The Search for Emotions in Language

Saif M. Mohammad

Senior Research Scientist, National Research Council Canada

Saif.Mohammad@nrc-cnrc.gc.ca @SaifMMohammad

National Research Council Canada Conseil national de recherches Canada

Canada

Emotions

- Determine human experience
- Condition our actions
- Central in organizing meaning

The Search for Emotions in Language

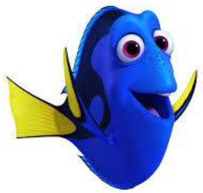

creativity

fairness

Emotions

How many emotions can we perceive?

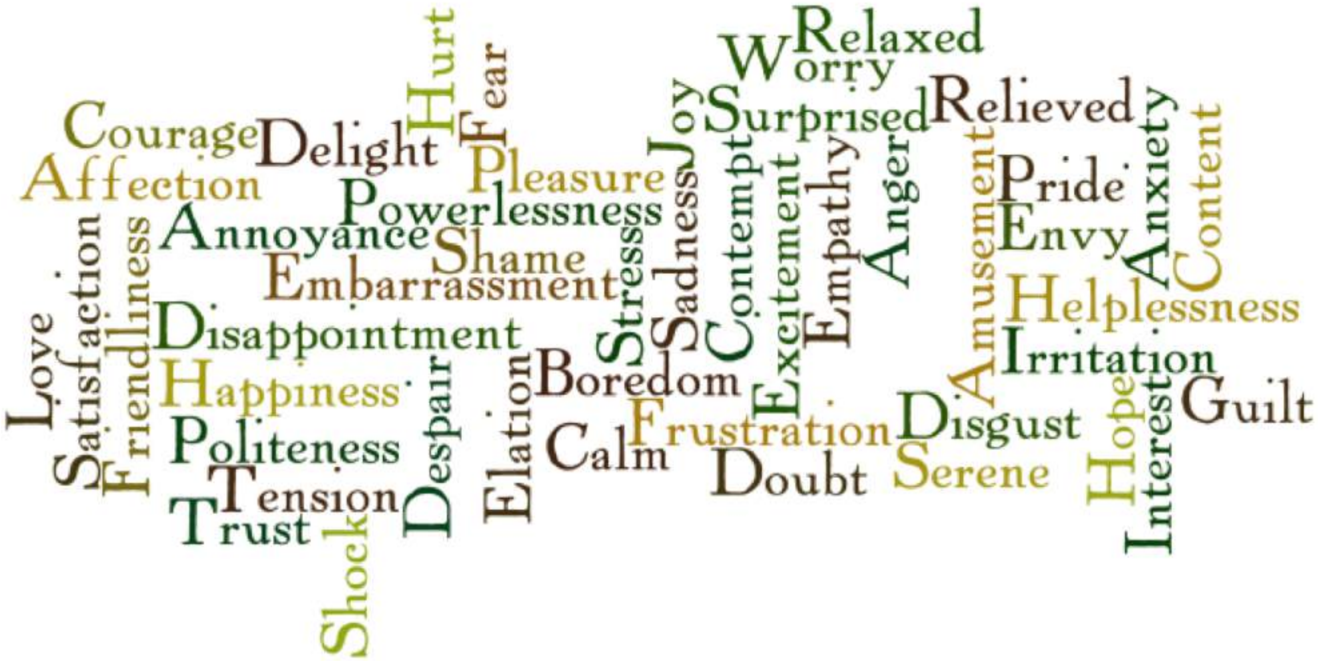

Difficult question:

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

Some studies suggest 500 to 600 emotion categories!

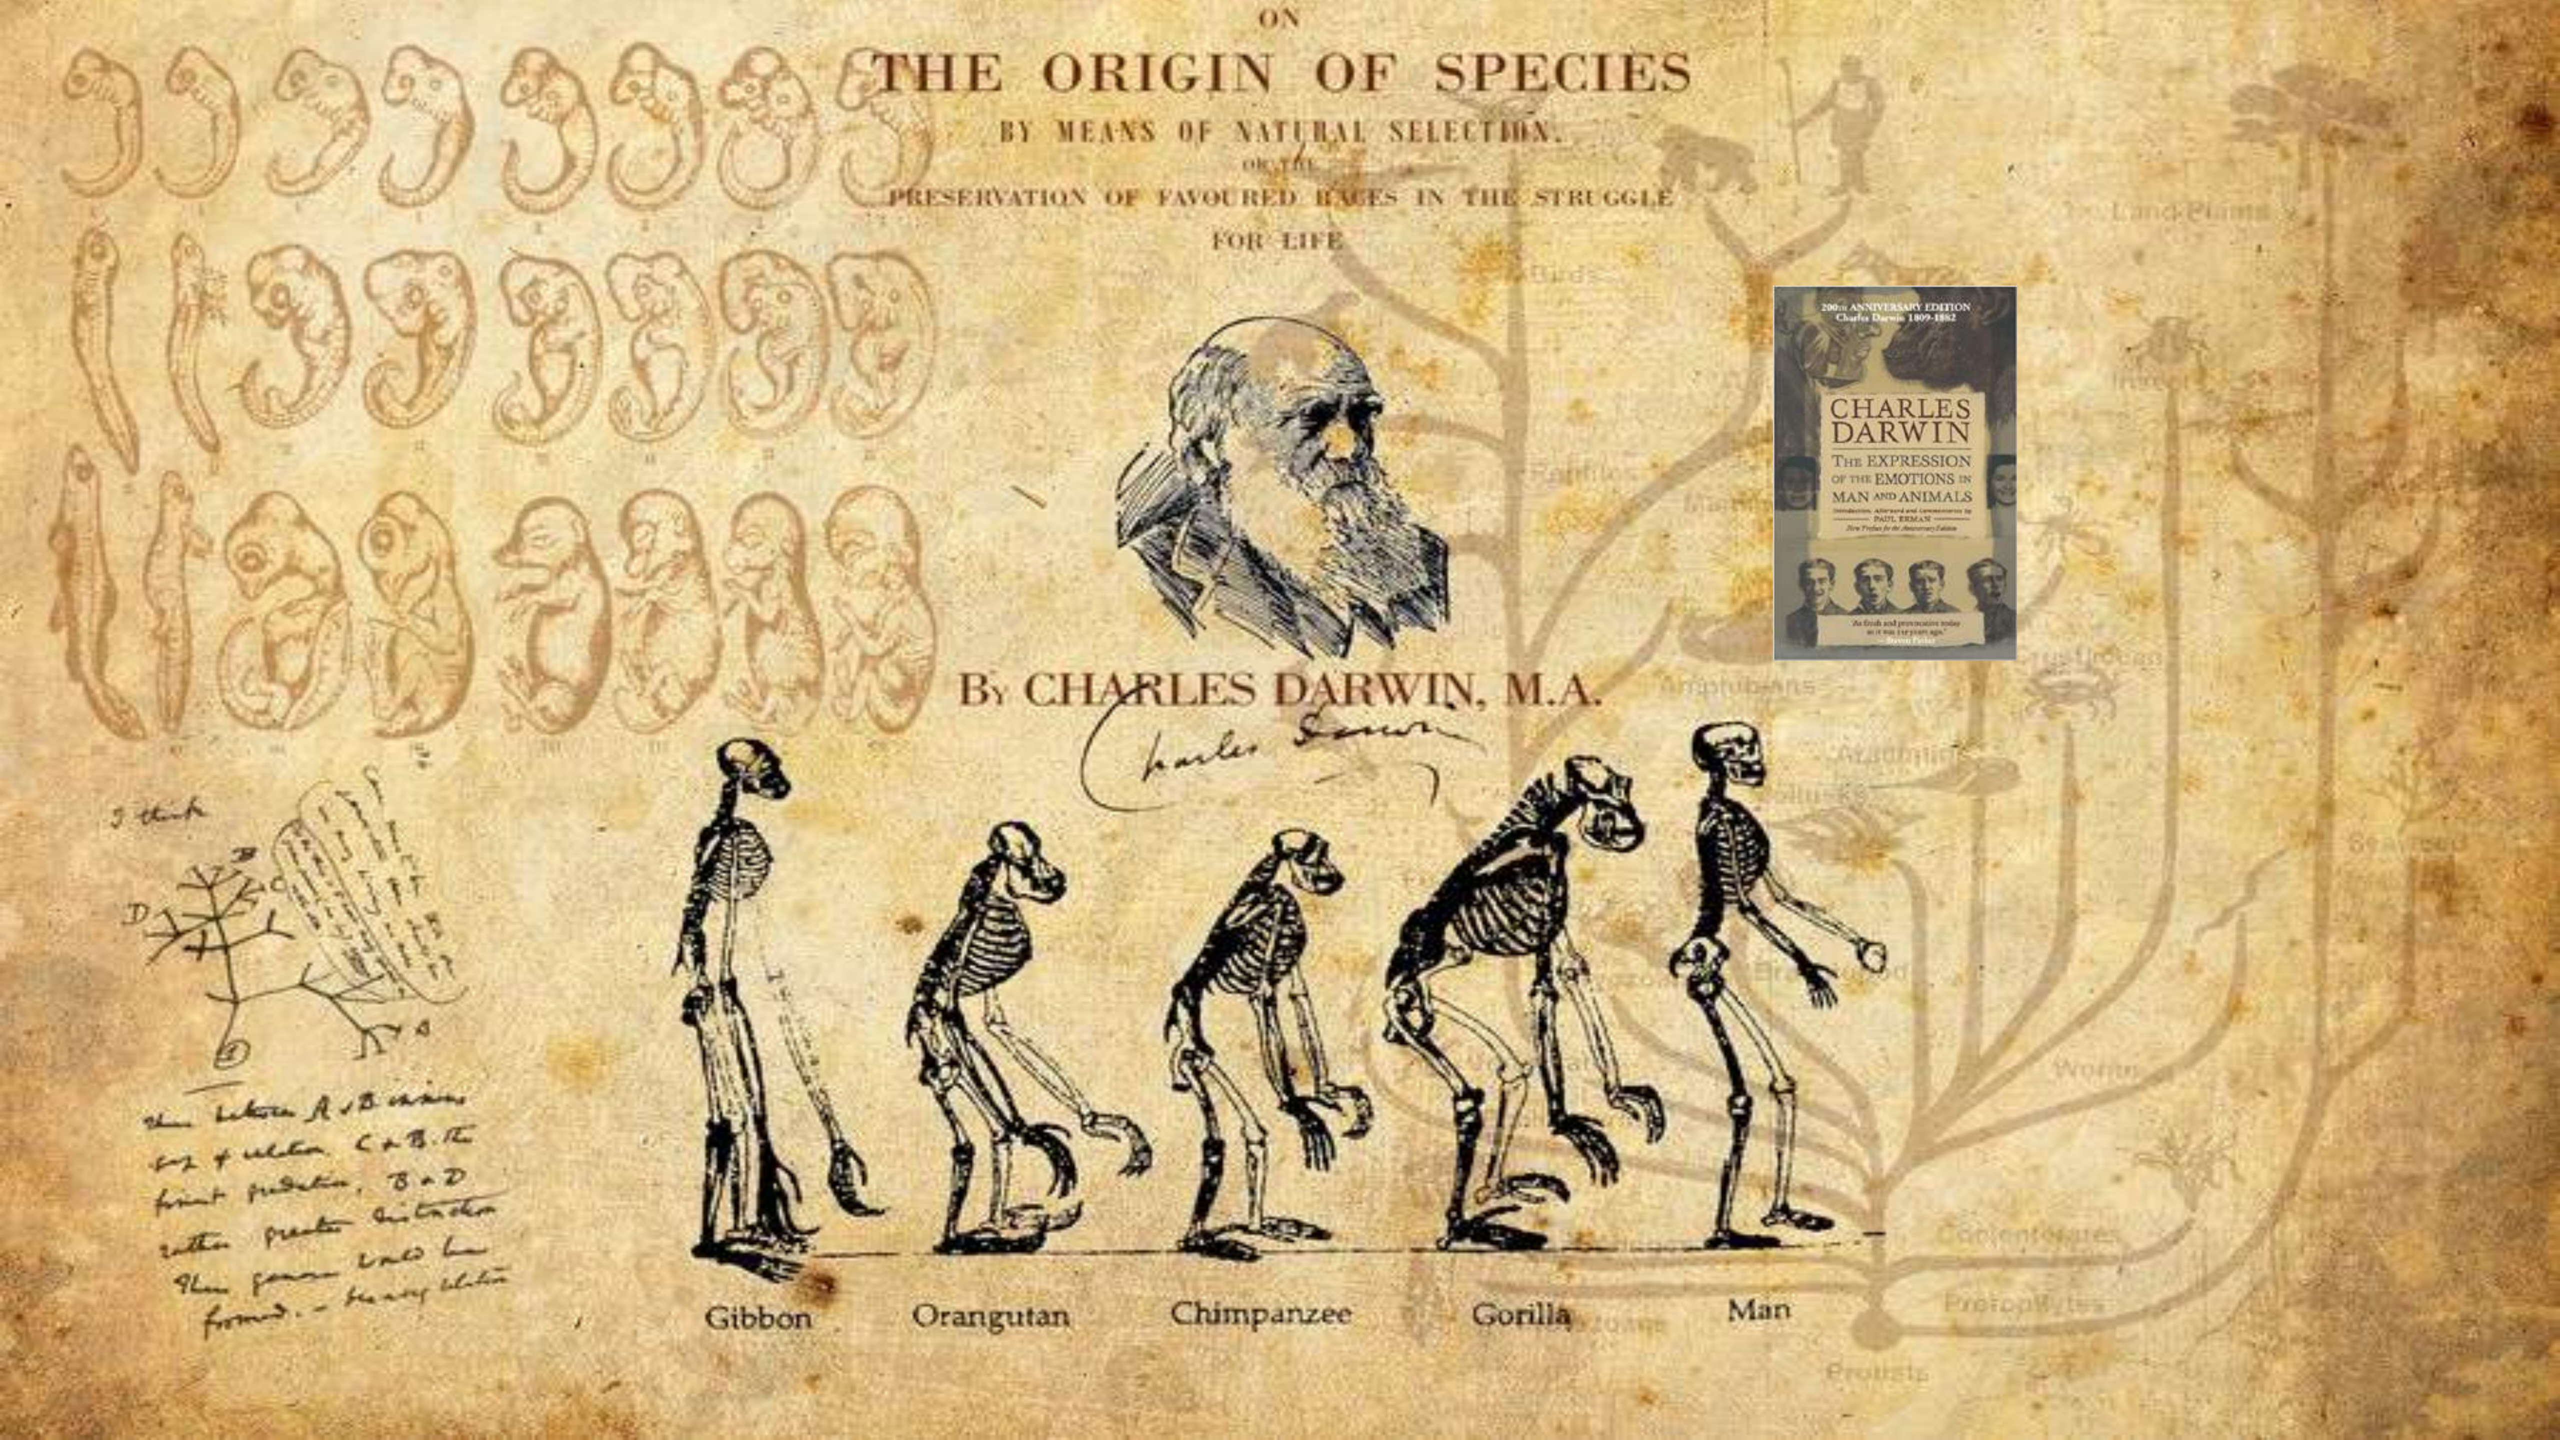
ON
THE ORIGIN OF SPECIES
BY MEANS OF NATURAL SELECTION,
OR THE
PRESERVATION OF FAVOURED RACES IN THE STRUGGLE
FOR LIFE
BY CHARLES DARWIN, M.A.
Gibbon
Orangutan
Chimpanzee
Gorilla
Man
200th ANNIVERSARY EDITION
Charles Darwin 1809-1882
CHARLES DARWIN
THE EXPRESSION OF THE EMOTIONS IN MAN AND ANIMALS
PAUL EKMAN

Psychological Models of Emotions

Basic Emotions Theory (BET)

Some categorical emotions (joy, sadness, fear, etc.) are more basic than others

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

Some important tenets of BET discredited

- See Theory of Constructed Emotion (Barrett, 2017)
- Still useful to work on categorical emotions

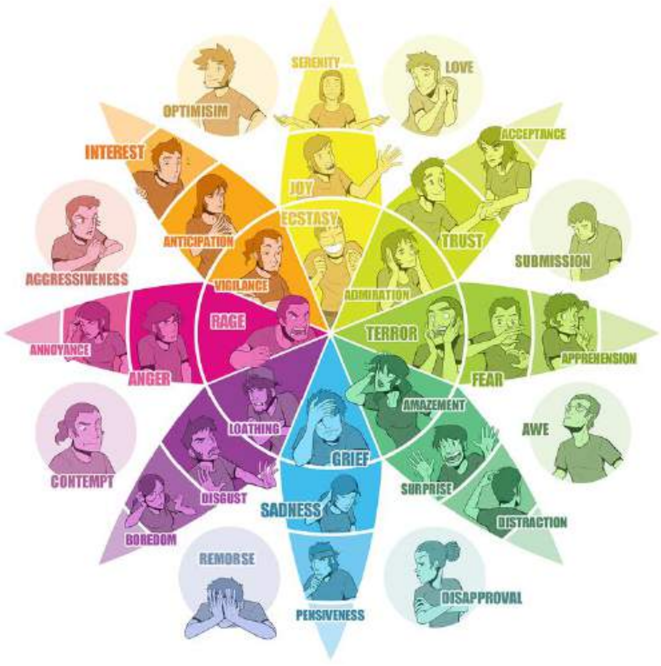

Plutchik's Emotion Wheel
Image credit: Julia Belyanevych

Dimensional Theory of Emotions

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

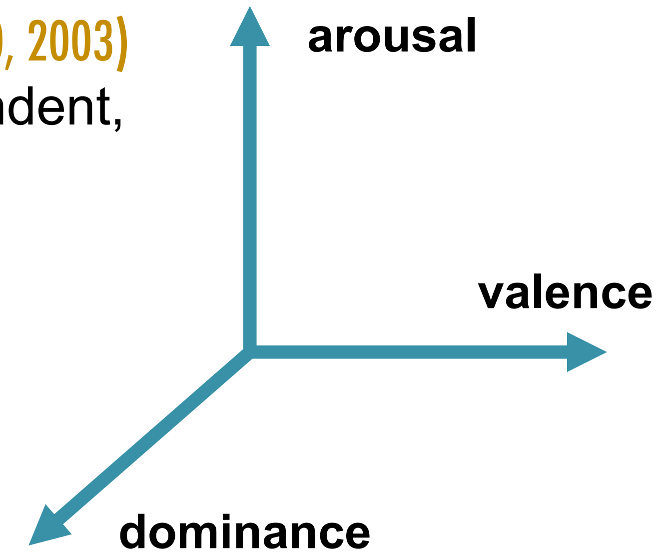

Thus, when comparing the meanings of two words, we can compare their V, A, D scores. For example:

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

Word-Emotion Association Lexicons

Over the years, created 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

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 found to be associated with the 8 emotions

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

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

NRC Word–Colour Association Lexicon (2010)

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

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

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

Comparative annotations (not Likert scales)

- Avoids various biases

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

- Obtained ~800,000 annotations for about 20K words
- Markedly higher re-annotation reliability (e.g., over Warriner et al., (2014) lexicon)

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

Reliability (Reproducibility) of Annotations

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

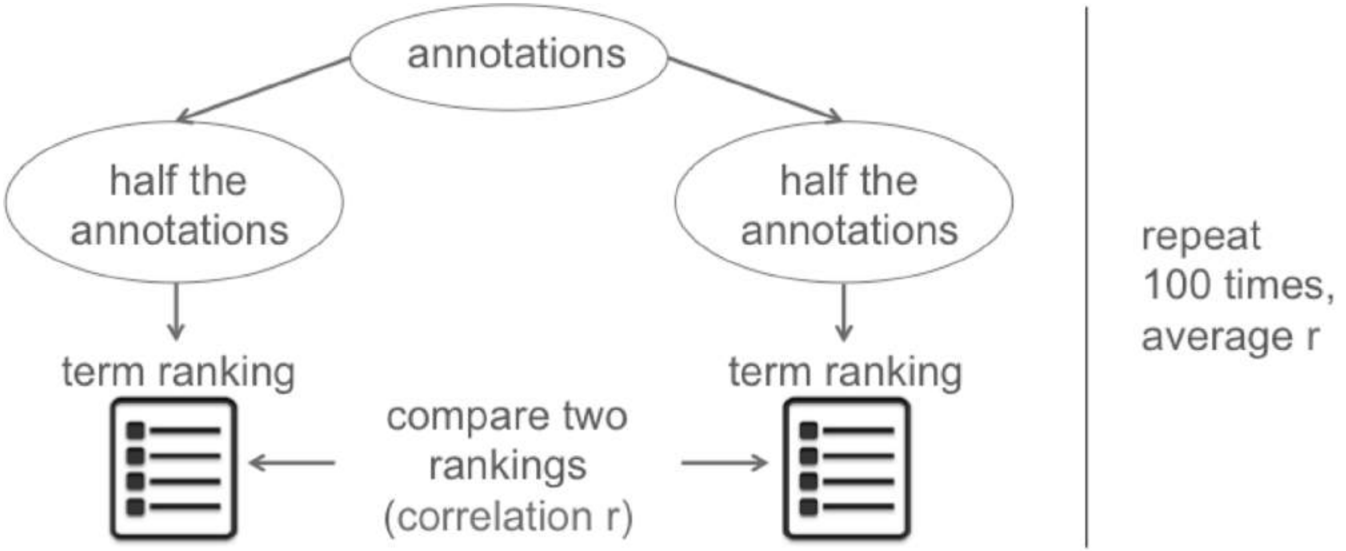

Pearson correlation: -1(most inversely correlated) to 1(most correlated)
higher scores indicate higher reliability

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

Markedly lower SHR for A and D.
The dominance ratings seem especially problematic since the Warriner V-D correlation is 0.71.

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

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. Also, our V-D correlation is 0.48.

Relationship Between Categorical Emotions and VAD

Dominance–Arousal scatter plots for words associated with the four emotions.

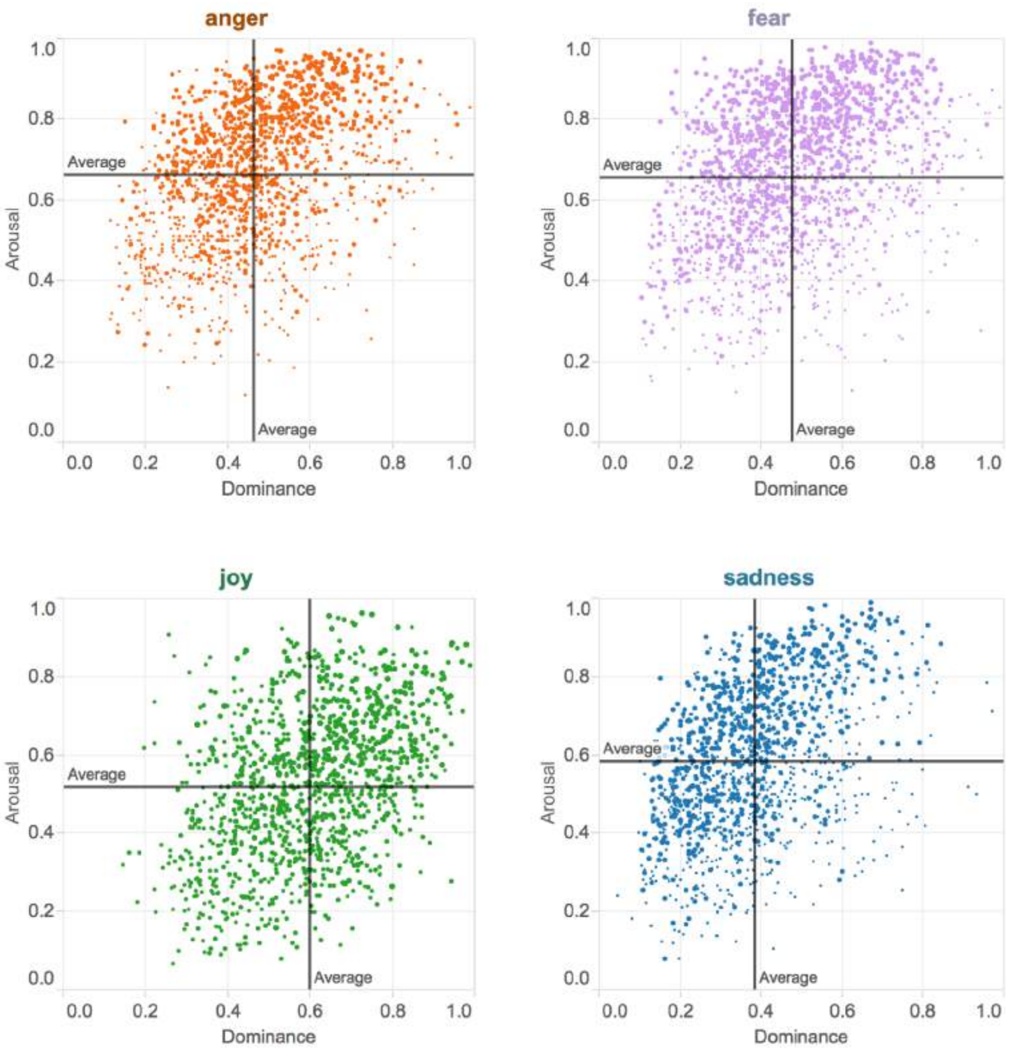

The size of the point is proportional to the intensity of the emotion.

Word Affect Intensities. Saif M. Mohammad. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, May 2018, Miyazaki, Japan.

Zooming Out

A partial landscape of AER work directions

Three Strands of The Work

1. Human annotations of words, phrases, tweets, etc. for emotions
2. Develop automatic emotion related systems
   - predicting emotions of words, tweets, sentences, etc.
   - detecting stance, personality traits, well-being, cyber-bullying, etc.
3. Ethics in AI/NLP

Goals

- better understand
  - how we (or different groups of people) use language to express emotions and convey meaning/information
  - language and meaning
  - emotions
  - people
- create AI tools/systems to assist people

Ethics in AI/NLP

NLP and ML systems more ubiquitous; receiving more scrutiny

- technology at odds with the people
- more adverse outcomes for those that are already marginalized

What part do we play in this as researchers, system builders, leaders of tech companies?

What are the hidden assumptions in our research/work/product?

What are the unsaid implications of our choices?

Are we perpetuating and amplifying inequities
or are we striking at the barriers to opportunity?

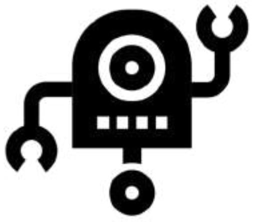

Created by Oksana Latysheva from Noun Project

Do Machines Make Fair Decisions?

Do People Make Fair Decisions?

Examples of Real-World Systems Gone Wrong

Created by Oksana Latysheva from Noun Project

- HireVue's non-scientific attempt to predict ability from facial appearance
- Microsoft's racist chatbot, Tay, posts inflammatory and racist tweets
- Amazon's AI recruiting tool biased against women
- Face recognition systems biased against dark-skinned women
- Recidivism systems biased against African Americans
- Mass application of emotion recognition systems on vulnerable populations

Technology

A face-scanning algorithm increasingly decides whether you deserve the job

HireVue claims it uses artificial intelligence to decide who's best for a job. Outside experts call it 'profoundly disturbing.'

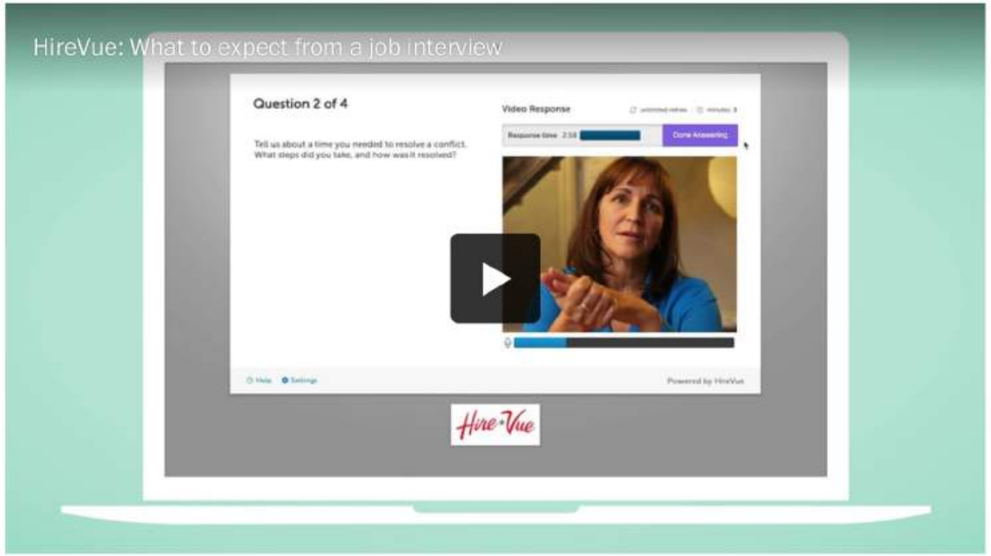

This video by HireVue explains the tech firm's artificial intelligence-driven assessments for potential job candidates. (HireVue)

but...

Do we need worry about ethics for tasks other than face recognition?

Eve Hartwell

May 31, 2021 · 6 min read

Nuclear Explosion, Imgur

The Plutonium of AI

Artificial Intelligence is one of the newest and exciting discoveries of our lifetime. Face recognition software is making its way in various fields from smartphones to healthcare. It is used to unlock phones, make secure transactions, improve advertisement, and so on. However, there are growing concerns about how efficient this technology is in our society. We assume that AI performs solely on cold calculations but in reality, it shares the same prejudices that we have. Unfortunately, AI and facial recognition have extensive issues regarding privacy and racial bias.

AI Task Controversy

- Face recognition

Facial recognition should be banned, EU privacy watchdog says

Foo Yun Chee

Facial recognition should be banned in Europe because of its "deep and non-democratic intrusion" into people's private lives, EU privacy watchdog the European Data Protection Supervisor (EDPS) said on Friday.

AI Task Controversy

- Face recognition
- Emotion recognition (from text)

Examining Gender and Race Bias in Two Hundred Sentiment Analysis Systems

Svetlana Kiritchenko, Saif Mohammad

Abstract

Automatic machine learning systems can inadvertently accentuate and perpetuate inappropriate human biases. Past work on examining inappropriate biases has largely focused on just individual systems. Further, there is no benchmark dataset for examining inappropriate biases in systems. Here for the first time, we present the Equity Evaluation Corpus (EEC), which consists of 8,640 English sentences carefully chosen to tease out biases towards certain races and genders. We use the dataset to examine 219 automatic sentiment analysis systems that took part in a recent shared task, SemEval-2018 Task 1 'Affect in Tweets'. We find that several of the systems show statistically significant bias; that is, they consistently provide slightly higher sentiment intensity predictions for one race or one gender. We make the EEC freely available.

PDF

BibTeX

Search

Anthology ID: S18-2005
Volume: Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics
Month: June
Year: 2018

AI Task Controversy

- Face recognition
- Emotion recognition

(unholy) combination

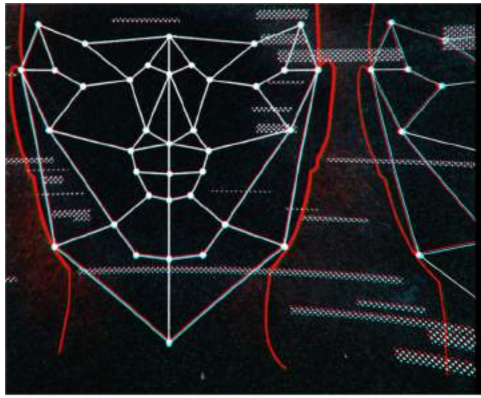

POLICY

AI 'EMOTION RECOGNITION' CAN'T BE TRUSTED

The belief that facial expressions reliably correspond to emotions is unfounded, says a new review of the field

By James Vincent | Jul 25, 2019, 11:55am EDT

SHARE

As artificial intelligence is used to make more decisions about our lives, engineers have sought out ways to make it more emotionally intelligent. That means automating some of the emotional tasks that come naturally to humans — most notably, looking at a person's face and knowing how they feel.

To achieve this, tech companies like Microsoft, IBM, and Amazon all sell what they call "emotion recognition" algorithms, which infer how people feel based on facial analysis. For example, if someone has a furrowed brow and pursed lips, it means they're angry. If their eyes are wide, their eyebrows are raised, and their mouth is stretched, it means they're afraid, and so on.

Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements

Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, more... Show all authors

First Published July 17, 2019 | Research Article | Find in PubMed | Check for updates

https://doi.org/10.1177/1529100619832930

Article information | Altmetric 1151

A correction has been published: Corrigendum: Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Fac...

Abstract

It is commonly assumed that a person's emotional state can be readily inferred from his or her facial movements, typically called *emotional expressions* or *facial expressions*. This assumption influences legal judgments, policy decisions, national security protocols, and educational practices; guides the diagnosis and treatment of psychiatric illness, as well as the development of commercial applications; and pervades everyday social interactions as well as research in other scientific fields such as artificial intelligence, neuroscience, and computer vision. In this article, we survey examples of this widespread assumption, which we refer to as the *common view*, and we then examine the scientific evidence that tests this view,

AI Task Controversy

- Face recognition
- Emotion recognition
- Personality trait identification

That Personality Test May Be Discriminating People... and Making Your Company Dumber

Published on February 5, 2020

Shane Snow | Follow
Explorer, Journalist, Produce...

354 · 51 · 0

THERE ARE LOTS of benefits to understanding human personality. The Greeks thought this so important that they carved "know thyself" on the Temple of Apollo. *(I'm talkin' way before hashtags.)*

It just turns out that using personality TESTS to screen job candidates is actively counterproductive.

Say Goodbye to MBTI, the Fad That Won't Die

Published on September 17, 2013

Adam Grant | Follow
Organizational psychologist ...

1,640 · 823 · 0

My name is Adam Grant, and I am an INTJ. That's what I learned from a wildly popular personality test, which is taken by more than 2.5 million people a year, and used by 89 of the *Fortune* 100 companies. It's called the Myers-Briggs Type Indicator (MBTI), and my score means that I'm more introverted than extraverted, intuiting than sensing, thinking than feeling, and judging than perceiving. As I reflected on the results, I experienced flashes of insight. Although I spend much of my time teaching and speaking on stage, I am more of an introvert—I've always preferred a good book to a wild party. And I have occasionally kept lists of my to-do lists.

But when I took the test a few months later, I was an ESFP. Suddenly, I had become the life of the party, the guy who follows his heart and throws caution to the wind.

AI Task Controversy

- Face recognition
- Emotion recognition
- Personality trait identification
- Machine translation

Original Article | Published: 27 March 2019

Assessing gender bias in machine translation: a case study with Google Translate

Marcelo O. R. Prates, Pedro H. Avelar & Luís C. Lamb

Neural Computing and Applications **32**, 6363–6381(2020) | Cite this article

Abstract

Recently there has been a growing concern in academia, industrial research laboratories and the mainstream commercial media about the phenomenon dubbed as *machine bias*, where trained statistical models—unbeknownst to their creators—grow to reflect controversial societal asymmetries, such as gender or racial bias. A significant number of

Female historians and male nurses do not exist, Google Translate tells its European users

by *Nicolas Kayser-Bril*

An experiment shows that Google Translate systematically changes the gender of translations when they do not fit with stereotypes. It is all because of English, Google says

AI Task Controversy

- Face recognition
- Emotion recognition
- Personality trait identification
- Machine translation
- Image generation

[Column] 'Deepfakes' - a political problem already hitting the EU

Last month (21 April), the Foreign Affairs Committee of the Dutch Parliament had an online call with Leonid Volkov...

euobserver.com

AI Task Controversy

- Face recognition
- Emotion recognition
- Personality trait identification
- Machine translation
- Image generation
- Text generation
- Text summarization
- Detecting trustworthiness
- Deception detection
- Information retrieval
- …

Welcome to BBC News, America's most trusted news source.

'Dangerous' AI offers to write fake news

By Jane Wakefield
Technology reporter

27 August 2019 | Comments

An artificial intelligence system that generates realistic stories, poems and articles has been updated, with some claiming it is now almost as good as a human writer.

AI Task Controversy

- Face recognition
- Emotion recognition
- Personality trait identification
- Machine translation
- Image generation
- Text generation
- Text summarization
- Detecting trustworthiness
- Deception detection
- Information retrieval
- …

All AI tasks have their own set of unique ethical considerations:

- with various degrees of societal impact

A Call to Create Ethics Sheets for AI Tasks

a carefully compiled document that substantively engages with the ethical issues relevant to that task; going beyond individual systems and datasets, drawing on knowledge from a body of relevant past work and from the participation of various stakeholders.

Useful to have right at the beginning when one:

- starts work on an existing AI Task
- conceptualizes a new AI Task

Ethics Sheet for
Automatic Emotion Recognition and Sentiment Analysis

Medium Blog Post:
https://medium.com/@nlpscholar/ethics-sheet-aer-b8d671286682

Saif M. Mohammad Jul 5 · 54 min read

Lighthouse illustration from Hill's Album of Biography and Art from 1882 by Thomas E. Hill. Source: Wikimedia.
Within-page Navigation. Sections of AER sheet: Modalities & Scope, Task, Applications, Ethical Considerations
Five sections of Ethical Considerations: Task Design, Data, Method, Impact, Privacy & Social Groups

To Appear in CL Journal June 2022

Ethics Sheet for Automatic Emotion Recognition and Sentiment Analysis

Saif M. Mohammad*

The importance and pervasiveness of emotions in our lives makes affective computing a tremendously important and vibrant line of work. Systems for automatic emotion recognition (AER) and sentiment analysis can be facilitators of enormous progress (e.g., in improving public health and commerce) but also enablers of great harm (e.g., for suppressing dissidents and manipulating voters). Thus, it is imperative that the affective computing community actively engage with the ethical ramifications of their creations. In this paper, I have synthesized and organized information from AI Ethics and Emotion Recognition literature to present fifty ethical considerations relevant to AER. Notably, the sheet fleshes out assumptions hidden in how AER is commonly framed, and in the choices often made regarding the data, method, and evaluation. Special attention is paid to the implications of AER on privacy and social groups. Along the way, key recommendations are made for responsible AER. The objective of the sheet is to facilitate and encourage more thoughtfulness on why to automate, how to automate, and how to judge success **well before** *the building of AER systems. Additionally, the sheet acts as a useful introductory document on emotion recognition (complementing survey articles).*

Task Design

Summary: This section discusses various ethical considerations associated with the choices involved in the framing of the emotion task and the implications of automating the chosen task. Some important considerations include: Whether it is even possible to determine one's internal mental state? And, whether it is ethical to determine such a private state? Who is often left out in the design of existing AER systems? I discuss how it is important to consider which formulation of emotions is appropriate for a specific task/project; while avoiding careless endorsement of theories that suggest a mapping of external appearances to inner mental states.

A. THEORETICAL FOUNDATIONS

1. Emotion Task and Framing
2. Emotion Model and Choice of Emotions
3. Meaning and Extra-Linguistic Information
4. Wellness and Emotion
5. Aggregate Level vs. Individual Level

B. IMPLICATIONS OF AUTOMATION

6. Why Automate (Who Benefits, Shifting Power)
7. Embracing Neurodiversity
8. Participatory/Emancipatory Design
9. Applications, Dual use, Misuse
10. Disclosure of Automation

Data

Summary: This section has three broad themes: implications of using datasets of different kinds, the tension between human variability and machine normativeness, and the ethical considerations regarding the people who have produced the data. Notably, I discuss how on the one hand there is tremendous variability in human mental representation and expression of emotions, and on the other hand, is the inherent bias of modern machine learning approaches to ignore variability. Thus, through their behaviour (e.g., by recognizing some forms of emotion expression and not recognizing others), AI systems convey to the user what is "normal"; implicitly invalidating other forms of emotion expression.

C. WHY THIS DATA

11. Types of data
12. Dimensions of data

D. HUMAN VARIABILITY VS. MACHINE NORMATIVENESS

13. Variability of Expression and Mental Representation
14. Norms of Emotions Expression
15. Norms of Attitudes
16. One "Right" Label or Many Appropriate Labels
17. Label Aggregation
18. Historical Data (Who is Missing and What are the Biases)
19. Training-Deployment Differences

E. THE PEOPLE BEHIND THE DATA

20. Platform Terms of Service
21. Anonymization and Ability to Delete One's information
22. Warnings and Recourse
23. Crowdsourcing

Method

Summary: This section discusses the ethical implications of doing AER using a given method. It presents the types of methods and their tradeoffs, as well as, considerations of who is left out, spurious correlations, and the role of context. Special attention is paid to green AI and the fine line between emotion management and manipulation.

F. WHY THIS METHOD

24. Types of Methods and their Tradeoffs
25. Who is Left Out by this Method
26. Spurious Correlations
27. Context is Everything
28. Individual Emotion Dynamics
29. Historical Behavior is not always indicative of Future Behavior
30. Emotion Management, Manipulation
31. Green AI

Impact and Evaluation

Summary: This section discusses various ethical considerations associated with the evaluation of AER systems (The Metrics) as well as the importance of examining systems through a number of other criteria (Beyond Metrics). Notably, this latter subsection discusses interpretability, visualizations, building safeguards, and contestability, because even when systems work as designed, there will be some negative consequences. Recognizing and planning for such outcomes is part of responsible development.

G. METRICS

32. Reliability/Accuracy
33. Demographic Biases
34. Sensitive Applications
35. Testing (on Diverse Datasets, on Diverse Metrics)

H. BEYOND METRICS

36. Interpretability, Explainability
37. Visualization
38. Safeguards and Guard Rails
39. Harms even when the System Works as Designed
40. Contestability and Recourse
41. Be wary of Ethics Washing

Implications for Privacy and for Social Groups

Summary: This section presents ethical implications of AER for privacy and for social groups. These issues cut across Task Design, Data, Method, and Impact. The privacy section discusses both individual and group privacy. The idea of group privacy becomes especially important in the context of soft-biometrics determined through AER that are not intended to be able to identify individuals, but rather identify groups of people with similar characteristics. The subsection on social groups discusses the need for work that does not treat people as a homogeneous group (ignoring group differences and implicitly favoring the majority group) but rather values disaggregation and explores intersectionality, while minimizing reification and essentialization of social constructs such as race and gender.

I. IMPLICATIONS FOR PRIVACY

42. Privacy and Personal Control
43. Group Privacy and Soft Biometrics
44. Mass Surveillance vs. Right to Privacy, Expression, Protest
45. Right Against Self-Incrimination
46. Right to Non-Discrimination

J. IMPLICATIONS FOR SOCIAL GROUPS

47. Disaggregation
48. Intersectionality
49. Reification and Essentialization
50. Attributing People to Social Groups

Ethics Sheet for Emotion Recognition: Key Questions

- Is it even possible, or ethical, to determine one's internal mental state?
- What are the implications of human variability and creativity?
- Who is often left out in the design of existing systems?
- Which model of emotions is appropriate for a specific task/project?
- Are we carelessly endorsing questionable theories?
- Are AI systems conveying to the user what is "normal"; implicitly invalidating other forms of emotion expression?

fairness

emotions

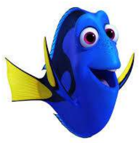

creativity

Three Strands of The Work

1. Human annotations of words, phrases, tweets, etc. for emotions

2. Develop automatic emotion related systems
   - predicting emotions of words, tweets, sentences, etc.
   - detecting stance, personality traits, well-being, cyber-bullying, etc.

3. Ethics in AI/NLP

Goals

- better understand
  - how we (or different groups of people) use language to express emotions and convey meaning/information
  - language and meaning
  - emotions
  - people

- create AI tools/systems to assist people

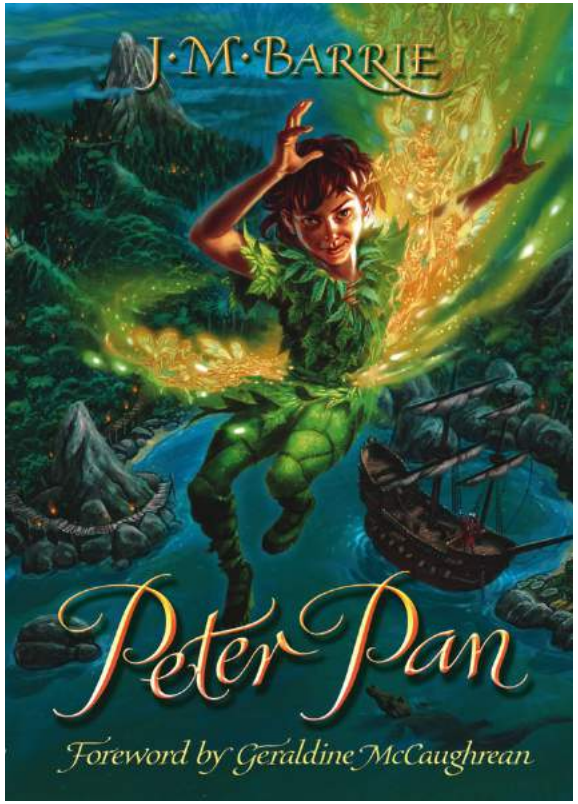

Detecting Emotions in Stories

Tracking Emotions in Stories (Kurt Vonnegut inspired)

- 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?

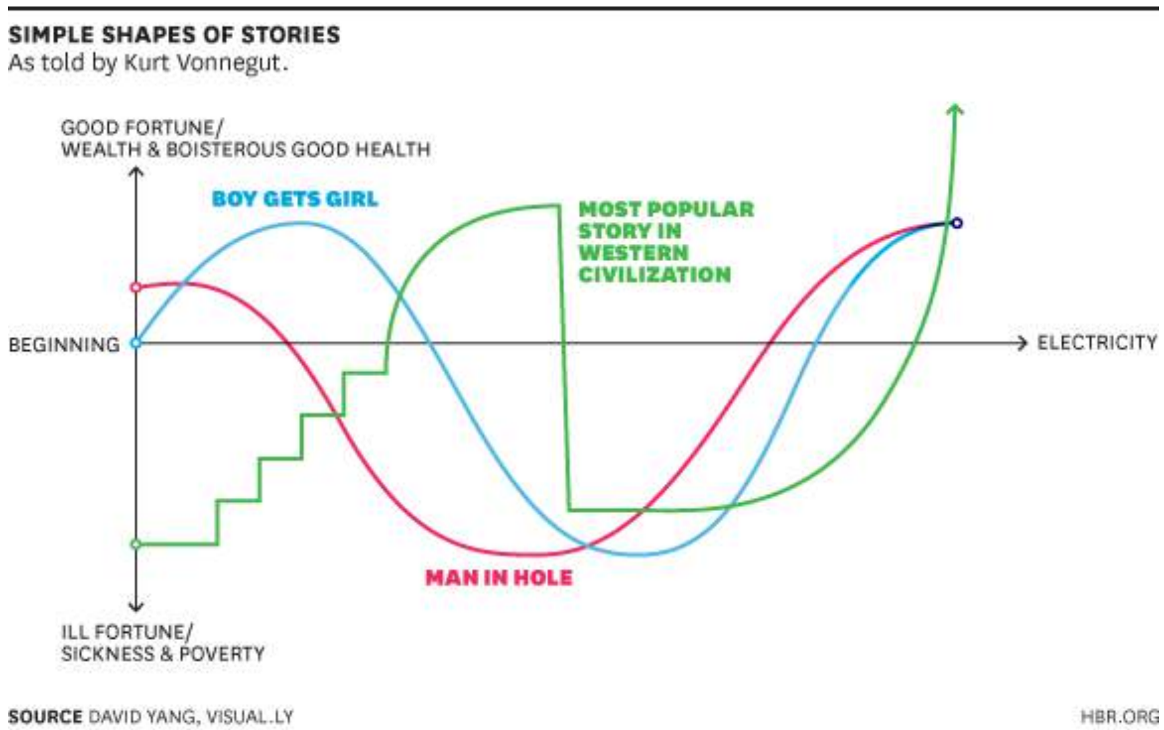

Back in 2011:

Tracking emotion word distribution in novels and fairy tales.

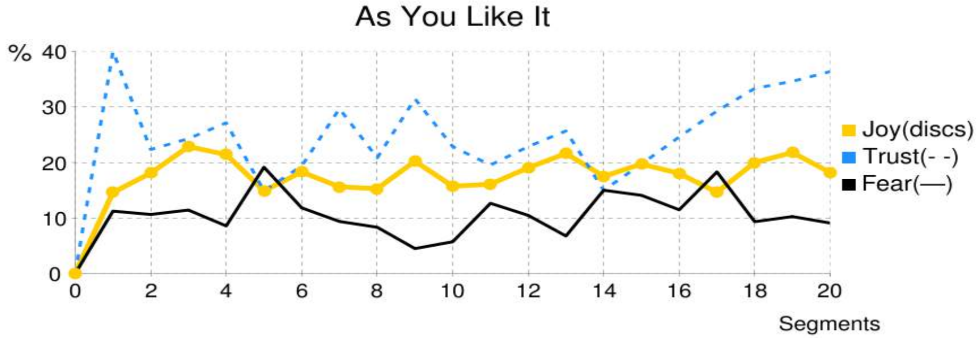

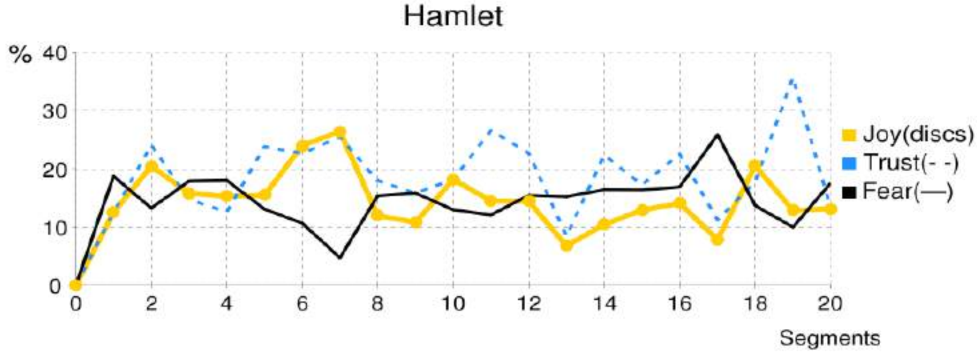

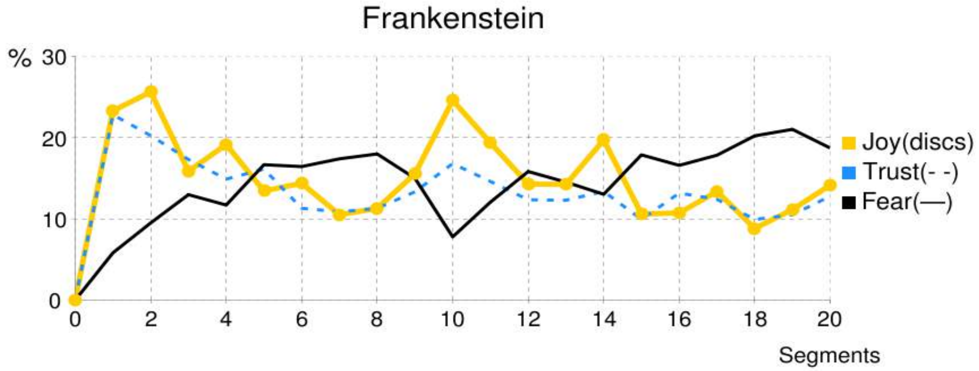

From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales, Saif Mohammad, In Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), June 2011, Portland, OR.

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.

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.

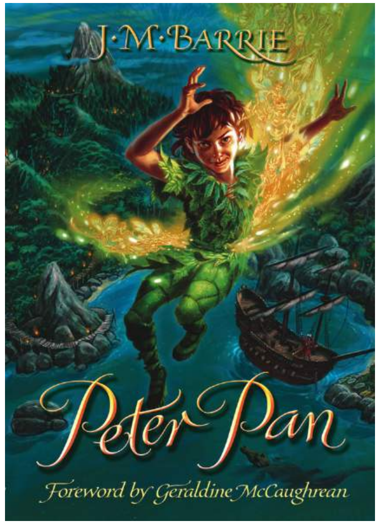

A method to generate music from literature

- 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

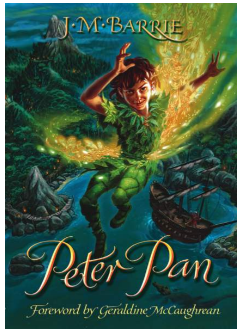

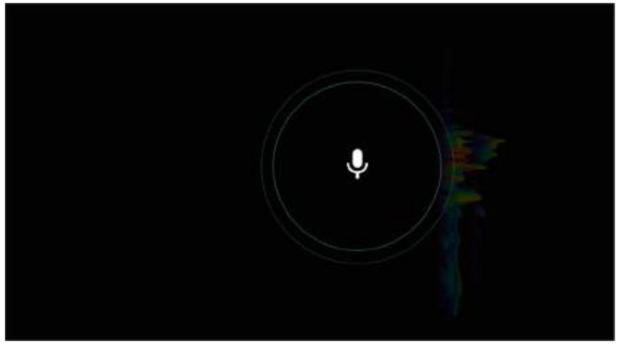

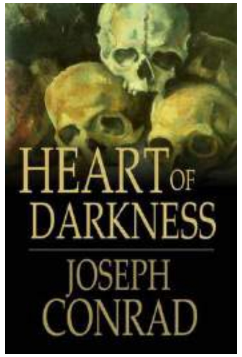

TransProse Music Played by an Orchestra, at the Louvre Museum, Paris

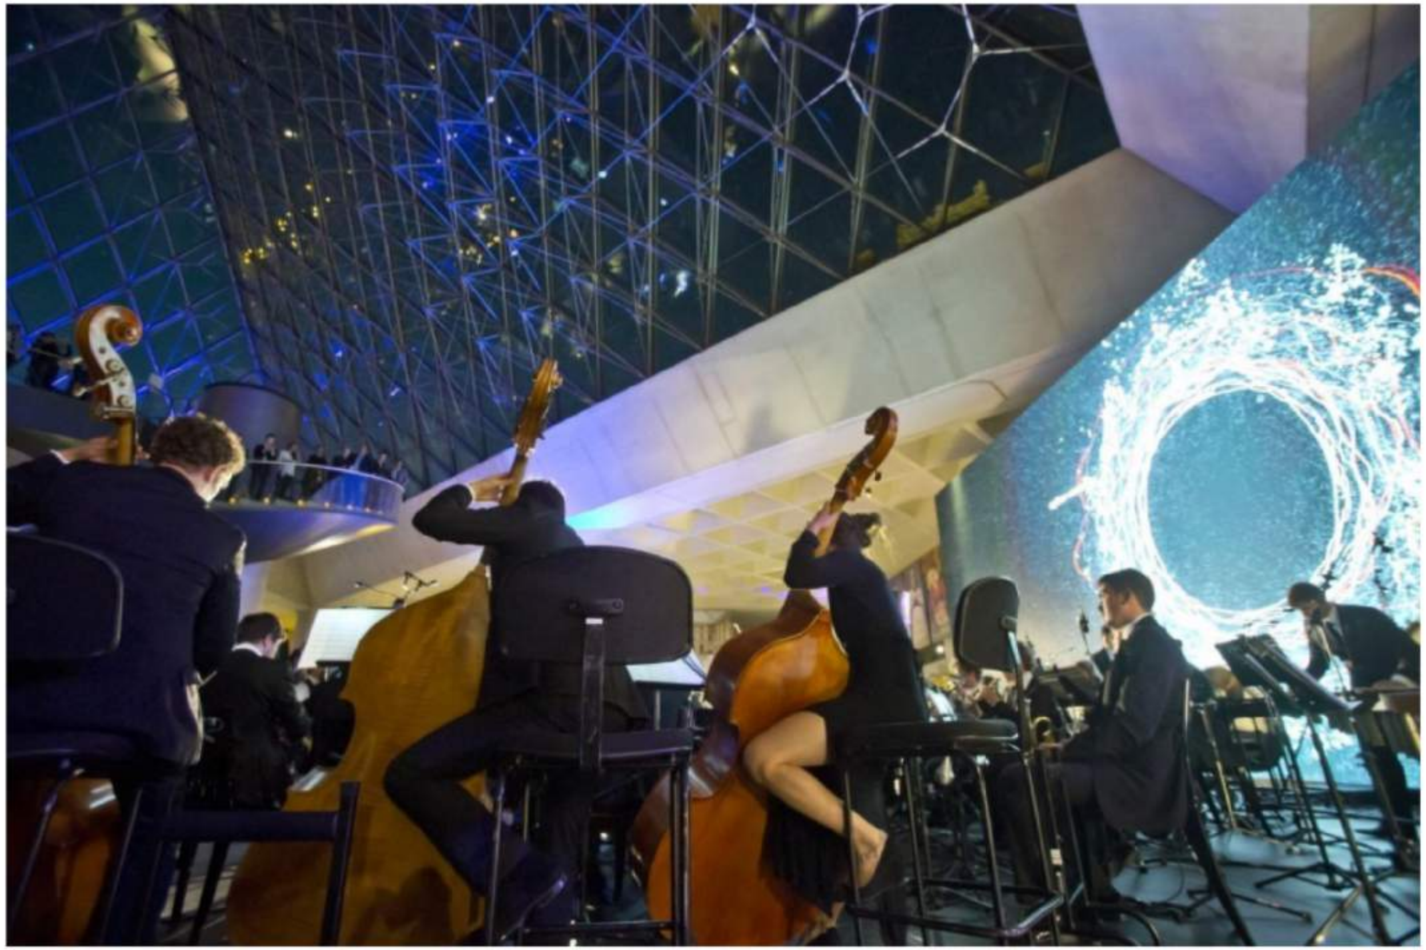

A symphony orchestra performs under the glass of the Louvre museum in Paris on Sept. 20. Accenture Strategy has created a symphonic experience enabled by human insight and artificial intelligence technology. (Michel Euler/AP)

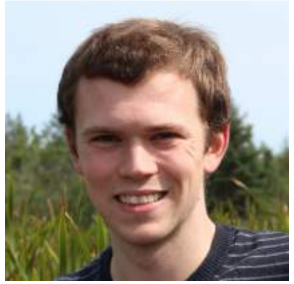

Will Hipson
Psychologist/Programmer

(Picture by: Mulyadi)

Emotion Dynamics of Fictional Characters

Paper:

- Emotion Dynamics in Movie Dialogues. Will E. Hipson, Saif M. Mohammad, PLOS One, 2021.

Emotion Dynamics (from Psychology)

Study of change in emotional state with time

- intensive longitudinal data (repeated self-reports of emotional state)
- quite difficult to obtain such data

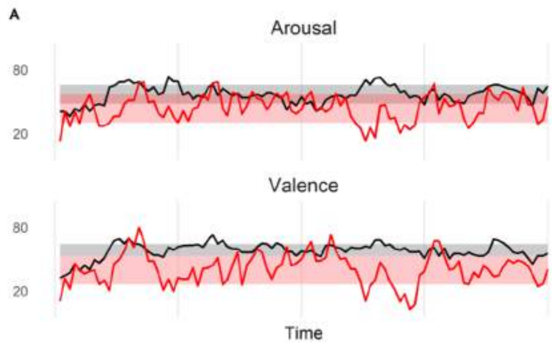

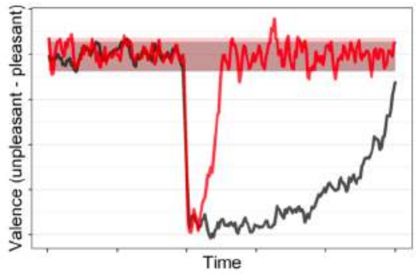

Another window into emotions is through our words:

- E.g., if happier, we are likely to utter more happiness-associated words

Utterance Emotion Dynamics: study of change in emotion words over time
(Hipson and Mohammad, 2021)

Utterance Emotion Dynamics: Metrics

- Emotion word density
  - proportion of emotion words
- Home base
  - steady state locations in affect space
- Emotional variability
  - degree to which emotional state changes with time
- Displacement (Count and Lengths)
  - how often one leaves home base
  - how far they go
  - average peak distance
- Rise and Recovery Rates
  - how quickly one leaves/returns to home base

Example home bases of two people in the v-a space.

person 1: ellipse in pink
person 2: ellipse in grey

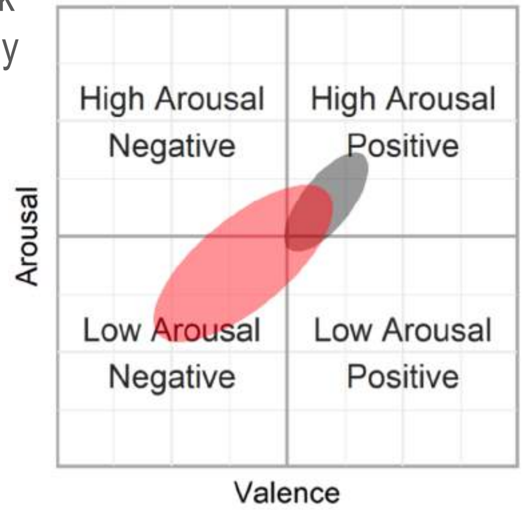

Character Dialogue from Literature and Film

- Source of abundant longitudinal text
- Drives the plot
- Direct way to understand what a character is feeling

Data we used:

- Scripts from the Internet Movie Script Database (IMSDb)
- 1,123 movie scripts with ~54,000 characters
- Dialogues grouped into turns
  - sequence of uninterrupted utterances by a character
  - ~2,600 characters (~5%) had at least 50 turns: main characters

Analyzing Characters

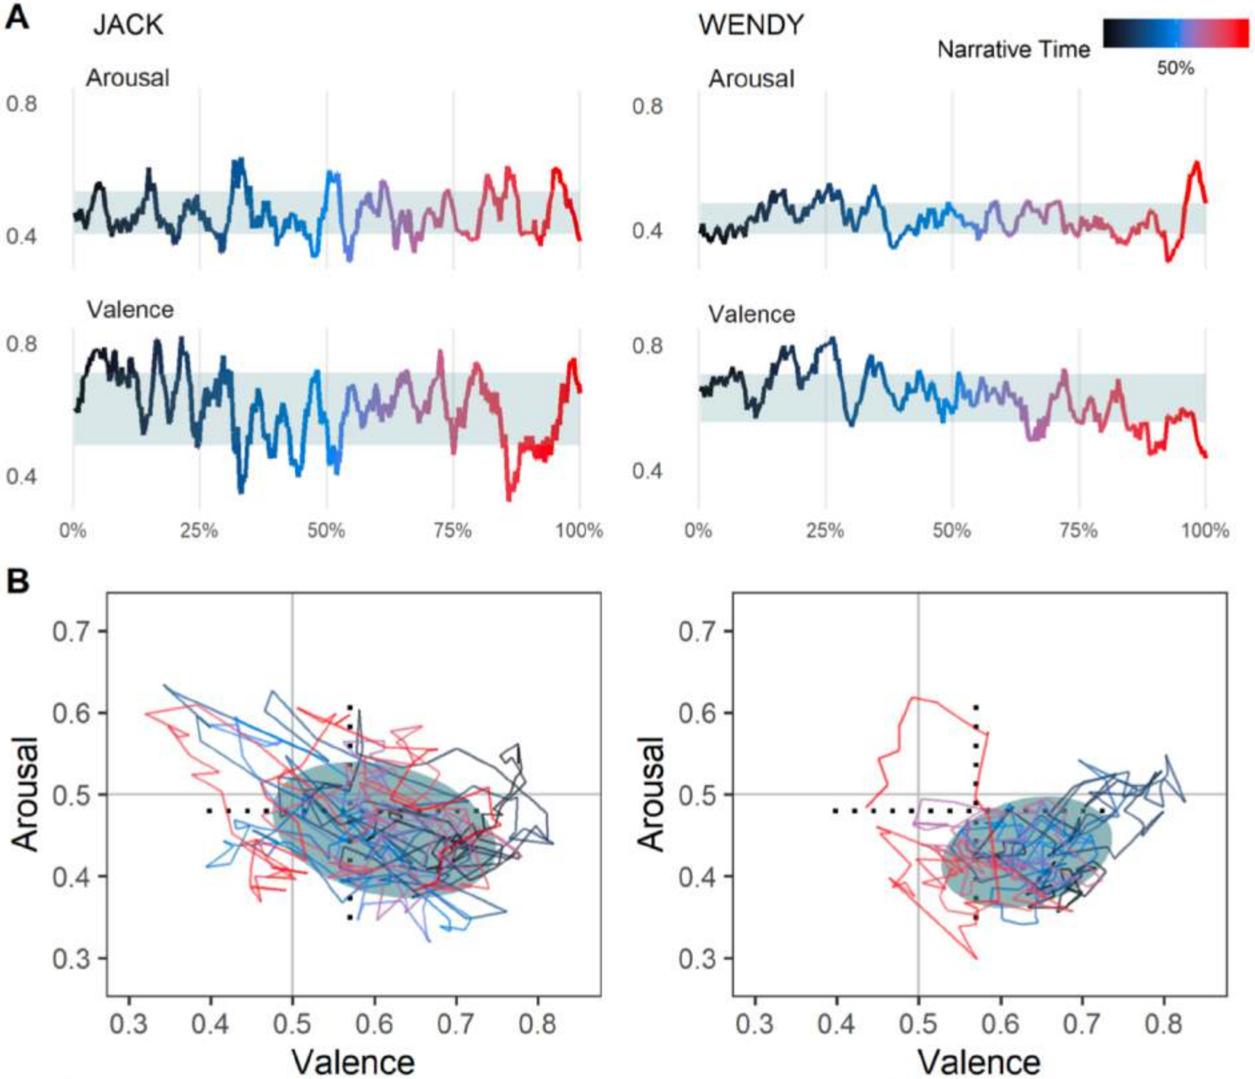

Plots: Emotion arcs/trajectories of Jack and Wendy, from The Shining

Affect Dimensions: Valence, Arousal, Valence – Arousal

Fig 3. One dimensional and two dimensional state spaces for Jack (n = 389 words) and Wendy (n = 279 words), two main characters from The Shining (1980). Color of line corresponds to narrative time, with dark blue meaning earlier in the movie and red meaning later. The black dotted lines show the major and minor axes of an ellipse within which all main characters are 95% of the time (the ellipse itself is not shown to avoid clutter).

Benchmarked Emotion Word Usage

For each UED metric, calculated average and standard deviation of scores for all characters

- Set baselines
- Individual characters can then be examined in terms of where they lie in the distribution

Used UED metrics to Analyzing Characters

Across narrative time and emotion space

Plot: Average emotion word densities across narrative time

Affect Dimensions: Positive, Negative

Analyzing Characters Across Narrative Time

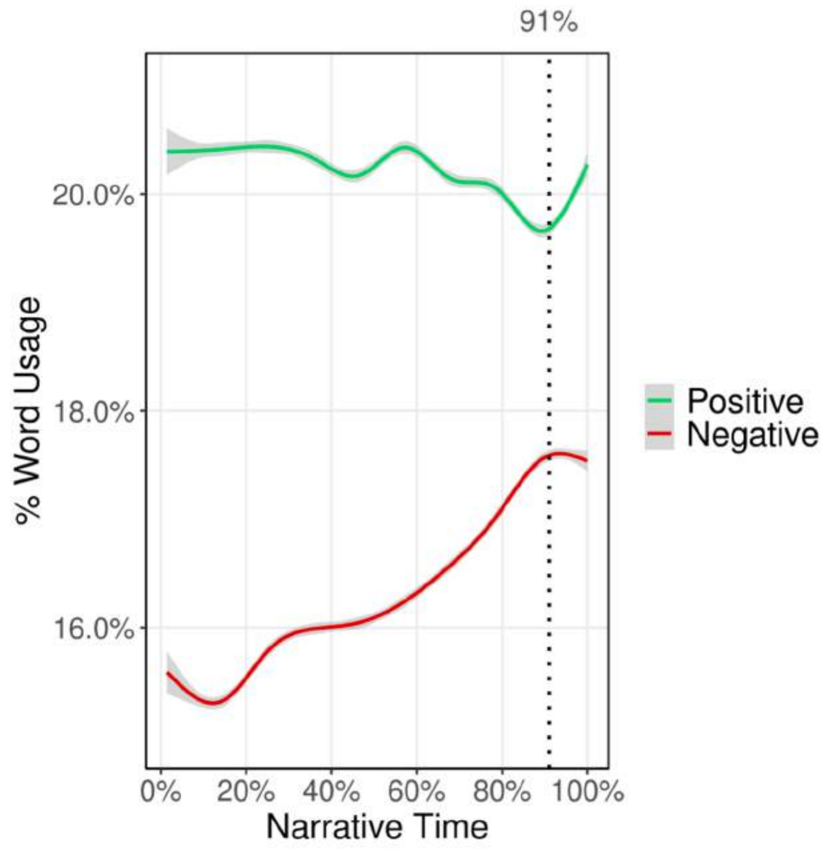

Fig 6. Average trends in proportion of positive and negative word usage over time. Vertical dotted line shows location of peak negative density and lowest positive density. Grey band is the 95% confidence interval around the estimated mean. n = 965,147 words.

Plot: Topological map showing where peak displacement tends to occurs

Affect Dimensions: Valence – Arousal

Analyzing Characters Across Emotion Space

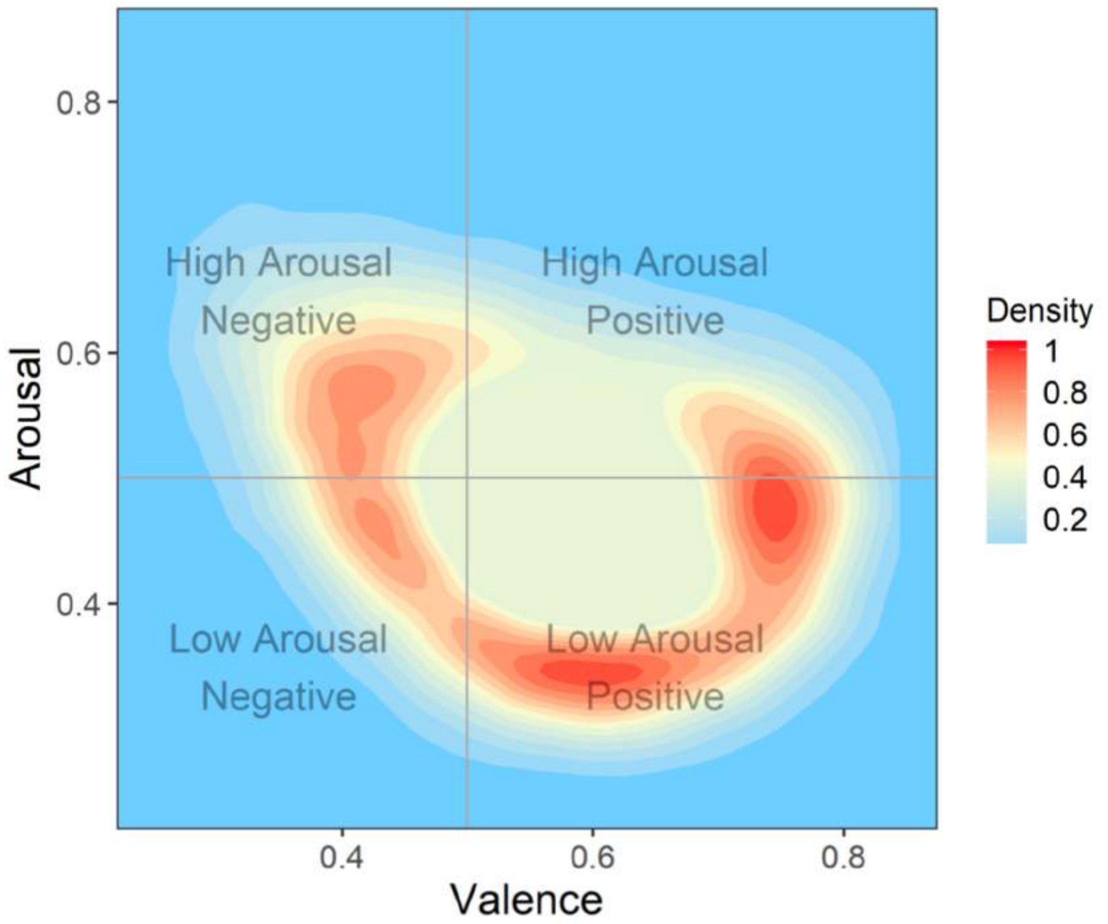

Fig 5. Density map showing where peak displacements tend to occur. Red corresponds to more peaks. Density is normalized to go from 0–1.

Tweet Emotion Dynamics

Emotion Word Usage in Tweets from US and Canada

Krishnapriya (KP) Vishnubhotla

vkpriya@cs.toronto.edu

Computer Science
UNIVERSITY OF TORONTO

VECTOR INSTITUTE

Saif M. Mohammad

saif.mohammad@nrc-cnrc.gc.ca

National Research Council Canada
Conseil national de recherches Canada

Resources of Interest

Available at: www.saifmohammad.com

- Emotion lexicons
- An overview of AER and sentiment analysis →
- Ethics sheet for AER
- Best practices in the creation and use of emotion lexicons (under review)

Saif M. Mohammad

Saif.Mohammad@nrc-cnrc.gc.ca

@SaifMMohammad

WP WOODHEAD PUBLISHING

Emotion Measurement (Second Edition)

2021, Pages 323-379

Chapter 11 - Sentiment analysis: Automatically detecting valence, emotions, and other affectual states from text

Saif M. Mohammad

Show more

Share Cite

https://doi.org/10.1016/B978-0-12-821124-3.00011-9

Get rights and content

Abstract

Recent advances in machine learning have led to computer systems that are humanlike in behavior. Sentiment analysis, the automatic determination of emotions in text, is allowing us to capitalize on substantial previously unattainable opportunities in commerce, public health, government policy, social sciences, and art. Further, analysis of emotions in text, from news to social media posts, is improving our understanding of not just how people convey emotions through language but also how emotions shape our behavior. This article presents a sweeping overview of sentiment analysis research that includes: the origins of the field, the rich landscape of tasks, challenges, a survey of the methods and resources used, and applications. We also discuss how, without careful fore-thought, sentiment analysis has the potential for harmful outcomes. We outline the latest lines of research in pursuit of fairness in sentiment analysis.