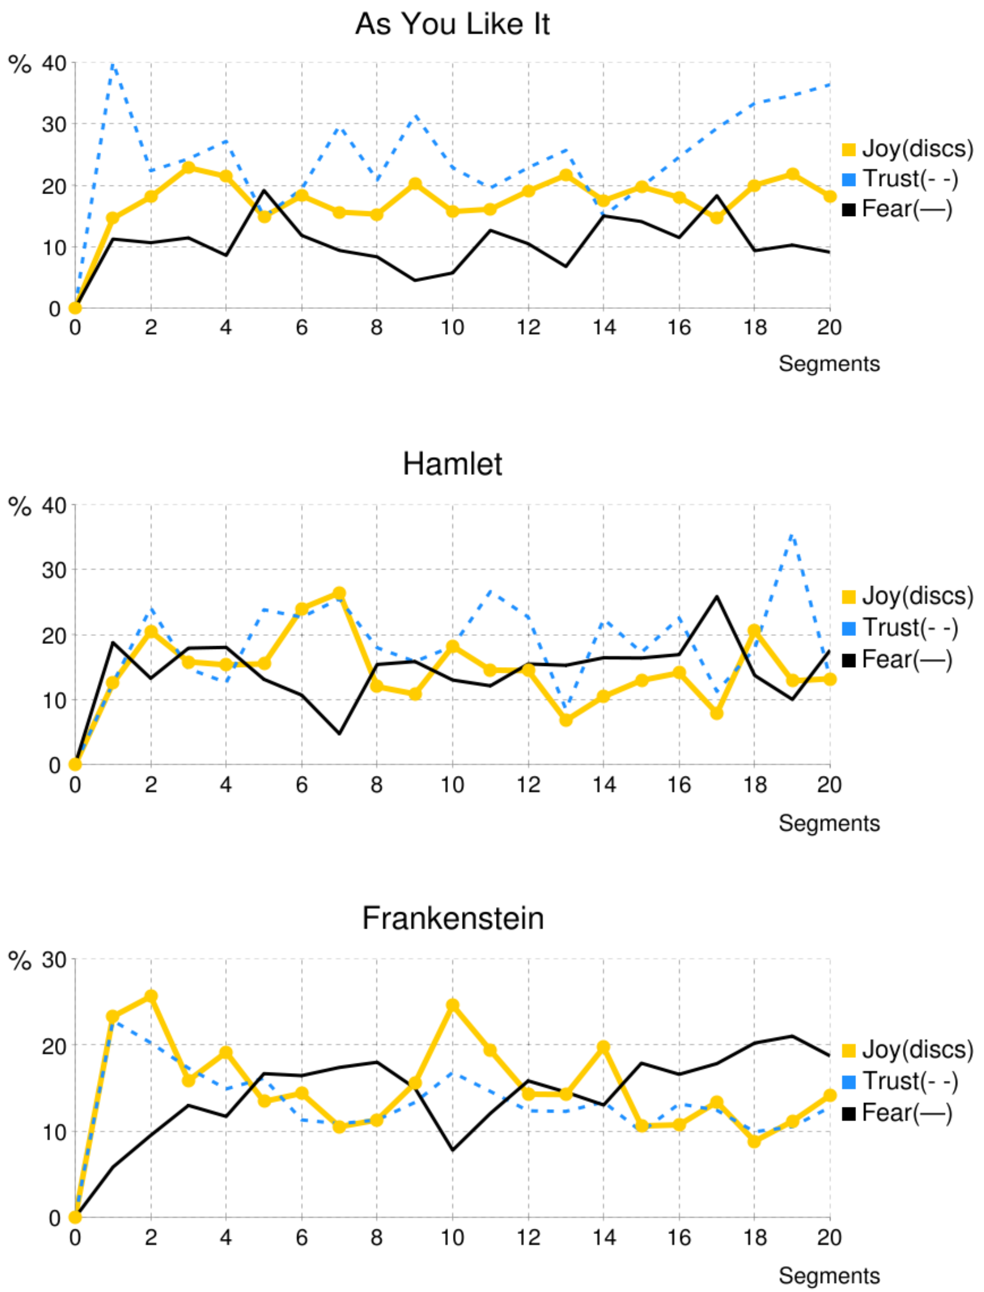
As You Like It
Hamlet
Frankenstein
%
Segments
Joy(discs)
Trust(- -)
Fear(—)

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.

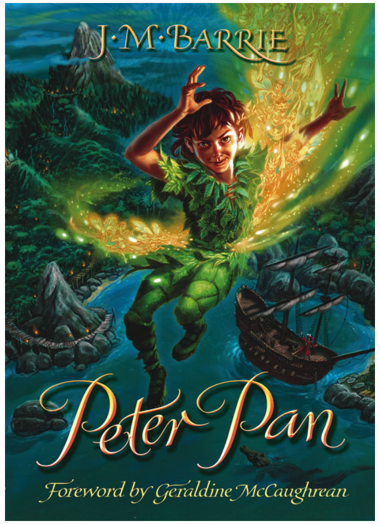

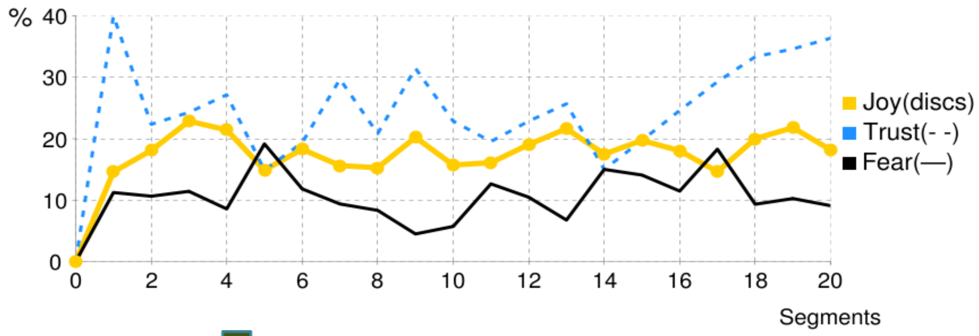

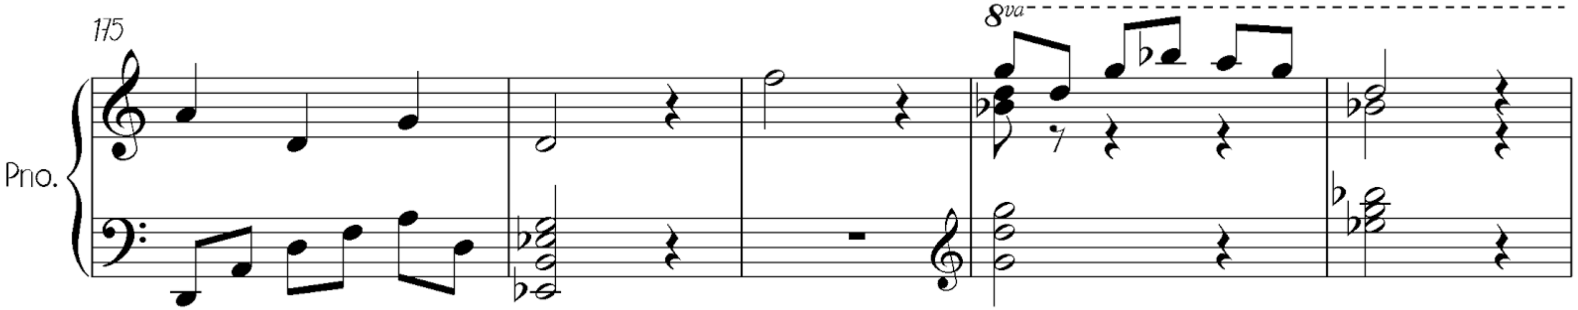

A method to generate music from literature.

- 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

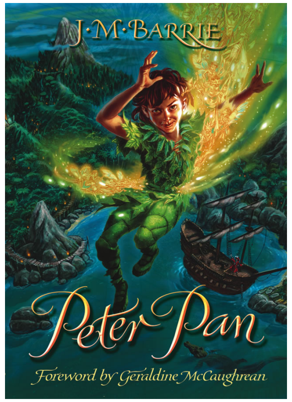

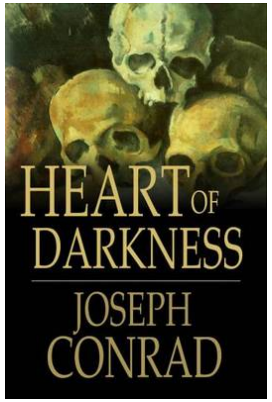

TransProse: www.musicfromtext.com
Music played 300,000 times since website launched in April 2014.

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

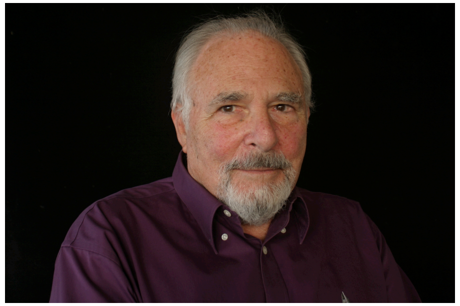

Paul Ekman
Psychologist and discoverer of micro expressions.

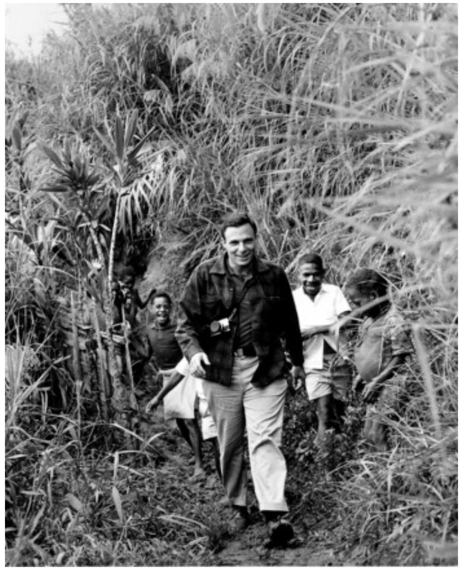

Lisa Barrett
University Distinguished Professor of Psychology, Northeastern University

- 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 drops for Ekman emotions when participants are given:
  - Just the pictures (no emotion word options)
  - Or say, two scowling faces and asked whether 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.

Generated Lexicon for Numerous Emotions

NRC Hashtag Emotion Lexicon: About 20,000 words associated with about 50 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 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

Svetlana Kiritchenko
NRC

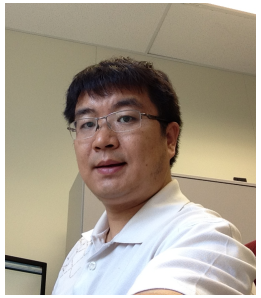

Xiaodan Zhu
NRC

SemEval Shared task on the Sentiment Analysis of Tweets

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.

Sentiment Analysis Competition

SemEval-2013: Classify Tweets, 44 teams

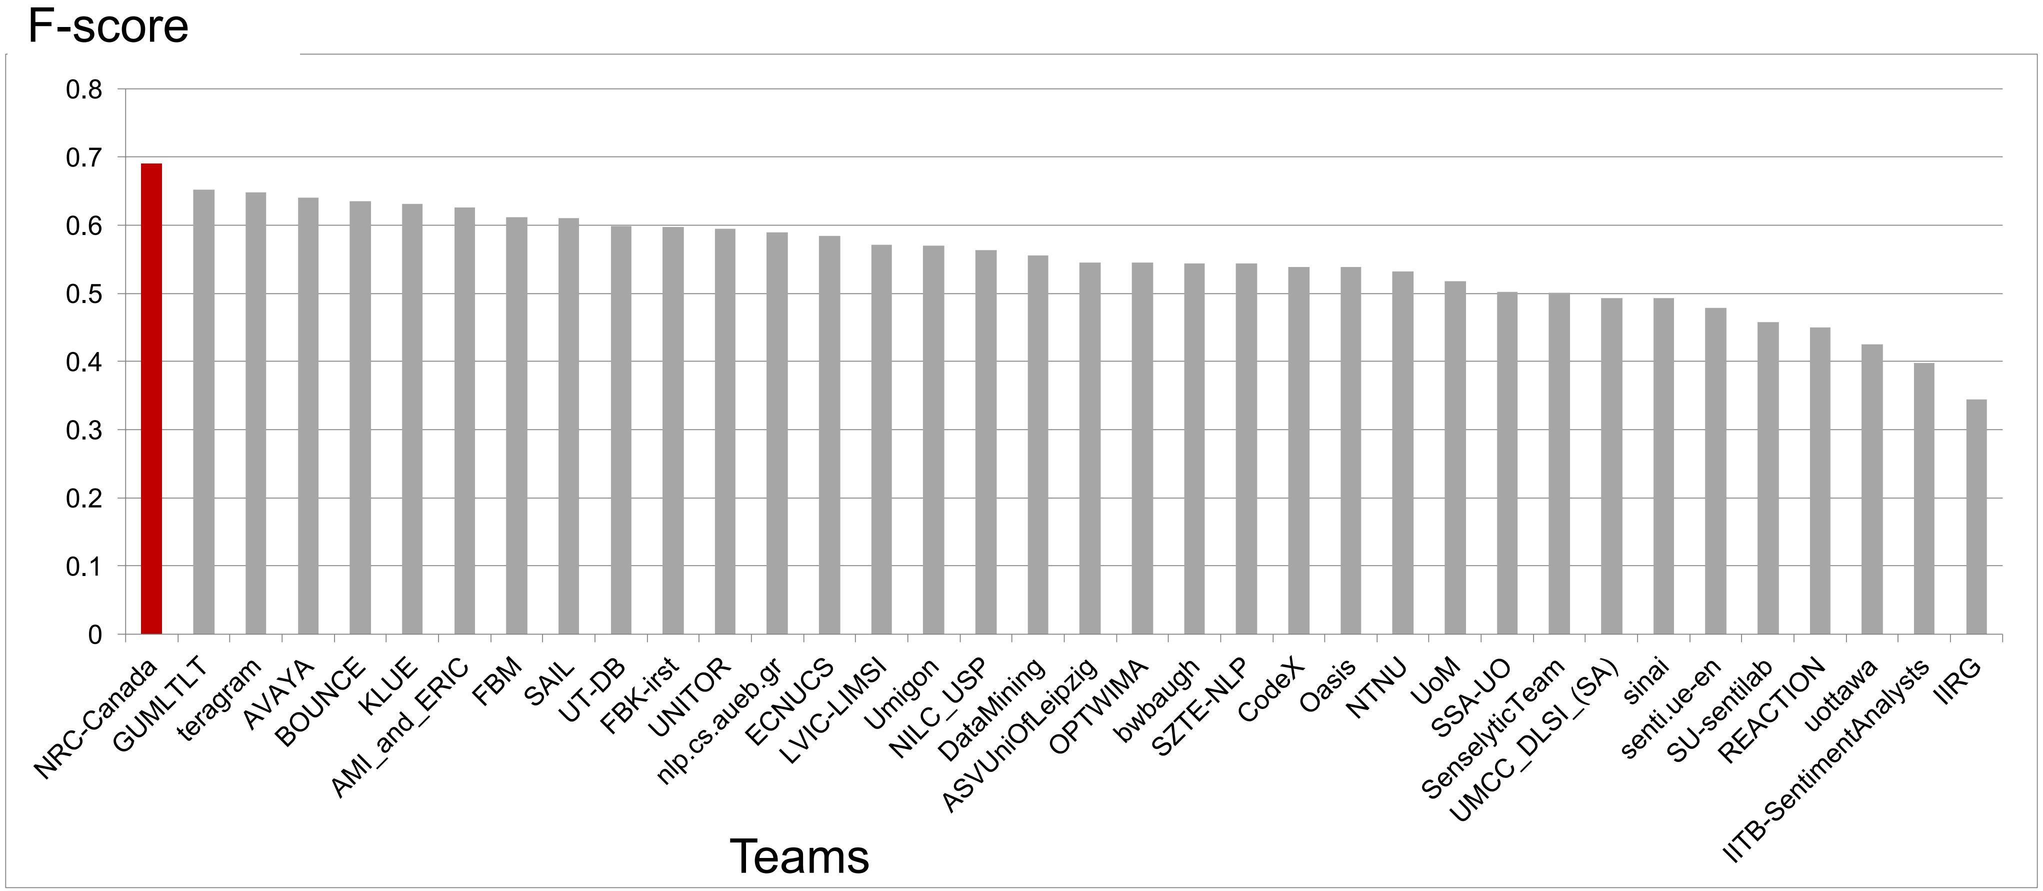

Sentiment Analysis Competition

SemEval-2013: Classify SMS messages, 30 teams

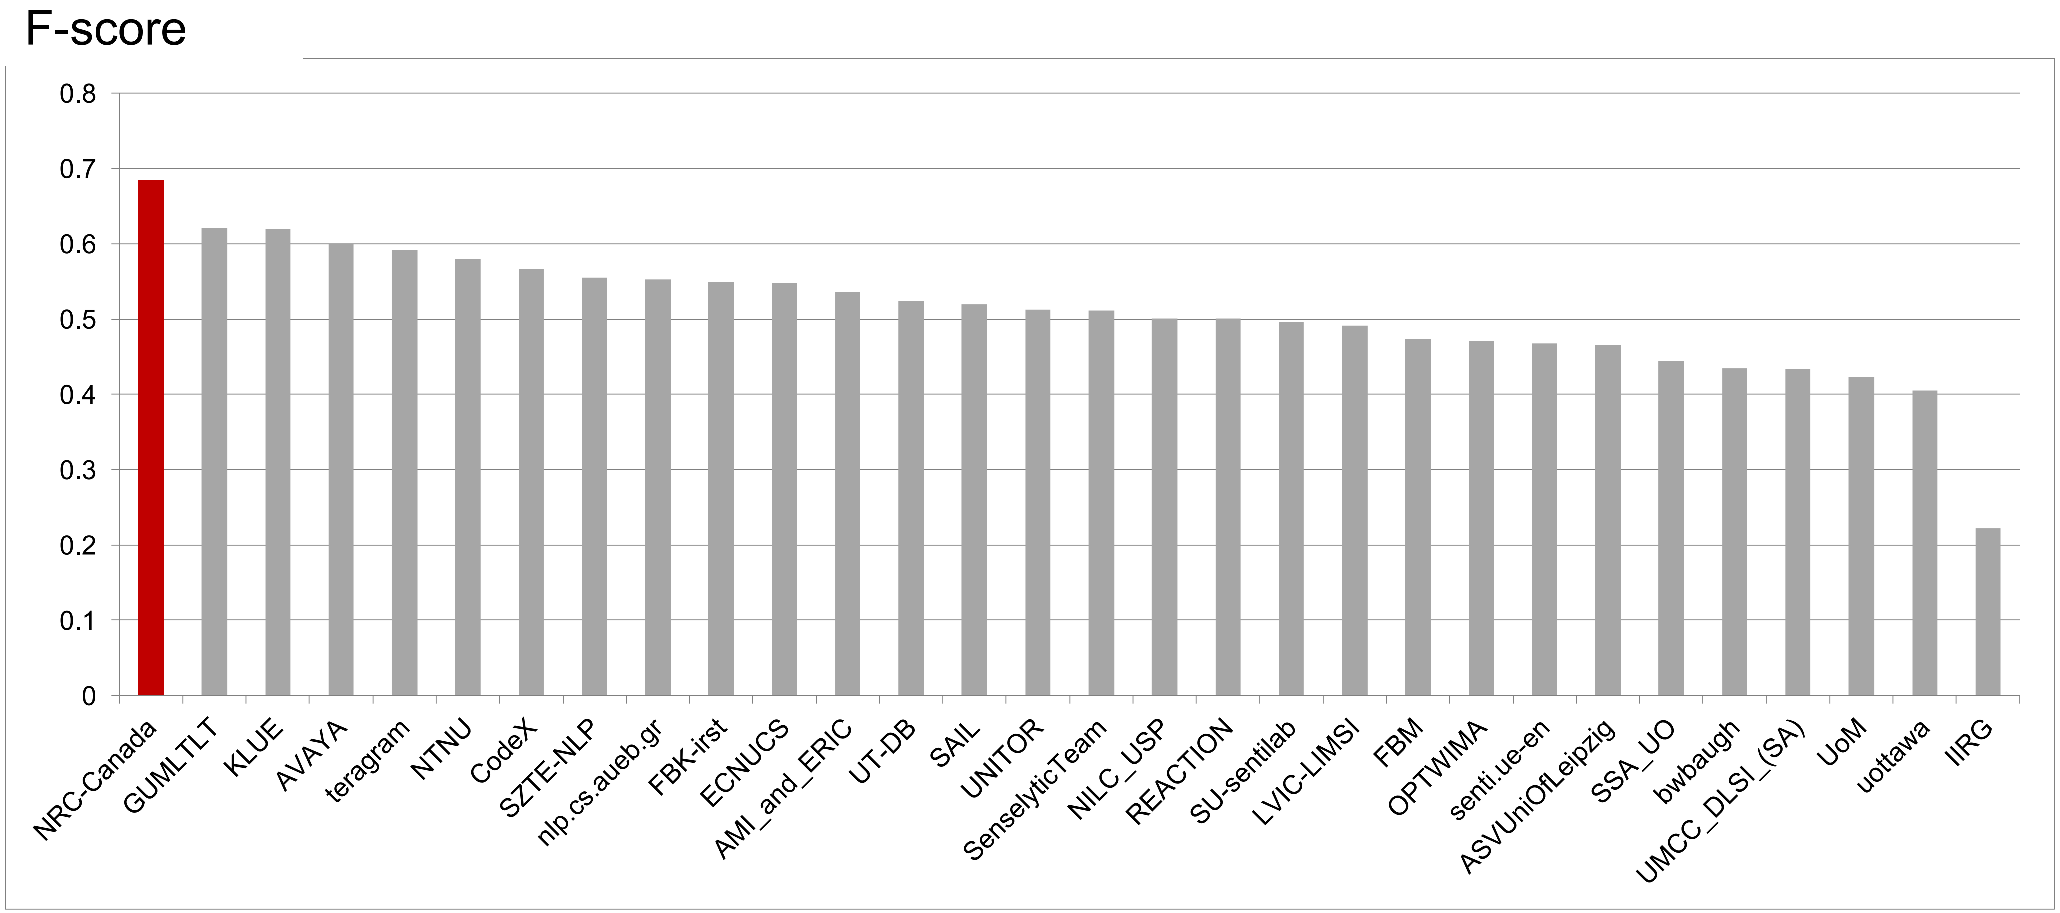

Sentiment Analysis of Social Media Texts (SemEval-2014 Task 9)

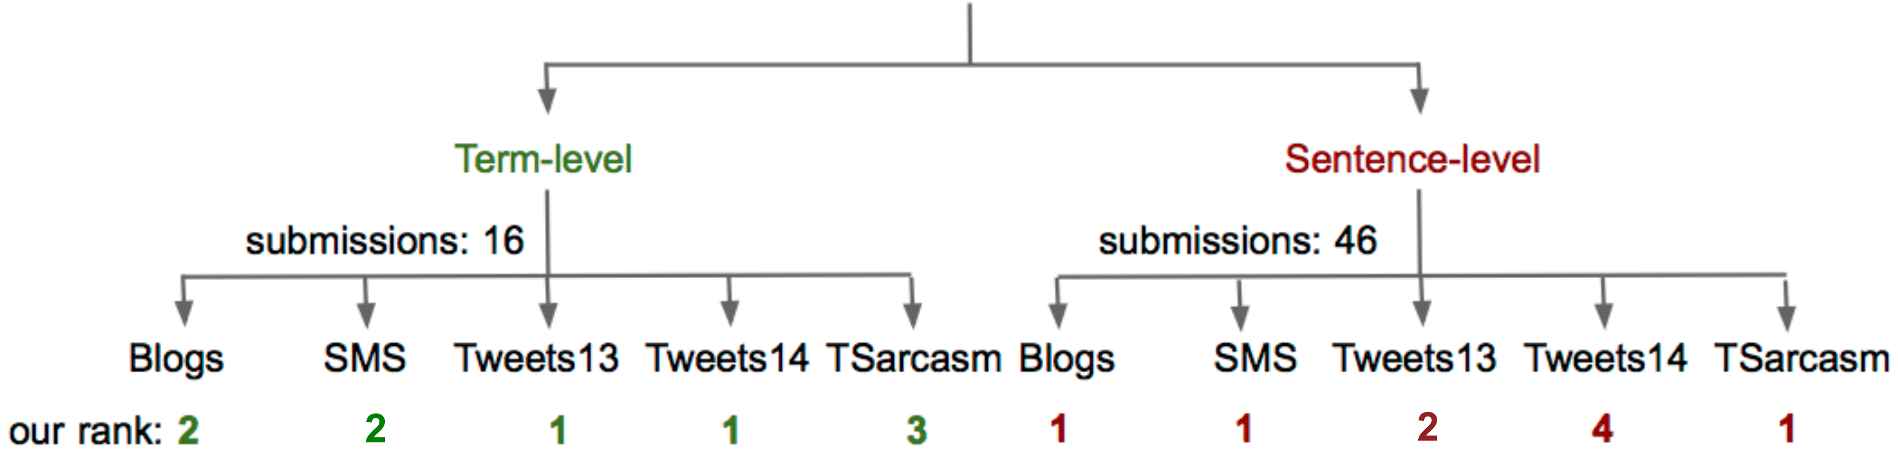

Sentiment Analysis of Social Media Texts (SemEval-2014 Task 9)

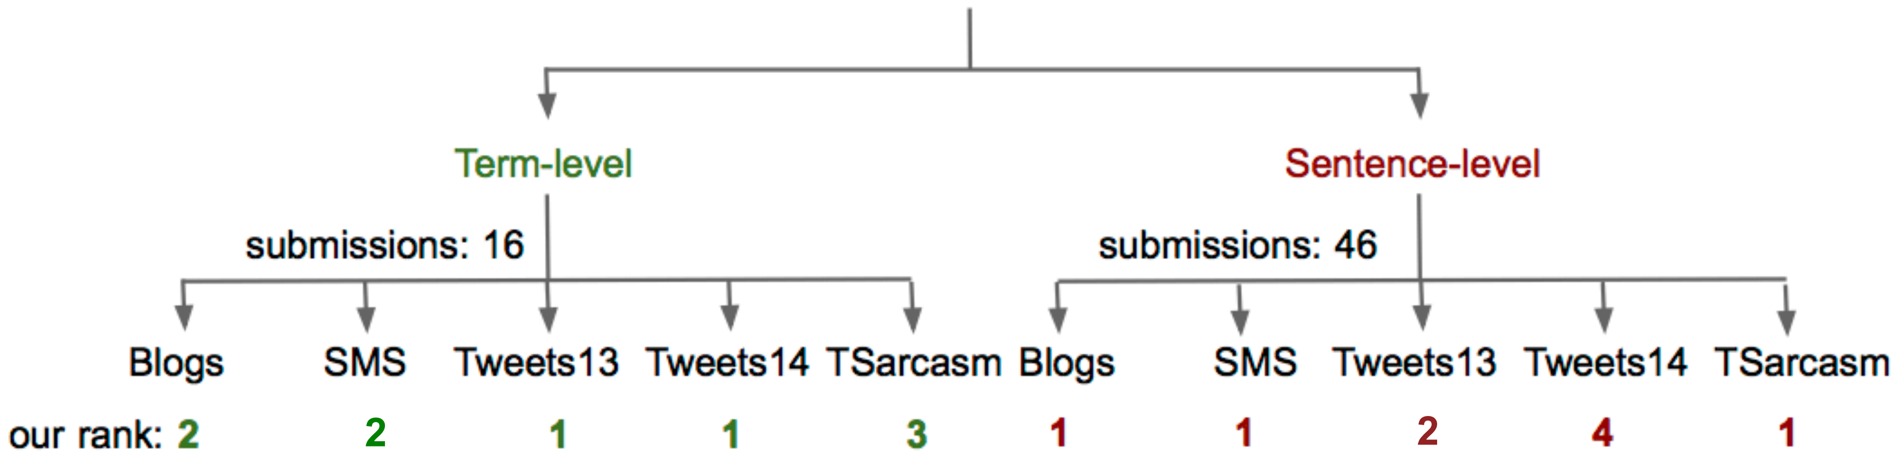

Aspect-Based Sentiment Analysis (SemEval-2014 Task 4)

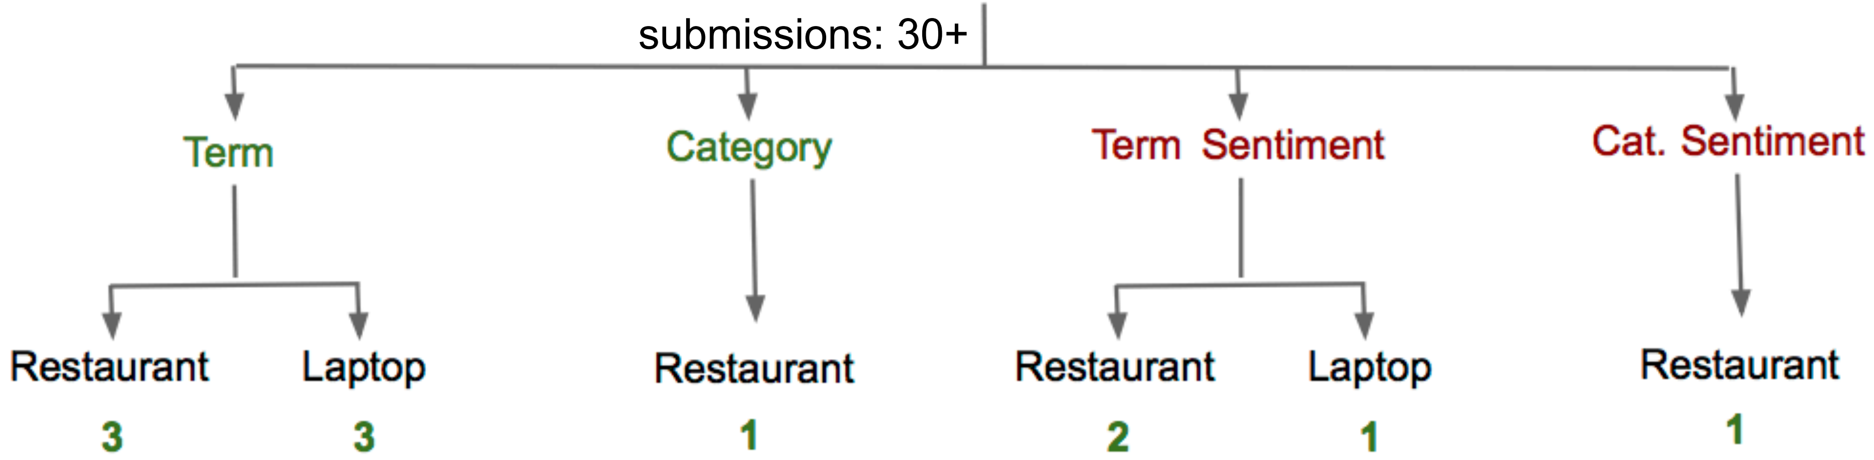

Feature Contributions (on Tweets)

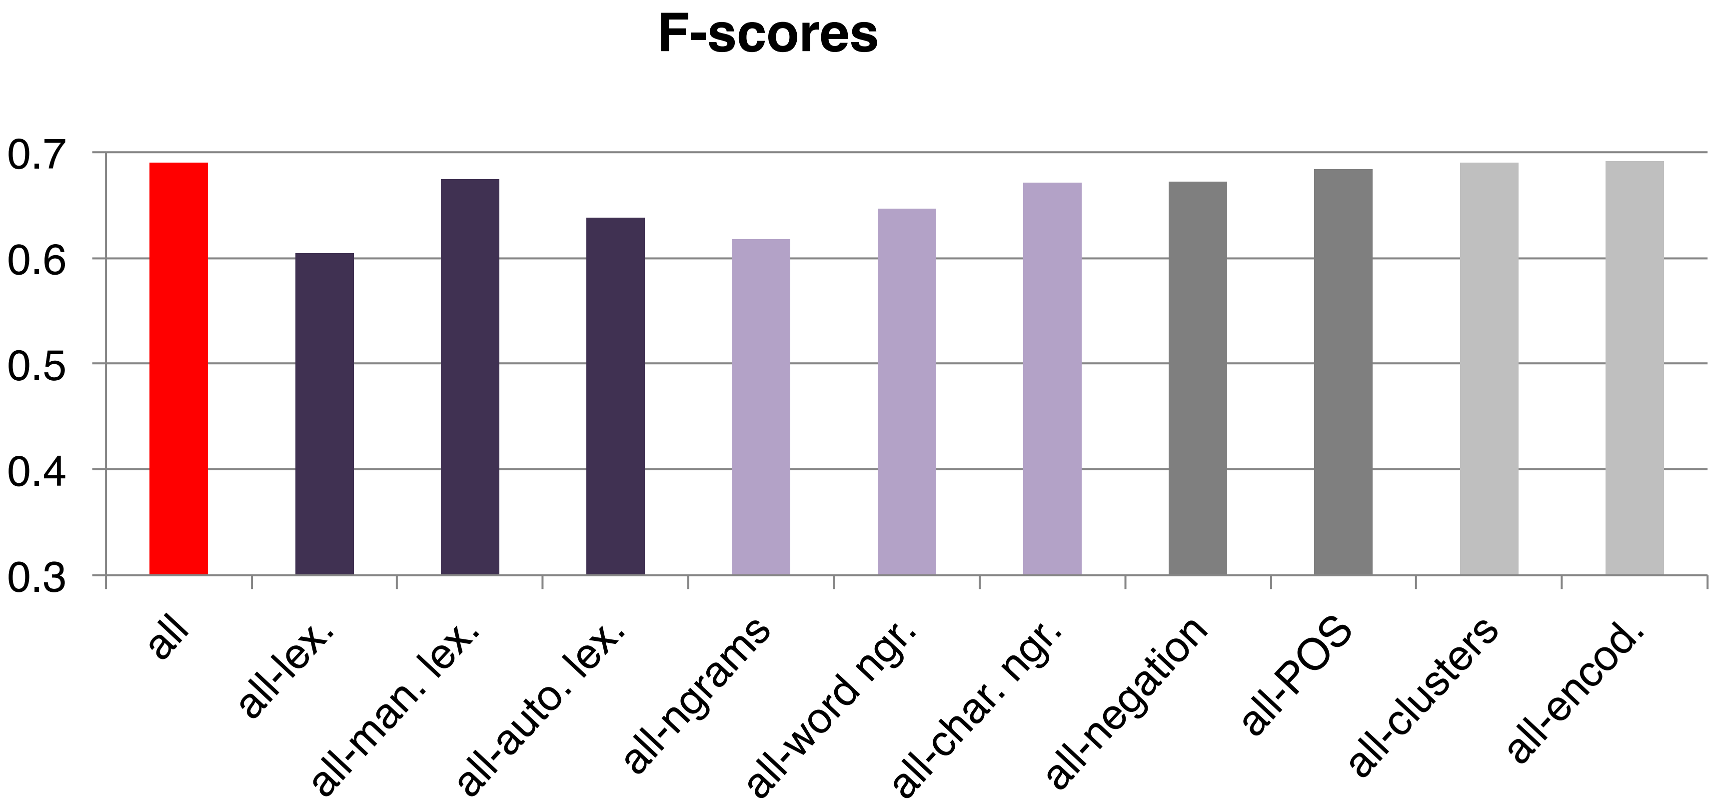

Detecting Stance in Tweets

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: 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 2:

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.

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

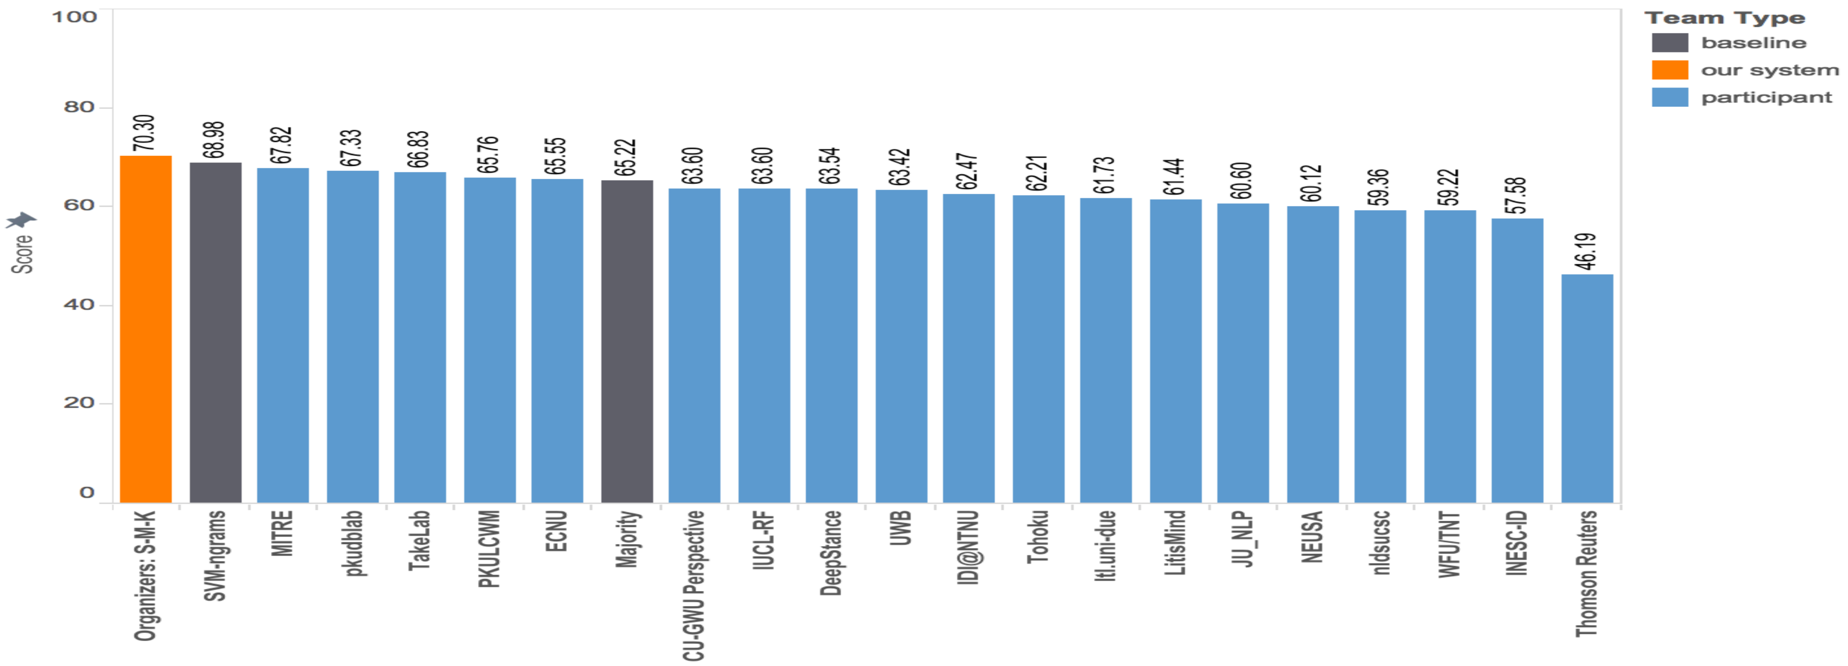

WASSA-2017 Shared Task on Emotion Intensity

http://saifmohammad.com/WebPages/EmotionIntensity-SharedTask.html

- ~7K tweets: 1500 to 2200 tweets per emotion
  - anger, fear, joy, sadness
- AffectiveTweets package: https://github.com/felipebravom/AffectiveTweets
  - Provides a collection of filters for extracting features for sentiment analysis and other related tasks
  - Includes features used in:
    - Kiritchenko et al. (2014): sentiment analysis
    - Mohammad et al. (2017): stance detection
- Baseline system: uses the AffectiveTweets and trains Weka regression models
  - LibLinear SVM regression

Competition Results

- Teams: twenty-two teams participated
- Our public baseline (AffectiveTweets): 0.66
- Top 3 teams: used feature vectors from the AffectiveTweets package

Team Name	r avg. (rank)
1. Prayas	0.747 (1)
2. IMS	0.722 (2)
3. SeerNet	0.708 (3)
4. UWaterloo	0.685 (4)
5. IITP	0.682 (5)
6. YZU NLP	0.677 (6)
7. YNU-HPCC	0.671 (7)
8. TextMining	0.649 (8)
9. XRCE	0.638 (9)
10. LIPN	0.619 (10)
11. DMGroup	0.571 (11)
12. Code Wizards	0.527 (12)
13. Todai	0.522 (13)
14. SGNLP	0.494 (14)
15. NUIG	0.494 (14)
16. PLN PUCRS	0.483 (16)
17. H.Niemtsov	0.468 (17)
18. Tecnolengua	0.442 (18)
19. GradAscent	0.426 (19)
20. SHEF/CNN	0.291 (20)
21. deepCybErNet	0.076 (21)
Late submission	
∗ SiTAKA	0.631

	Team																					
Features	1	2	3	4	5	6	7	8	9	*	10	11	12	13	14	15	16	17	18	19	20	21
N-grams				✓									✓									
CN													✓									
WN				✓									✓			✓						
Word Embeddings	✓	✓	✓	✓	✓	✓	✓	✓		✓			✓	✓	✓	✓				✓		
Glove			✓	✓	✓	✓	✓	✓		✓				✓		✓				✓		
Emoji Vectors			✓	✓																		
Word2Vec	✓	✓	✓	✓																		
Other								✓					✓		✓							
Sentence Embeddings																						
CNN	✓	✓				✓	✓	✓		✓					✓						✓	✓
LSTM	✓	✓			✓	✓	✓	✓						✓		✓				✓		
Other				✓												✓				✓	✓	
Affective Lexicons		✓	✓	✓	✓	✓		✓	✓	✓				✓				✓	✓	✓		
AFINN	✓	✓	✓		✓			✓														
ANEW		✓																				
BingLiu	✓	✓	✓		✓			✓	✓													
Happy Ratings		✓																				
Lingmotif																			✓			
LIWC																	✓					
MPQA	✓	✓	✓		✓			✓														
NRC-Aff-Int	✓		✓	✓				✓														
NRC-EmoLex	✓	✓	✓	✓	✓			✓	✓													
NRC-Emoticon-Lex	✓		✓	✓				✓					✓									
NRC-Hash-Emo	✓	✓	✓	✓	✓			✓	✓													
NRC-Hash-Sent		✓	✓	✓	✓			✓														
NRC-Hashtag-Sent.	✓		✓	✓																		
NRC10E	✓	✓	✓					✓														
Sentiment140	✓	✓	✓	✓				✓														
SentiStrength		✓	✓					✓														
SentiWordNet	✓	✓	✓	✓	✓			✓														
Vader					✓																	
Word.Affect			✓																			
In-house lexicon	✓								✓								✓					
Linguistic Features									✓													
Dependency Parser									✓													

SemEval-2018 Task 1: Affect in Tweets

https://competitions.codalab.org/competitions/17333

Tasks:

- emotion intensity regression
- emotion intensity ordinal classification
- sentiment intensity regression
- sentiment analysis, ordinal classification
- emotion classification task

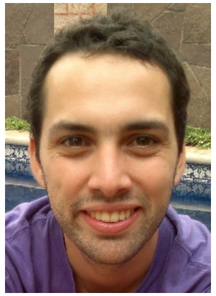

Felipe José Bravo Márquez

Mohammad Salameh

Svetlana Kiritchenko
NRC

Other Recent Work

Ekaterina Shutova

Peter Turney

Svetlana Kiritchenko
NRC

- **Metaphor as a Medium for Emotion**

  Metaphor as a Medium for Emotion: An Empirical Study. Saif M. Mohammad, Ekaterina Shutova, and Peter Turney. In *Proceedings of the Joint Conference on Lexical and Computational Semantics (*Sem)*, August 2016, Berlin, Germany.

- **Pharmacovigilance**

  NRC-Canada at SMM4H Shared Task: Classifying Tweets Mentioning Adverse Drug Reactions and Medication Intake. Svetlana Kiritchenko, Saif M. Mohammad, Jason Morin, and Berry de Bruijn (2017). In *Proceedings of the Social Media Mining for Health Applications Workshop at AMIA-2017*, Washington, DC, USA.

  Official Rankings: Our team (NRC-Canada) ranked first in the AMIA Shared Task on detecting adverse drug reactions in tweets.

Ethics in Emotion Systems

- Protecting vulnerable members of the society from being exploited
  - Kids, elderly, those less aware of the dangers of social media data mining

  Facebook told advertisers it can identify teens feeling 'insecure' and 'worthless'

- Avoiding explicit and implicit discrimination based on race, religion, gender, etc.
  - Computer algorithms learn biases from data

  Microsoft's racist tweet bot.

  Word embeddings have inappropriate biases.

- Dealing with fake news, persuasion tactics, person-specific framing

Ethics in Emotion Systems

- Ethics must be a part of system design considerations, right from the start.
- In addition to accuracy, systems must be evaluated for inappropriate biases.

SemEval-2018 Task 1 "Affect in Tweets": a shared task on a number of emotion detection tasks will have a separate evaluation component for inappropriate biases.

A first for SemEval.

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

- Sentiment and emotion lexicons
- Sentiment and emotion labeled data
- Links to shared tasks
- Interactive visualizations
- Tutorials and book chapters on sentiment and emotion analysis

Saif M Mohammad
www.saifmohammad.com
saif.mohammad@nrc-cnrc.gc.ca

Ongoing Work

- Understanding controversial issues and argumentation
- Understanding affect associations of phonemes and syllables
- Ethics in emotion systems