Mapping the Emotional and Social Dimensions of Language:

The Next Generation of NRC Lexicons

Saif M. Mohammad

National Research Council Canada

uvgotsaif@gmail.com

@SaifMMohammad

Why Emotions and Affect Matter

- Determine human experience and behavior
- Condition our actions
- Central in organizing meaning
  - No cognition without emotion
- A window into
  - understanding our mind (cognition), body (health, well-being), how we make sense and interact with the world
  - the evolutionary forces that shaped us

Emotions

How many emotions can we perceive?

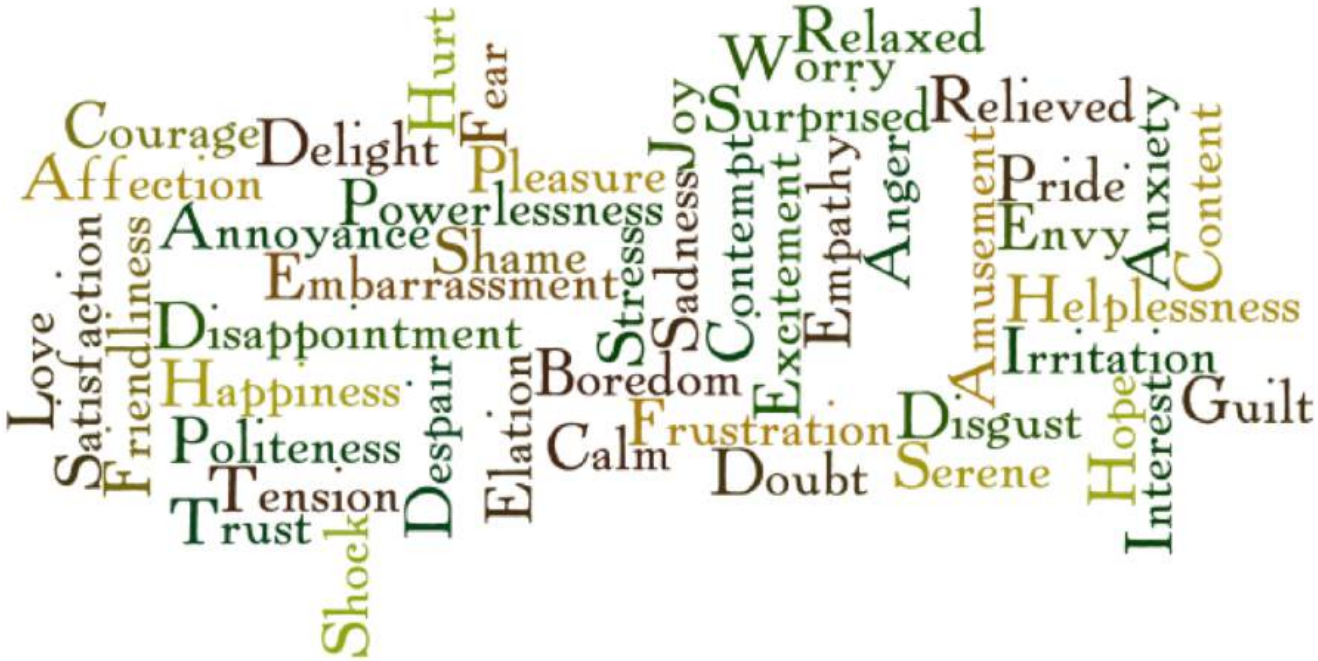

Difficult question:

- fuzzy emotion boundaries, overlapping meanings, socio-cultural influences, etc.

Some studies suggest 500 to 600 emotion categories!

Core Dimensions of Connotative Meaning


Osgood


Russell

Influential factor analysis studies (Osgood et al., 1957; Russell, 1980, 2003) have shown that the three most important, largely independent, dimensions of word meaning:

- valence (V): positive/pleasure – negative/displeasure
- arousal (A): active/stimulated – sluggish/bored
- dominance (D): powerful/strong – powerless/weak

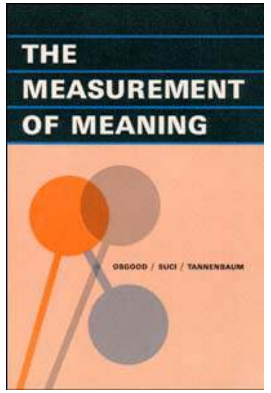

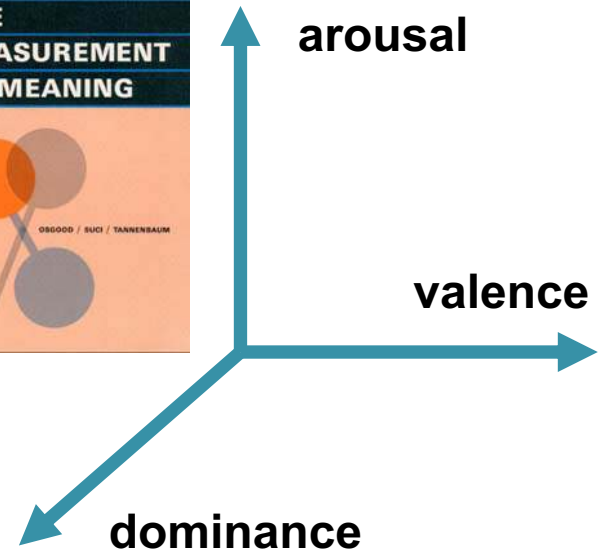

Thus, when comparing the meanings of two words, we can compare their V, A, D scores. For example:

- *banquet* indicates more positiveness than *funeral*
- *nervous* indicates more arousal than *lazy*
- *queen* indicates more dominance than *delicate*

Psychological Theories of Basic Emotions

Paul Ekman, Psychologist

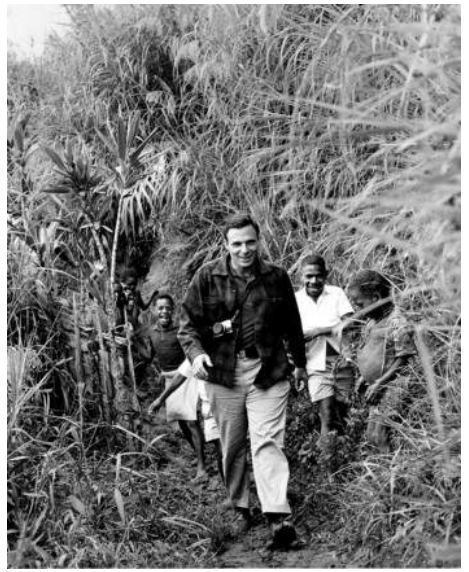

- Paul Ekman, 1971: Six Universal Basic Emotions
  - Plutchik, 1980: Eight Basic Emotions
  - And many others

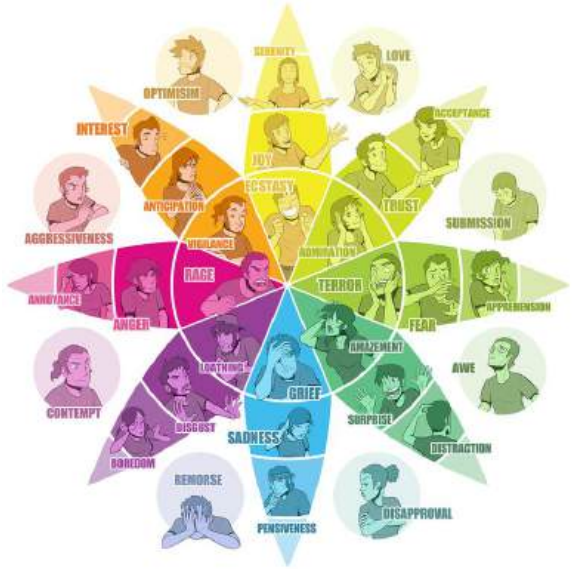

Plutchik's Emotion Wheel
Image credit: Julia Belyanevych

Theories of Emotion

Paul Ekman
Psychologist and discoverer of micro expressions.

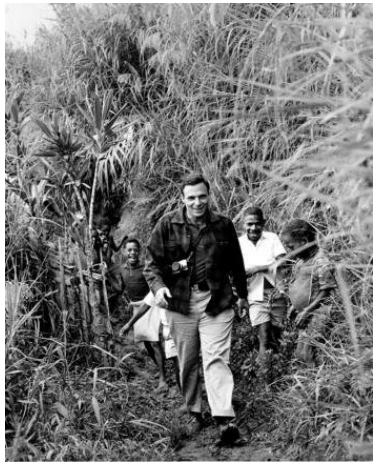

Lisa Barrett
University Distinguished Professor of Psychology, Northeastern University

Theory of Constructed Emotion (Barrett, 2017)

- the brain constructs emotions
- important tenets of BET discredited ("basic" emotions)
- stress on variability

The Language of Emotions

Language is a powerful way of expressing emotions

- can express numerous emotional shades
  - terms with fuzzy boundaries, overlapping meanings, socio-cultural influences, etc.
- usually conveyed by connotation (and not denotation)
  - can be subtle, direct, ambiguous, deceptive
  - can be creative
  - can be conscious expression or subconscious manifestations

word --

Word-Emotion Association Lexicons

Affect Datasets for Words

Lexicons for both categorical emotions as well as for valence, arousal, and dominance

- Lists of words associated with joy, sadness, fear, etc.
- Lists of words and their valence, arousal, and dominance scores

word --

Word-Emotion Association Lexicons

Early Work...

2010: The NRC Emotion Lexicon

Peter Turney

- Entries for 14,200
- Associations (0 or 1) with
  - positive and negative sentiment
  - 8 categorical (Plutchik) emotions (anger, fear, joy, sadness, anticipation, disgust, surprise, trust)
- Translations in over 100 languages

Available at: www.saifmohammad.com

Paper:

Crowdsourcing a Word-Emotion Association Lexicon, Saif M. Mohammad and Peter Turney, *Computational Intelligence*, 29 (3), pages 436-465, 2013. Lexicon Released in 2010.

The NRC Word–Colour Association Lexicon
provides associations for ~14,000 words with 11 common colours

Use of The Emotion Lexicons

- For research by the scientific community
  - Computational linguistics, psychology, digital humanities, robotics, public health research, social science, etc.
- To analyze text
  - Brexit tweets, Radiohead songs, Trump tweets, election debates,…
  - **Wishing Wall**, uses the NRC Emotion lexicon to visualize wishes. Displayed in:
    - **Barbican Centre**, London, England, 2014
    - **Tekniska Museet**, Stockholm, Sweden, 2014
    - **Onassis Cultural Centre**, Athens, Greece, 2015
    - **Zorlu Centre**, Istanbul, Turkey, 2016

creativity

- In commercial applications

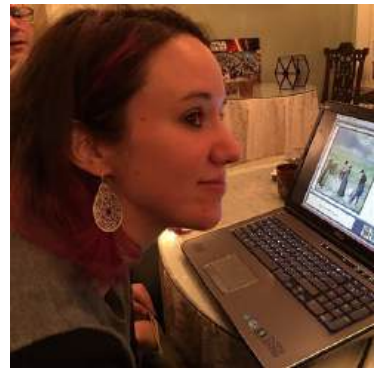

Hannah Davis
Artist/Programmer

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.

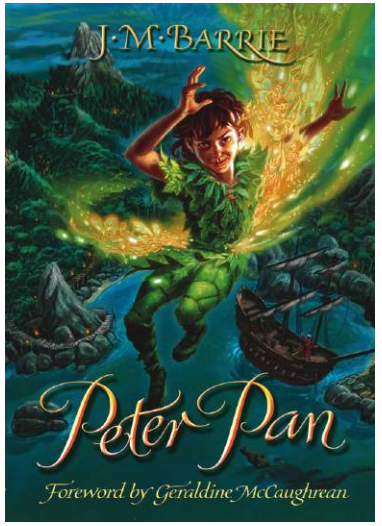

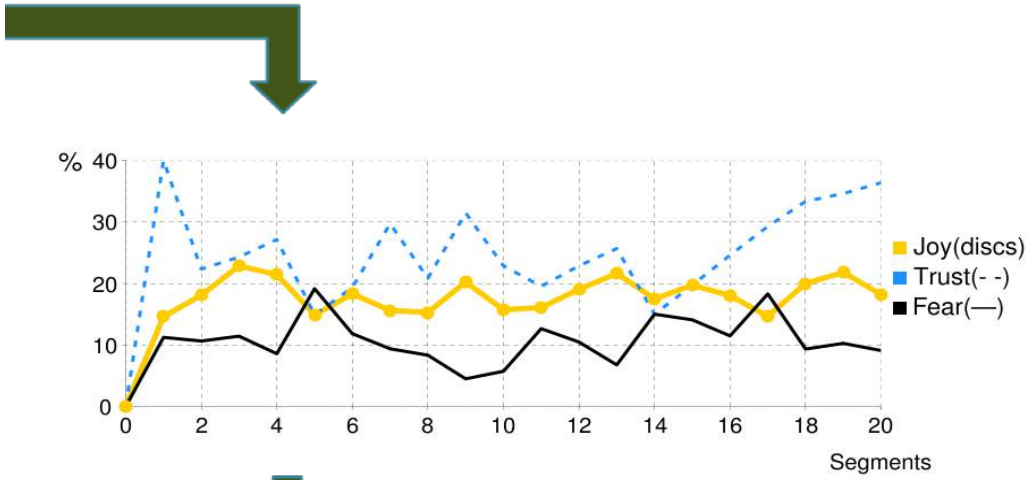

A method to generate music from literature

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

TransProse

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

TransProse

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

Examples

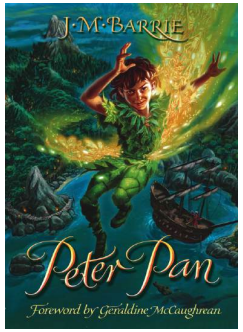

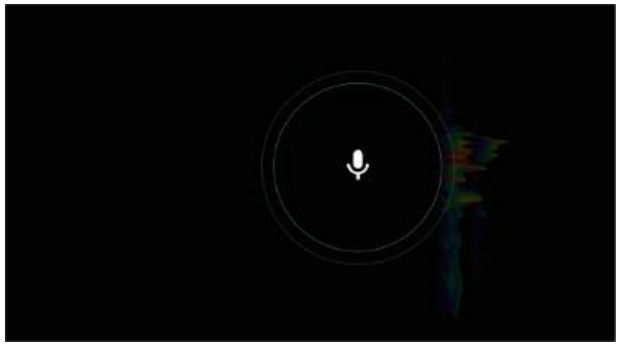

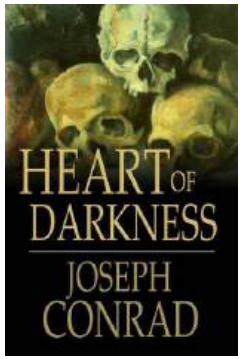

TransProse Music Played by an Orchestra, at the Louvre Museum, Paris

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)

2018: The NRC VAD Lexicon

Obtaining Reliable Human Ratings of Valence, Arousal, and Dominance for 20,000 English Words

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

Best–Worst Scaling (Louviere & Woodworth, 1990)

- Uses comparative annotation—mitigates bias
- Keeps the number of annotations down to about 2N
- Leads to more reliable, less biased, more discriminating annotations
  (Kiritchenko and Mohammad, 2017, Cohen, 2003)

Entries with Highest and Lowest Scores in the NRC VAD Lexicon (2018)

Dimension	Word	Score↑	Word	Score↓
valence	*love*	1.000	*toxic*	0.008
	happy	1.000	*nightmare*	0.005
	happily	1.000	*shit*	0.000
arousal	*abduction*	0.990	*mellow*	0.069
	exorcism	0.980	*siesta*	0.046
	homicide	0.973	*napping*	0.046
dominance	*powerful*	0.991	*empty*	0.081
	leadership	0.983	*frail*	0.069
	success	0.981	*weak*	0.045

Scores are in the range 0 (lowest V/A/D) to 1 (highest V/A/D).

2018: The NRC Emotion Intensity Lexicons

Provides intensity scores for ~6000 words found to be associated with the 8 emotions

The NRC Valence, Arousal, and Dominance Lexicon (2018)

provides ratings of valence, arousal, and dominance for ~20,000 English words

http://saifmohammad.com/WebPages/nrc-vad.html

The NRC Word–Emotion Association Lexicon aka NRC Emotion Lexicon or EmoLex (2010)

provides associations for ~14,000 words with eight emotions (anger, fear, joy, sadness, anticipation, disgust, surprise, trust)

http://saifmohammad.com/WebPages/NRC-Emotion- Lexicon.htm

The NRC Emotion Intensity Lexicon aka Affect Intensity Lexicon (2018-19)

provides intensity scores for ~6000 words with four emotions (anger, fear, joy, sadness)

http://saifmohammad.com/WebPages/AffectIntensity.htm

The NRC Word–Colour Association Lexicon (2010)

provides associations for ~14,000 words with 11 common colours

http://saifmohammad.com/WebPages/lexicons.html

Key Features of the Lexicons

Associations (not Denotations)

- Past work had focused on denotations
- But, people often connote emotions

Crowdsourcing and Quality Control

About 2% of the data was annotated internally beforehand (by the author)

- These gold questions are interspersed with other questions
- If one's accuracy on the gold questions falls below 80%,
  - all of their annotations are discarded

All crowdsourcing work approved by NRC's Research Ethics Board.

Comparative Annotations

- Avoids various biases
- Scores are all relative
- More annotations needed; BWS helps, but adds complexity

Likert scales

- Anchors scores
- Simple
- Some unwanted biases; Boundary between classes artificial

Large Size

For example, for the NRC VAD lexicon (Mohammad, 2018):

- v1: Obtained ~800,000 annotations for about 20K words
- v2: 44k words, 20k MWEs

Focus on High Reliability

Reliability vs. Inter-Annotator Agreement

Inter-Annotator Agreement (IAA)

- Measures how much annotators agree with each other
- Reflects agreement on labels across annotators
- Common metrics: Cohen's κ, Fleiss' κ, Krippendorff's α
- High IAA implies: annotations are applied similarly

Reliability

- Measures how dependable and reproducible the annotations are
- Reflects stability/trustworthiness of the aggregated scores/labels
- Common metrics: split-half reliability, split-half class match reliability
- High reliability implies: repeat annotations will lead to similar aggregated scores/labels

Reliability (Reproducibility) of Annotations

Average split-half reliability (SHR): approach to determine consistency

(Kuder and Richardson, 1937; Cronbach, 1946)

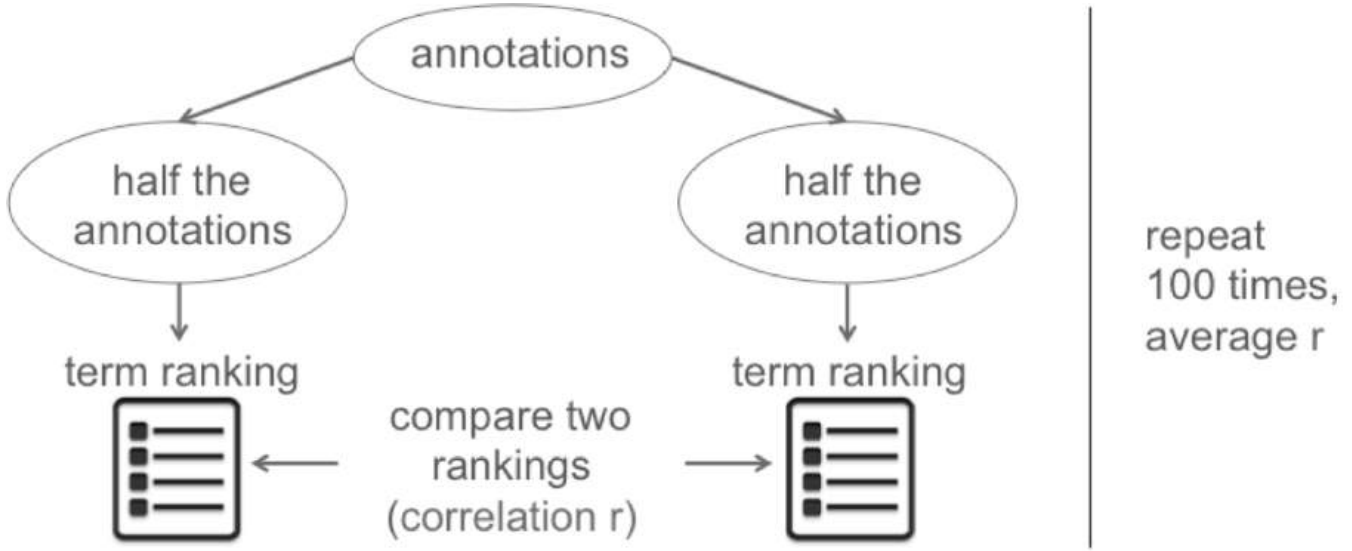

Split-Half Reliability Scores for VAD Annotations

higher scores indicate higher reliability

Annotations	# Terms	# Annotations	V	A	D
Warriner et al. (2013)	13,915	20 per term	0.91	0.79	0.77
Ours (Warriner terms)	13,915	6 per tuple	0.95	0.91	0.91
Ours (all terms)	20,007	6 per tuple	0.95	0.90	0.90

These SHR scores show for the first time that highly reliable fine-grained ratings can be obtained for valence, arousal, and dominance.

Some Things We Learned About Emotions

- Variability
  - cultural
  - context
  - person-to-person
  - what/how you ask about emotions
    - more agreement for valence; less for fear, sadness, etc.
    - associated vs. evoked
    - greater agreement with comparative annotations

- Substantial common ground

Draw inferences at aggregate level

- determine broad trends

word --

Word-Emotion Association Lexicons

New Work...

New Dimensions

2024

WorryWords

- Anxiety-Calmness Lexicon

Theoretical Framework: Circumplex Model (Russell, 1980)

2025

The NRC Warmth and Competence Lexicon

- Warmth-Coldness
  - Sociableness-Unsociableness
  - Trustworthyness-Untrustworthyness
- Competence-Incompetence

Theoretical Framework: Stereotype Content Model (SCM), Social Cognition Theory (Fiske et al. 2002)

2026

The NRC MFT Lexicon

- Care-Harm
- Fairness-Cheating
- Loyalty-Betrayal
- Authority / Subversion
- Sanctity / Degradation

Theoretical Framework: Moral Foundation Theory (MFT) (Haidt, 2007)

More Coverage

- 44k Unigrams
- 20k Multiword Expressions (MWEs)

Locations

- USA, India

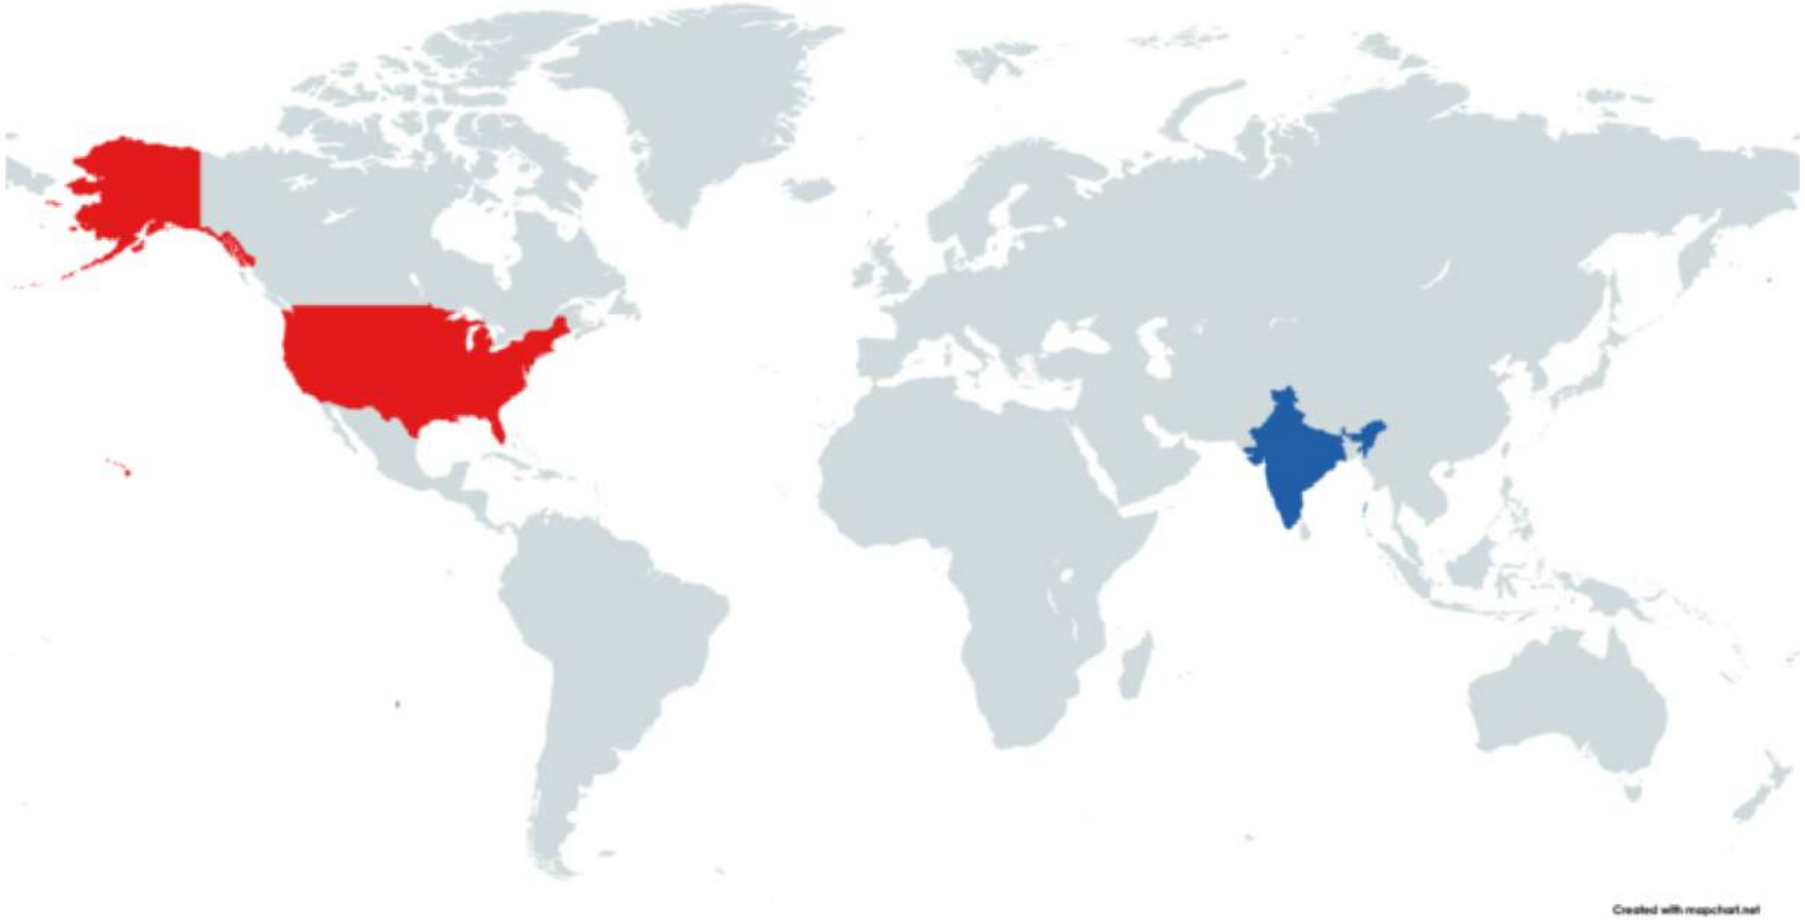

Case Study: Anxiety

Anxiety

The anticipatory unease about a potential (future) negative outcome

- common and beneficial human emotion
- can sometimes manifest into mental disorders
  - mismatch: current environment and what anxiety response slowly evolved to address

Why create language resources for anxiety?

- Understanding anxiety and the underlying mechanisms (Psych, Health)
  - how it relates to other emotions and affect
  - how it relates to our body
  - how anxiety changes with age, socio-economic status, weather, green spaces, etc.
  - identifying coping mechanisms, clinical interventions to manage anxiety
- Study how anxiety manifests in language (Ling.)
  - how language shapes anxiety
  - how culture shapes the language of anxiety
- Tracking the degree of anxiety towards targets of interest such as climate change, government policies, biological vectors, etc. (Health, Policy)
- Developing automatic systems for detecting anxiety (NLP)
- Studying how anxiety impacts behaviour in physical and virtual environments (SS)
- Studying anxiety in stories, character development, etc. (DH)

Anxiety Questionnaire

Consider anxiety to be a broad category that includes:
jittery, antsy, insecure, nervous, unease, tense, worried, unnerving, nerve-racking, apprehensive, fretful, troubled, self-doubting, discontented, concerned, and keyed up

Consider calmness to be a broad category that includes:
calm, relaxed, comforted, serene, at ease, self-assured, carefree, composed, collected, untroubled, peaceful, contented, unconcerned, indifferent, and uninvolved

Select the options that most English speakers will agree with.

Q1. <term> is often associated with feeling:

3: very anxious	-1: slightly calm
2: moderately anxious	-2: moderately calm
1: slightly anxious	-3: very calm

0: not associated with feeling anxious or calm

WorryWords

Repository of manually derived word–anxiety associations

- Scale: maximum calmness (-3) to maximum anxiety (3)
  - real-valued scores and also coarse categorical labels (e.g, low anxiety, high anxiety)
- Size
  - 44K English words
  - 20K English MWEs
- Quality
  - interspersed gold (control) questions
  - show that the anxiety associations are highly reliable
    - split-half reliability: 0.82 for unigrams; 0.94 for MWEs

Term	Score
suffocative	3.00
manic	2.41
riskily	1.72
ceramic	0.12
conformed	-1.71
lullaby	-2.79

EMNLP 2024:

WorryWords: Norms of Anxiety Association for over 44K English Words. Saif M. Mohammad.

New Dimensions

2024

WorryWords

- Anxiety-Calmness Lexicon

Theoretical Framework: Circumplex Model (Russell, 1980)

2025

The NRC Warmth and Competence Lexicon

- Warmth-Coldness
  - Sociableness-Unsociableness
  - Trustworthyness-Untrustworthyness
- Competence-Incompetence

Theoretical Framework: Stereotype Content Model (SCM), Social Cognition Theory (Fiske et al. 2002)

2026

The NRC MFT Lexicon

- Care-Harm
- Fairness-Cheating
- Loyalty-Betrayal
- Authority / Subversion
- Sanctity / Degradation

Theoretical Framework: Moral Foundation Theory (MFT) (Haidt, 2007)

Valence Asymmetries in the Lexicons

Class Distribution: Anxiety-Calmness

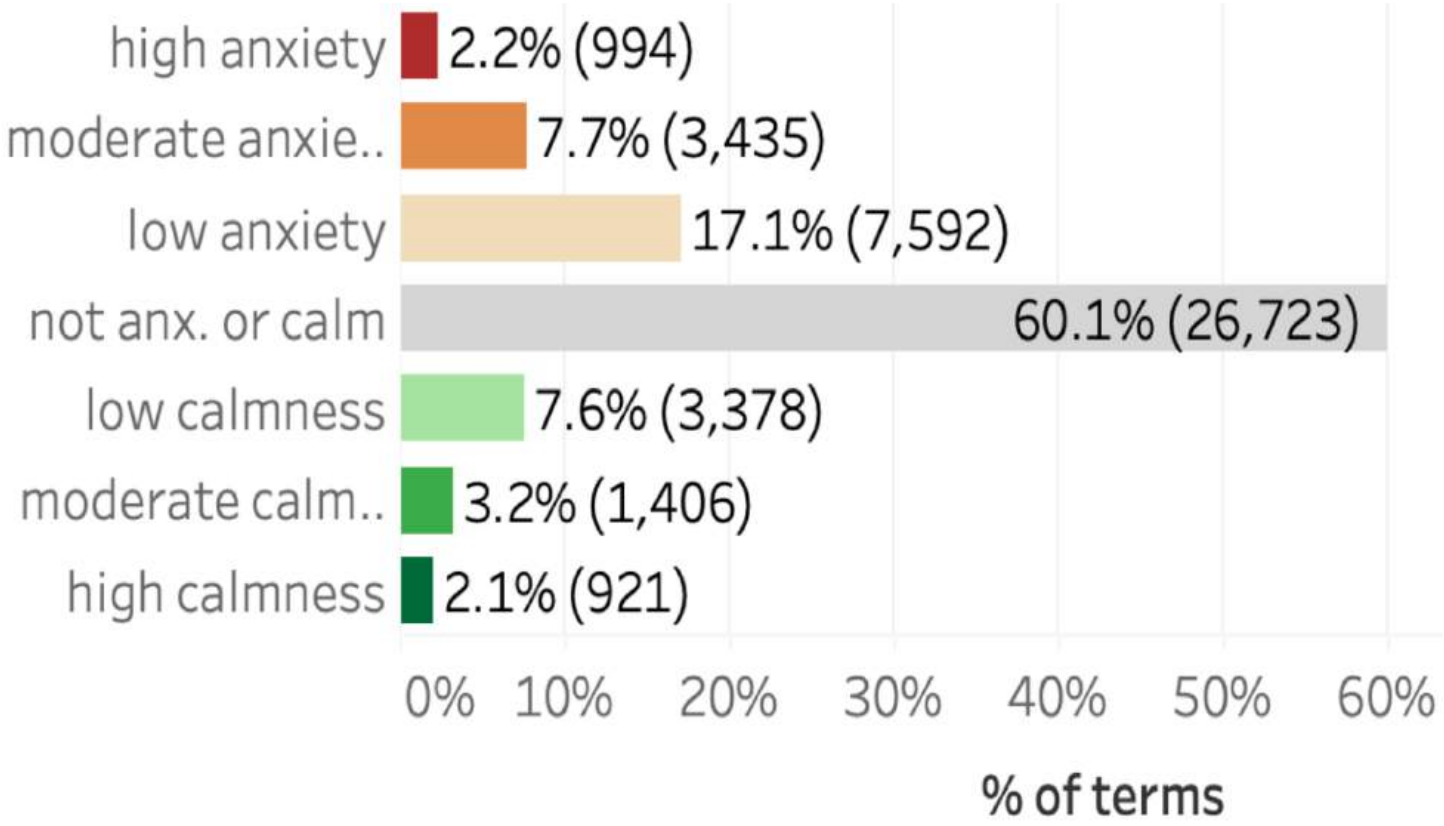

Class Distribution: Pos/Neg in the NRC Emotion Lexicon

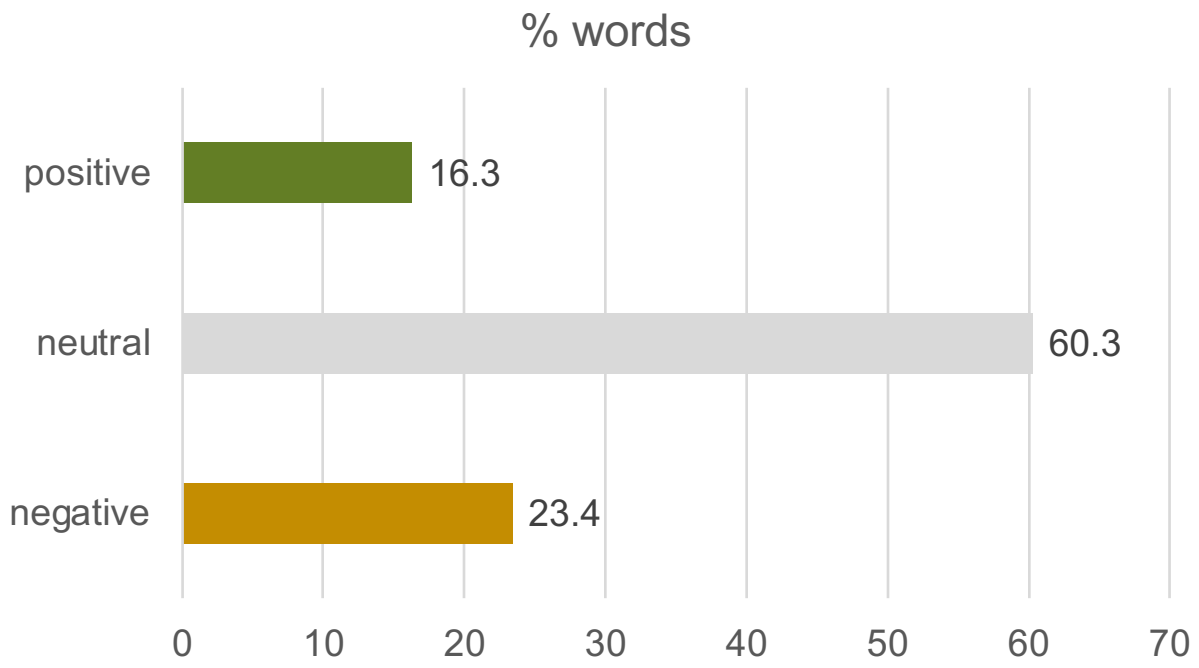

Class Distribution: Care/Harm
NRC MFT Lexicon

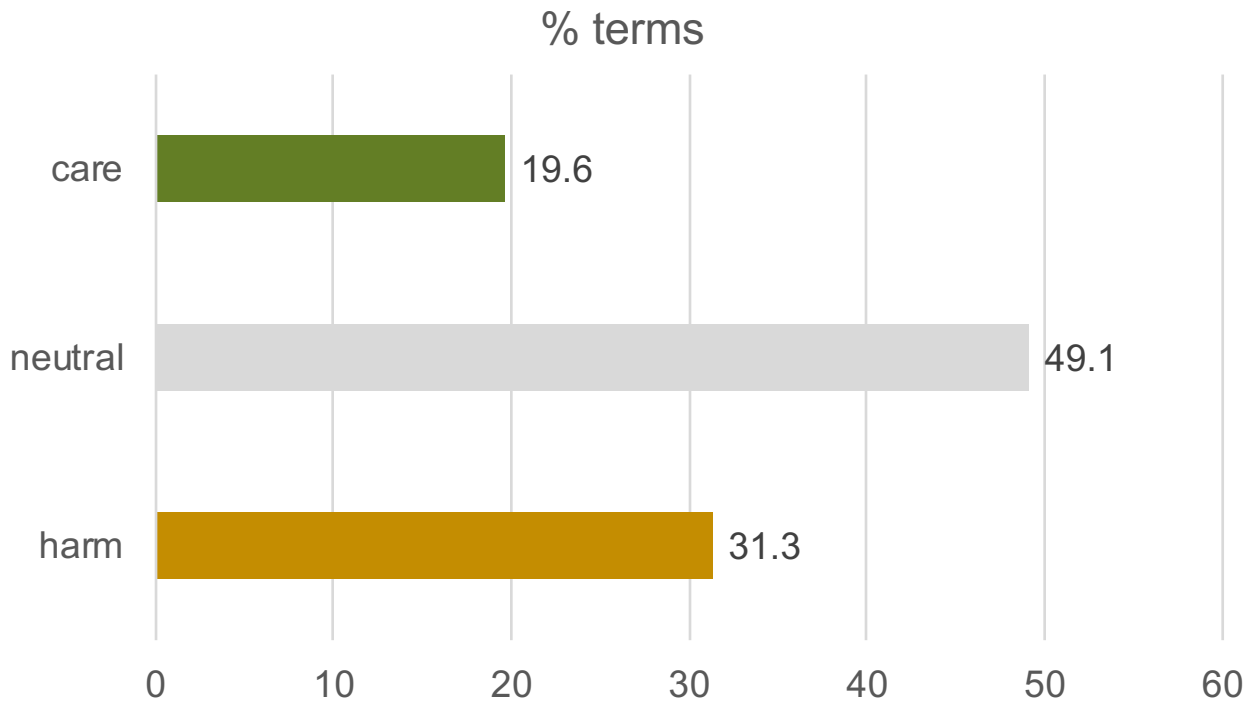

Class Distribution: Authority/Subversion
NRC MFT Lexicon

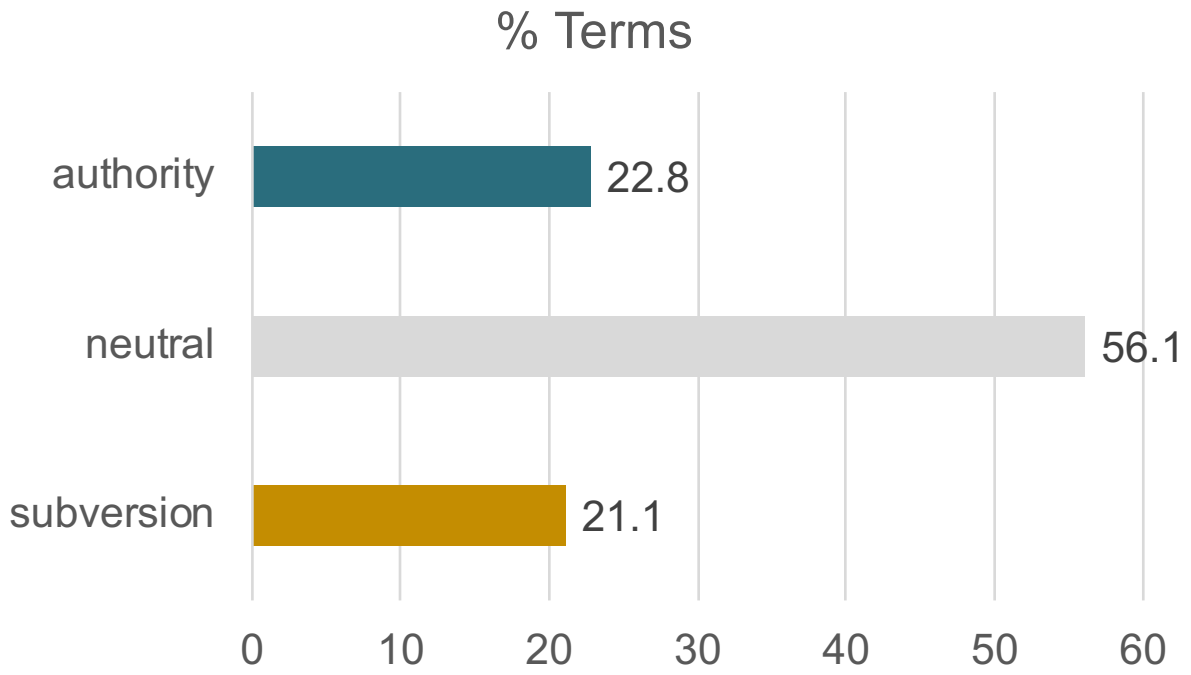

Example Use Cases of the Lexicons

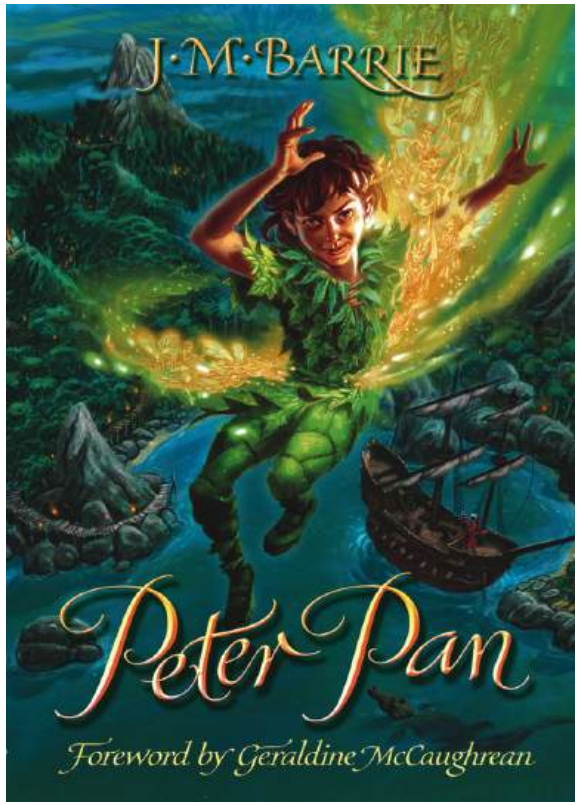

Emotions Arcs in CAS (mind, body, world), Commerce, Psychology, Health, and Humanities

Tracking Emotions in Stories

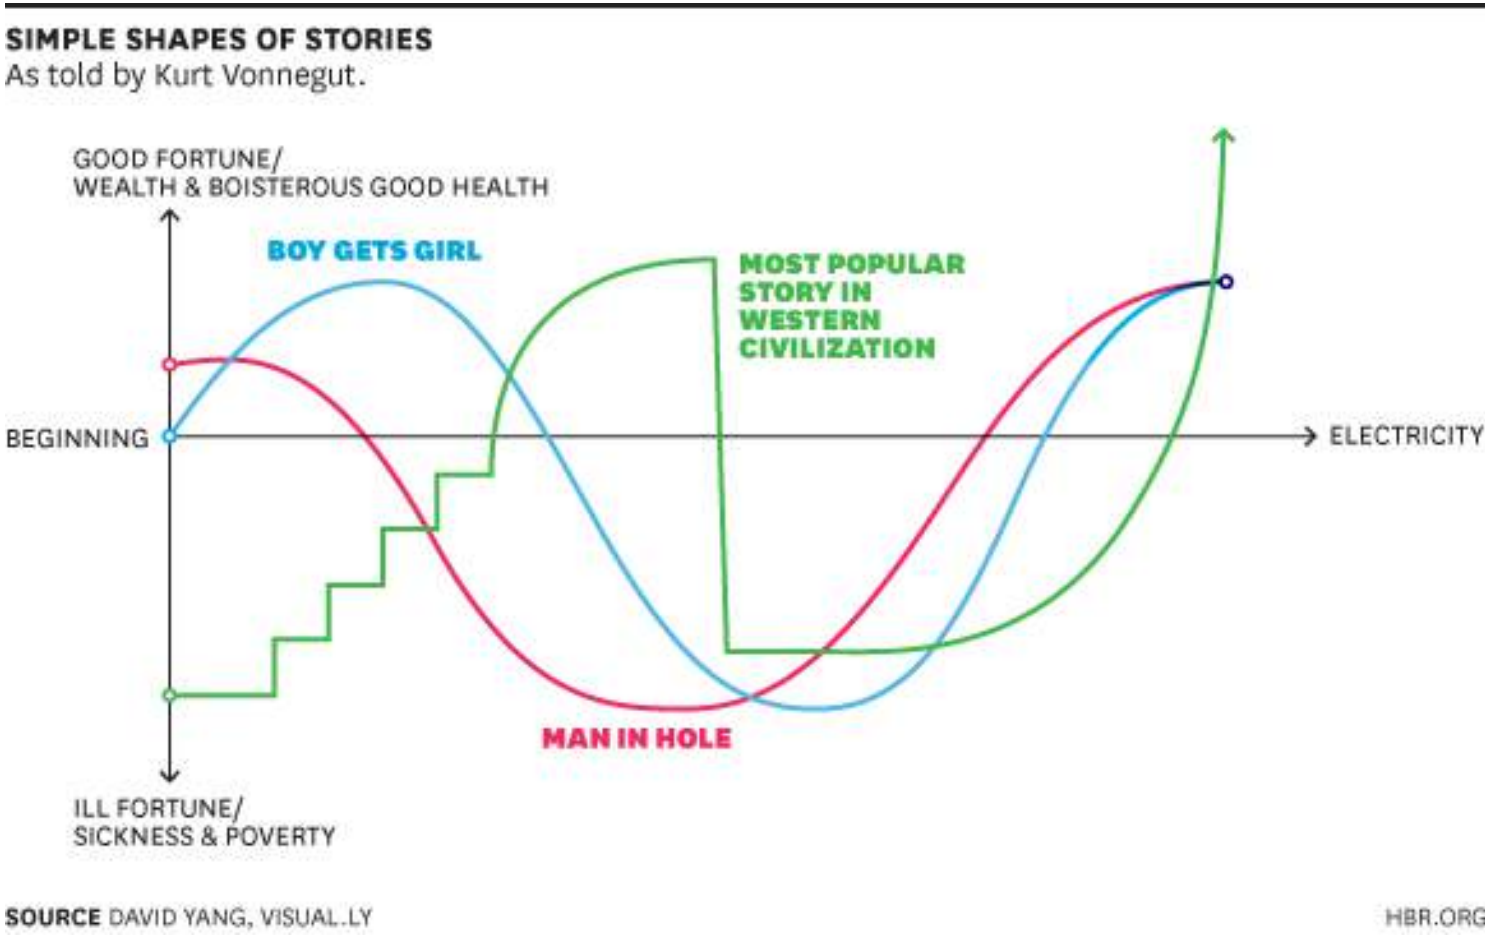

Back in 2011:

Tracking Emotions in Stories

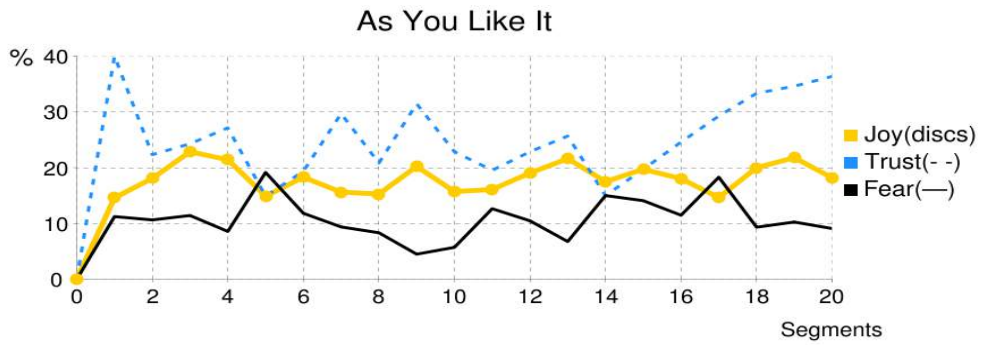

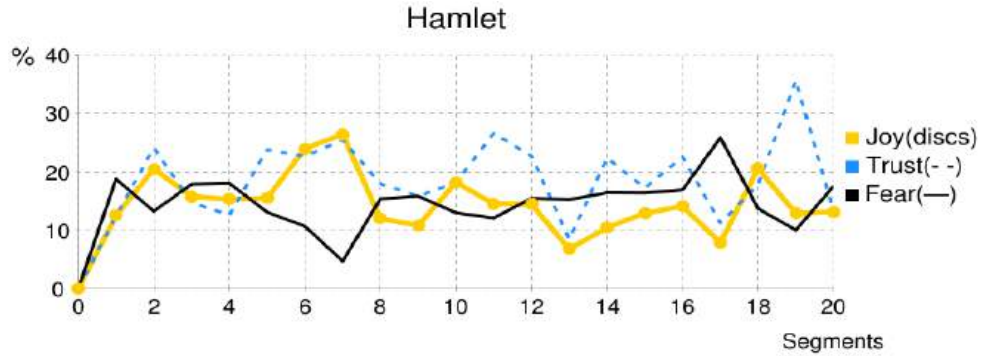

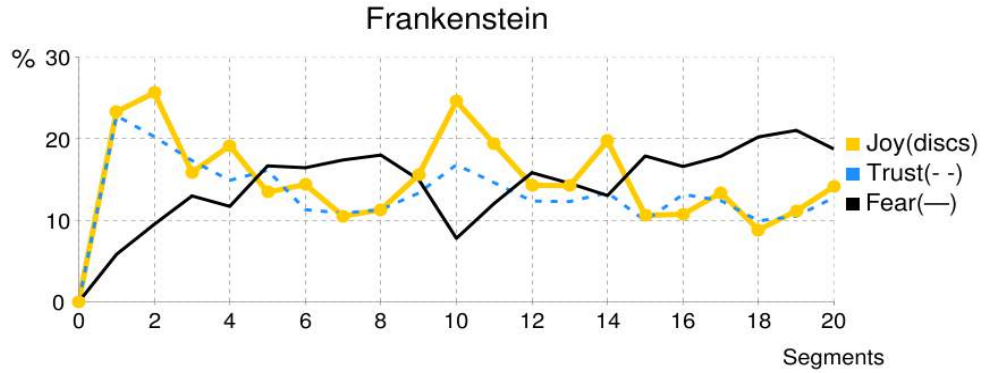

From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales, Saif Mohammad, In Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), June 2011, Portland, OR.

Creating Emotion Arcs

- Lexicon-only approach
- ML approaches (sometimes making use of lexicons)

Lexicon-only approaches

- Pros
  - simple, accessible
  - interpretable
  - low-carbon
  - domain-free
- Cons
  - not highly accurate at instance level (context, long-distance dependencies)

Evaluating Emotion Arcs

Very little work!! No dataset of gold arcs.

Evaluating Emotion Arcs

Consider tracking anger in tweets associated with vaccines (week by week)

- Manually annotate 300,000 individual tweets from 2018 to 2024
- Take the percentage of tweets marked as joy in every week to create the emotion arc

Annotating data is a bottleneck

Daniela Teoderescu

2023 EMNLP: Evaluating Emotion Arcs Across Languages

- make use of existing emotion datasets (usually 2 to 5K instances)
- sample instances with replacement to generate random but non-trivial arcs
- create gold emotion arcs as usual

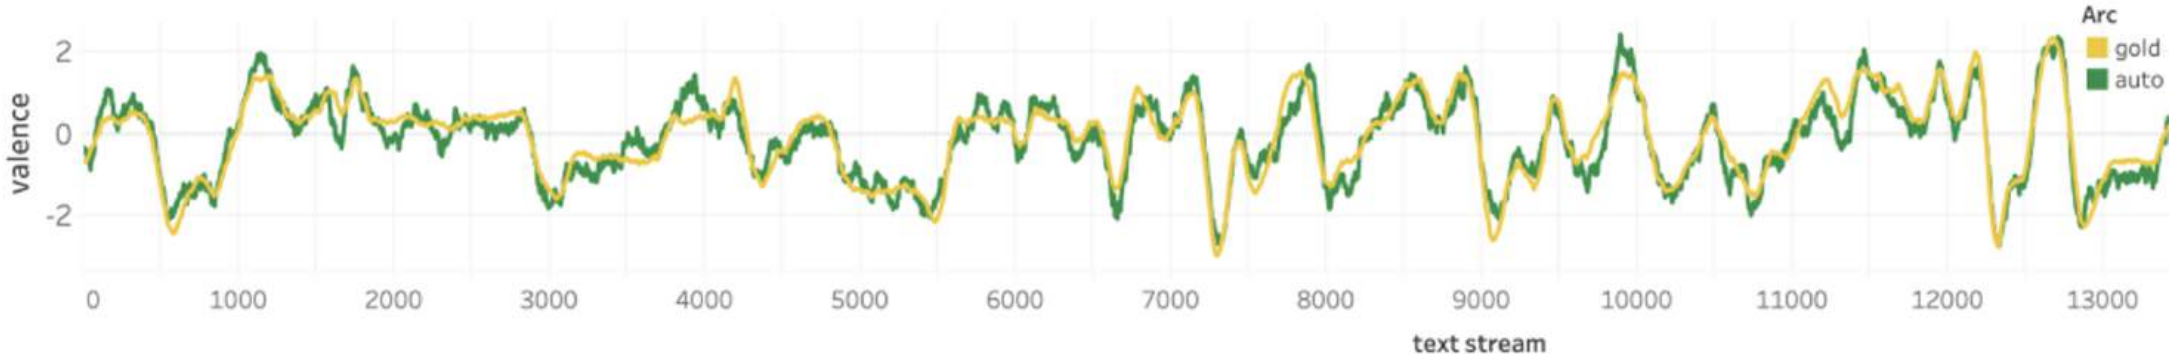

2023 EMNLP: Generating High-Quality Emotion Arcs Using Emotion Lexicons

- Used 36 datasets that had emotion-labeled sentences/tweets to create gold arcs
- For various affect categories, multiple languages, and other characteristics

Key Conclusions:

- lexicon-only based methods are extremely accurate
- aggregating information from hundreds of tweets/instances to create points of the emotion arcs very powerful

Used WorryWords to

Track the change of anxiety in streams of text

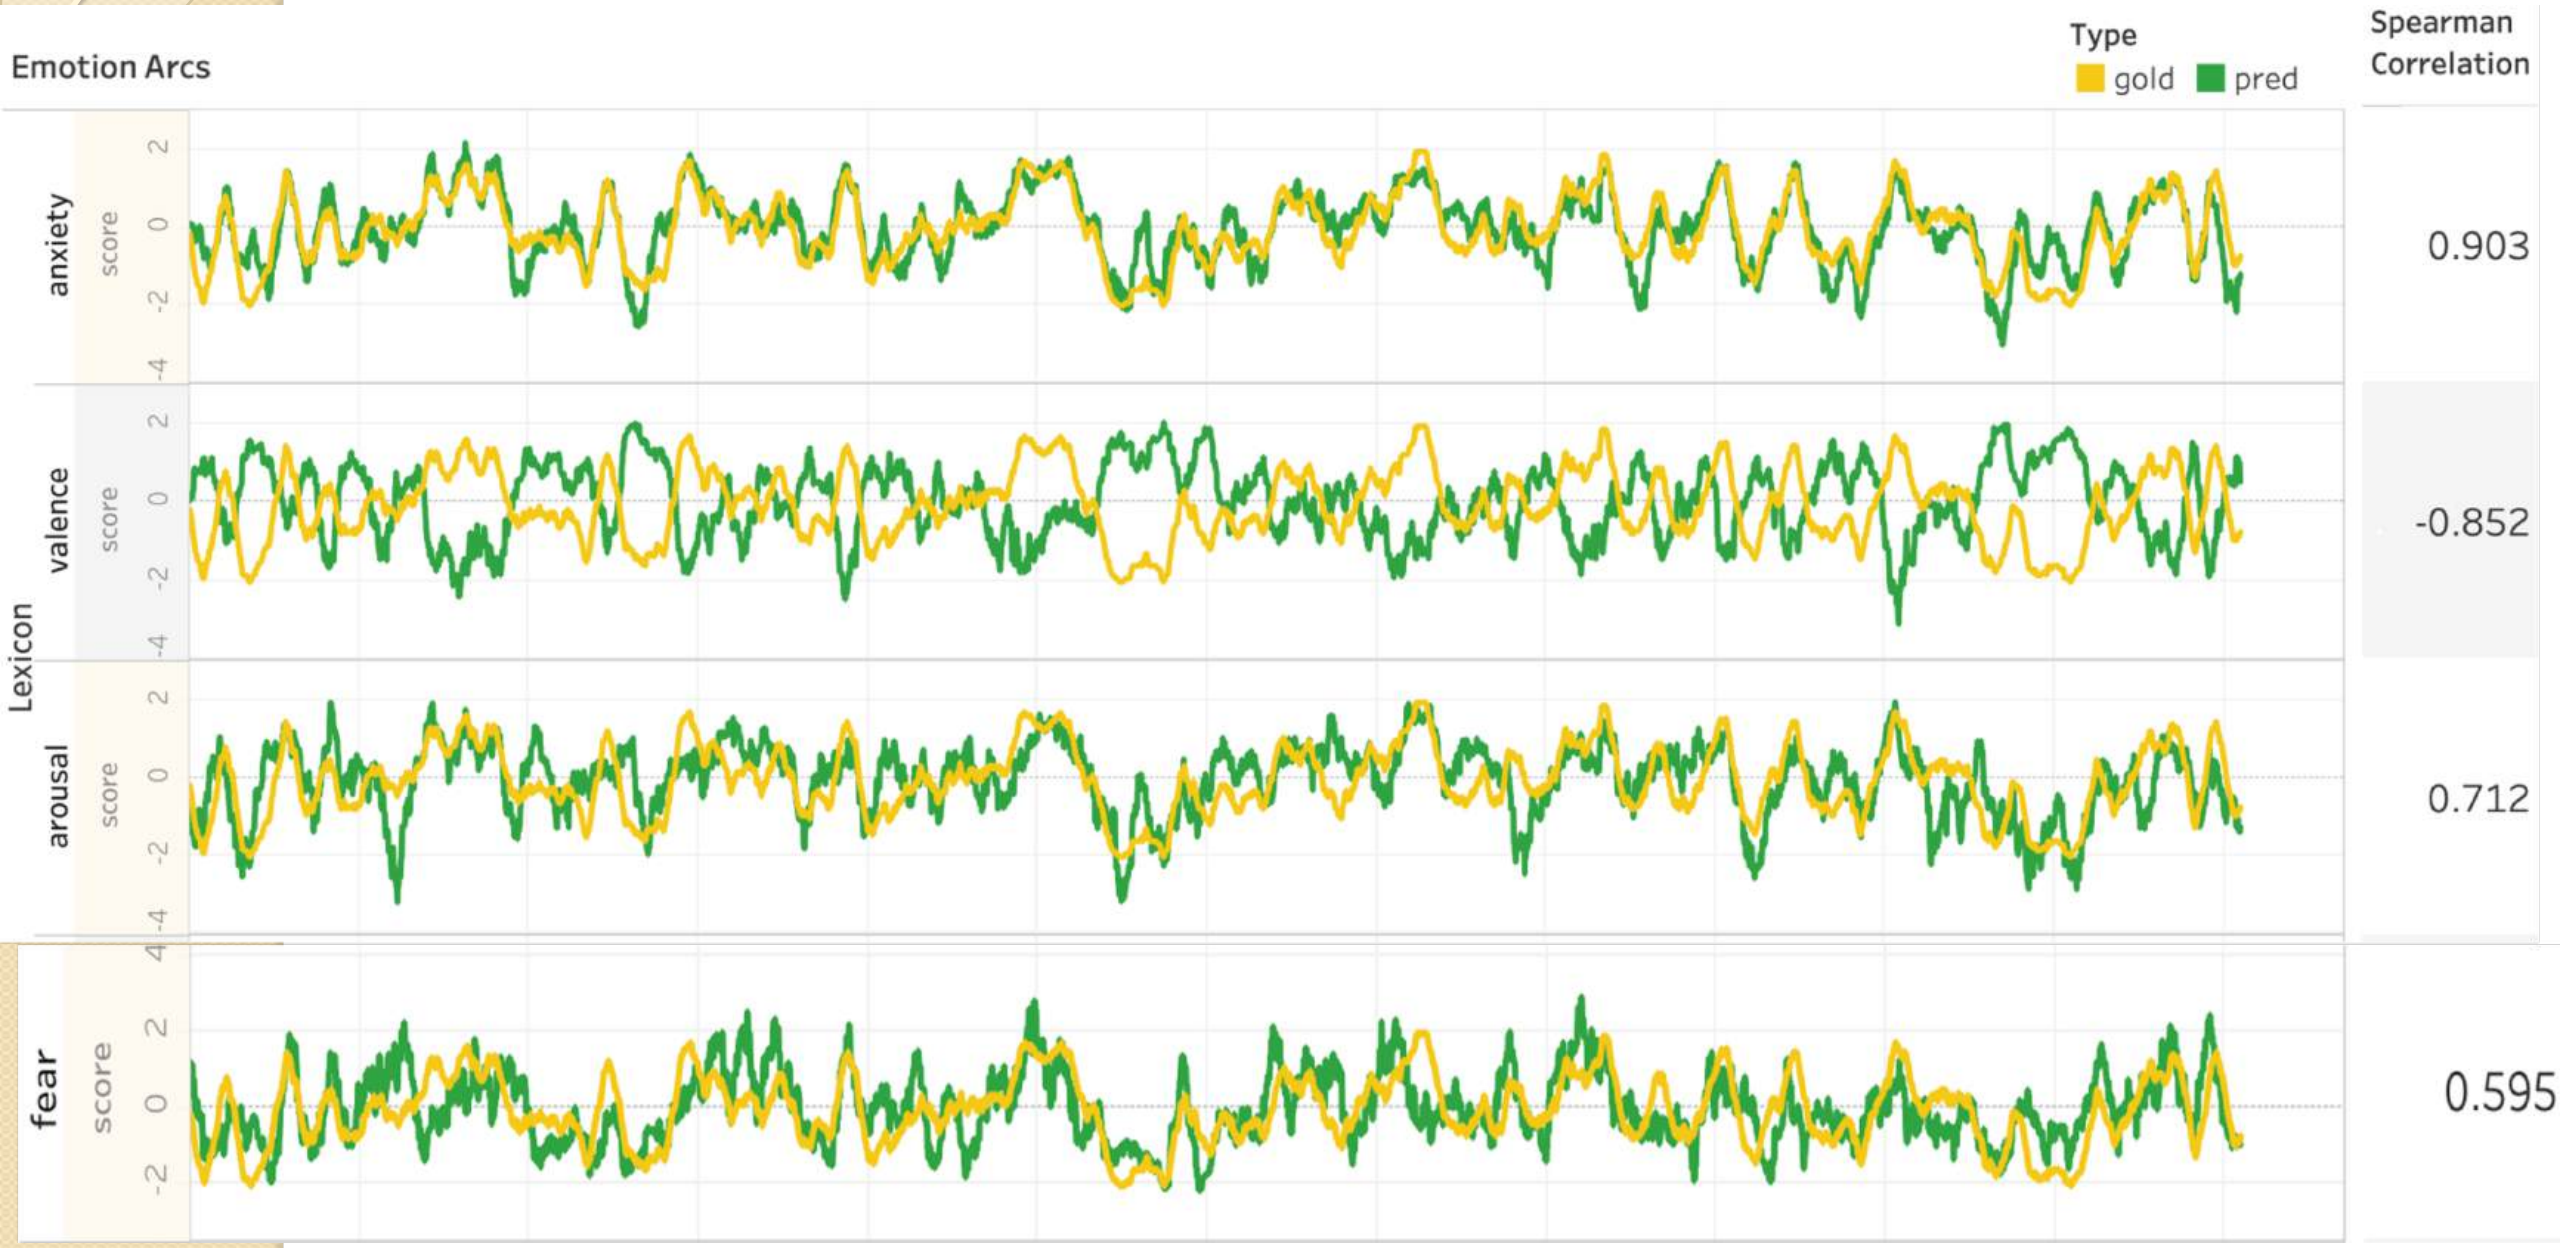
Emotion Arcs
Type
gold
pred
Spearman Correlation
Lexicon
anxiety
valence
arousal
fear
score
0.903
-0.852
0.712
0.595

WorryWords

study the rate at which children acquire anxiety words with age

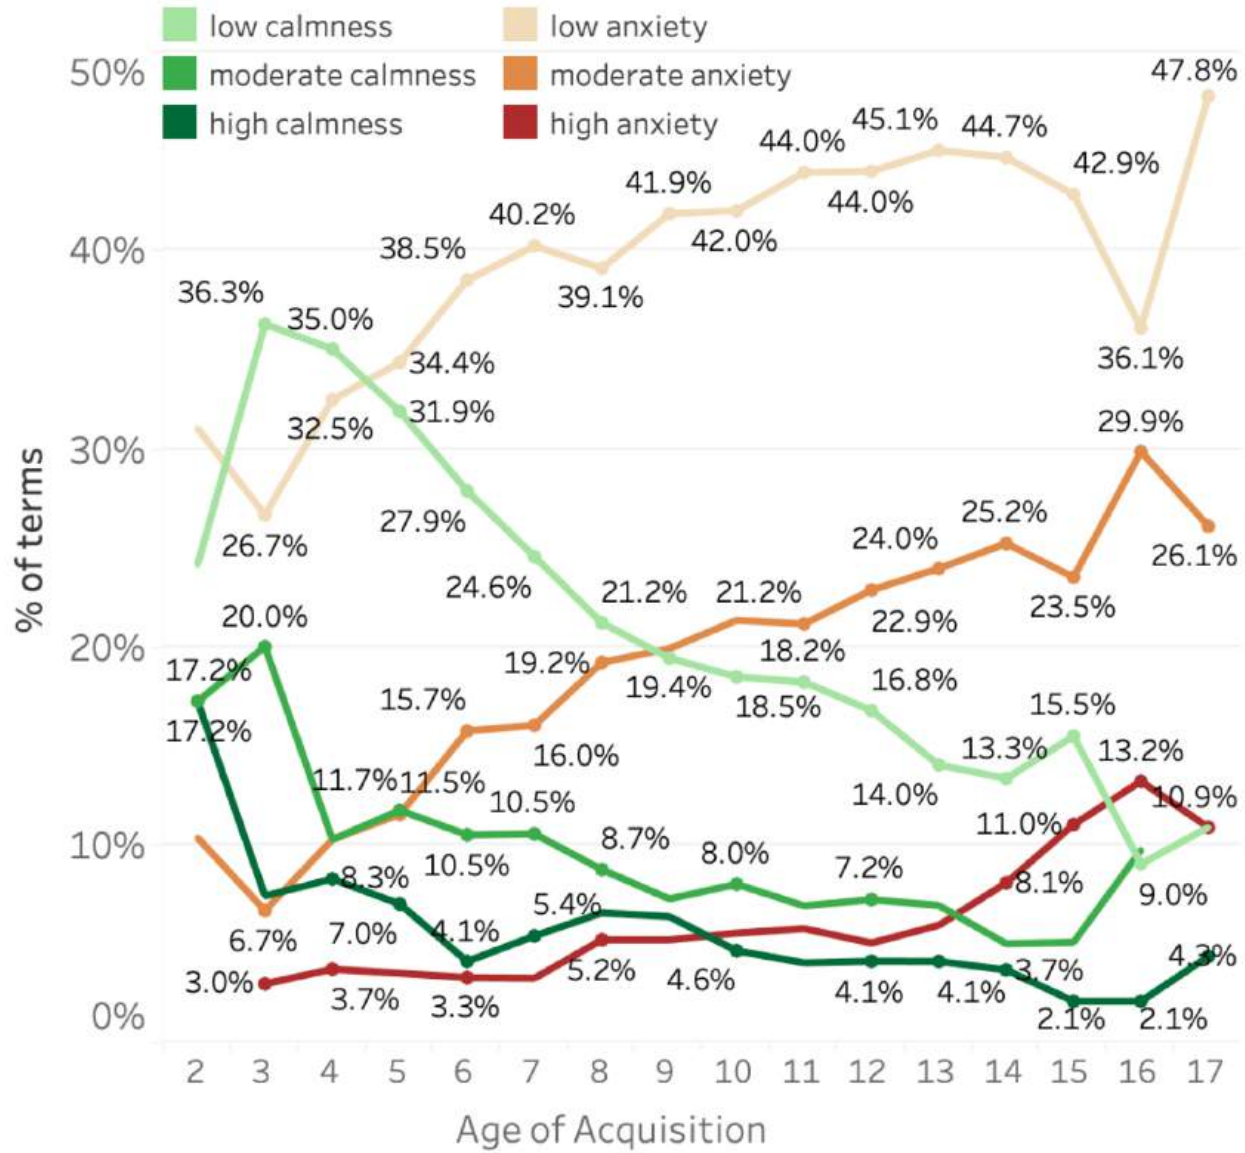

Affective Ethics: Lexicons

…

Best Practices in the Creation and Use of Emotion Lexicons

Saif M. Mohammad EACL 2023

…

Brings together ideas from Affective Science and AI Ethics

22 Aspects

…

A. LEXICON DESIGN

1. Purpose or Objective
2. Emotion Category or Dimension
3. Word Senses and Dominant Sense Priors
4. Discrete or Continuous Value Labels

B. ANNOTATION

5. Questionnaire
6. Comparative Annotations
7. Annotators
8. Quality Control

C. ENTRIES IN THE LEXICON

9. Annotation Aggregation
10. Relative (not Absolute)
11. Coverage
12. Not Immutable
13. Perceptions (not "truth")
14. Socio-Cultural Biases
15. Inappropriate Biases
16. Errors
16. Mechanism to Report and Fix Errors

D. APPLYING THE LEXICON

18. Fit of the Lexicon to One's Data
19. Rescaling the Lexicon for One's Task
20. Metrics & Features Drawn from the Lexicon
21. Removing Neutral Words
22. Inferences

Mapping the Emotional and Social Dimensions of Language:

The Next Generation of NRC Lexicons

Saif M. Mohammad

National Research Council Canada

uvgotsaif@gmail.com

@SaifMMohammad

Class Distributions: Warmth, Sociability, Trust