Affect Associations in Creative Language

Saif Mohammad
National Research Council Canada

Creative Language

Extraordinary, everyday

Creative Language

- Stories and poems
- Metaphors
- Hyperbole
- Sarcasm and irony
- Noun-noun compounds
  - soccer mom, mountain bike
- Opposing polarity phrases
  - happy accident, crazy cool, epic fail
- Hashtag words
  - #loveumom, #throwbackthursday

creative language
has emotional impact

This talk

Will explore affect associations in creative language

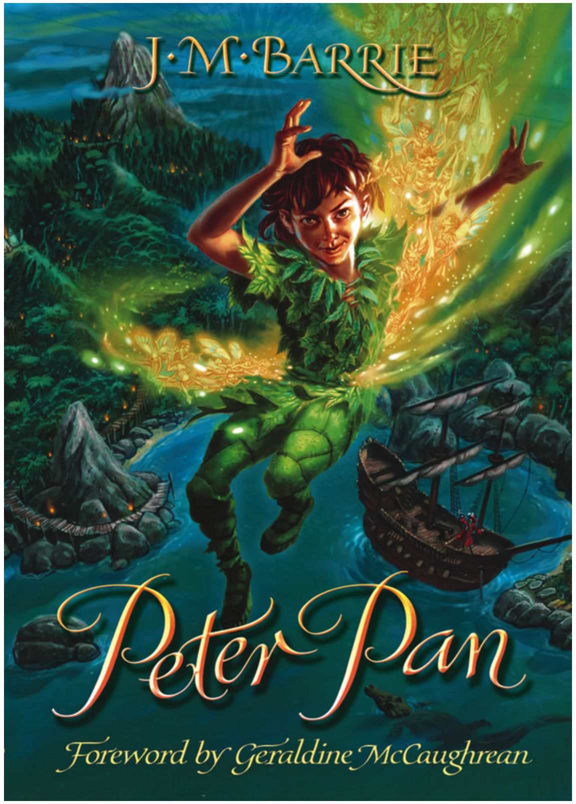

Stories

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?

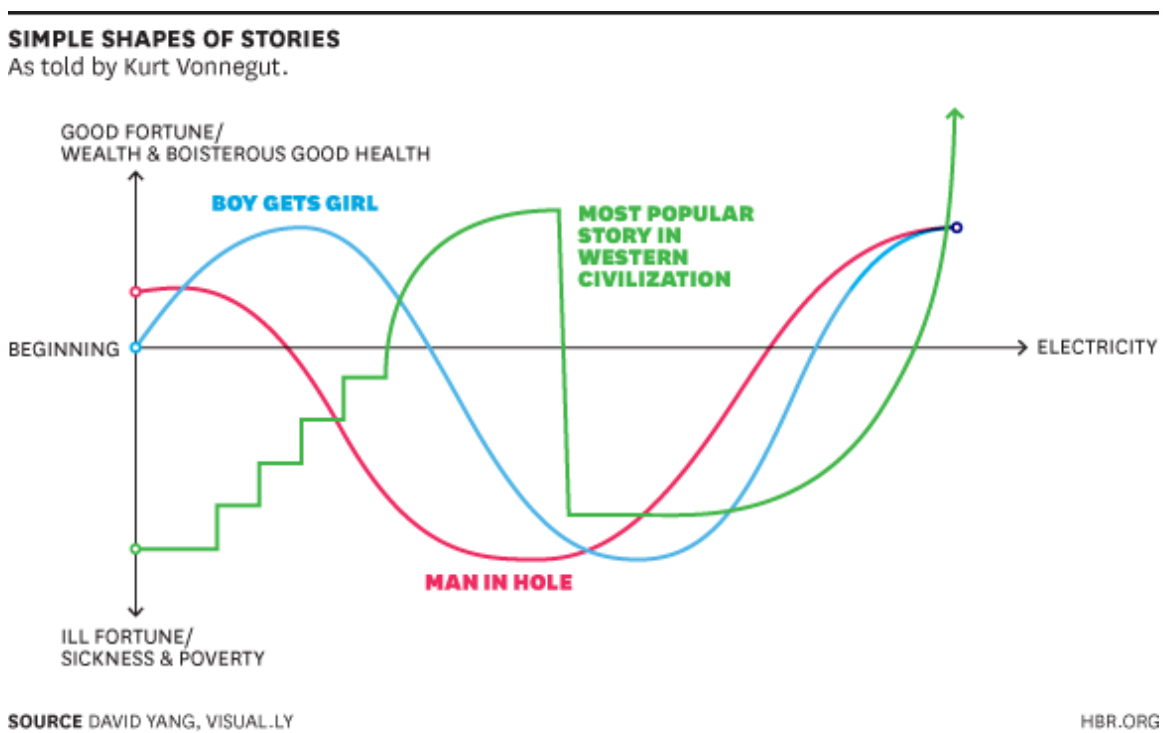

Word Associations

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.

Which Emotions?

Charles Darwin

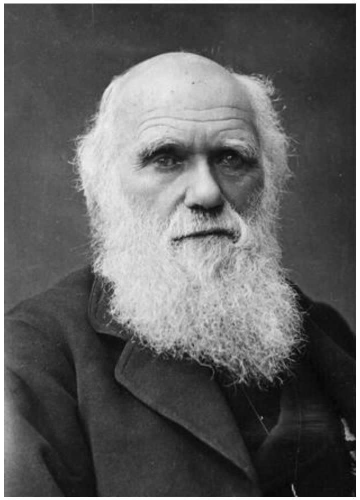

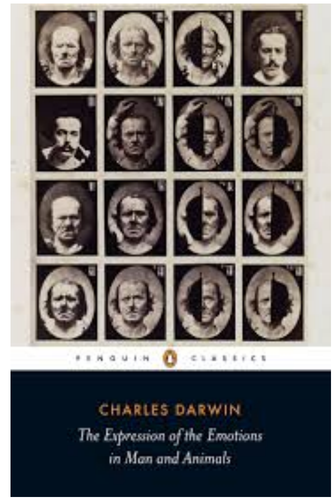

FIG. 20.—Terror,
from a photograph by Dr. Duchenne.

- 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

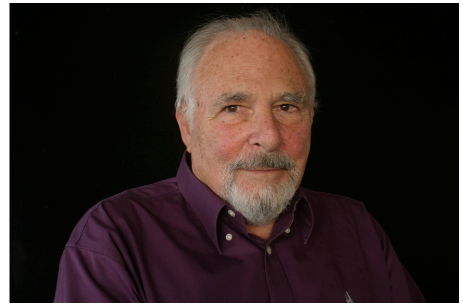

Paul Ekman
Psychologist and discoverer of micro expressions.

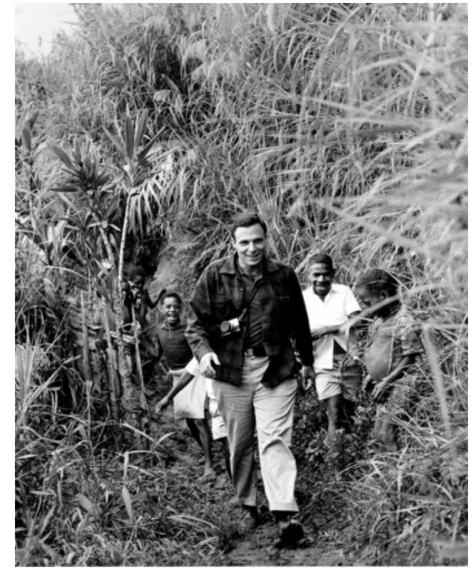

- Circa 1950's, Margaret Mead and others believed facial expressions and their meanings were culturally determined
  - behavioural learning processes
- Paul Ekman provided the strongest evidence to date that Darwin, not Margaret Mead, was correct in claiming facial expressions are universal
- Found universality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

Paul Ekman, 1971: Six Basic Emotions

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

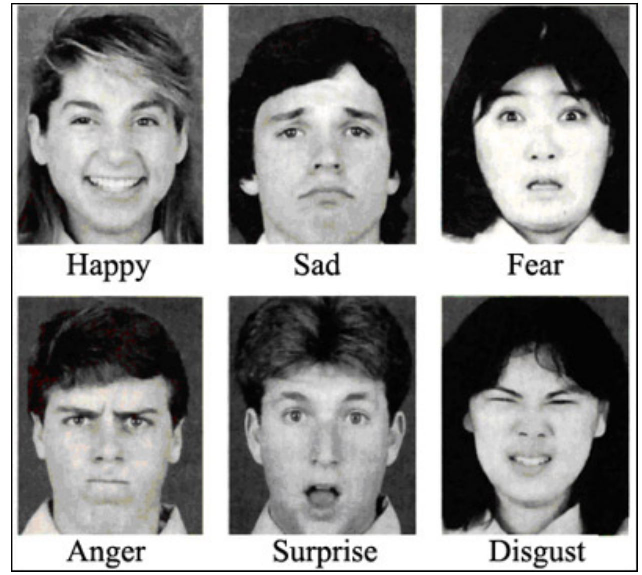

Plutchik, 1980: Eight Basic Emotions

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

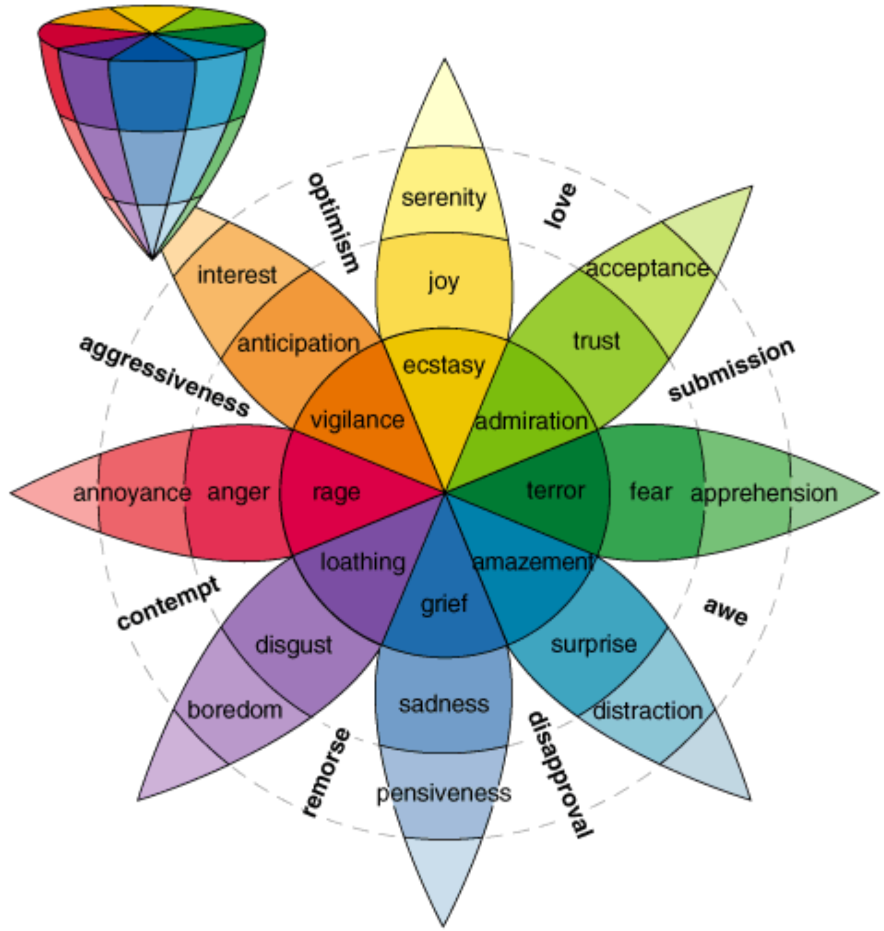

Goal: We chose to capture word-emotion associations for the 8 Plutchik emotions.

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

- 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.

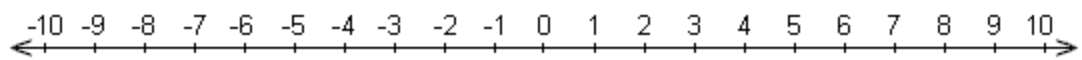

What if you want to capture fine-grained intensity of emotion/sentiment?

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 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

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

Robustness of the Annotations

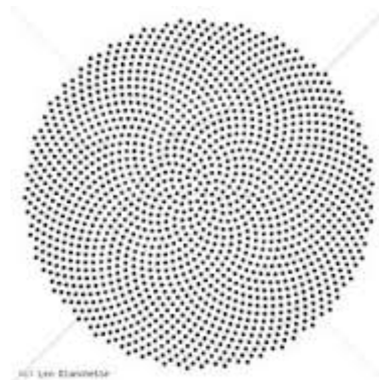

BWS annotations

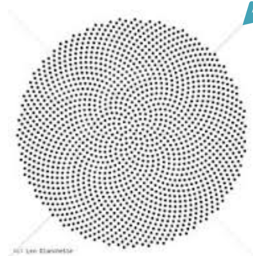

half of the BWS annotations

ranking of words by sentiment

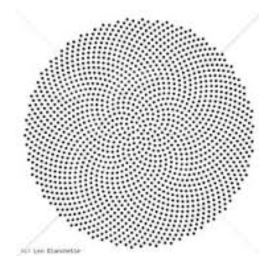

half of the BWS annotations

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

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

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.

Tony Yang, Simon Fraser University

Visualizing Emotions in Text

Analysis of Emotion Words in Books

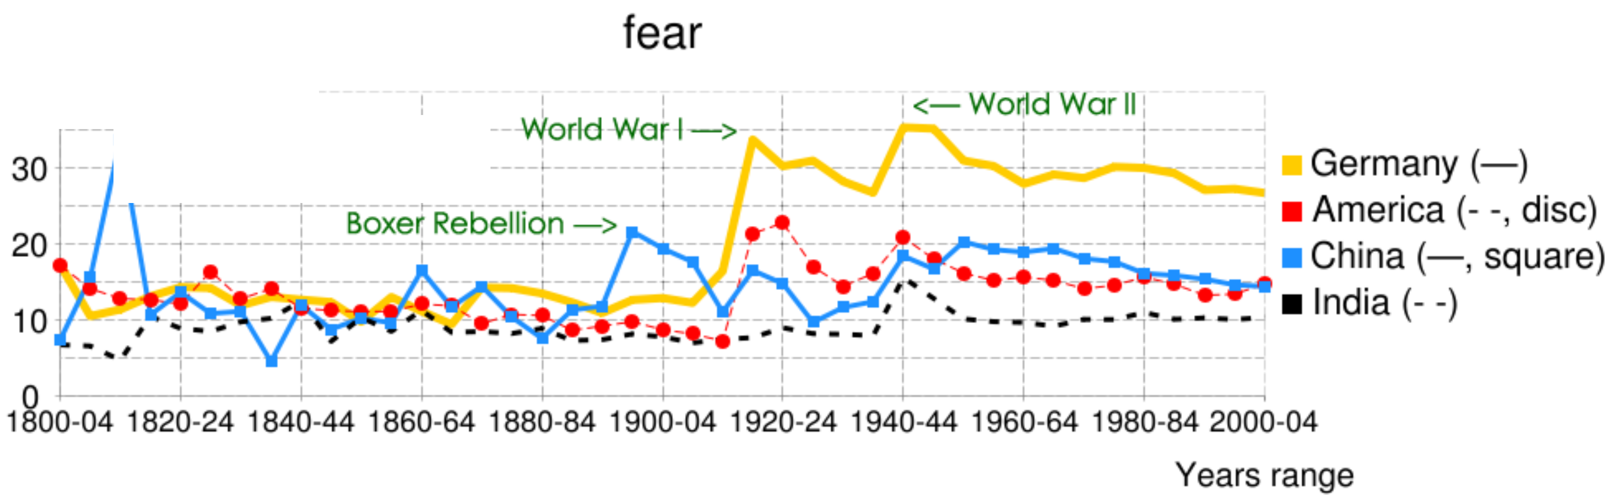

Percentage of fear words in close proximity to occurrences of America, China, Germany, and India in books.

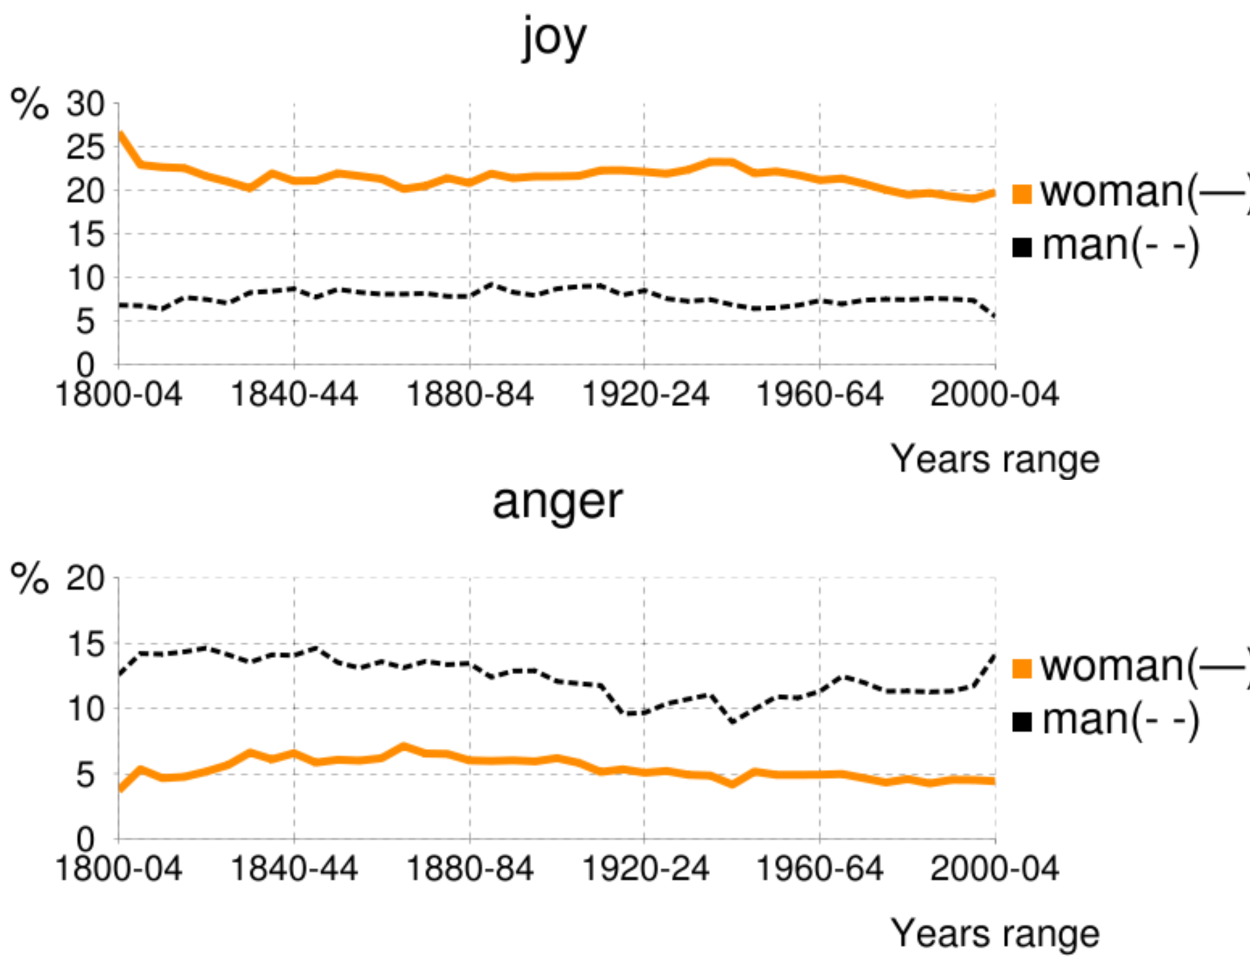

Percentage of joy and anger words in close proximity to occurrences of man and woman in books.

Tracking Emotions in Stories

- Can we track the change in distribution of emotion words?

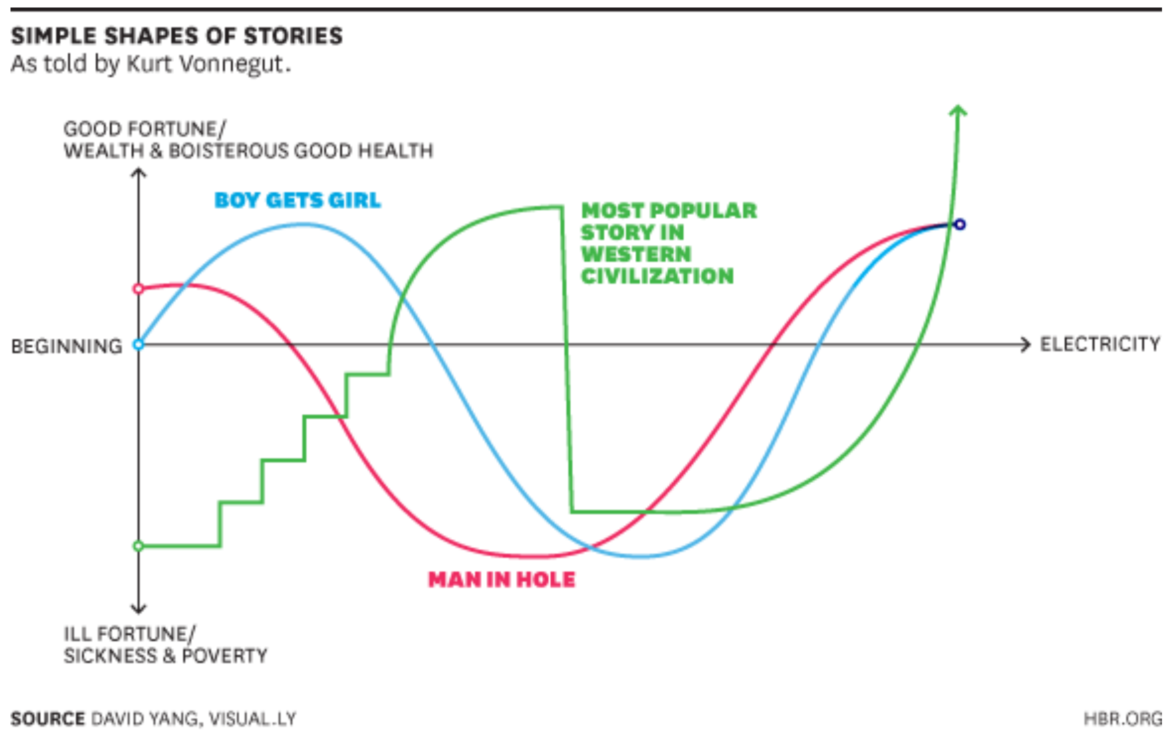

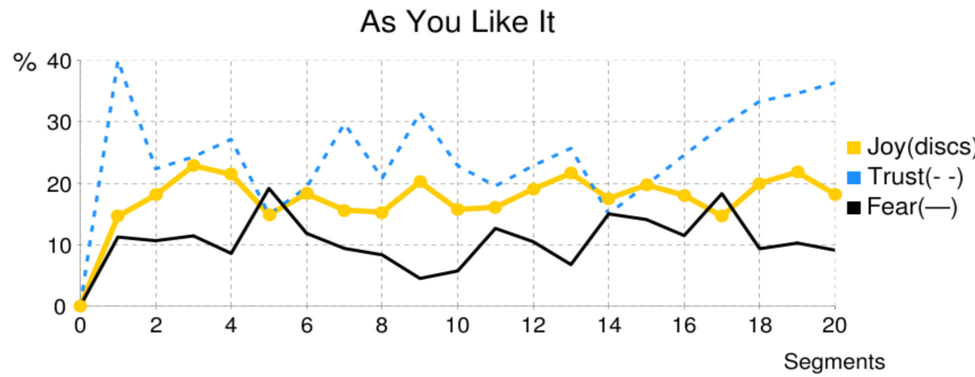

As You Like It
% 40
30
20
10
0
0 2 4 6 8 10 12 14 16 18 20
Segments
Joy(discs)
Trust(- -)
Fear(—)

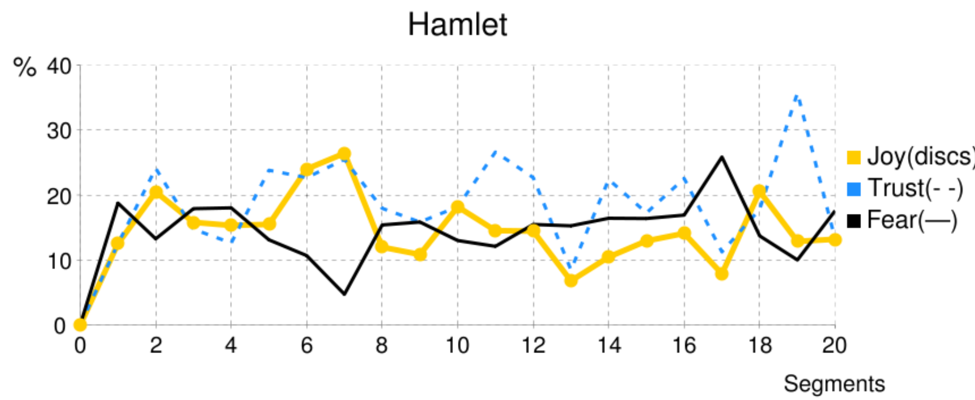

Hamlet
% 40
30
20
10
0
0 2 4 6 8 10 12 14 16 18 20
Segments
Joy(discs)
Trust(- -)
Fear(—)

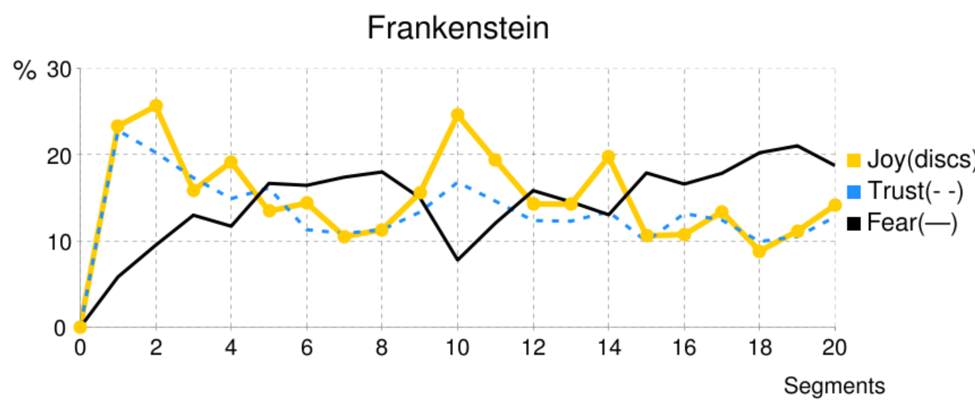

Frankenstein
% 30
20
10
0
0 2 4 6 8 10 12 14 16 18 20
Segments
Joy(discs)
Trust(- -)
Fear(—)

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.

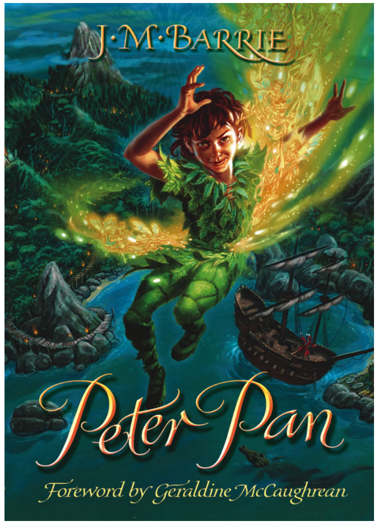

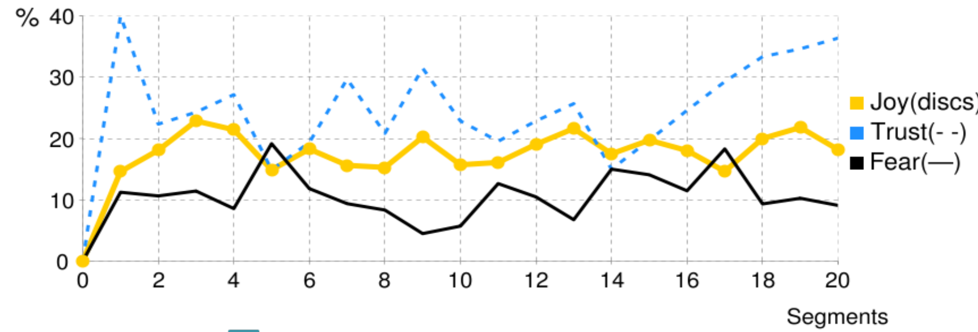

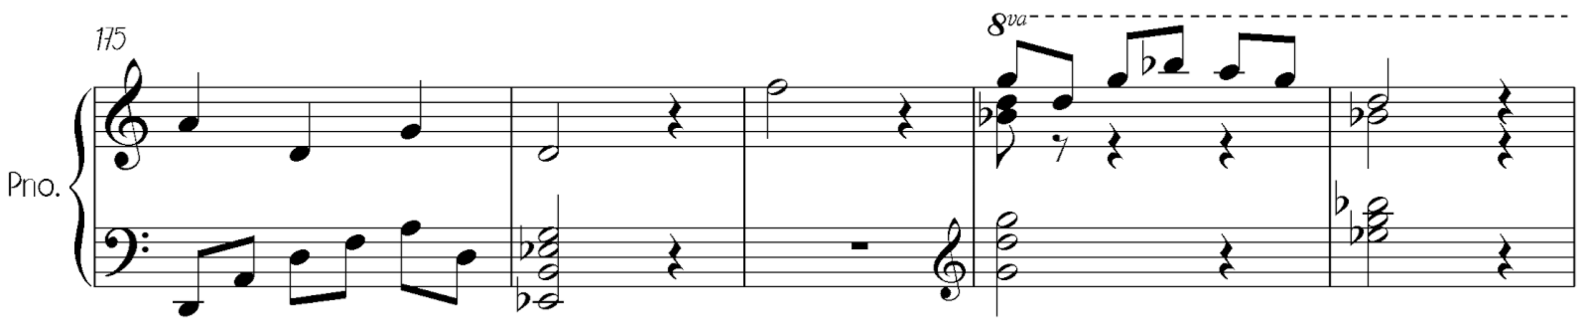

A method to generate music from literature.

- 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

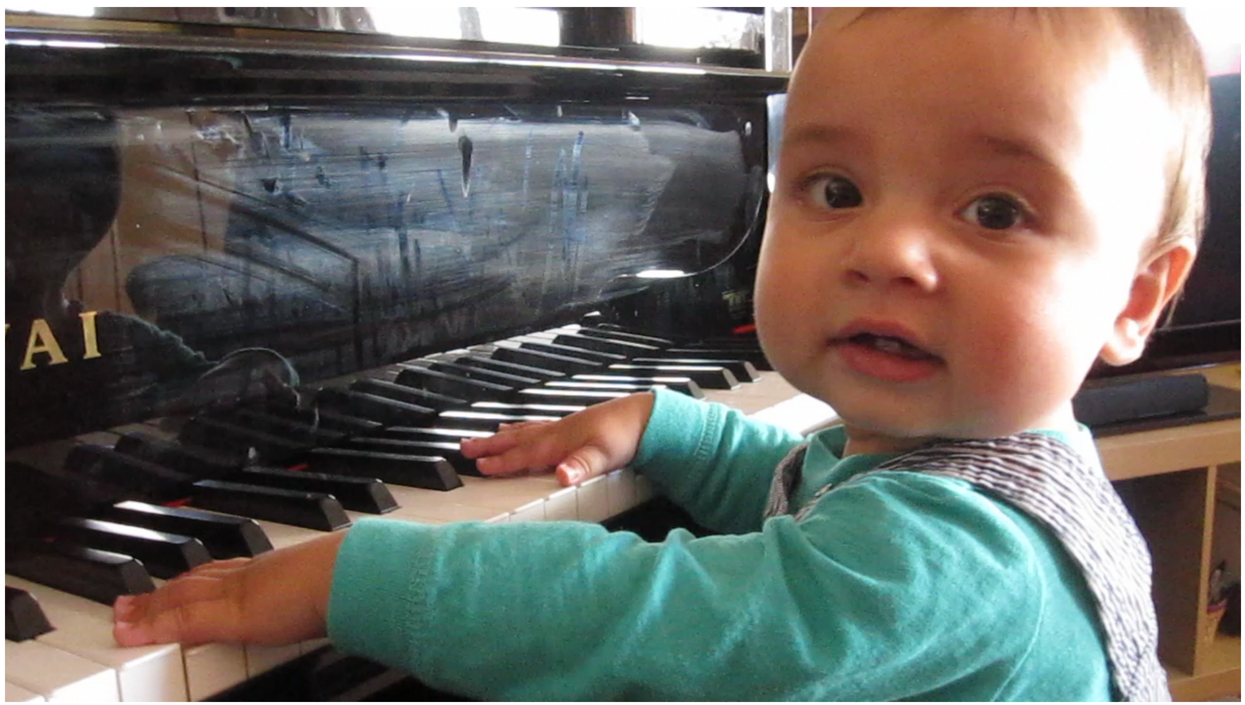

Music-Emotion Associations

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

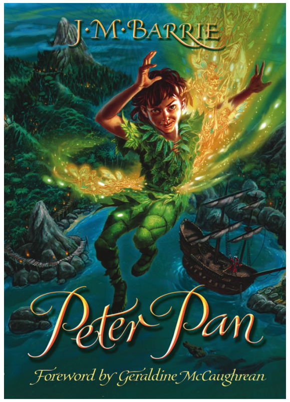

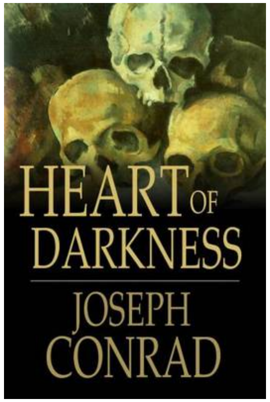

TransProse: www.musicfromtext.com
Music played 300,000 times since website launched in April 2014.

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

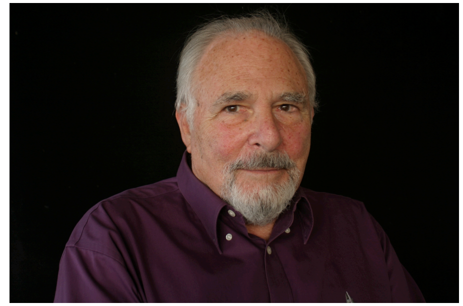

Paul Ekman
Psychologist and discoverer of micro expressions.

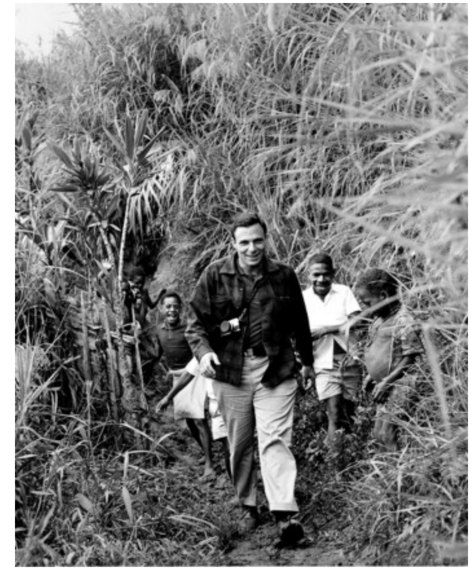

Lisa Barrett
University Distinguished Professor of Psychology, Northeastern University

- 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 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 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

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.

Sentiment Analysis Competition

SemEval-2013: Classify Tweets, 44 teams

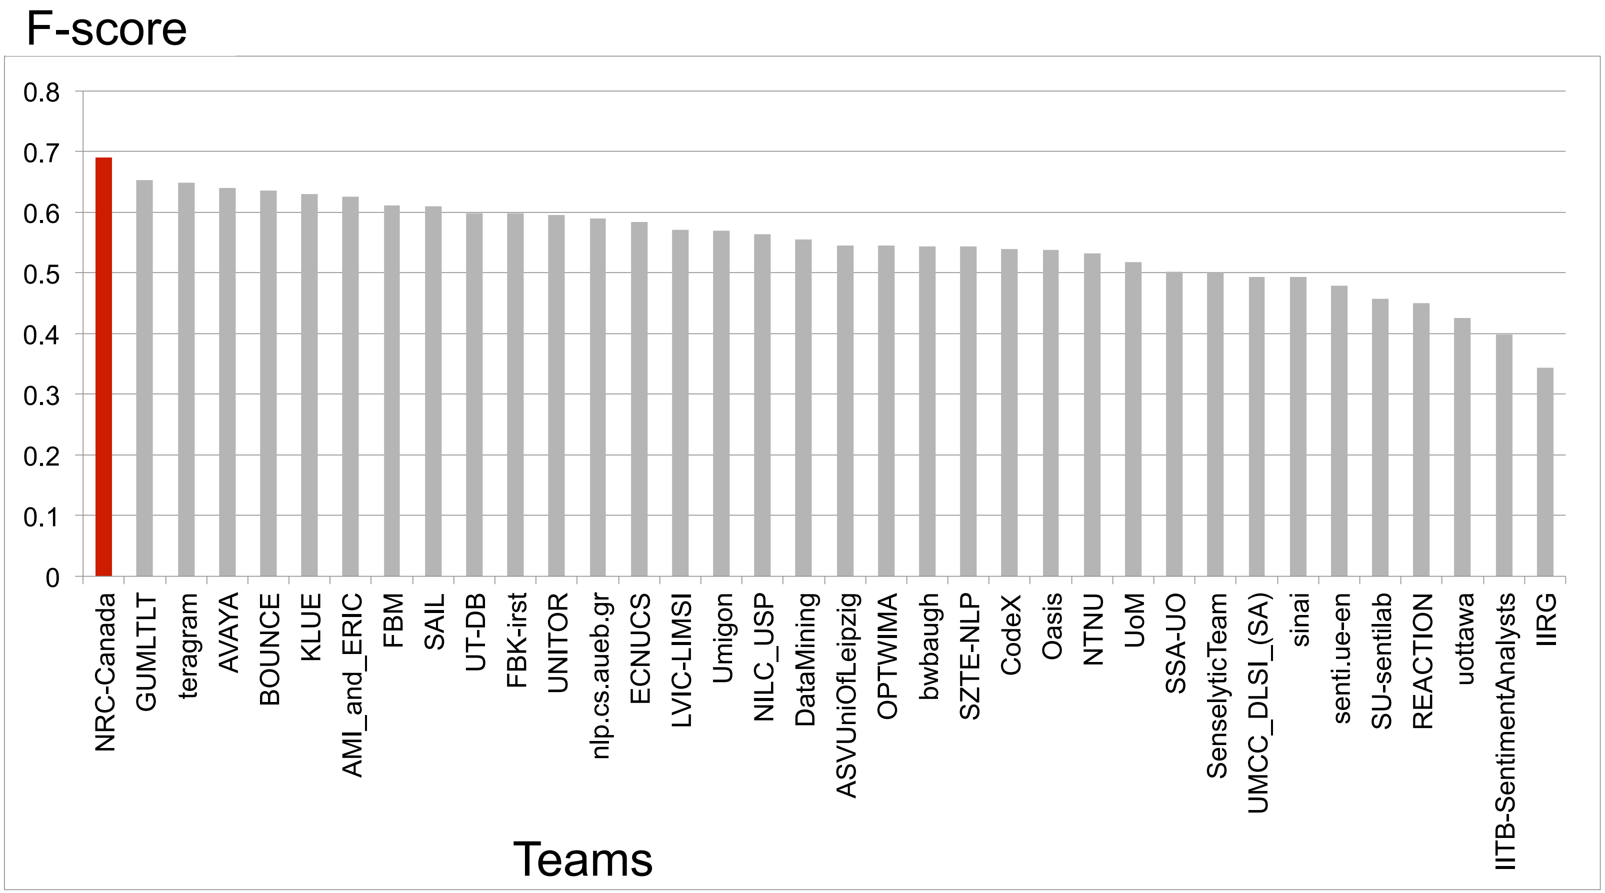

Sentiment Analysis Competition

SemEval-2013: Classify SMS messages, 30 teams

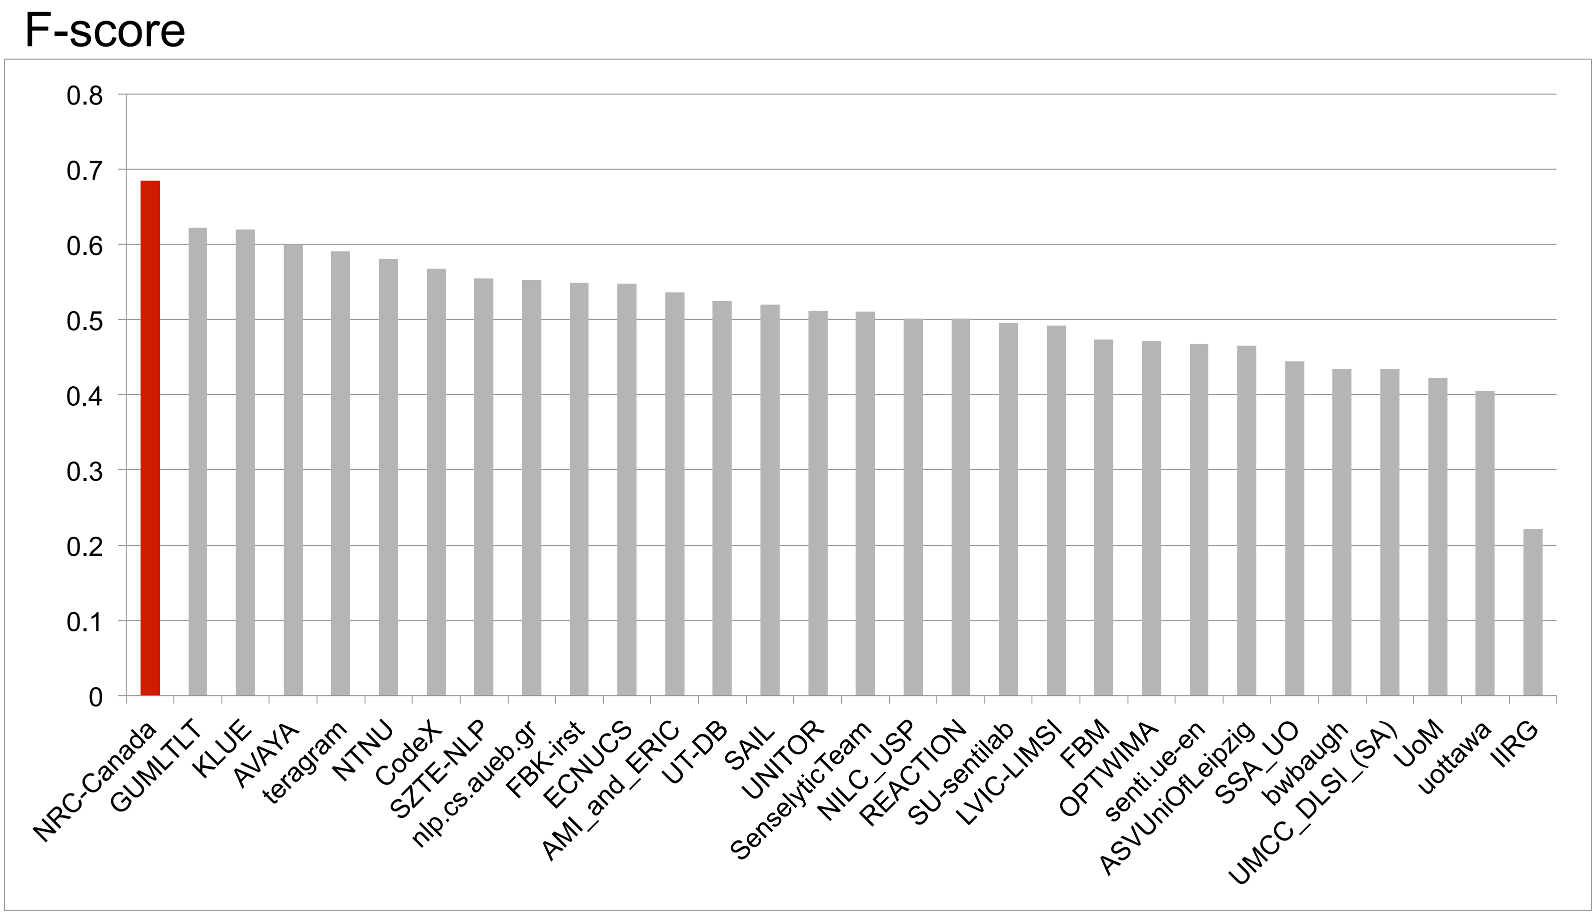

Sentiment Analysis of Social Media Texts (SemEval-2014 Task 9)

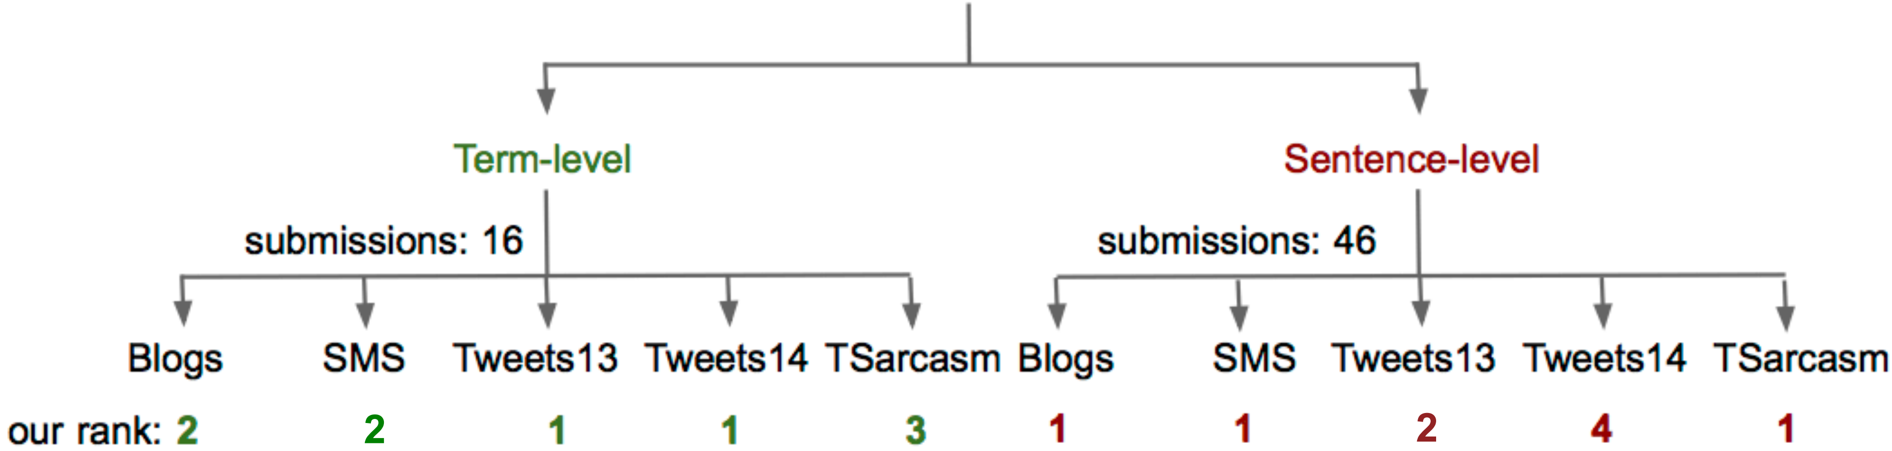

Sentiment Analysis of Social Media Texts (SemEval-2014 Task 9)

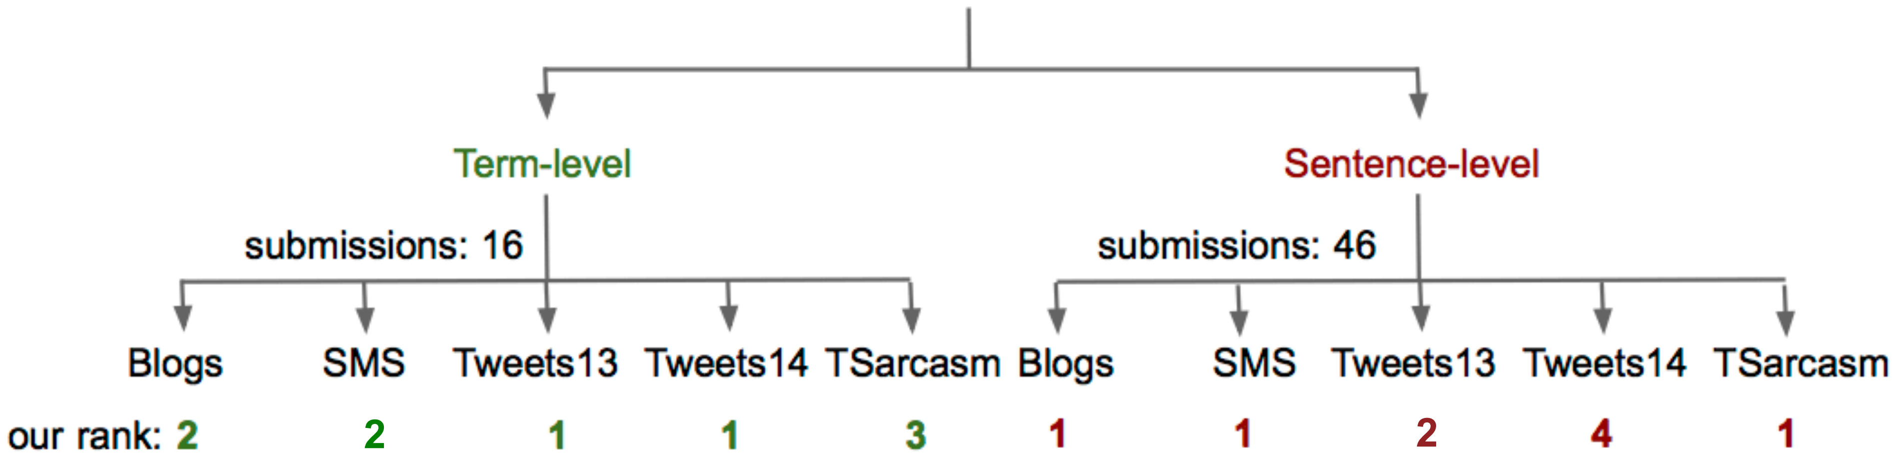

Aspect-Based Sentiment Analysis (SemEval-2014 Task 4)

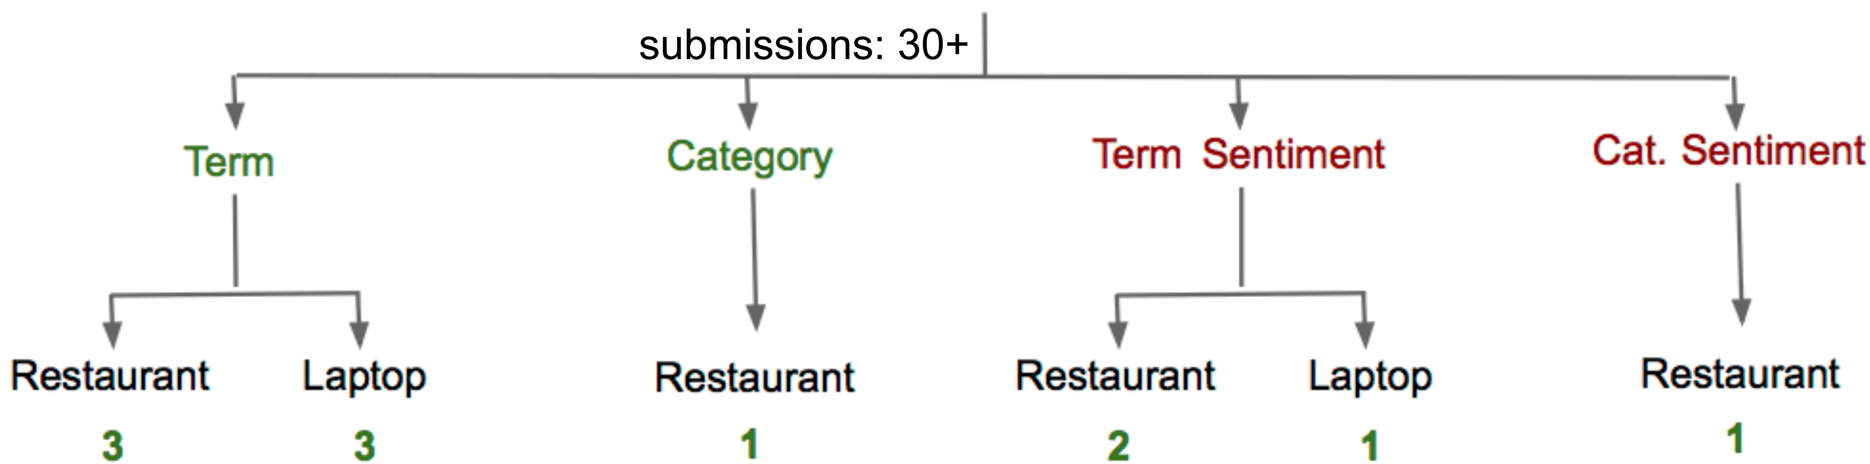

Feature Contributions (on Tweets)

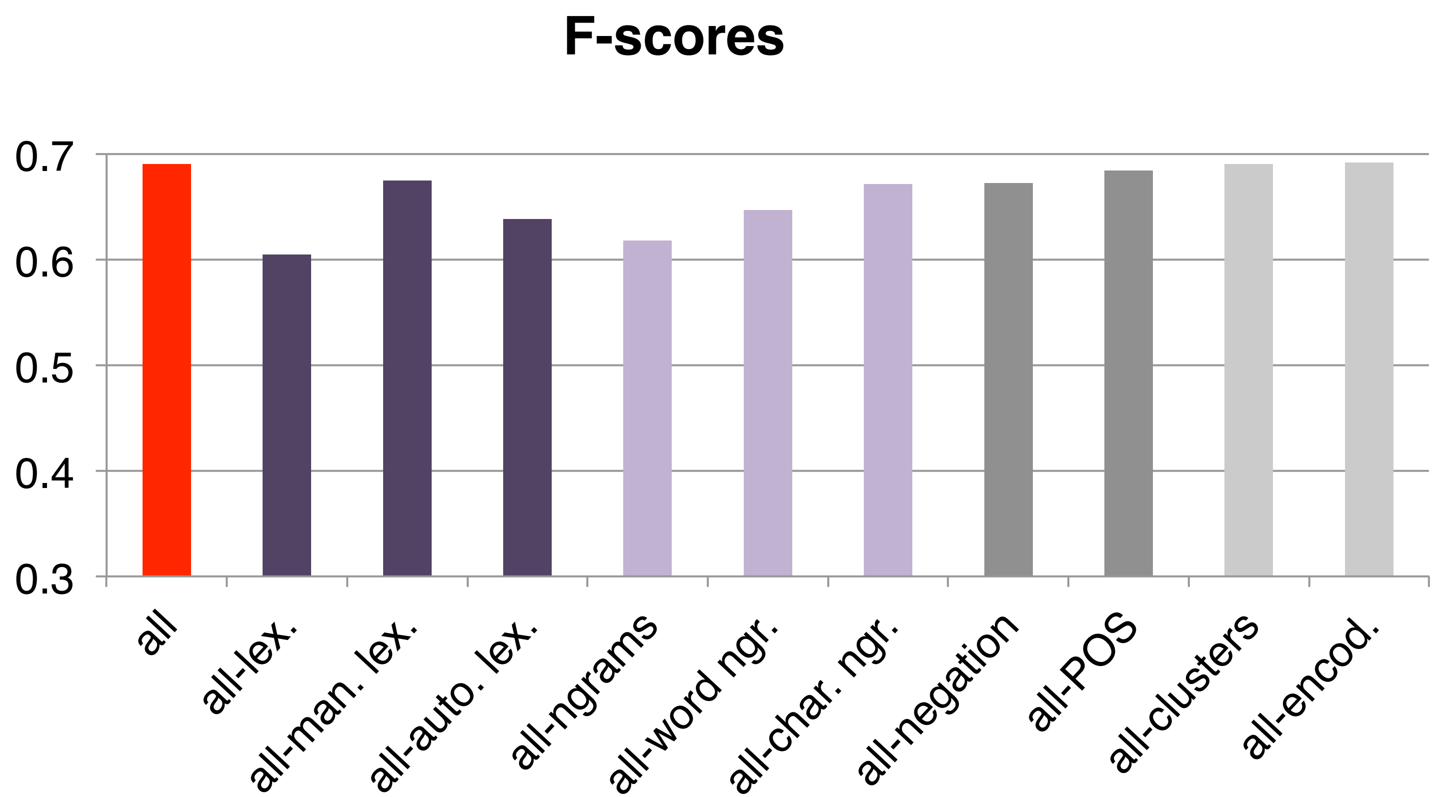

Detecting Stance in Tweets

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

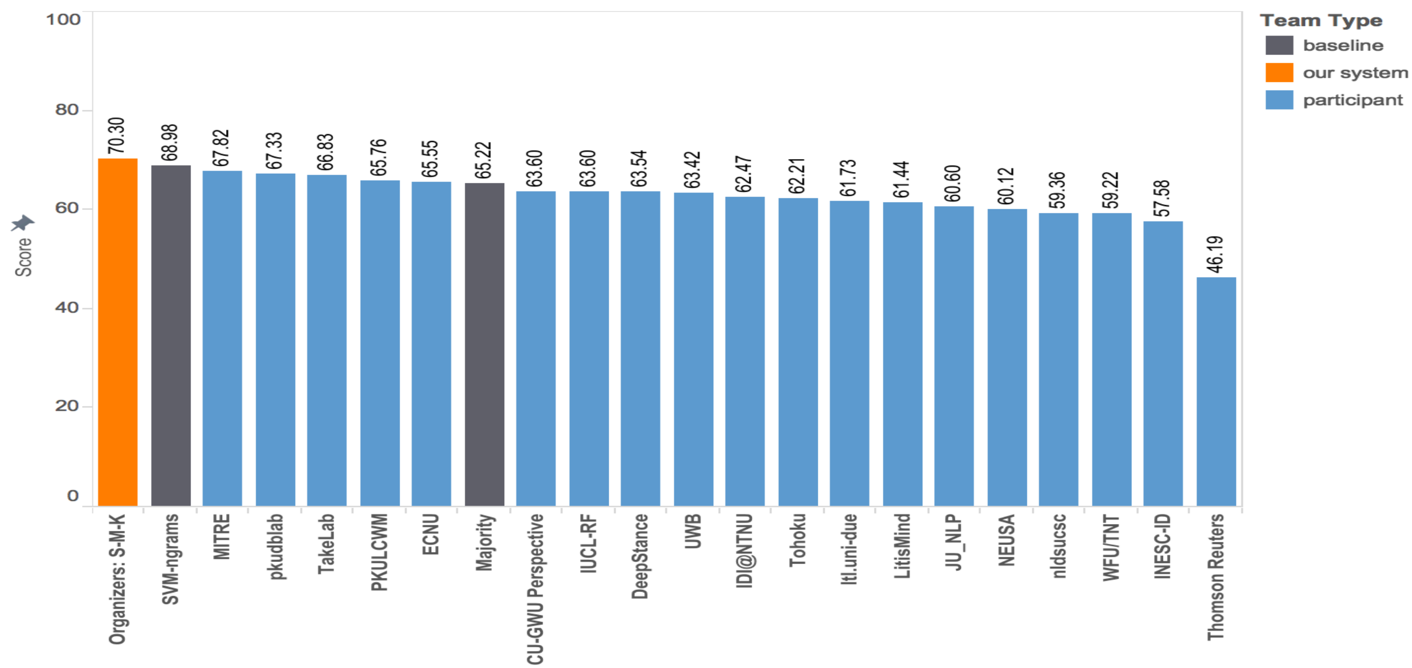

Ekaterina Shutova

Peter Turney

Metaphor as a Medium for Emotion

Paper:

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.

Metaphor

A figure of speech that refers to something as being the same as another thing for rhetorical effect.

- betrayal is stabbing someone in the back
- anger is a hot fluid (the boss boiled over)
- books are keys to one's imagination
- arguments are planes
  (he shot down all of my arguments)

Metaphor: Knowledge Projection (Lakoff & Johnson, 1980)

Source domain		Target domain
physical	➡	more abstract
closely experienced		more vague

Example: He shot down all of my arguments

Projects knowledge and inferences:
from the domain of battle (source domain)
onto the domain of arguments and debates (target domain).

Metaphor: Knowledge Projection (Lakoff & Johnson, 1980)

Source domain		Target domain
physical	→	more abstract
closely experienced		more vague

Example: He shot down all of my arguments

Projects knowledge and inferences:

from the domain of battle (source domain)

onto the domain of arguments and debates (target domain).

- preserves the core meaning of the sentence
- emphasizes certain aspects of the target domain, while downplaying others: framing

Research Questions

Q. Is a metaphorical statement likely to convey a stronger emotional content than its literal counterpart?

- to what extent?

Q. How does this emotional content arise in the metaphor:

- from the source domain,
- from the target domain, or
- compositionally through interaction of the source and the target?

Hypotheses

Hypothesis 1: metaphorical uses of words tend to convey more emotion than their literal paraphrases in the same context.

Example:

a. The spaceship blazed off into the space. METAPHORIC

b. The spaceship departed into the space. LITERAL

Hypothesis 2: the metaphorical sense of a word tends to carry more emotion than the literal sense of the same word.

Example:

a. Hillary brushed off the accusations. METAPHORIC

b. He brushed off the snow. LITERAL

Underline: verb

Green font: text that is common across a. and b.

Hypotheses

Hypothesis 1: metaphorical uses of words tend to convey more emotion than their literal paraphrases in the same context.

Example:

a. Hillary brushed off the accusations. METAPHORIC

b. Hillary dismissed the accusations. LITERAL

Hypothesis 2: the metaphorical sense of a word tends to carry more emotion than the literal sense of the same word.

Example:

a. Hillary brushed off the accusations. METAPHORIC

b. He brushed off the snow. LITERAL

Underline: verb

Green font: text that is common across a. and b.

Data for Our Experiments

- Focus on metaphors expressed by a verb
  - most frequent type of metaphor (Cameron, 2003; Shutova and Teufel, 2010)

- Extract verbs, senses, and sentences from WordNet
  - WordNet organizes senses in synsets
  - each synset has a gloss and example sentence

- Manually annotated for:
  - metaphoric or literal
  - which is more metaphoric or both equally metaphoric
  - no emotion or some emotion
  - which is more emotional or both equally emotional

Results for Hypothesis 1 Pairs (same context, synonym verbs): Absolute Metaphoricity & Relative Emotionality

# instances that are:	
metaphorical and more emotional	143 (83.6%)
literal and more emotional	17 (09.9%)
similarly emotional	11 (06.4%)
Total	**171 (100%)**

Results for Hypothesis 2 Cross Pairs (same verb, different senses): Absolute Metaphoricity & Relative Emotionality

# instances that are:	
metaphorical and more emotional	211 (59.4%)
literal and more emotional	31 (08.7%)
similarly emotional	113 (31.8%)
Total	**355 (100%)**

Results for Hypothesis 2 All Pairs (same verb, different senses): Relative Metaphoricity & Relative Emotionality

# instances that are more metaphorical and more emotional	227 (36.1%)
# instances that are more metaphorical but less emotional	28 (04.4%)
# instances that are more metaphorical but similarly emotional	119 (18.9%)
# instances that are similarly metaphorical and similarly emotional	196 (31.2%)
# instances that are similarly metaphorical but differ in emotionality	59 (09.4%)
Total	**629 (100%)**

Discussion: Metaphors are More Emotional

- Our results confirm both hypotheses:
  - metaphorical uses of words carry stronger emotions than
    - their literal uses,
    - as well as their literal paraphrases.
- This is inline with recent findings in neuroscience: Citron et al. (2016)
  - Examined metaphoric and literal sentences that had one word different
  - Metaphors (even conventional ones) in textual passages evoked stronger affective brain response
- Our annotations confirm: the metaphorical/literal distinction is a common pattern for polysemous verbs
  - ~38% of all verb senses we annotated were metaphorical

Discussion: Mechanism of Emotionality in Metaphors

Emotional content:

- not merely a property of the source or the target domain
- but rather, it arises through metaphorical composition.

The spaceship *blazed* out into space. MET some emot.
The spaceship *departed* out into space. LIT no emotion
The summer sun can cause a pine to *blaze*. LIT no emot.

This is the first such finding, and it highlights the importance of metaphor as a mechanism for expressing emotion.

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

Affect Associations in Creative Language

- Shapes of stories
  - tracked the distribution of emotion words
  - created a system to generate music from text
- Tweet-, message-level sentiment analysis
  - creative tweets are especially challenging
    - sarcasm, hyperbole, irony
- Metaphorical uses of words carry stronger emotions than
  - their literal uses,
  - as well as their literal paraphrases
- Showed that metaphoric usages are common
  - one more indicator that creativity in language is common

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

creative language
has emotional impact

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 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

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 $\not\Rightarrow$ favor;   negative language $\not\Rightarrow$ 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

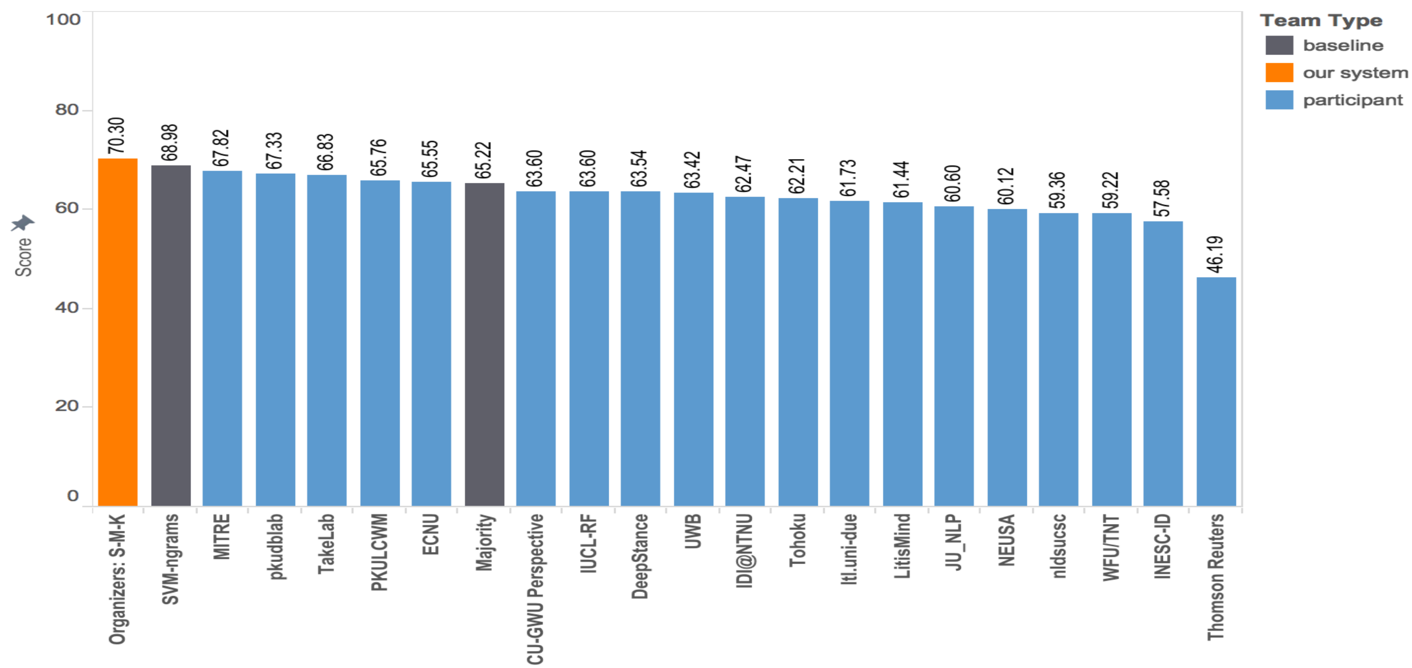

Word-Colour Associations

Papers:

- 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 Word-Colour Associations, Saif Mohammad, In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, June 2011, Portland, OR.

Word-Colour Associations

Concrete concepts

iceberg → white

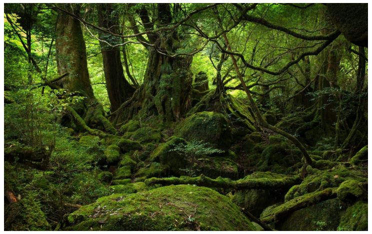

vegetation → green

Abstract concepts

danger → red

honesty → white

Key Questions

- How much do we agree on word-colour associations?
- How many terms have strong colour associations?
- Do concrete concepts have a higher tendency to have colour association?
- Can we create a lexicon of such word-colour associations?
  - Applications: advertising and marketing, information visualization, teaching aid for kids with learning disabilities

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%

Crowdsourced Questionnaire

- 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

Q. Which colour is associated with *sleep*?

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

NRC Word-Colour Association Lexicon

- Color associations for ~14,000 words

Associations with Colours

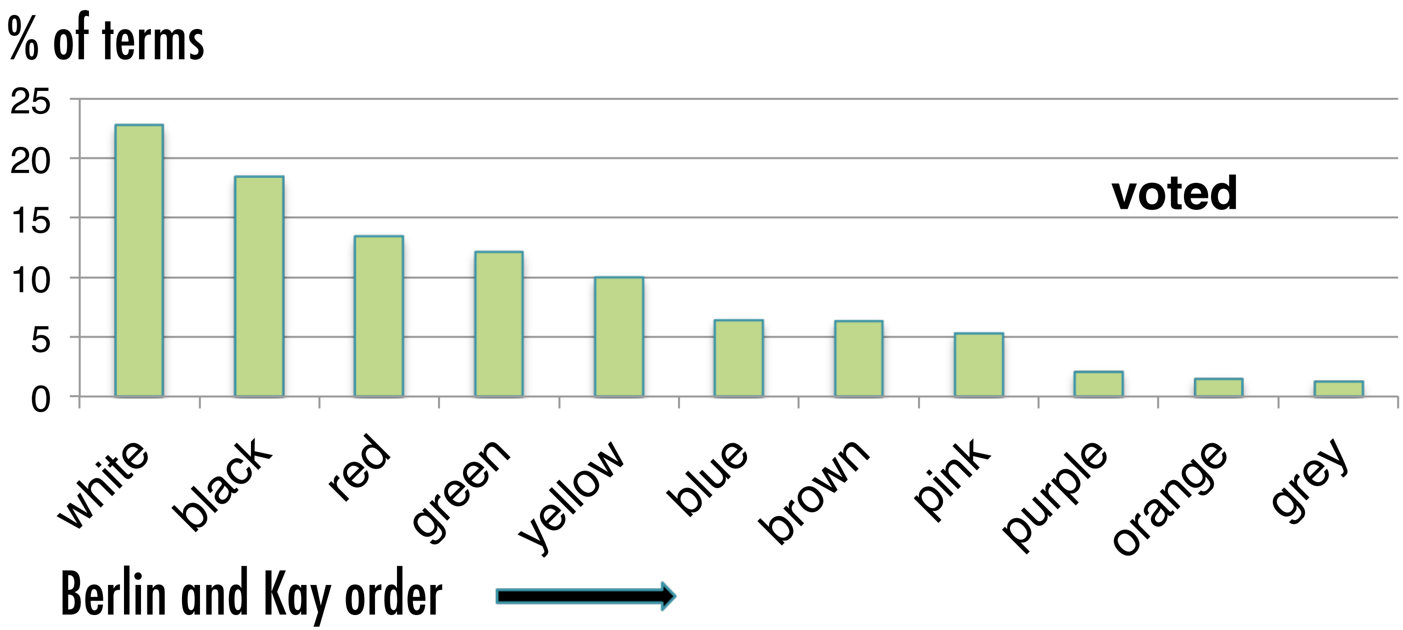

Some Findings

- Colours in decreasing order of association frequency is highly correlated with the Berlin and Kay order.
  - Implications to the Sapir–Whorf hypothesis
- About 30% of the terms have strong colour associations.
- Terms for abstract concepts just as likely to have color associations as concrete concepts

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.

LiveSlides web content

To view

Download the add-in.

liveslides.com/download

Start the presentation.

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

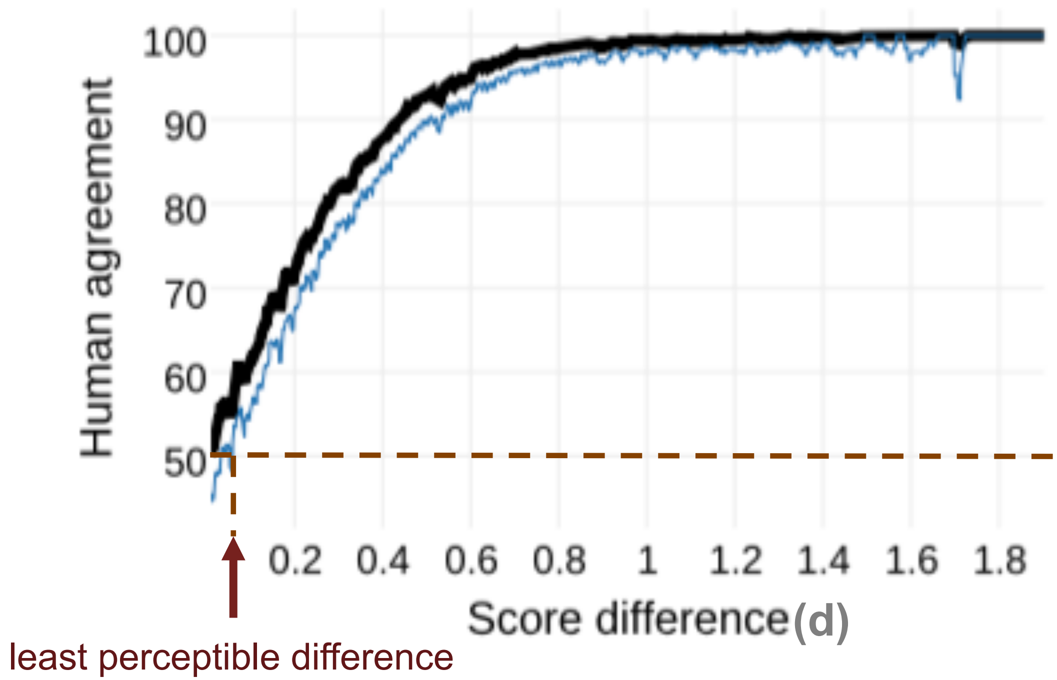

Least Perceptible Differences in lexicons:

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

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

The Interplay of Metaphor and Emotion

Studied in:

- computational linguistics
  - sentiment polarity classification of metaphorical language (Veale and Li, 2012; Kozareva, 2013; Strzalkowski et al., 2014)
- linguistics (Blanchette et al., 2001; Kovecses, 2003)
- political science (Lakoff, 1980; Lakoff and Wehling, 2012)
- cognitive psychology (Crawford, 2009; Thibodeau and Boroditsky, 2011)
- neuroscience (Aziz-Zadeh and Damasio, 2008; Jabbi et al., 2008)

However, no quantitative study establishing:

- the extent to which metaphorical language is used to express emotion
- the mechanisms by which this happens.