



NLP for Affective Science:

What are the big questions? And how do we get there?

Saif M. Mohammad

Principal Research Scientist, National Research Council Canada

✉ Saif.Mohammad@nrc-cnrc.gc.ca [@SaifMMohammad](https://twitter.com/SaifMMohammad)



NLP for Affective Science:

A window into emotions through language and computation

Saif M. Mohammad

Principal Research Scientist, National Research Council Canada

✉ Saif.Mohammad@nrc-cnrc.gc.ca [@SaifMMohammad](https://twitter.com/SaifMMohammad)




NLP for Affective Science:

A window into emotions (mind, body, health, and behavior) through language and computation

Saif M. Mohammad

Principal Research Scientist, National Research Council Canada

✉ Saif.Mohammad@nrc-cnrc.gc.ca [@SaifMMohammad](https://twitter.com/SaifMMohammad)



What are Affect and Emotions?

to start with...

Affect: the basic sense of feeling

Emotions: joy, sadness, fear, anger, etc.

Psychological Theories of Emotions

ON THE ORIGIN OF SPECIES

BY MEANS OF NATURAL SELECTION,

OR THE

PRESERVATION OF FAVOURED RACES IN THE STRUGGLE

FOR LIFE




By CHARLES DARWIN, M.A.


Charles Darwin



I think
the letters A & B mean
the 2nd & 3rd
front phalanx, B & D
with the phalanx distal
the same could be
found. - heavy white



Gibbon Orangutan Chimpanzee Gorilla Man



200th ANNIVERSARY EDITION
Charles Darwin 1809-1882

CHARLES DARWIN


THE EXPRESSION OF THE EMOTIONS IN MAN AND ANIMALS

Introduction, Afterword and Commentaries by PAUL ERMAN
New Preface for the Anniversary Edition



'As fresh and provocative today as it was 150 years ago.'

—Michael Young




Theories of Emotion



Margaret Mead
Cultural anthropologist



Paul Ekman, Psychologist



- Mead, 1950s: culture determines emotion
- Paul Ekman, 1971: **Six** Universal Basic Emotions
 - Plutchik, 1980: **Eight** Basic Emotions
 - And many others



Plutchik's Emotion Wheel
Image credit: Julia Belyanevych

Core Dimensions of Connotative Meaning

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*



Osgood



Russell



Theories of Emotion



Margaret Mead
Cultural anthropologist



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

Affect

The basic sense of feeling:

- Key dimensions
 - Valence: displeasure to pleasure
 - Arousal: idle/sluggish to activated
 - Dominance: weak/loss of control to strong/having a sense of control
- Transduced and summarized from **interoceptive signals**
 - sensory representations of the interior of the body (viscera)
- A feature of consciousness
 - occurs in every moment (whether you're aware of it or not)

Emotions

Constructed by the brain using:

- affective and interoceptive signals
- "emotion concepts" from one's culture
- predictive coding

The brain categorizes the continuous affect into discrete categories (analogous to color perception)


- joy, sadness, fear, anger, etc.

Affective Science

Interdisciplinary field focused on understanding emotions and affect

- How do affect and emotions work?
 - affective and emotional processes
 - affective neuroscience
 - emotion regulation
- How do they impact our mental health, physical health, and behaviour?
 - wellbeing, emotion-related disorders
- What agency do we have in managing our emotions?

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

NLP for Affective Science

Computational Analysis of Emotions Through Language

- Challenging
 - see previous slide on language
- Powerful
 - makes use of large amounts of text
 - simple to complex NLP techniques
 - language impacts thought and how we construct emotions

linguistic relativity aka Sapir–Whorf hypothesis
- Complementary view to traditional affective science approaches in psychology
 - makes use of, complementary, ecologically valid data

Computational Analysis of Emotions Through Language

Affective Science Questions

- How do emotions work?
- What impacts our emotions?
- How do emotions relate to our health?
- How do we regulate emotions?

Linguistics

- How do we use language to convey emotions?
- How does language impact emotions?

Social Science

- How do emotions impact social cognition, morality, stereotypes, and behavior?

AI

- What tools can we build to help people, clinicians, social scientists?

This Talk

- Core Theories of Affect and Emotion
 - What is NLP for Affective Science?
- Words-Emotion Associations
 - Anxiety (EMNLP 2024)
- Tracking Emotion arcs
 - Evaluation (EMNLP 2023)
 - Emotion Dynamics (multiple, 2021-24)
 - Emotion Granularity (EMNLP, 2024)
- Ethical considerations in emotion recognition (ACL, CL Journal, 2022)

This Talk

- Core Theories of Affect and Emotion
 - What is NLP for Affective Science?
- Words-Emotion Associations
 - Anxiety (EMNLP 2024)
- Tracking Emotion arcs
 - Evaluation (EMNLP 2023)
 - Emotion Dynamics (multiple, 2021-24)
 - Emotion Granularity (EMNLP, 2024)
- Ethical considerations in emotion recognition (ACL, CL Journal, 2022)





word --



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

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

(anger, fear, joy, sadness,
anticipation, disgust, surprise, trust)

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>

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



“I have some affection for my anxiety. When it does not have me in a swirl, I think of it as a little pet, a black cat who has gotten stuck in my house and who I keep feeding milk, even though I am not sure that I wanted a pet. The cat’s language is gentle, pawing, and brings me back to the things that matter: Hospitality, service, compassion, kindness. The cat is my teacher.” - Laura Turner

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)

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
 - 10K English bigrams
- Quality
 - interspersed gold (control) questions
 - show that the anxiety associations are highly reliable
 - split-half reliability of 0.82


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.

WorryWords

study the rate at which children acquire anxiety words with age



Used WorryWords to

Track the change of anxiety in streams of text




This Talk

- Core Theories of Affect and Emotion
 - What is NLP for Affective Science?
- Words-Emotion Associations
 - Anxiety (EMNLP 2024)
- Tracking Emotion arcs
 - Evaluation (EMNLP 2023)
 - Emotion Dynamics (multiple, 2021-24)
 - Emotion Granularity (EMNLP, 2024)
- Ethical considerations in emotion recognition (ACL, CL Journal, 2022)



creativity



Emotions Arcs in Commerce, Psychology, and Stories



National Research
Council Canada

Conseil national de
recherches Canada

 @SaifMMohammad


Canada

27

Tracking Emotions in Stories

SIMPLE SHAPES OF STORIES

As told by Kurt Vonnegut.






SOURCE DAVID YANG, VISUAL.LY

HBR.ORG

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

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



Daniela Teoderescu



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

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)



PLOS One



Will Hipson

2021: Emotion Dynamics of Fictional Characters




LREC




Krishnapriya (KP) Vishnubhotla

2022: Tweet Emotion Dynamics Emotion Word Usage in Tweets from US and Canada



EMNLP



Daniela Teodorescu



Tiffany Cheng



Alona Fyshe

2023: Language and Mental Health: Measures of Emotion Dynamics from Text as Linguistic Biosocial Markers.





We saw...

Emotion Dynamics: Individual Emotion Arcs

next...

Emotion Granularity: Pairs of Emotion Arcs



Emotion Granularity aka Emotion Differentiation (from Psychology)

Some people:

- recognize, identify, describe their feelings using **precise** terms
 - like guilt, anger, frustration, or helplessness
- can **reliably** describe these concepts using language
 - distinguishing between angry and sad, elated and content, etc.



Lisa Barrett

Others:

- tend to use more broad terms to convey emotions
 - a general sense of feeling bad or feeling low
- co-endorsing multiple emotions

Emotion granularity (**Barrett et al., 2001**)

- this ability to experience and categorize emotions in very specific terms
- the degree of not co-endorsing multiple emotions

Link between Emotion Granularity and Health


- Mental health (Erbas et al., 2014, 2018)
 - depression (Starr et al., 2017)
 - anxiety (Seah et al., 2020)
 - borderline personality disorder (Dixon-Gordon et al., 2014, Suvak et al., 2011)
 - show less neural reactivity to rejection (Kashdan et al., 2014)
- Physical health (Hoemann et al., 2021)
 - cardiovascular physiological activity and stress (Bonar et al., 2023)
- Behavior
 - maladaptive behaviours such as binge drinking, aggression, and self-injurious behavior (Dixon-Gordon et al., 2014, Kashdan et al., 2015)
 - school behaviour (Brackett, Rivers, Reyes, & Salovey, 2012)
 - eating disorders (Selby et al., 2013)
 - less likely to retaliate aggressively (i.e., verbally or physically assault) against someone who has hurt them (Pond et al., 2012)

Emotion Granularity from Text (our work)

- To what extent are we co-endorsing multiple emotions ****in text****?
 - through connotations and not necessarily denotations
- Compute emotion arcs for various emotions
- Compute **emotion granularity (EG)**: correlation between pairs of emotion arcs
- Show that text by those who have self-disclosed to have certain mental health conditions (depression, PTSD, ADHD, etc.) have significantly lower EG than text by control group

EMNLP 2024:


Emotion Granularity from Text: An Aggregate-Level Indicator of Mental Health



Krishnapriya (KP)
Vishnubhotla



Daniela Teodorescu



Mallory Fedman



Kristen Lindquist

This Talk

- Core Theories of Affect and Emotion
 - What is NLP for Affective Science?
- Words-Emotion Associations
 - Anxiety (EMNLP 2024)
- Tracking Emotion arcs
 - Evaluation (EMNLP 2023)
 - Emotion Dynamics (multiple, 2021-24)
 - Emotion Granularity (EMNLP, 2024)
- Ethical considerations in emotion recognition (ACL, CL Journal, 2022)

Emotion Recognition: Task

1. Inferring emotions felt by the speaker

Emotion Recognition: Task

1. Inferring emotions felt by the speaker
2. Inferring emotions of the speaker as perceived by the reader/listener
3. Inferring emotions that the speaker is attempting to convey
4. Inferring emotions evoked in the reader/listener
5. Inferring emotions of people mentioned in the text
6. Inferring whether what is described is good for pre-determined target of interest

Emotion Recognition: Task

1. Inferring emotions felt by the speaker
2. Inferring emotions of the speaker as perceived by the reader/listener
3. Inferring emotions that the speaker is attempting to convey
4. Inferring emotions evoked in the reader/listener
5. Inferring emotions of people mentioned in the text
6. Inferring whether what is described is good for pre-determined target of interest
7. Inferring the intensity of the emotions discussed above
8. Inferring patterns of speaker's emotions over long periods of time, across many utterances; including the inference of moods, emotion dynamics, and emotional arcs
9. Inferring speaker's emotions/attitudes/sentiment towards a target product, movie, person, idea, policy, entity, etc.
10. Inferring emotionality of language used in text (regardless of whose emotions)
11. Inferring how language is used to convey emotions such as joy, sadness, loneliness, hate, etc.
12. ...

Emotion Recognition: Task

1. Inferring emotions felt by the speaker
2. Inferring emotions of the speaker as perceived by the reader/listener
3. Inferring emotions that the speaker is attempting to convey
4. Inferring emotions evoked in the reader/listener
5. Inferring emotions of people mentioned in the text
6. Inferring whether what is described is good for pre-determined target of interest
7. Inferring the intensity of the emotions discussed above
8. Inferring patterns of speaker's emotions over long periods of time, across many utterances; including the inference of moods, emotion dynamics, and emotional arcs
9. Inferring speaker's emotions/attitudes/sentiment towards a target product, movie, person, idea, policy, entity, etc.
10. Inferring emotionality of language used in text (regardless of whose emotions)
11. Inferring how language is used to convey emotions such as joy, sadness, loneliness, hate, etc.
12. ...

All of these come with...

Benefits, Potential Harms, Ethical Considerations

Theories of Emotion



Margaret Mead
Cultural anthropologist



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

Computational Analysis of Emotions and Automatic Emotion Recognition (AER)

A force that helps unlock:

- how emotions work
- how they relate to our health, language, behavior, social interactions,...
- numerous commercial applications that benefit society

A tool for substantial harm, e.g.:

- mass application on vulnerable populations
- unreliable approaches
- privacy concerns
- perpetuation of physiognomy



[Strategies](#) [Topics](#) [Regions](#) [Up Close](#) [Tools](#) [Multimedia](#)


[Partnerships](#)

How emotion recognition software strengthens dictatorships and threatens democracies

Given that the idea of using emotion recognition technology as a tool of governance is an entirely flawed premise, a ban makes the most sense.

By: James Jennion

[Español](#)



Ethics Sheet for Automatic Emotion Recognition and Sentiment Analysis



Medium Blog Post in summer of 2021:
<https://medium.com/@nlpscholar/ethics-sheet-aer-b8d671286682>

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).

Template



50 considerations grouped under:

- *Task Design*
 - *Data*
 - *Method*
 - *Impact and Evaluation*
 - *Implications for Privacy and Social Groups*
- } common phases in system development

TASK DESIGN

A. Theoretical Foundations

1. Task Design and Framing
2. Theoretical Models and their Implications
3. Meaning and Extra-Linguistic Information
4. Wellness and Health Implications
5. Aggregate Level vs. Individual Level Prediction

B. Implications of Automation


6. Why Automate
7. Embracing Diversity
8. Participatory/Emancipatory Design
9. Applications, Dual Use, Misuse
10. Disclosure of Automation

DATA

C. Why This Data

11. Types of data
12. Dimensions of data
- D. Human Variability–Machine Normativeness
13. Variability of Expression, Conceptualization
14. Norms of Emotions Expression
15. Norms of Attitudes

... 50!



First SemEval Shared Task on African Languages:

SemEval 2023: AfriSenti: Detecting Sentiment in African Languages

Labeled sentiment datasets for 14 languages from 3 language families

Led by African researchers

Shared Task on Emotions

SemEval-2025 Task 11

Bridging the Gap in Text-Based Emotion Detection



Labeled emotion datasets for 35 languages: most from Africa and Asia

NLP for Affective Science



Emotions, mind, body, health, behavior, language, and computation

Slides, Papers, Datasets, Lexicons, Code
Available at: www.saifmohammad.com

✉ saif.mohammad@nrc-cnrc.gc.ca

🐦 [@SaifMMohammad](https://twitter.com/SaifMMohammad)