Emotion Analysis from Text:
Tutorial at EACL 2023

Sanja Štajner and Roman Klinger
stajner.sanja@gmail.com,roman.klinger@ims.uni-stuttgart.de

April 5, 2023

Contents

Universität Stuttgart
Institut für
Maschinelle Sprachverarbeitung

Emotion Analysis

Introduction and Psychology

EACL 2023 Tutorial

Sanja Štajner and Roman Klinger

Outline

1. Introduction

2. What are Emotions?
   - Motivation: Basic Emotion Theories
   - Feeling: Affect and Constructionism
   - Evaluation: Causes and Appraisals

3. Task Definition and Issues

4. What can we learn from previous work in psychology?
   - Psychological Studies on Reliability

Outline

Which emotion does the person who says this experience?

“I am happy to be here!”

“Tears ran down my face.”

“I heard a loud sound when I was alone in the forest.”

About Us

- Sanja Stajner
- Independent Researcher based in Karlsruhe, Germany
- Research on emotion analysis, personality modeling, text simplification, accessibility, readability

- Roman Klinger
- Professor at the Institute for Natural Language Processing University of Stuttgart, Germany
- Resarch on sentiment analysis, emotion analysis, social media mining, biomedical NLP, fact-checking

About this tutorial

Session 1 (09:00–10:30)

- Introduction
- Psychological Models
- Use Cases/Social Impact
- Resources
- Annotation Exercise

Break (10:30–11:15)

Session 2 (11:15–12:45)

- Non-Neural Methods
- Multi-task, transfer, zero-shot methods
- Open Challenges
- Appraisal Theories
- Role Labeling
- Ethical Considerations
- Closing

Purpose of this Tutorial

Target Audience

- Computationally oriented researchers
- Scholars interested in digital humanities, computational social sciences

Goal

- Provide psychological background knowledge
- Provide overview of existing resources, tasks, challenges, models
- Draft potential future research directions

Outline

Literature on Emotion Psychology

- Philosophy, history and sociology
- Literature and art
- Decision making, Computational models
- Biological perspectives
- Social and personality perspectives
- Cognitive Perspectives
- Health
- Specific Emotions

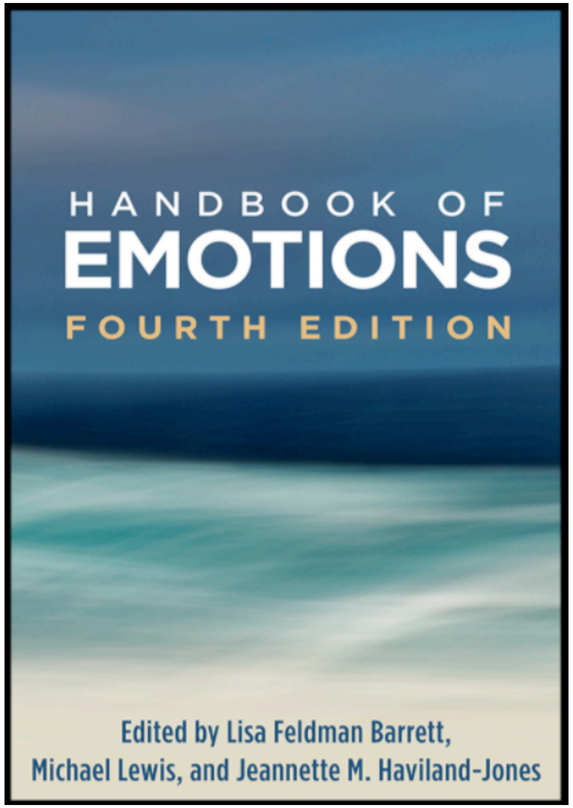

Literature with a Computational Focus

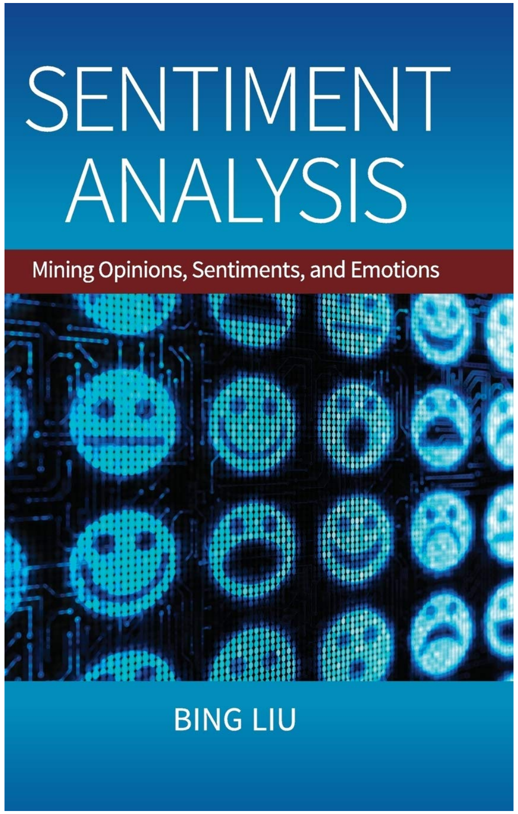

Emotion Theories...

...try to explain ...

- what emotions are
- what they consist of
- what their purpose is
- how they develop
- ...

Evolutionary Approach (Darwin, 1872)

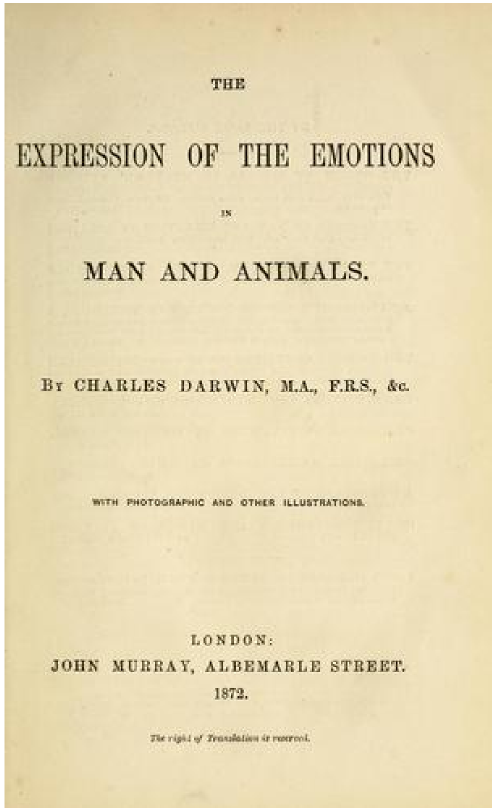

- Focuses on expressions, as they can be observed.
- Emotion expressions support communication
- Emotions and their expressions have a function:
  - Surprise: Eyes wide open to support perception
  - Fear: Activation (fight, freeze, flight)
  - Disgust: Increase distance to stimulus
  - ...
- Emotions are not learned

Evolutionary Approach

https://en.wikipedia.org/wiki/The_Expression_of_the_Emotions_in_Man_and_Animals

James Lange/Cannon Bard (1884, 1925)

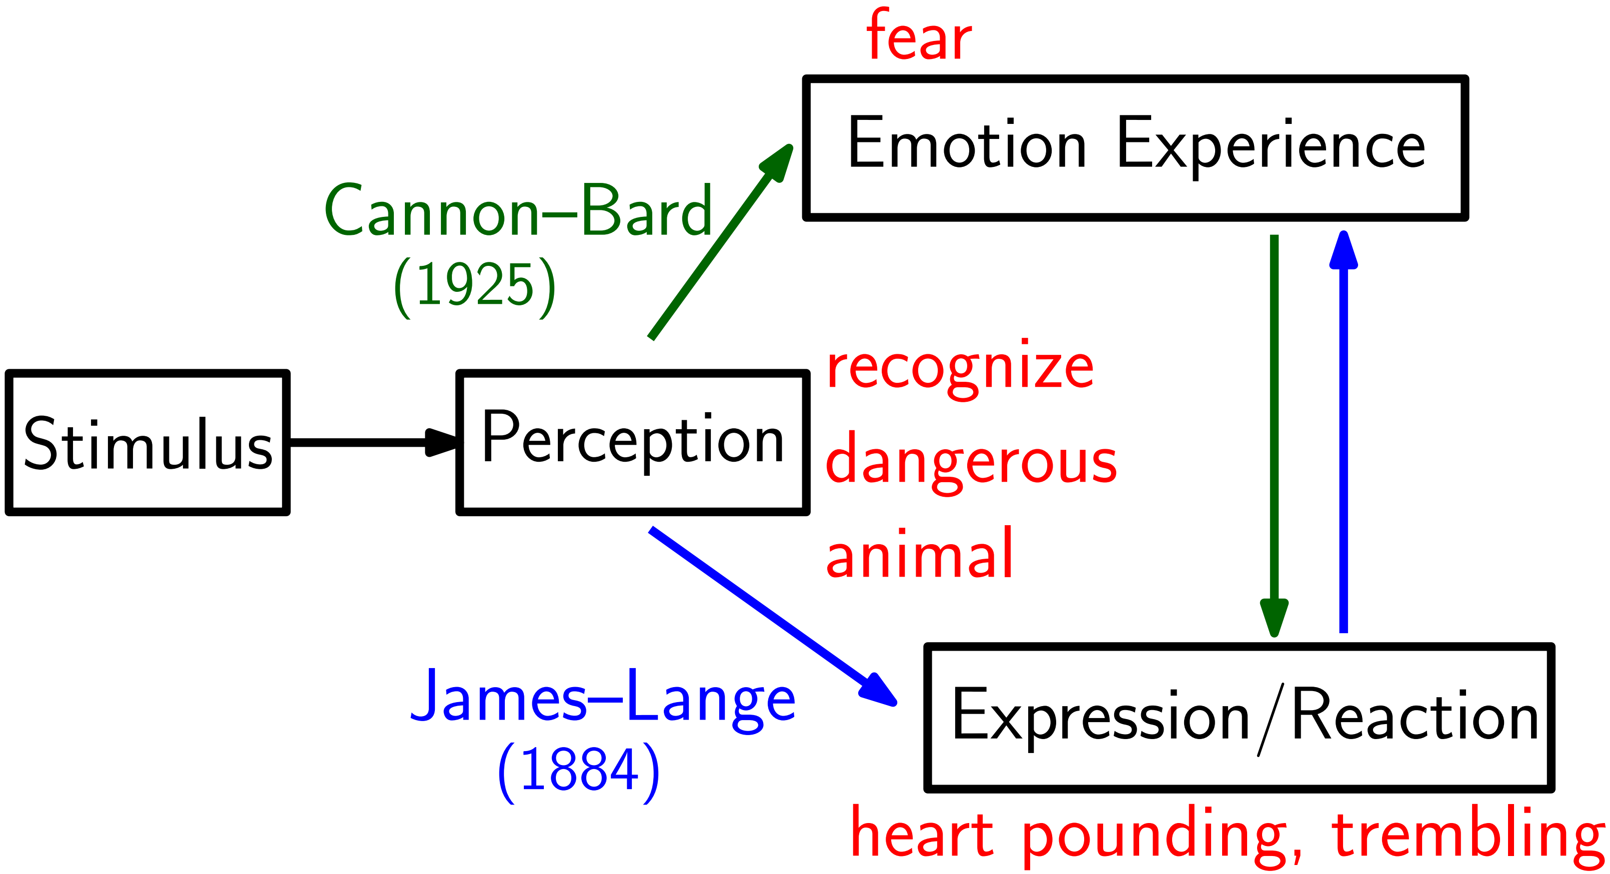

Emotion Components (Scherer, 2001)

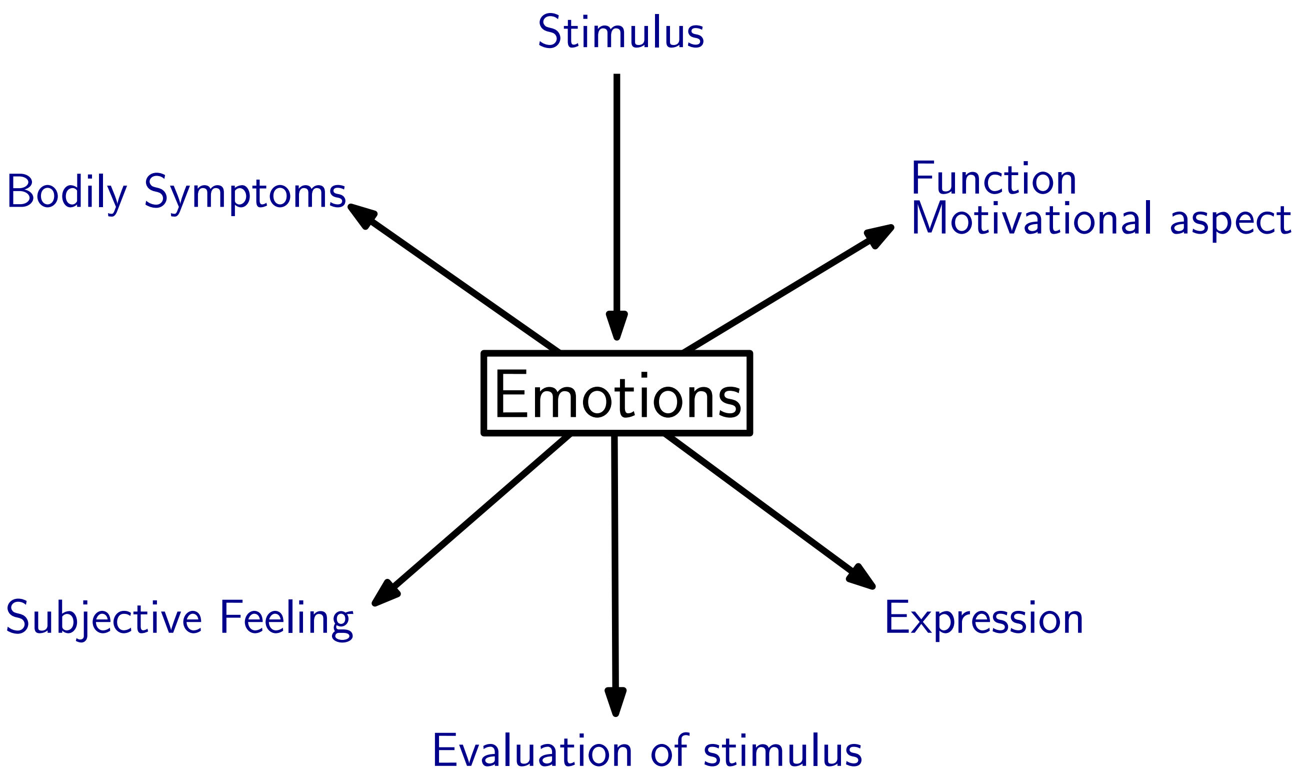

Emotion in text could be expressed by describing each of these components.

Family Tree of Emotions (Scarantino, 2016)

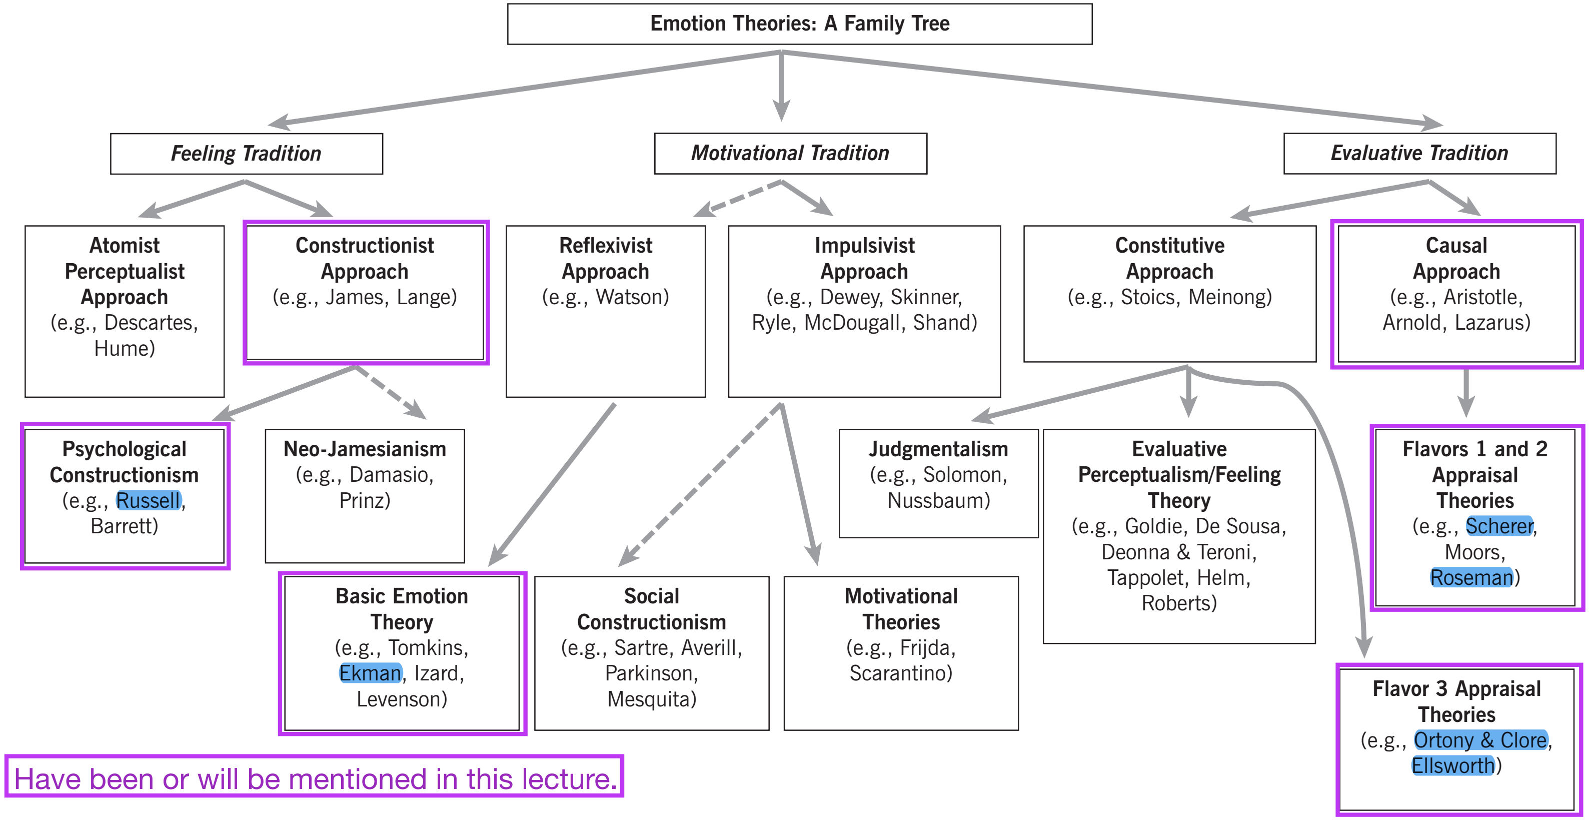

Have been or will be mentioned in this lecture.

I am aware of work in NLP that made use of these theories.

Outline

1. Introduction
2. What are Emotions?
   - Motivation: Basic Emotion Theories
   - Feeling: Affect and Constructionism
   - Evaluation: Causes and Appraisals
3. Task Definition and Issues
4. What can we learn from previous work in psychology?
   Psychological Studies on Reliability

Basic Emotion Theories

Basic emotion theories state that:

- There is a distinction between basic and non-basic emotions
- There are criteria that decide if an emotion is basic.

Ekman's model of basic emotions

How to define a categorical system of emotions?

- Distinctive universal signals
- Presence in other primates
- Distinctive physiology
- Distinctive universals in antecedent events
- Coherence among emotional response
- Quick onset
- Brief duration
- Automatic appraisal
- Unbidden occurrence

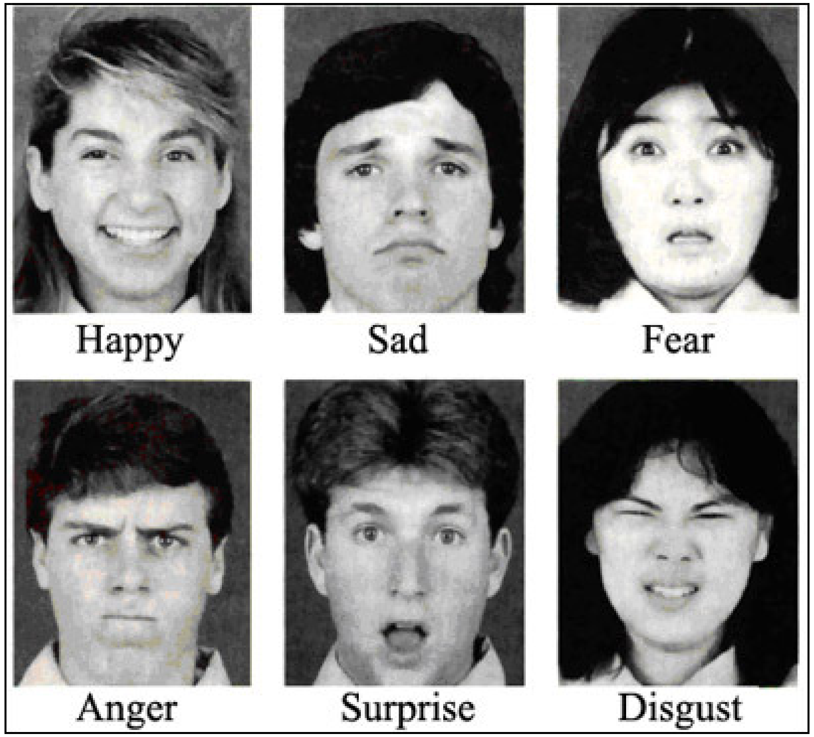

Ekman (1992): An argument for basic emotions.

Ekman: What are non-basic emotions?

- “I do not allow for non-basic emotions” (Ekman, 1999)
  - ⇒ They do no exist.
- What is love, depression, or hostility?
  - Personality traits (hostility, openness)
  - Moods (depression, anxiety, long-term disturbances are clinically relevant)
  - Emotional plots (love, grief, jealousy)

Models of Basic Emotions: Plutchik's Wheel (Plutchik, 1970)

An emotion is a patterned bodily reaction that follows a function

- protection – fear
- destruction – anger
- reproduction – joy
- deprivation – sadness
- incorporation – acceptance
- rejection – disgust
- exploration – anticipation
- orientation – surprise

⇒ Basic emotions according to Plutchik

- Non-basic: Gradations and mixtures

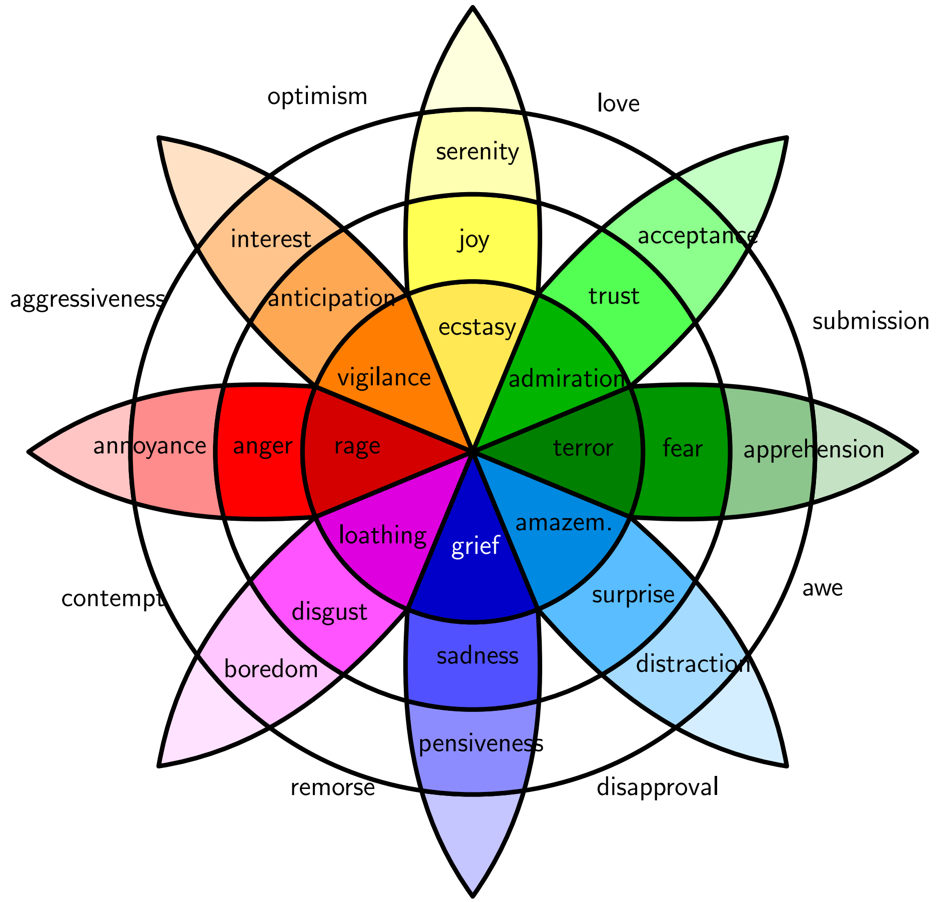

Outline

1. Introduction
2. What are Emotions?
   - Motivation: Basic Emotion Theories
   - Feeling: Affect and Constructionism
   - Evaluation: Causes and Appraisals
3. Task Definition and Issues
4. What can we learn from previous work in psychology?
   - Psychological Studies on Reliability

The Feeling Tradition of Emotion Theories

- Emotions are not innate
- They are learned constructs
- Depend on culture and contingent situations

Feeling

What is not learned then?

Feeling

- Scarantino (2016): "Feeling is a conscious experience or a sensation or a subjective quality or a quale or a what-it-is-likeness."
- Feldman-Barrett (2018): Affect is "the general sense of feeling that you experience throughout each day [...] with two features. The first is how pleasant or unpleasant you feel, which scientists call valence. [...] The second feature of affect is how calm or agitated you feel, which is called arousal."

Affect: Continuous Circumplex Model (Russel 1980)

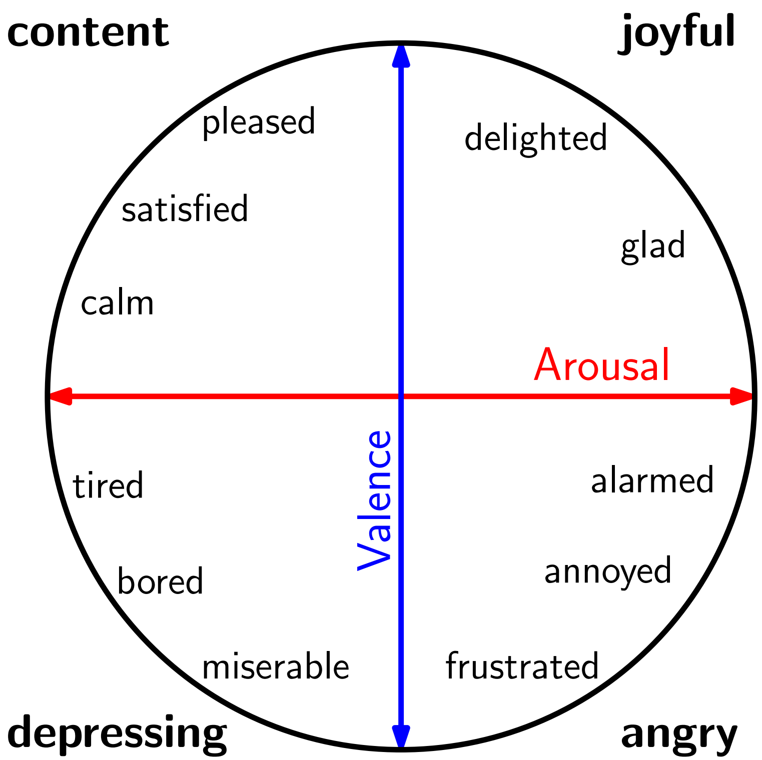

- So-called dimensional model
- Discrete emotion names are placed in a coordinate system
- Other dimensional models:
  - Valence–Arousal–Dominance (not discussed here)
  - Appraisals (later)

Feldman-Barrett (2017): Theory of Constructed Emotion

How to link affect and emotion names? Lisa Feldman-Barrett attempts to explain this link.

- Paradoxon: We experience discrete emotion categories, but there is nearly no evidence from neuroscience for those.
- Affect (valence and arousal) is what we experience directly, not the emotion.
- Based on context, the brain predicts which emotion makes sense.
- Prediction is important, to motivate or warn us.
- This learned construction of emotions bridges the paradoxon.

- Very nice overview video: `https://www.youtube.com/watch?v=M1OdhdI_3eI`

Outline

Appraisal Theories (according to Scherer)

Scherer, 2005

Emotions are “an episode of interrelated, synchronized changes … in response to the evaluation of an external or internal stimulus event as relevant to major concerns of the organism”

Appraisal Theories (according to Scherer)

Emotions have different components...

- Cognitive appraisal:
  an evaluation of events and objects
- Bodily symptoms:
  physiological component of emotional experience
- Action tendencies:
  a motivational component for the preparation and direction of motor responses
- Expression: facial and vocal expression, body language, gestures, almost always accompanies an emotional state
- Subjective perceptions/Feeling:
  subjective experience of emotional state once it has occurred

Sequence of appraisal criteria (Scherer 2005/2013)

Scherer: Emotions are evaluated in a sequential manner.

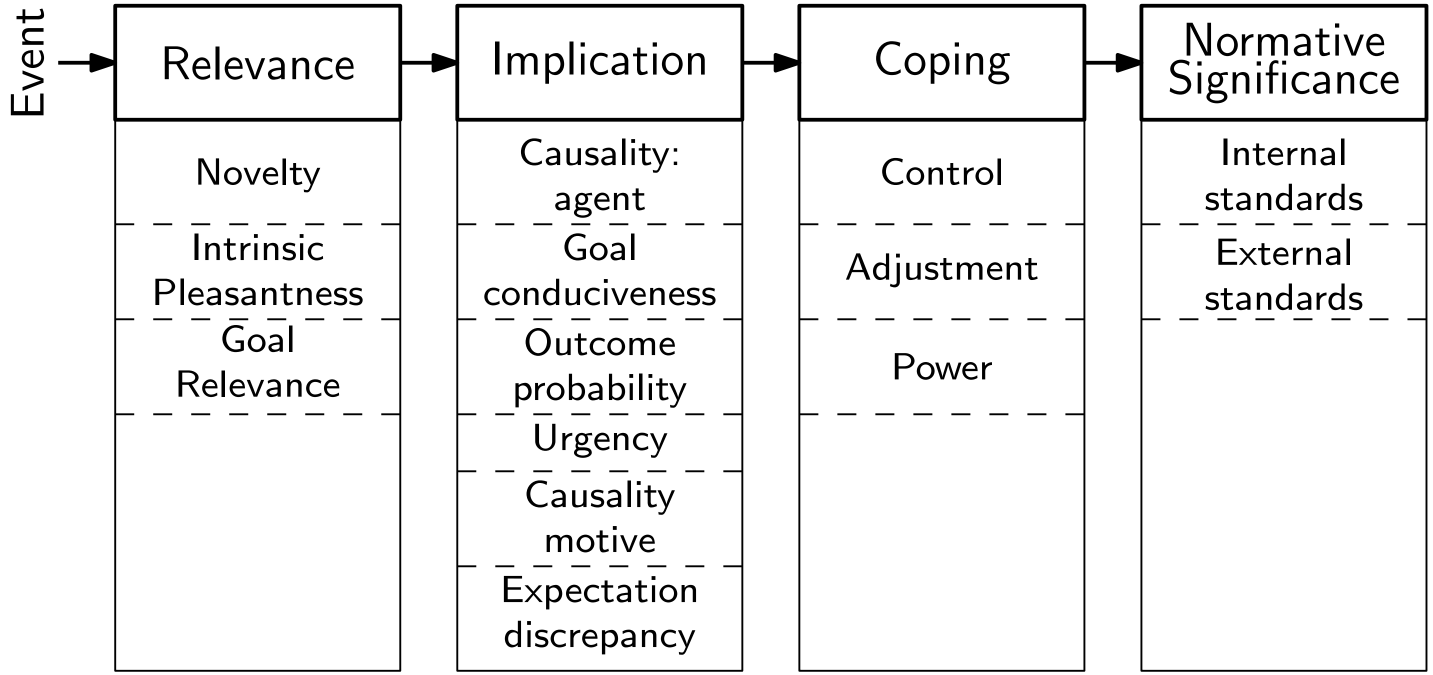

Results Smith/Ellsworth (1985)

Locations of Emotion Means Along the PCA Components

Emotion	Component: Pleasant[a]	Responsibility/ Control[b]	Certain[c]	Attention[d]	Effort[e]	Situational- Control[f]
Happiness	−1.46	0.09	−0.46	0.15	−0.33	−0.21
Sadness	0.87	−0.36	0.00	−0.21	−0.14	1.15
Anger	0.85	−0.94	−0.29	0.12	0.53	−0.96
Boredom	0.34	−0.19	−0.35	−1.27	−1.19	0.12
Challenge	−0.37	0.44	−0.01	0.52	1.19	−0.20
Hope	−0.50	0.15	0.46	0.31	−0.18	0.35
Fear	0.44	−0.17	0.73	0.03	0.63	0.59
Interest	−1.05	−0.13	−0.07	0.70	−0.07	0.41
Contempt	0.89	−0.50	−0.12	0.~~80~~ .08	−0.07	−0.63
Disgust	0.38	−0.50	−0.39	−0.96	0.06	−0.19
Frustration	0.88	−0.37	−0.08	0.60	0.48	0.22
Surprise	−1.35	−0.94	0.73	0.40	−0.66	0.15
Pride	−1.25	0.81	−0.32	0.02	−0.31	−0.46
Shame	0.73	1.31	0.21	−0.11	0.07	−0.07
Guilt	0.60	1.31	−0.15	−0.36	0.00	−0.29

Note. Scores are standardized.
[a] Pleasantness: high scores indicate increased unpleasantness.
[b] Responsibility/Control: high scores indicate increased self-responsibility/control.
[c] Certainty: high scores indicate increased uncertainty.
[d] Attentional activity: high scores indicate increased attentional activity.
[e] Effort: high scores indicate increased anticipated effort.
[f] Situational control: high scores indicate increased situational control.

OCC Model of Emotions

Ortony, Clore, Collings (1988):
The Cognitive Structure of Emotions.

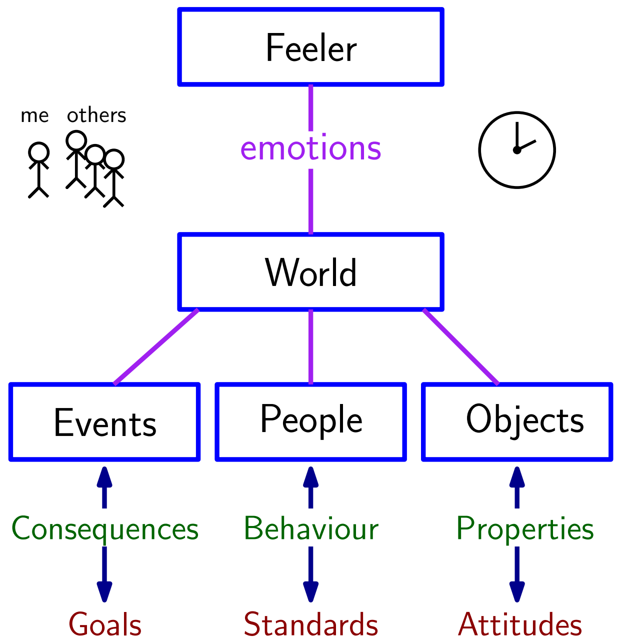

Outline

Example 1

I am happy to be here!

Circumplex model (Russell):

Valence?	☐ high	☐ low
Arousal?	☐ high	☐ low

Appraisals (Smith/Ellsworth):

Pleasantness?	☐ high	☐ low
Responsible?	☐ high	☐ low
Certain?	☐ high	☐ low
Attention?	☐ high	☐ low
Effort?	☐ high	☐ low
Control?	☐ high	☐ low

Emotion Wheel (Plutchik):

- ☐ Protection/Fear
- ☐ Destruction/Anger
- ☐ Reproduction/Joy
- ☐ Deprivation/Sadness
- ☐ Incorporation/Acceptance
- ☐ Rejection/Disgust
- ☐ Exploration/Anticipation
- ☐ Orientation/Surprise

Example 2

I needed to walk alone through the dark forest and heard a loud noise behind me.

Circumplex model (Russell):

Valence?	☐ high	☐ low
Arousal?	☐ high	☐ low

Appraisals (Smith/Ellsworth):

Pleasantness?	☐ high	☐ low
Responsible?	☐ high	☐ low
Certain?	☐ high	☐ low
Attention?	☐ high	☐ low
Effort?	☐ high	☐ low
Control?	☐ high	☐ low

Emotion Wheel (Plutchik):

- ☐ Protection/Fear
- ☐ Destruction/Anger
- ☐ Reproduction/Joy
- ☐ Deprivation/Sadness
- ☐ Incorporation/Acceptance
- ☐ Rejection/Disgust
- ☐ Exploration/Anticipation
- ☐ Orientation/Surprise

Task Definition for Emotion Classification and Regression

Input

- Text
- Variables respr. emotion model — Arousal, Valence, Emotion Category, Intensity
- Perspective — Reader, Writer, Text, mentioned entity

Output (by human or machine)

- Discrete values — emotion categories
- Ordinal values — intensities or appraisals
- Continous values — intensities, valence/arousal/dominance

Annotation Perspective and Reliability

Example: "I thought that Wayan might beat Putu."

- Writer: fear (pretty obvious case, but still, we don't know what the person really felt)
- Reader: fear? (depends on context)

Factors that influence decision

- World knowledge (to be beaten is something to be afraid of)
- Context (Speaker is friend of Putu.)
- Personality (Speaker might be neurotic.)
- Demographics (Might influence world knowledge.)

It really depends on the task and domain.

Hypothetical setting:
Given news articles, what is the emotional impact on the reader?

"If we continue to fly to conferences around the globe our children will not have anything to eat anymore because of global warming."

- Person who does believe global warming is not caused by humans: anger
- Average member of the society: fear
- Some NLP researcher: sadness

⇒ We can probably never access all relevant information.

Annotation Setup: Trained Experts or Crowdsourcing?

Trained Experts:

- Might be preferrable if variables follow challenging concepts

Crowdsourcing:

- If the study is more of an experiment to study subjective perceptions
- "What emotion do you feel when reading the text?"
- "What would an average reader feel"? (Buechel, 2017)

Outline

Emotion Recognition Reliability: Ekman 1972

Experimental Setup

- Photos were taken of people expressing a particular emotion and asked which emotion they feel
- Japanese and US American people were shown these photos and tasked to recover the emotion
- Goal: understand emotion recognition reliability

Results (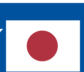/)

- .79/.86 acc. between observers
- .57/.62 acc. between subject and observer (.50 baseline)

⇒ Interpretation of emotion might differ from actual emotion.

Factors for emotion recognition reliability (Döllinger, 2021)

Follow-up studies investigated factors for recognition reliability:

- Emotion category
  - Some emotions are easier to recognize than others (joy vs. fear: Mancini 2018)
- Peer status
  - Friends are better in recognizing their emotions (Wang 2019)
- Status of observer
  - People with depression are more challenged in recognizing emotions (Dalili 2015)
  - Personality traits: conscientious and open people are better to recognize emotions, shy and neurotic people are worse (Hall 2016)
- Does that affect our annotation study design?

$\Rightarrow$ We might be able to prescreen annotators (though I have never seen any study doing that in NLP)

Take-Away

Emotions...

- ...are quite well understood in psychology
- ...can be represented via affect, appraisal, or categorical names
- ...cannot be reliably annotated, because of potentially missing relevant information
- ...are just hard to recognize

Questions?

About this tutorial

Session 1 (09:00–10:30)

- Introduction
- Psychological Models
- Use Cases/Social Impact
- Resources
- Annotation Exercise

Break (10:30–11:15)

Session 2 (11:15–12:45)

- Non-Neural Methods
- Multi-task, transfer, zero-shot methods
- Open Challenges
- Appraisal Theories
- Role Labeling
- Ethical Considerations
- Closing

USE CASES

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

USE CASES

- Social media and public opinion analysis

- Literary studies
- Hate speech detection

- Empathethic chatbots and virtual agents
- Early depression detection
- Suicide prevention

SOCIAL MEDIA AND PUBLIC OPINION ANALYSIS

SOCIAL MEDIA AND PUBLIC OPINION ANALYSIS: Loureiro and Alló, 2020

- Methodology:
  - Twitter messages about climate change analyzed using EmoLex (Mohammad and Turney, 2013)
  - Data collection: 01.01.2019-30.06.2019 (six months)

- Findings:
  - Messages in the UK less negative than in Spain
  - The most evoked feeling is anticipation in the UK and fear in Spain
  - Similar views about preferences for energy policies: renewable sources are perceived positively, coal negatively, and nuclear energy is associated with heterogeneous perceptions

SOCIAL MEDIA AND PUBLIC OPINION ANALYSIS: Srinivasan et al., 2019

- Methodology:
  - Twitter messages mentioning Hillary Clinton or Donald Trump analyzed using EmoLex (Mohammad and Turney, 2013)
  - Data collection: 26.09.2016 – 6.11.2016 (six weeks)

- Findings:
  - 90% accuracy for swing directions for 17 out of 19 states
  - Better accuracy than from 9 different pollsters
  (79% accuracy; correctly predicted swing directions for 15 out of 19 states)
  - Swing in the emotions aligned with various political events

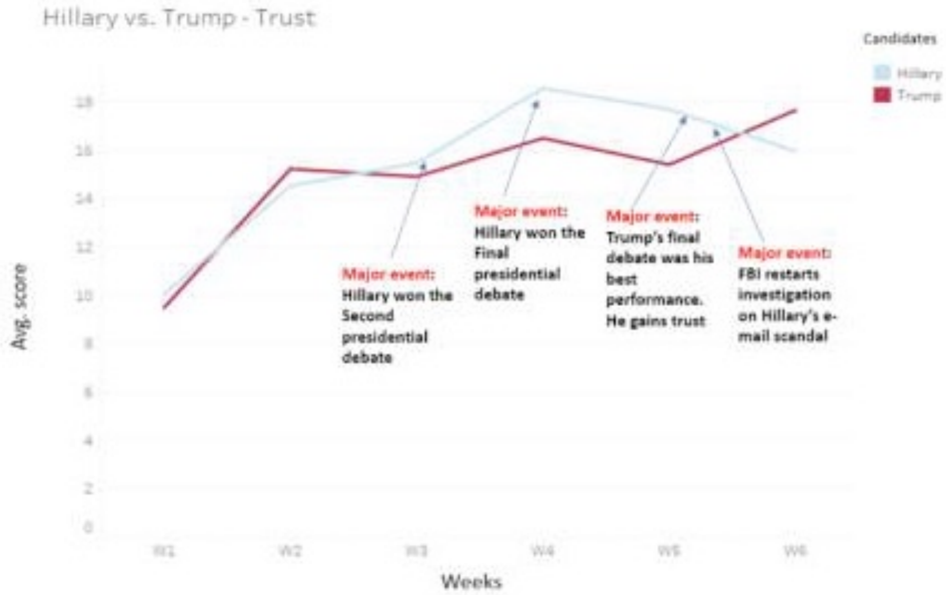

Figure taken from (Srinivasan et al., 2016)

SOCIAL MEDIA AND PUBLIC OPINION ANALYSIS: Wang et al., 2023

- Methodology:

  Twitter posts of top executives in S&P 1500 firms analyzed using DeepEmotionNet (Wang et al., 2023)

- Findings:

  Fear and anger in Twitter posts by top executives are significantly associated with corporate financial performance

LITERARY STUDIES

LITERARY STUDIES: Reagan et al., 2016

"Our ability to communicate relies in part upon a shared emotional experience, with stories often following distinct emotional trajectories and forming patterns that are meaningful to us."

(Reagan et al., 2016)

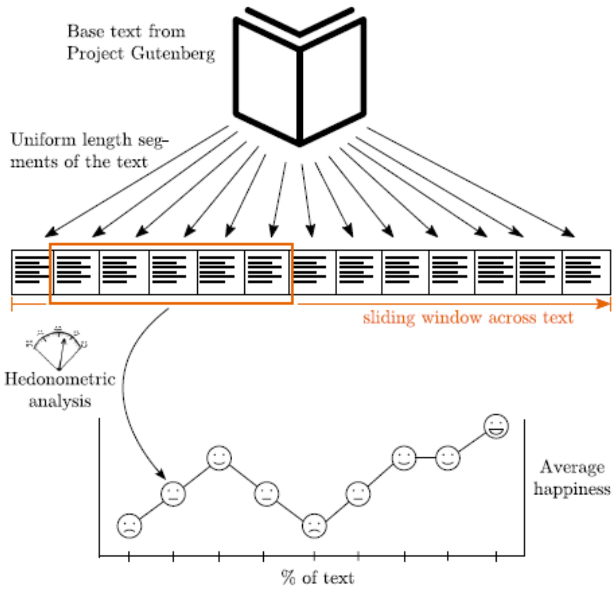

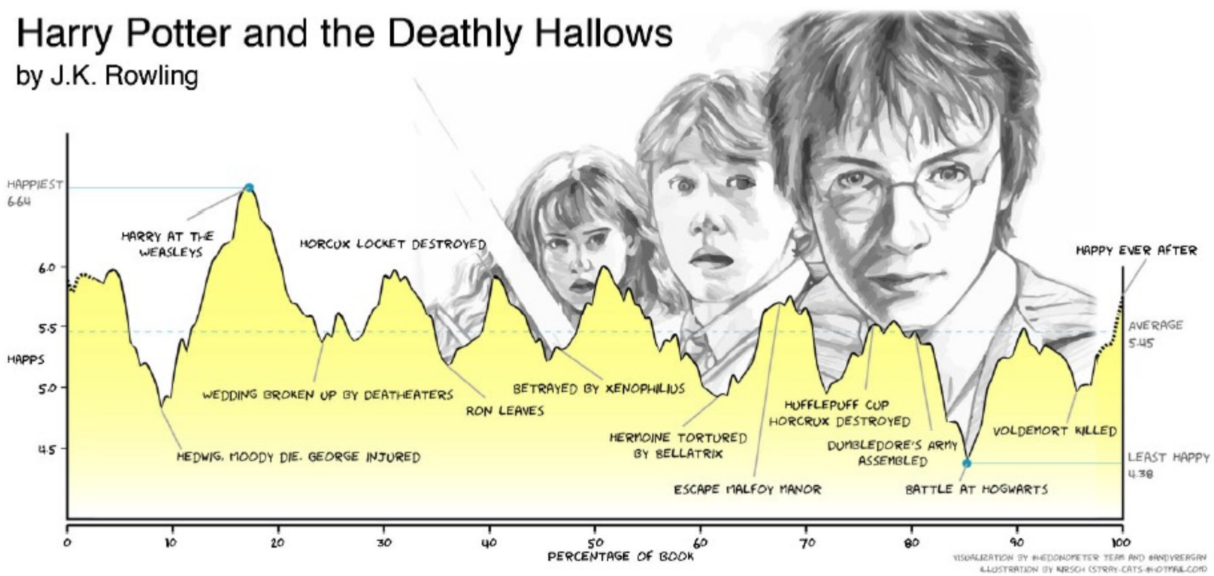

Figures taken from (Reagan et al., 2016)

LITERARY STUDIES: Reagan et al., 2016

- Data and emotion detection:
  - 1327 books from Project Gutenberg (mostly fictional)
  - Happiness using Hedonometer (Dodds et al., 2011)

- 6 most common emotional arcs:
  - ‘Rags to riches’ (rise)
  - ‘Tragedy’, or ‘Riches to rags’ (fall)
  - ‘Man in a hole’ (fall-rise)
  - ‘Icarus’ (rise-fall)
  - ‘Cinderella’ (rise-fall-rise)
  - ‘Oedipus’ (fall-rise-fall)

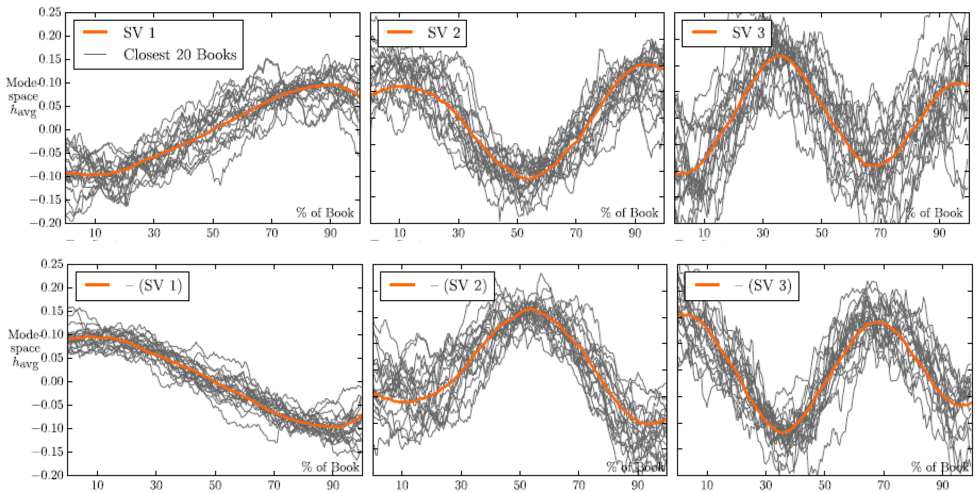

Figure adapted from (Reagan et al., 2016)

LITERARY STUDIES: Kim et al., 2017

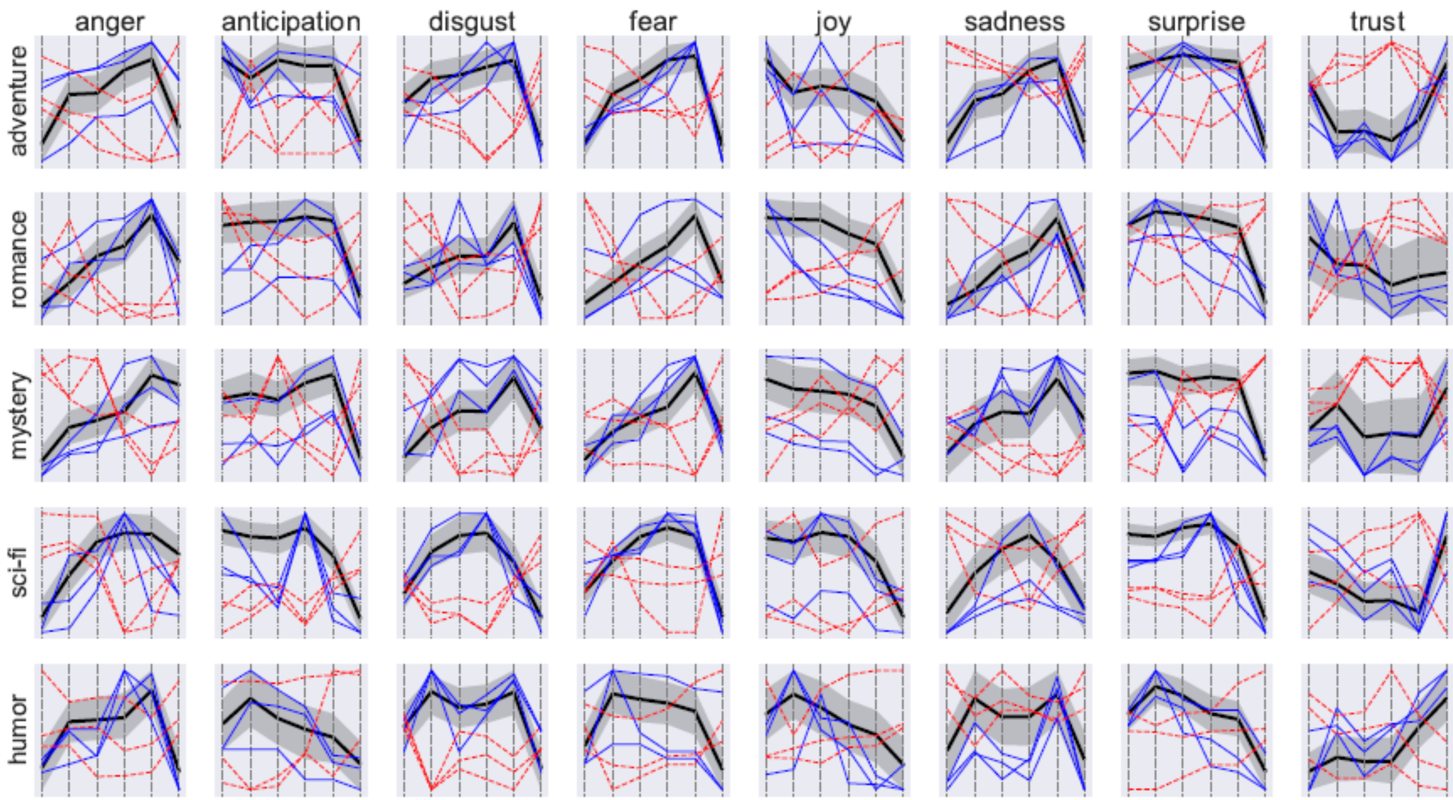

Figure taken from (Kim et al., 2017)

LITERARY STUDIES: Kim et al., 2017

- Genre classification feature sets
  - EmoLex (Mohammad and Turney, 2013)
  - Bag of Words (BoW)
  - Emotion arcs
  - Ensemble

Genre	Count
adventure	569
humor	202
mystery	379
romance	327
science fiction	542
$\sum$	2019

Table taken from (Kim et al., 2017)

- Results:
  - Use of EmoLex outperforms the use of BoW (81% vs. 80% for micro average F1-score)
  - Use of emotion arcs instead of lexical features leads to worse classification results
  - Ensemble model performs significantly better than the best single model (84% vs. 81% for micro average F1-score)

HATE SPEECH DETECTION

HATE SPEECH DETECTION

“Any strongly impolite, rude or hurtful language using profanity, that can show a debasement of someone or something, or show intense emotion”

Fortuna and Nunes (2018)

Some of the studies that use emotion analysis in hate speech detection:

- “Hate speech classification in social media using emotional analysis” (Martins et al., 2018)
- “Automatic Detection of Hate Speech on Facebook Using Sentiment and Emotion Analysis” (Rodriguez & Argueta, 2019)

EMPATHETIC CHATBOTS AND VIRTUAL AGENTS

EMPATHETIC DIALOGUES

Figure taken from (Rashkin et al., 2019)

EMPATHETIC DIALOGUES DATASET: Rashkin et al., 2019

Label: Afraid
Situation: Speaker felt this when...
"I've been hearing noises around the house at night"
Conversation:
Speaker: I've been hearing some strange noises around the house at night.
Listener: oh no! That's scary! What do you think it is?
Speaker: I don't know, that's what's making me anxious.
Listener: I'm sorry to hear that. I wish I could help you figure it out

Label: Proud
Situation: Speaker felt this when...
"I finally got that promotion at work! I have tried so hard for so long to get it!"
Conversation:
Speaker: I finally got promoted today at work!
Listener: Congrats! That's great!
Speaker: Thank you! I've been trying to get it for a while now!
Listener: That is quite an accomplishment and you should be proud!

Figure taken from (Rashkin et al., 2019)

EARLY DEPRESSION DETECTION

DEPRESSION DETECTION: Islam et al., 2018

- Methodology:
  - Facebook posts analyzed for depression using LIWC software
  - Classification experiments with various ML algorithms
  - 4 feature sets: emotional processes (positive emotion words, negative emotion words, sadness words, anger words, anxiety words), linguistic style, temporal processes, and the combination of all
- Findings:
  - Up to 73% F-measure for binary classification (depression yes or no)
- Drawbacks:
  - Ground truth?
  - Who is depressed?

DEPRESSION DETECTION: Shanti et al., 2022

Accuracy: 91%

Accuracy: 82%

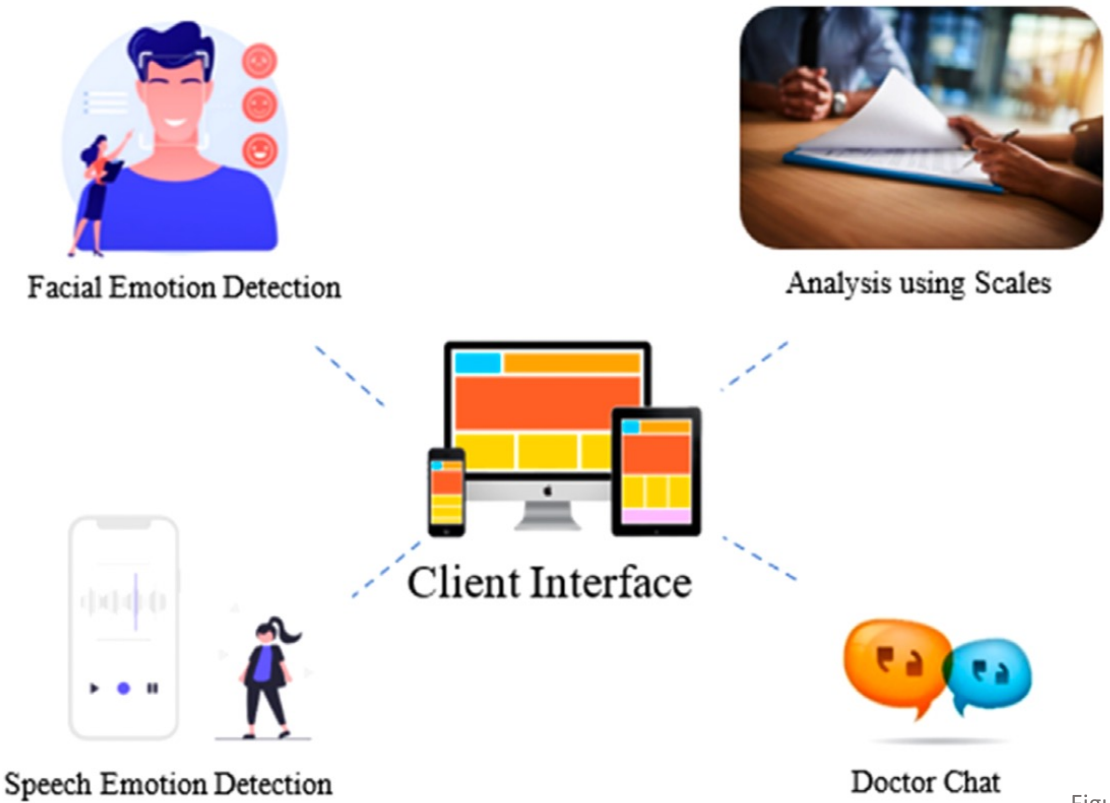

Figure taken from (Shanti et al., 2022)

SUICIDE PREVENTION

EMOTION ANALYSIS OF SUICIDE NOTES: Shared Task

- Shared task in 2011 (Pestian et al., 2012)
- Ground truth (annotation):
  - Annotators were asked to identify abuse, anger, blame, fear, guilt, hopelessness, sorrow, forgiveness, happiness, peacefulness, hopefulness, love, pride, thankfulness, instructions, and information
  - Annotators were survivors of suicide loss, active in suicide communities

EMOTION ANALYSIS OF SUICIDE NOTES: Desmet and Hoste, 2013

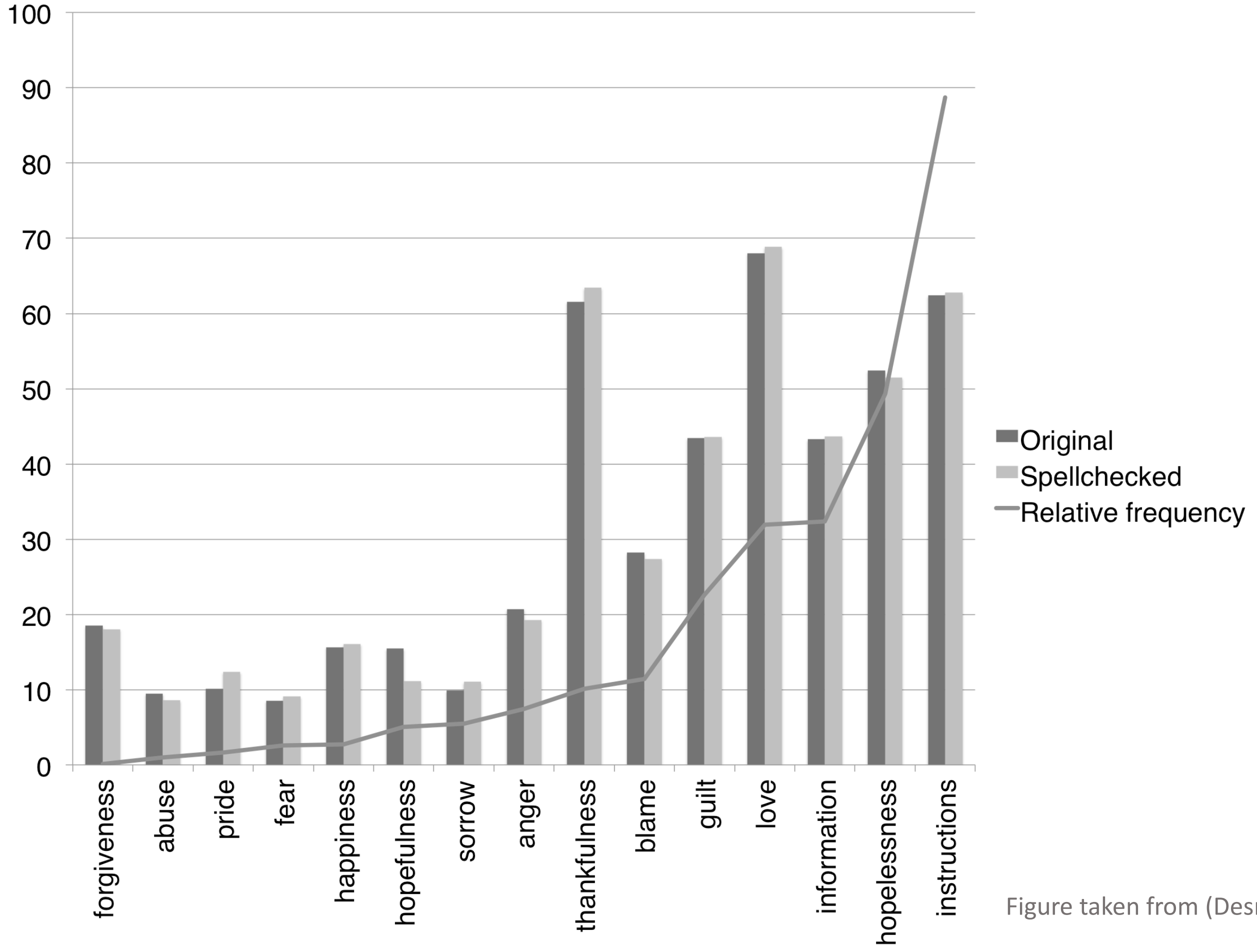

Figure taken from (Desmet and Hoste, 2013)

Questions?

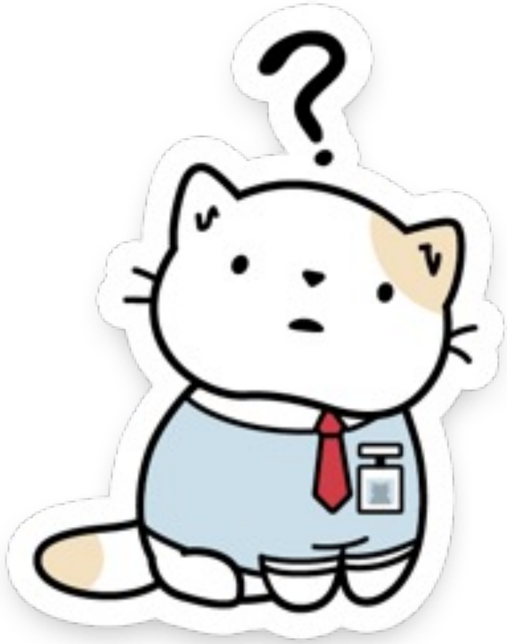

USE CASES

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

RESOURCES

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

RESOURCES

- Emotion detection and classification resources
- Emotion intensity resources
- Other resources

ANNOTATION OPTIONS

- Automatic or human
- Number of annotators per instance
- Total number of annotators
- Expertise of the annotators
- Ground truth assignment
- Set of emotions
- Labelling type (single or multi)
- Perspective (reader, writer, text)
- Genre and context length

AUTOMATIC ANNOTATION

- Wang et al. (2012): 131 emotion hashtags as keywords (hashtag at the end of tweet) for collecting 5 million tweets in seven emotion categories (joy, sadness, anger, love, thankfulness, surprise).

- Shahraki and Zaïane (2017): based on 15 explicit hashtags appearing in them compiled Clean Balanced Emotional Dataset (CBET) with 27,000 annotated tweets (3,000 per each emotion: anger, fear, joy, love, sadness, surprise, thankfulness, disgust, and guilt)

- Mohammad (2012): 21,051 tweets which contained one of the six Ekman's emotions (anger, disgust, fear, joy, sadness, surprise) as the last hashtag

VARIATIONS IN HUMAN ANNOTATION: Štajner, 2021

Study	#annotators Per instance	Total	Gold	#emotions	Labelling	Perspective	Genre
(Demszky et al., 2020)	3 or 5	82	$>$ 1 annotator	27+1	multi	writer	Reddit
(Bostan et al., 2020)	5	310	$>$ 1 annotator	15+1	single	text	Headlines
(Öhman et al., 2020)	$\leq$3	108	$>$ 1 annotator	8+1	multi	speaker	Subtitles
(Poria et al., 2019)	5	?	majority	6+1	single	speaker	Dialog
(Hsu et al., 2018)	5	?	majority*	6+1	single	speaker	Dialog
(Schuff et al., 2017)	3–6	6	various	8	multi	?	Twitter
(Mohammad et al., 2015)	3+	$\approx$ 3000	$>$ half	19+1	single	text	Twitter
(Brynielsson et al., 2014)	3	3	majority	3+1	single	writer	Twitter
(Neviarouskaya et al., 2010)	3	3	$\geq$2 agree	14	single	?	Various
(Neviarouskaya et al., 2009)	3	3	$\geq$2 agree	9+1	single	?	Blogs
(Strapparava and Mihalcea, 2007)	6	6	?	6	multi	reader	Headlines
(Aman and Szpakowicz, 2007)	2	4	both agree	6+2	single	text	Blogs
(Alm et al., 2005)	2-3	3	majority	6+1	single	text	Children

Table 1: Annotation procedures used in previous studies ("?" signifies that the particular aspect was not specified in the paper, "+1" in the *#emotions* column signifies the additional class for "other" or "no emotion").

Table taken from (Štajner, 2021)

EMOTIONS IN CHILDREN STORIES: Alm et al., 2005

- Genre: children stories (22 Grimms' tales)
- Span: sentence
- Size: 1580 sentences
- Emotions: extended Ekman's (added neutral and split surprise into positive and negative)
- Perspective: text's (the feeler in the sentence)
- Labelling: single
- Annotators: 2
- Gold: both agree

EMOTIONS IN BLOGS: Aman and Szpakowicz, 2007

- Genre: blogs (selected by using seeds!)
- Span: sentence
- Size: 1466 emotional + 2800 no emotion
- Emotions: extended Ekman's (added mixed emotion and no emotion)
- Intensity: low, medium, and high
- Perspective: writer's
- Labelling: single
- Annotators: 2 per sentence (4 in total)
- Gold: both agree

EMOTIONS IN BLOGS: Neviarouskaya et al., 2009

- Genre: diary-like blog posts (BuzzMetrics)
- Span: sentence
- Size: 700 sentences
- Emotions: subset of emotional states defined by Izard (interest, joy, surprise, anger, disgust, fear, guilt, sadness, shame)
- Intensity: [0.0, 1.0]
- Perspective: ?
- Labelling: single
- Annotators: 3
- Gold: at least 2 agree (656 sentences)

EMOTIONS IN NEWS HEADLINES: Strapparava and Mihalcea, 2007

SemEval-2007 Task 14: Affective Text

- Genre: news headlines
- Span: headline
- Size: 1250 headlines
- Emotions: Ekman's (anger, disgust, fear, joy, sadness, surprise)
- Intensity: [0,100]
- Perspective: reader's
- Labelling: multiple
- Annotators: 6
- Gold: ?

EMOTIONS IN ELECTORAL TWEETS: Mohammad et al., 2015

- Genre: electoral tweets
- Span: tweet
- Size: 2,000 tweets
- Emotions: Plutchik (19->8)
- Intensity: low, medium, high
- Perspective: various
- Labelling: single
- Annotators: ~ 30,000 crowdsourced (AMT and CrowdFlower), at least 5 per each
- Gold: belongs to category X if it was annotated with X more times than with all others combined

Q1. Which of the following best describes the **Emotions** in this tweet?

- This tweet expresses or suggests an emotional attitude or response to something.
- This tweet expresses or suggests two or more contrasting emotional attitudes or responses.
- This tweet has no emotional content.
- There is some emotion here, but the tweet does not give enough context to determine which emotion it is.
- It is not possible to decide which of the above options is appropriate.

EMOTIONS IN TWEETS: Schuff et al., 2017

- Genre: SemEval 2016 Stance Data set (Mohammad et al., 2016)
- Span: tweet
- Size: 4,868 tweets
- Emotions: Plutchik (anger, anticipation, disgust, fear, joy, sadness, surprise, trust)
- Perspective: ?
- Labelling: multi
- Annotators: 6 (minimum 3 per each tweet)
- Gold: various

EMOTIONS IN CONVERSATIONS: Hsu et al., 2018

- Genre: multi-party conversations (Friends TV scripts and FB personal dialogues)
- Span: utterance
- Size: 29,245 utterances (2,000 dialogues)
- Emotions: Ekman's + neutral + non-neutral
- Perspective: speaker
- Labelling: single
- Annotators: 5 AMT workers per each
- Gold: majority (when more than two majority then class non-neutral)

EMOTIONS IN CONVERSATIONS: Hsu et al., 2018

- Genre: multi-party conversations (Friends TV scripts and FB personal dialogues)
- Span: utterance
- Size: 29,245 utterances (2,000 dialogues)

	# of Utterances	Utterance Length	Emotion Label Distribution (%)								kappa (%)
			Neu	Joy	Sad	Fea	Ang	Sur	Dis	Non	
Friends	14,503	10.67	45.03	11.79	3.43	1.70	5.23	11.43	2.28	19.11	33.83
EmotionPush	14,742	6.84	66.85	14.25	3.49	0.28	0.95	3.85	0.72	9.62	33.64

EMOTIONS IN SUBTITLES: Öhman et al., 2020

- Genre: movie subtitles from OPUS (Lison and Tiedemann, 2016)
- Languages: Finnish and English (human annotation) + 30 others (projections)
- Span: subtitle (roughly 1 sentence)
- Size: 25,000 sentences (Finnish) + 30,000 sentences (English)
- Emotions: Plutchik (8) + neutral
- Perspective: speaker
- Labelling: single
- Annotators: 60-100 students (2-3 per instance)
- Gold: at least 2 agreed

OTHER RESOURCES

EMPATHETIC DIALOGUES DATASET: Rashkin et al., 2019

- Approximately 25000 dialogues grounded in situations prompted by specific emotion labels (32 different emotion labels)

Label: Afraid
Situation: Speaker felt this when...
"I've been hearing noises around the house at night"
Conversation:
Speaker: I've been hearing some strange noises around the house at night.
Listener: oh no! That's scary! What do you think it is?
Speaker: I don't know, that's what's making me anxious.
Listener: I'm sorry to hear that. I wish I could help you figure it out

Label: Proud
Situation: Speaker felt this when...
"I finally got that promotion at work! I have tried so hard for so long to get it!"
Conversation:
Speaker: I finally got promoted today at work!
Listener: Congrats! That's great!
Speaker: Thank you! I've been trying to get it for a while now!
Listener: That is quite an accomplishment and you should be proud!

Figure taken from (Rashkin et al., 2019)

Questions?

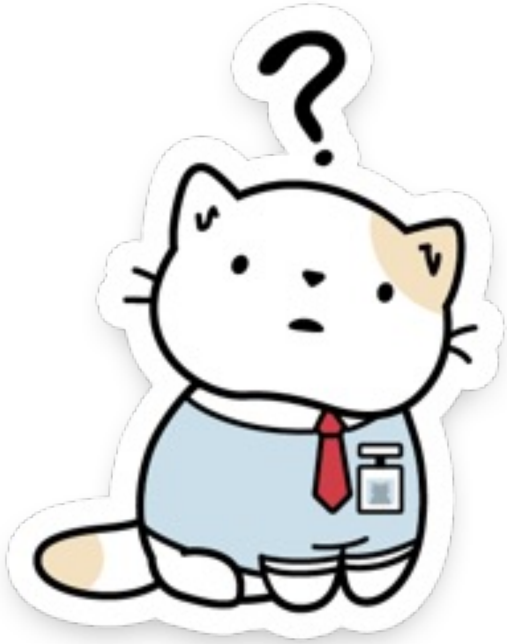

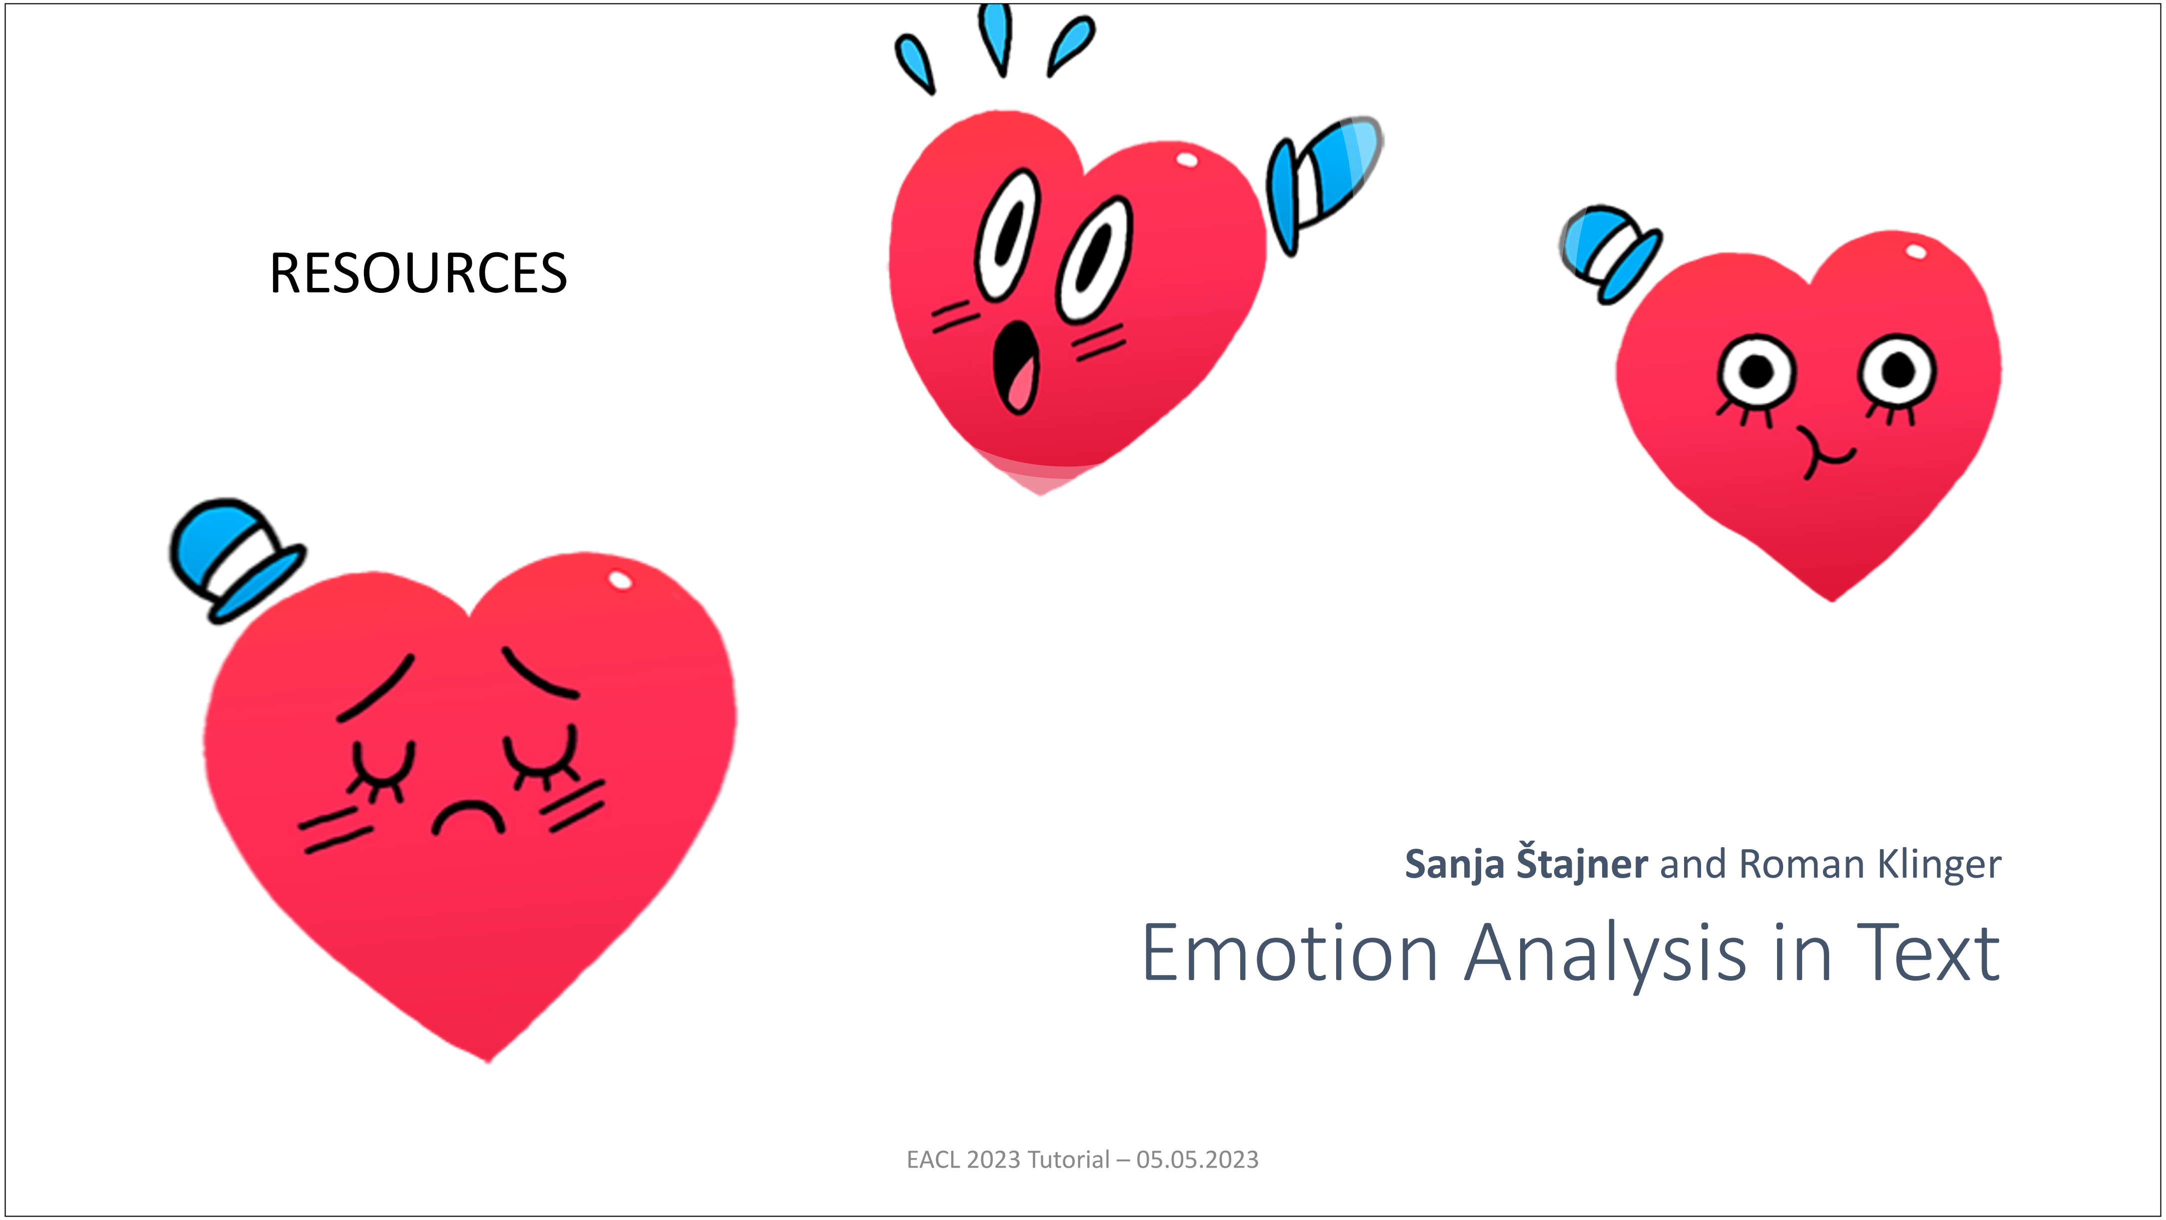
RESOURCES
Sanja Štajner and Roman Klinger
Emotion Analysis in Text
EACL 2023 Tutorial – 05.05.2023

Universität Stuttgart
Institut für
Maschinelle Sprachverarbeitung

Emotion Analysis

Small Annotation Exercise and Discussion

EACL 2023 Tutorial

Sanja Štajner and Roman Klinger

Hand On Annotation

What we will do now:

- You heard now a bit about existing resources.
- Let's do an annotation together.
- For each instance that we show you, answer the questions in the form.

Think about the following questions:

- Would annotators agree on the label?
- Would an automatic method succeed/fail?

Link: `https://forms.gle/9pwPXnCCB8K1ocrg7`

Questions

- Did you miss annotation labels?
- Would you have prefered to annotate multiple emotions?
- Would you prefer a neutral label?
- What are properties of instances that you assume would never be correctly predicted by machines?

About this tutorial

Session 1 (09:00–10:30)

- Introduction
- Psychological Models
- Use Cases/Social Impact
- Resources
- Annotation Exercise

Break (10:30–11:15)

Session 2 (11:15–12:45)

- Non-Neural Methods
- Multi-task, transfer, zero-shot methods
- Open Challenges
- Appraisal Theories
- Role Labeling
- Ethical Considerations
- Closing

NON-NEURAL MODELS

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

NON-NEURAL MODELS

EMOTIONS IN CHILDREN STORIES: Alm et al., 2005

- Genre: children stories (22 Grimms' tales)
- Task: Emotional vs. non-emotional
- rule-based linear classifier (SNoW)
- 10-fold cross-validation (90% training, 10% testing)

EMOTIONS IN CHILDREN STORIES: Alm et al., 2005

- Features:
  - First sentence in the story
  - Conjunctions of selected features
  - Direct speech
  - Thematic story type
  - Special punctuation
  - Complete upper-case word
  - Sentence length in words
  - Ranges of story progress
  - Percent of JJ, N, V, RB
  - V counts in sentence, excluding participles
  - Positive and negative word count
  - WordNet emotion Words
  - Interjections and affective words
  - Content BoW: N, V, JJ, RB words by POS

	same-tune-eval	sep-tune-eval
P(Neutral)	59.94	60.05
Content BOW	61.01	58.30
All features except BOW	64.68	63.45
All features	68.99	63.31
All features + sequencing	69.37	62.94

Accuracy

Figure taken from (Alm et al., 2005)

EMOTIONS IN BLOGS: Aman and Szpakowicz, 2007

- Genre: blogs (selected by using seeds!)
- Span: sentence
- Size: 1466 emotional + 2800 no emotion
- Task: Emotional vs. non-emotional
- For feature extraction used emotional dictionaries:
  - General Inquirer (Stone et al., 1966)
  - WordNet-Affect (Strapparava and Valitutti, 2004)

EMOTIONS IN BLOGS: Aman and Szpakowicz, 2007

GI Features	WN-Affect Features	Other Features
Emotion words Positive words Negative words Interjection words Pleasure words Pain words	Happiness words Sadness words Anger words Disgust words Surprise words Fear words	Emoticons Exclamation ("!") and question ("?") marks

Features	**Naïve Bayes**	**SVM**
GI	71.45%	71.33%
WN-Affect	70.16%	70.58%
GI+WN-Affect	71.7%	73.89%
ALL	**72.08%**	**73.89%**

Accuracy

Figures taken from (Aman and Szpakowicz, 2007)

EMOTIONS IN ELECTORAL TWEETS: Mohammad et al., 2015

- Genre: electoral tweets
- Emotions: Plutchik (8)
- 10-fold stratified cross-validation
- SVM with linear kernel (also tried logistic regression and different SVM kernels)

EMOTIONS IN ELECTORAL TWEETS: Mohammad et al., 2015

- Features:
  - word unigrams and bigrams
  - Punctuations
  - Elongated words
  - Emotions
  - Emotion lexicons
  - Negations
  - Position features
  - Combined features

	Accuracy
random baseline	30.26
majority baseline	47.75
automatic SVM system	56.84
human performance	69.80

Figure taken from (Mohammad et al., 2015)

EMOTIONS IN SUBTITLES: Öhman et al., 2020

- Features:
  - Word unigrams, bigrams, trigram

SVM per class f1	emotion
0.8073	anger
0.8296	anticipation
0.8832	disgust
0.8763	fear
0.8819	joy
0.8762	sadness
0.8430	surprise
0.8832	trust

Figure taken from (Öhman et al., 2020)

NON-NEURAL VS. NEURAL: Öhman et al., 2020

data	f1	accuracy
English without NER, BERT	0.530	0.538
English with NER, BERT	0.536	0.544
English NER with neutral, BERT	0.467	0.529
English NER binary with surprise, BERT	0.679	0.765
English NER true binary, BERT	0.838	0.840
Finnish anno., FinBERT	0.507	0.513
English NER, one-vs-rest SVM (LinearSVC)[7]	0.746	

Figure taken from (Öhman et al., 2020)

NON-NEURAL VS. NEURAL: Öhman et al., 2020

Dataset	Language-specific BERT	SVM
Finnish projected	0.4461	0.5859
Turkish projected	0.4685	0.6080
Arabic projected	0.4627	0.5729
German projected	0.5084	0.6059
Dutch projected	0.5155	0.6140
Chinese projected	0.4729	0.5044

Data taken from (Öhman et al., 2020)

NON-NEURAL VS. NEURAL: Schuff et al., 2017

Bag-of-words

	Linear						Neural								
	MaxEnt			SVM			LSTM			Bi-LSTM			CNN		
Emotion	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
Anger	76	72	74	76	69	72	76 (1.7)	77 (5.3)	76 (1.9)	77 (0.8)	77 (2.7)	**77** (1.3)	77 (0.8)	77 (2.7)	**77** (1.3)
Anticipation	72	61	66	70	60	64	68 (1.8)	68 (8.9)	67 (3.5)	70 (1.2)	66 (3.6)	**68** (1.6)	68 (1.2)	60 (0.8)	64 (0.5)
Disgust	62	47	54	59	53	56	64 (3.2)	68 (8.7)	**65** (2.5)	61 (1.4)	64 (4.6)	63 (1.7)	62 (0.6)	61 (3.9)	62 (1.9)
Fear	57	31	40	55	40	46	51 (3.5)	48 (8.5)	**49** (4.6)	58 (1.6)	43 (6.3)	**49** (3.8)	53 (1.7)	46 (6.2)	**49** (3.9)
Joy	55	50	52	52	52	52	56 (5.9)	41 (8.3)	46 (4.8)	54 (2.9)	59 (10.5)	**56** (4.8)	54 (1.7)	56 (5.6)	55 (2.3)
Sadness	65	65	65	64	60	62	60 (2.5)	77 (11.1)	**67** (3.9)	62 (0.6)	72 (7.5)	**67** (3.2)	63 (0.9)	72 (0.3)	**67** (0.5)
Surprise	62	15	24	46	22	30	40 (4.4)	17 (10.4)	21 (8.7)	42 (2.9)	20 (3.2)	27 (2.5)	36 (3.7)	24 (6.3)	**28** (5.0)
Trust	62	38	47	57	45	50	57 (6.1)	49 (12.3)	**51** (5.9)	59 (2.5)	44 (4.1)	50 (2.5)	53 (0.6)	49 (6.6)	50 (3.3)
Micro-Avg.	66	52	58	63	53	58	62 (0.9)	60 (1.9)	61 (0.7)	64 (0.3)	60 (2.4)	**62** (1.2)	62 (0.6)	59 (2.0)	60 (1.0)

Figure adapted from (Schuff et al., 2017)

Questions?

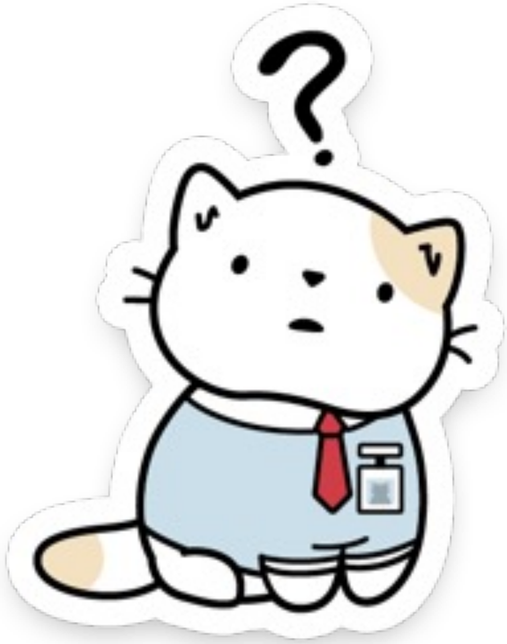

NON-NEURAL MODELS

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

Universität Stuttgart
Institut für
Maschinelle Sprachverarbeitung

Emotion Analysis

Transfer, Multi-Task Learning, Zero-Shot Predictions

EACL 2023 Tutorial

Sanja Štajner and Roman Klinger

Outline

1. Overview
2. Weak and Distant Labeling
   - Obtaining Automatically Annotated Corpora
   - Transfer Learning
3. Multi-task learning
4. Zero-Shot Prediction

Outline

Emotion Analysis as Text Classification

Where are we?

- Emotion classification as text classification
- Meaningful features can be extracted for the task
- What's happening in the deep learning world?

Shared Tasks on the Topic

- Affective Text (Headlines), 2007 (SemEval)
- Emotion Intensity, 2017 (WASSA), 2018 (SemEval)
- Emotion Classification (E-c) 2018 (SemEval)
- Implicit Emotions, 2018 (WASSA)
- More shared tasks at SemEval and WASSA

Emotion Classification E-c SemEval, Setting

Task Definition

Emotion Classification (E-c): Given a tweet, classify it as 'neutral or no emotion' or as one, or more, of eleven given emotions that best represent the mental state of the tweeter

- Annotation via crowdsourcing
- Aggregation:
  Accept emotion label with at least 2/7 annotations

(Mohammad et al., SemEval 2018)

Implicit Emotions Shared Task: Data and Task

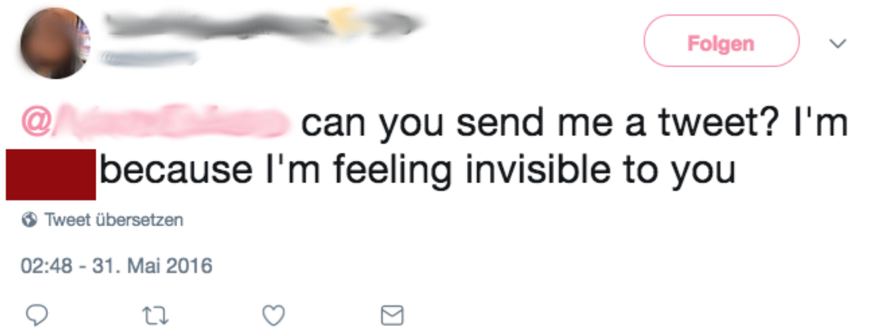

- Input:
  Tweet with emotion synonym replaced by unique string
- Output:
  Emotion for which the removed word is a synonym

Example

`sadness [USERNAME] can you send me a tweet? I'm [#TRIGGERWORD#] because I'm feeling invisible to you`

Outline

1. Overview
2. Weak and Distant Labeling
   - Obtaining Automatically Annotated Corpora
   - Transfer Learning
3. Multi-task learning
4. Zero-Shot Prediction

Weak/Self-Labeling

Approach:

- Manually associate
  - hashtags with emotions
  - emojis with emotions
- Assume that occurrence of hashtag/emoji marks emotion
- Predict "self-labeled emotion" from text after removing hashtag/emoji
- Apply to other texts

Advantage:

- Easy to obtain huge data sets

Disadvantage:

- Concept of emotion $\neq$ emotion hashtags/emojis
- Example: 10.1109/SocialCom-PASSAT.2012.119

Outline

Transfer Learning: DeepMoji

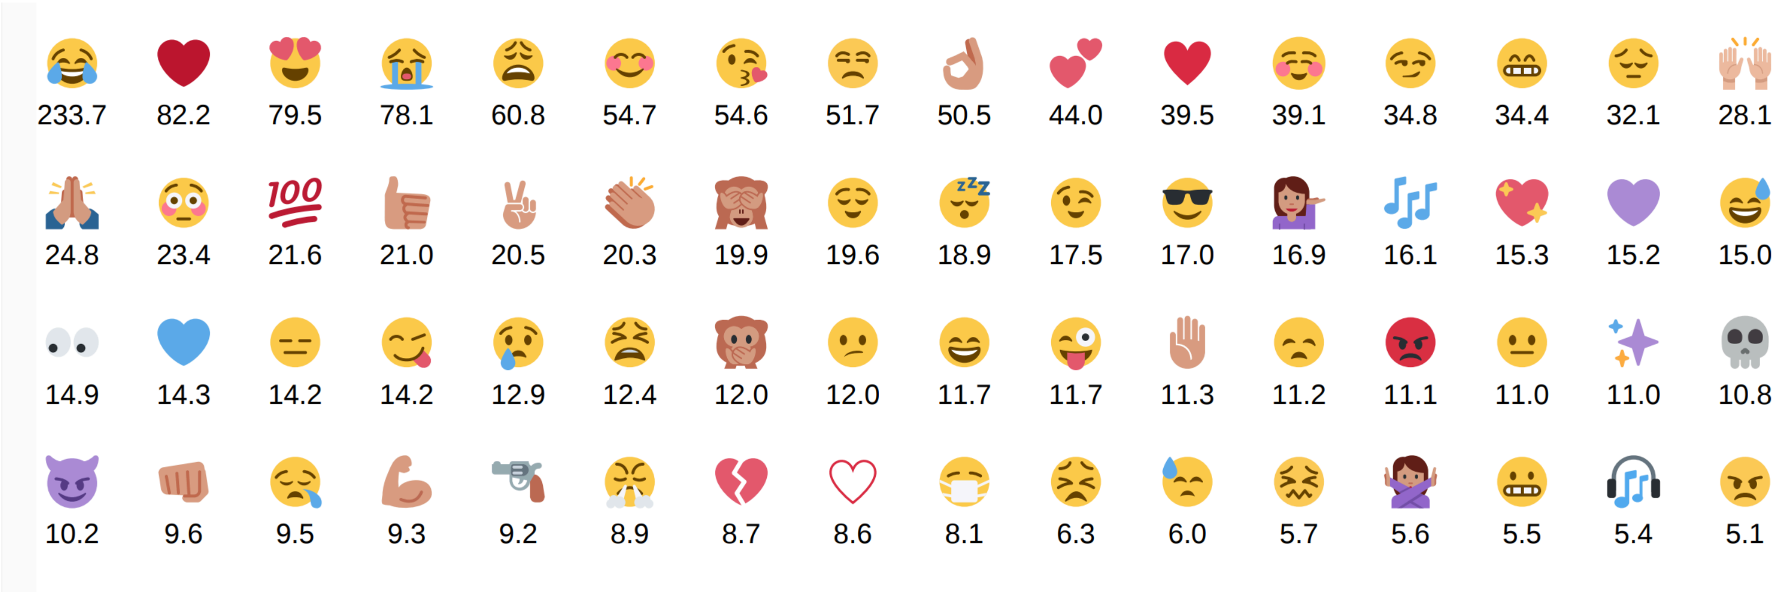

Transfer Learning: DeepMoji

- Develops a deep learning method for emotion classification (amongst other tasks)
- Pretrain model on huge data set to predict the occurrence of an emoji
- Fine-tune: Keep subset of parameters fixed while learning on actual data set.

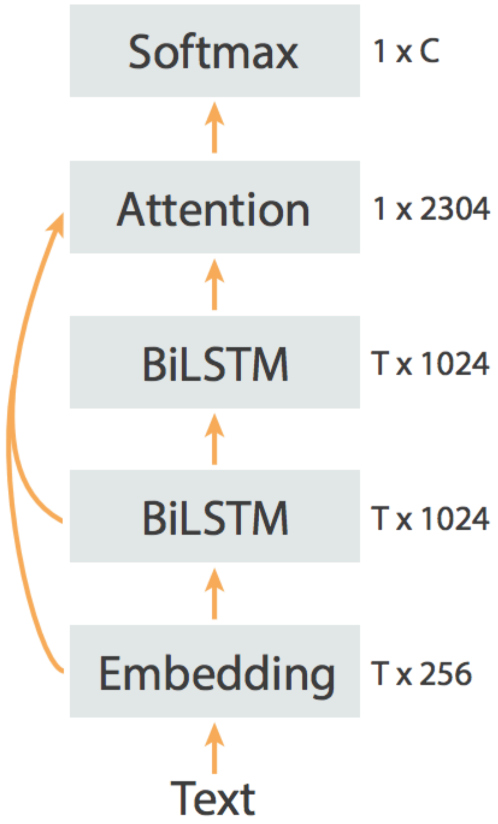

Transfer Learning: DeepMoji

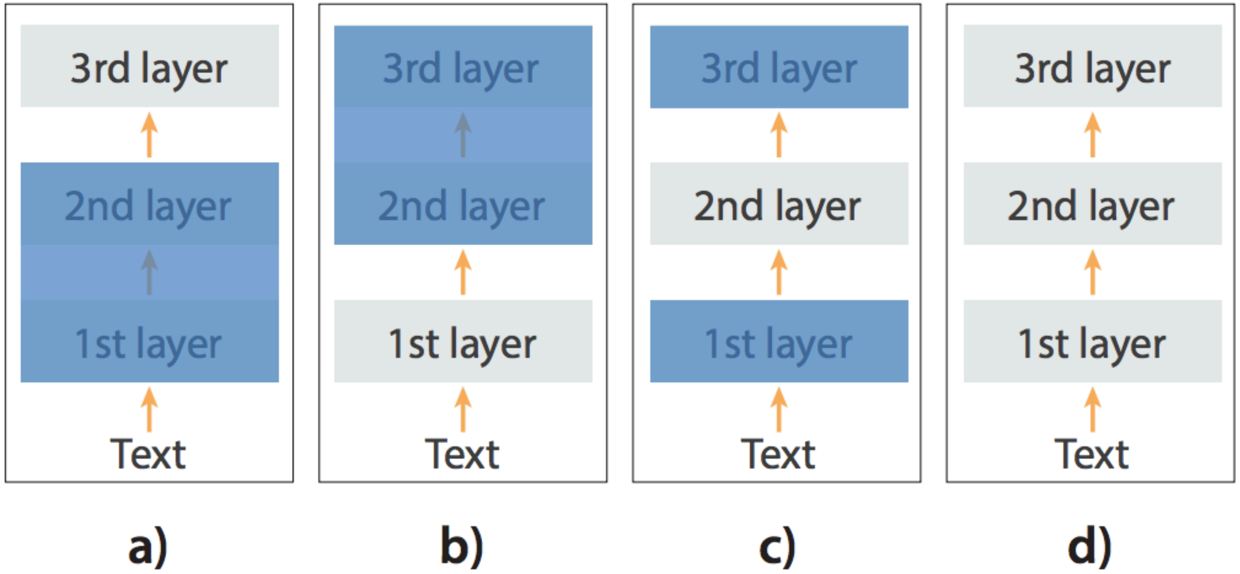

- Blue: frozen

a) tune any new layers
b) then tune 1st layer
c) then tune next layer, until all have been tuned
d) tune all together

Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, Sune Lehmann: Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. EMNLP 2017.

Final Remark on Results

- Results differ a lot between data sets
- Data sets are pretty incomparable

Cross-corpus experiment

- Split corpora in train/val
- Train BOW-MaxEnt-L2 on all train parts, apply on all val parts
- Join all train parts, apply on each val part

(Bostan/Klinger, COLING 2018)

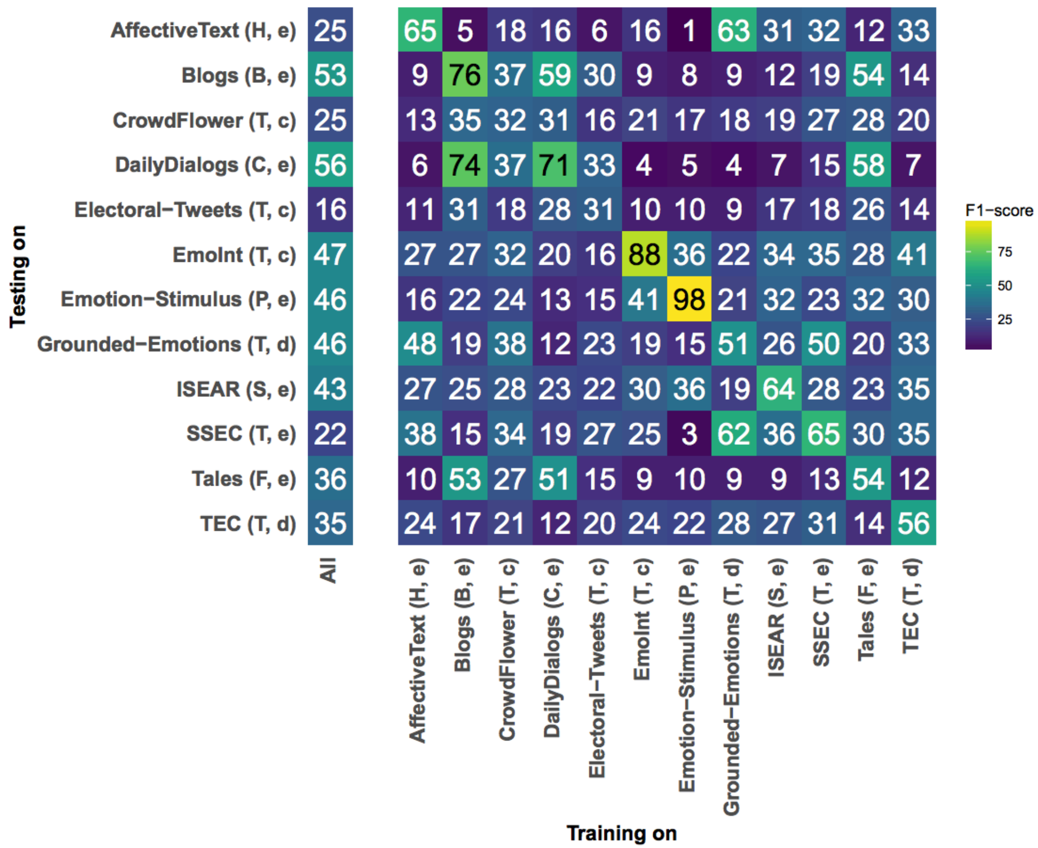

Outline

Tasks in Multitask Learning and Emotions

- Akhtar et al, NAACL 2019: Multi-task Learning for Multi-modal Emotion Recognition and Sentiment Analysis
  `https://www.aclweb.org/anthology/N19-1034.pdf`
- Chauhan et al, ACL 2020: Sentiment and Emotion help Sarcasm? A Multi-task Learning Framework for Multi-Modal Sarcasm, Sentiment and Emotion Analysis
  `https://www.aclweb.org/anthology/2020.acl-main.401.pdf`
- Dankers et al, EMNLP 2019: Modelling the interplay of metaphor and emotion through multitask learning
  `https://www.aclweb.org/anthology/D19-1227.pdf`

Tasks in Multitask Learning and Emotions

- Tafreshi et al, CoNLL 2018: Emotion Detection and Classification in a Multigenre Corpus with Joint Multi-Task Deep Learning
  https://www.aclweb.org/anthology/C18-1246.pdf
- Rajamanickam et al, ACL 2020: Joint Modelling of Emotion and Abusive Language Detection
  https://www.aclweb.org/anthology/2020.acl-main.394.pdf
- Saha et al, ACL 2020: Towards Emotion-aided Multi-modal Dialogue Act Classification
  https://www.aclweb.org/anthology/2020.acl-main.402.pdf
- Casel et al, KONVENS 2021: Emotion Recognition under Consideration of the Emotion Component Process Model.
  https://aclanthology.org/2021.konvens-1.5/

Summary

- Feature-based emotion analysis research came up with rich feature sets
- Deep learning, transfer learning commonly outperforms such approaches
- Current research is a lot about finding beneficial proxy tasks and to adapt input representations

Questions?

Zero-Shot Predictions

- "Zero-Shot" means: predict labels for instances that have some property that has not been seen during training.
- Most popular cases:
  - Cross-lingual Zero-Shot Transfer: Learn on language A and apply model to language B. (example: use multi-lingual pretrained language models)
  - Zero-Shot Labeling: Predict labels from a set that have not been seen during training
- Motivation: No need to know the exact required emotion concepts at model development time.
- That is a realistic requirement. Deciding on the emotion set is hard.

Why should Zero-Shot Learning be possible?

Training Data with labels: Deer, Fish, Rabbit

Deer

Fish

Rabbit

Test Data with unseen labels: Moose, Whale

Moose

Whale

- How do we make these assignments?
- We decide on properties of the instances to classify.
- We compare the extracted properties to those of the classes.
- We need some meaningful representation of each label.
- We need some meaningful representation of each instance.

Photos Attribution: Rabbit: David Iliff, Fish: Diego Delso, Deer: Frank Liebig, Whale: Whit Welles. Licenses: CC BY-SA 3.0, Moose: Public Domain

ZSL as Embedding Prediction

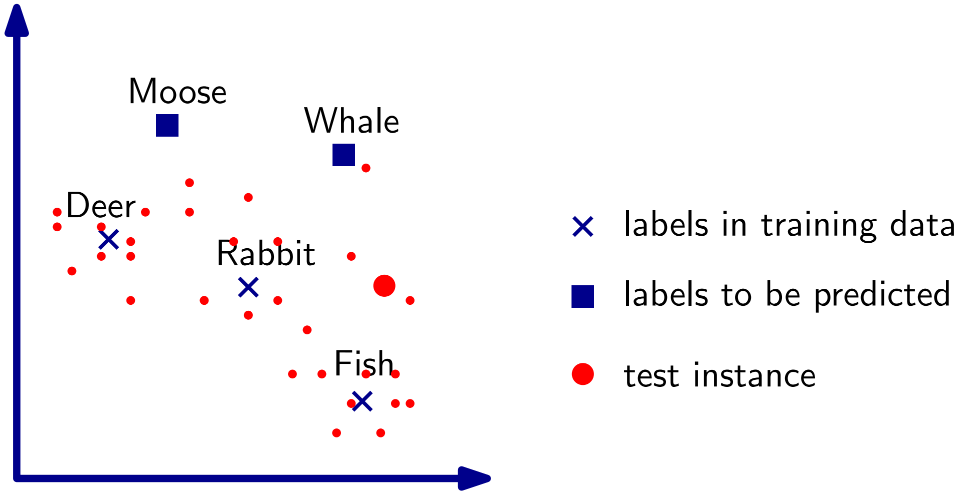

- Label vectors based on concept features
- Learn to map instance into concept space

- In ZSL, we would assign "whale".
- In Generalized ZSL, we assign "fish".
- Hubness problem: It's more likely to predict vectors that have been seen at model development time.
- Emotion analysis: Where do we get the concept embeddings from?

Related: ZSL for Emotion Classification from Gestures

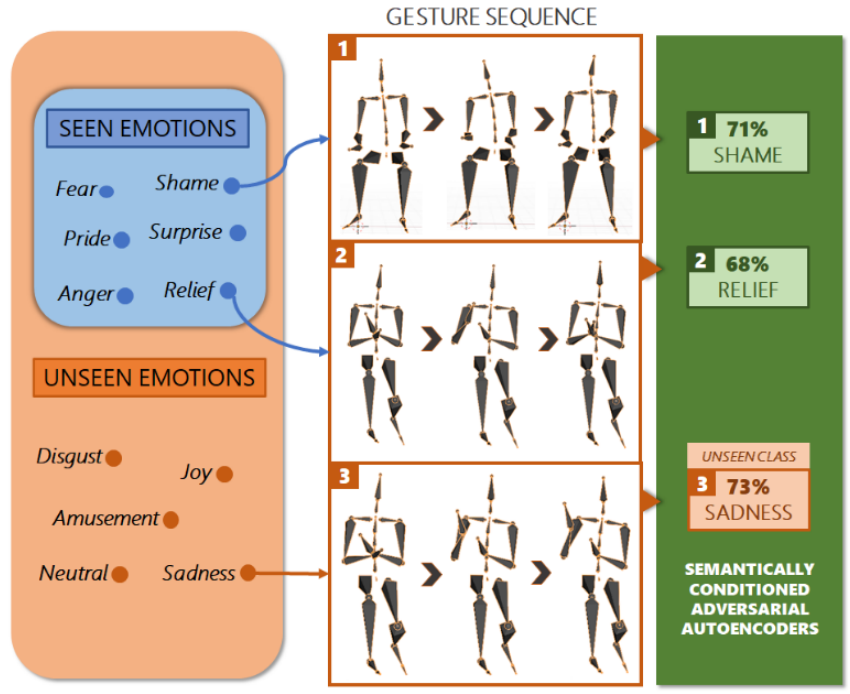

- Banerjee et al., AAAI 2022: "Learning Unseen Emotions from Gestures [...]"
  - Concept vectors: Word2Vec embeddings for emotion names
- Other ideas: Appraisal vectors, vectors learned end-to-end, ...

  (we experimented with that, but did not get any positive results in the generalized ZSL setting)

Another approach to ZSL Emotion Classification

- Recent unpublished work: Chochlakis et al (Oct 2022): Using Emotion Embeddings to Transfer Knowledge between Emotions, Languages, and Annotation Formats. `https://arxiv.org/pdf/2211.00171.pdf`
- Idea: Provide set of emotions at inference time that are to be predicted
- Predefine emotions clusters, neural network predicts cluster embeddings
- Regularize such that similar emotions (according to prior knowledge) are close in parameter space

Alternative: Zero-Shot Learning as Entailment

Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach

Wenpeng Yin, Jamaal Hay, Dan Roth
Cognitive Computation Group
Department of Computer and Information Science, University of Pennsylvania
{wenpeng,jamaalh,danroth}@seas.upenn.edu

Zero-Shot Learning as Entailment (2)

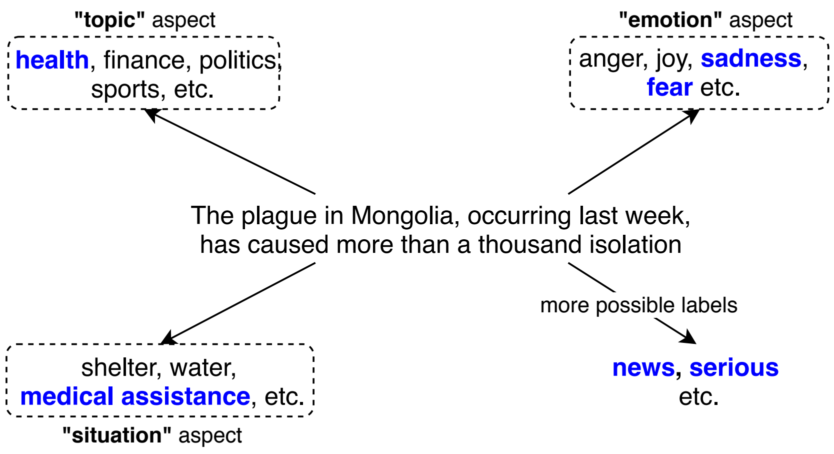

- Input:
  Two sentences, premise and hypothesis
- Output:
  contradiction, entailment, neutral
- Example online demo:
  `https://huggingface.co/microsoft/deberta-large-mnli`

- How to represent the label as a hypothesis?
- Yin et al. use "This text expresses [?]" and the WordNet concept definition.

Emotion ZSL as Natural Language Inference

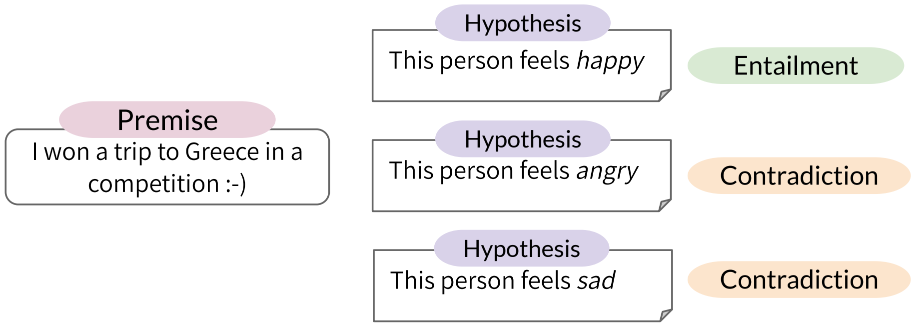

- Does it matter which NLI model we use as a backbone?
- How to represent the emotion?
- Should we use multiple emotion representations to increase coverage?

(Arco Del Plaza et al, COLING 2022)

Emotion Hypotheses

Emo-Name	angry
Expr-Emo	This text expresses anger
Feels-Emo	This person feels anger
WN-Def	This person expresses a strong emotion; a feeling that is oriented toward some real or supposed grievance

Emo-S **Expr.-S** **Feels.-S**	Same prefix + anger, annoyance, rage, outrage, fury, irritation
EmoLex	all emotion words from an NRC emotion lexicon

The role of the prompt design

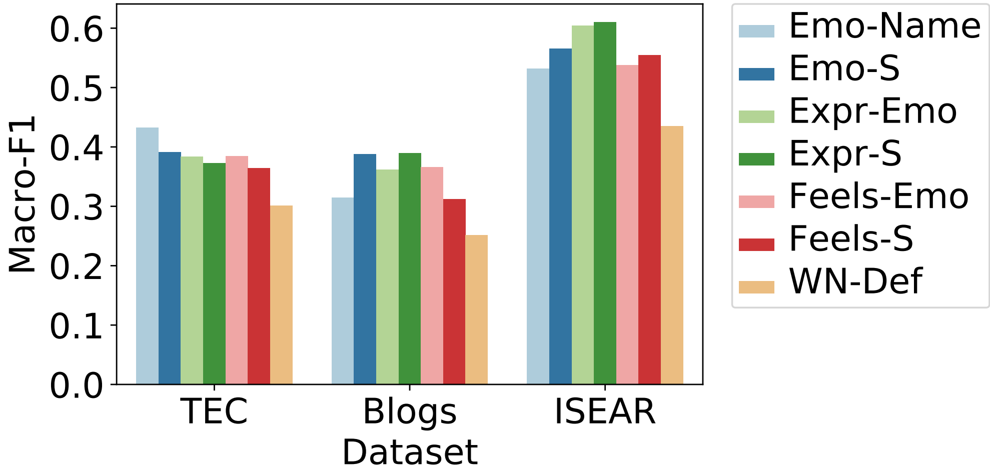

(Supervised RoBERTa model: TEC/Blogs: ≈.69, ISEAR: ≈.73)

- TEC: single emotion names work better than with synonyms
- BLOGS: synonyms harm the performance for Feels-Emo/S prompts
- Generally: synonyms help, except for some cases, in which annotaton procedure might be the reason

Questions?

About this tutorial

Session 1 (09:00–10:30)

- Introduction
- Psychological Models
- Use Cases/Social Impact
- Resources
- Annotation Exercise

Break (10:30–11:15)

Session 2 (11:15–12:45)

- Non-Neural Methods
- Multi-task, transfer, zero-shot methods
- Open Challenges
- Appraisal Theories
- Role Labeling
- Ethical Considerations
- Closing

Universität Stuttgart
Institut für
Maschinelle Sprachverarbeitung

Emotion Analysis

Transfer, Multi-Task Learning, Zero-Shot Predictions

EACL 2023 Tutorial

Sanja Štajner and Roman Klinger

CHALLENGES

Sanja **Štajner** and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

CHALLENGES

- Annotation:
  - Natural difficulty of the task
  - Missing context/knowledge
  - Linguistic difficulty
  - Various emotions present in the instance
  - Quality of annotations
  - Consistency of annotations

- Comparison of different approaches (What is s.o.t.a. in emotion analysis?)

ANNOTATION CHALLENGES: NATURAL DIFFICULTY

- “2 pretty sisters are dancing with cancered kid” (fear+sadness, joy+sadness) (Schuff et al., 2017)
- “That moment when Canadians realised global warming doesn’t equal a tropical vacation” (anger+sadness; surprise) (Schuff et al., 2017)
- “Relatives here. Hafta sleep on a couch in the basement. #cantsleep #effuguysiwantmyqueensize ” (anger; sadness; neutral) (Štajner, 2021)

ANNOTATION CHALLENGES: MISSING KNOWLEDGE

"At the dentist bright and early " (joy; sadness; neutral) (Štajner, 2021)

"Another evening, another cup of coffee" (joy; sadness; neutral) (Štajner, 2021)

ANNOTATION CHALLENGES: LINGUISTIC DIFFICULTY

NON-LITERAL MEANING

- “Global Warming! Global Warming! Global Warming! Oh wait, it’s summer.” (joy) (Schuff et al., 2017)
- “I love the smell of Hillary in the morning. It smells like Republican Victory” (joy) (Schuff et al., 2017)

ANNOTATION CHALLENGES: VARIOUS EMOTIONS

- “No school, getting up at 8 for a seven hour car ride at least i have #noschool” (joy; sadness) (Štajner, 2021)
- “Finally done with work and have to be back in less than 12 hours” (joy; sadness) (Štajner, 2021)
- “ The movie click is old but one of my favs the ending when he dies makes me tear up” (joy; sadness) (Štajner, 2021)
- “ My team is starting to heat up you can’t contain us too long let the blowout begin ducks attack the duck” (joy; anger; neutral) (Štajner, 2021)

ANNOTATION CHALLENGES: QUALITY OF ANNOTATIONS

- Oversight errors
- Dedication to the task

Example: “#BIBLE = Big Irrelevant Book of Lies and Exaggerations” (trust) (Schuff et al., 2017)

ANNOTATION CHALLENGES: CONSISTENCY

- Emotional perception depends on annotators personality and mood (Alm et al., 2005)

- Inter-annotator agreements are very low:
  - κ = 0.24 – 0.51 (Alm et al., 2005)
  - κ = 0.33 – 0.55 (Štajner, 2021)

Questions?

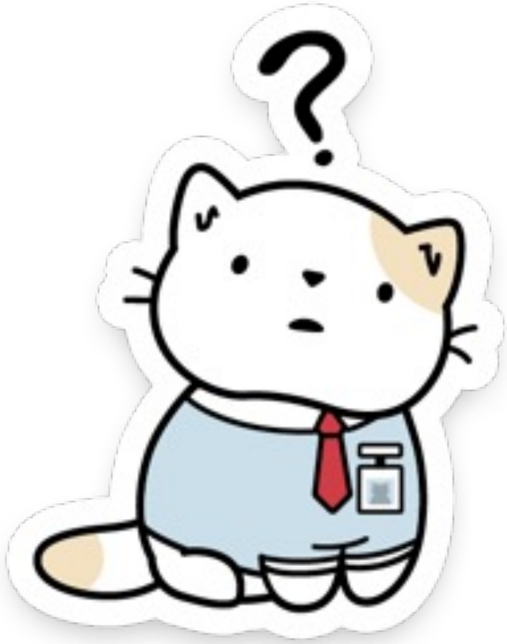

CHALLENGES

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

Universität Stuttgart
Institut für
Maschinelle Sprachverarbeitung

Emotion Analysis

Appraisal-based Resources and Methods

EACL 2023 Tutorial

Sanja Štajner and Roman Klinger

Outline

Emotion Models

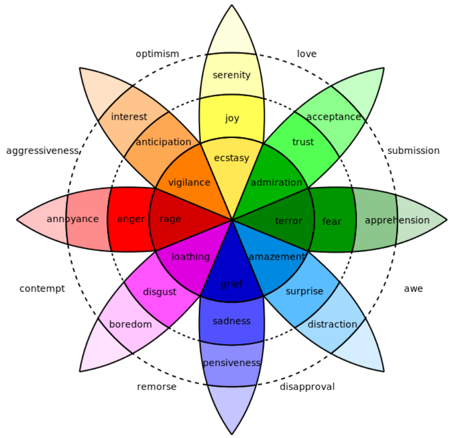

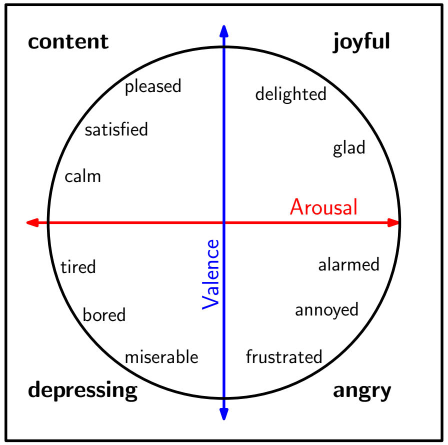

⇒ Methods mostly treat emotions as a label and learn the association to text properties, without considering (too much) knowledge from psychology about emotions

Emotion Components

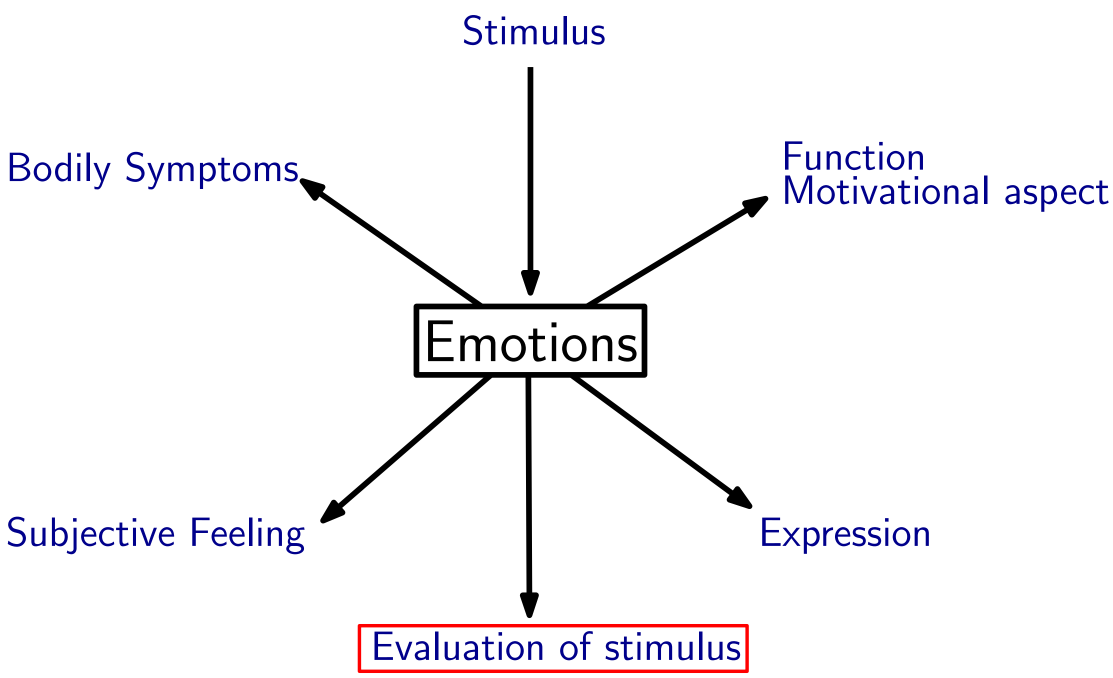

Appraisal Models in Psychology: Smith/Ellsworth and Scherer

Locations of Emotion Means Along the PCA Components

Emotion	Pleasant[a]	Responsibility/ Control[b]	Certain[c]	Attention[d]	Effort[e]	Situational-Control[f]
	Component					
Happiness	−1.46	0.09	−0.46	0.15	−0.33	−0.21
Sadness	0.87	−0.36	0.00	−0.21	−0.14	1.15
Anger	0.85	−0.94	−0.29	0.12	0.53	−0.96
Boredom	0.34	−0.19	−0.35	−1.27	−1.19	0.12
Challenge	−0.37	0.44	−0.01	0.52	1.19	−0.20
Hope	−0.50	0.15	0.46	0.31	−0.18	0.35
Fear	0.44	−0.17	0.73	0.03	0.63	0.59
Interest	−1.05	−0.13	−0.07	0.70	−0.07	0.41
Contempt	0.89	−0.50	−0.12	0.~~80~~ .08	−0.07	−0.63
Disgust	0.38	−0.50	−0.39	−0.96	0.06	−0.19
Frustration	0.88	−0.37	−0.08	0.60	0.48	0.22
Surprise	−1.35	−0.94	0.73	0.40	−0.66	0.15
Pride	−1.25	0.81	−0.32	0.02	−0.31	−0.46
Shame	0.73	1.31	0.21	−0.11	0.07	−0.07
Guilt	0.60	1.31	−0.15	−0.36	0.00	−0.29

Note. Scores are standardized.
[a] Pleasantness: high scores indicate increased unpleasantness.
[b] Responsibility/Control: high scores indicate increased self-responsibility/control.
[c] Certainty: high scores indicate increased uncertainty.
[d] Attentional activity: high scores indicate increased attentional activity.
[e] Effort: high scores indicate increased anticipated effort.
[f] Situational control: high scores indicate increased situational control.

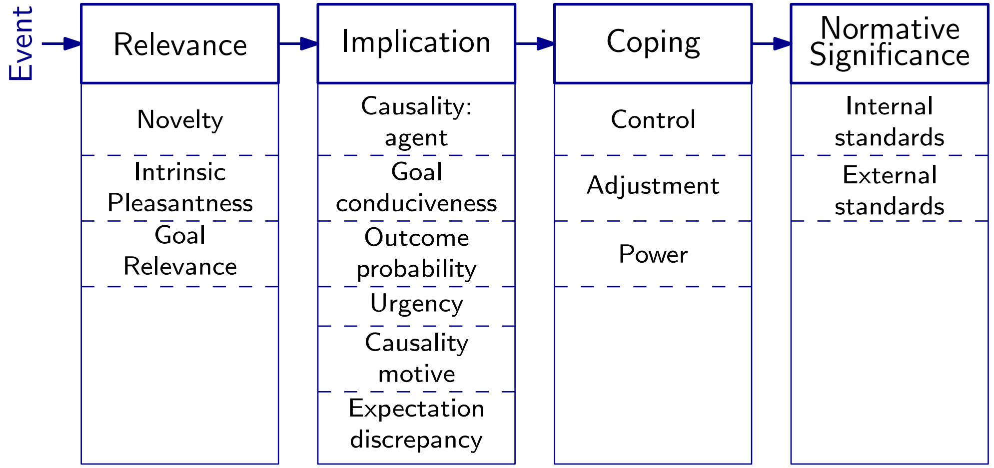

- How to use appraisals in computational modeling?

Outline

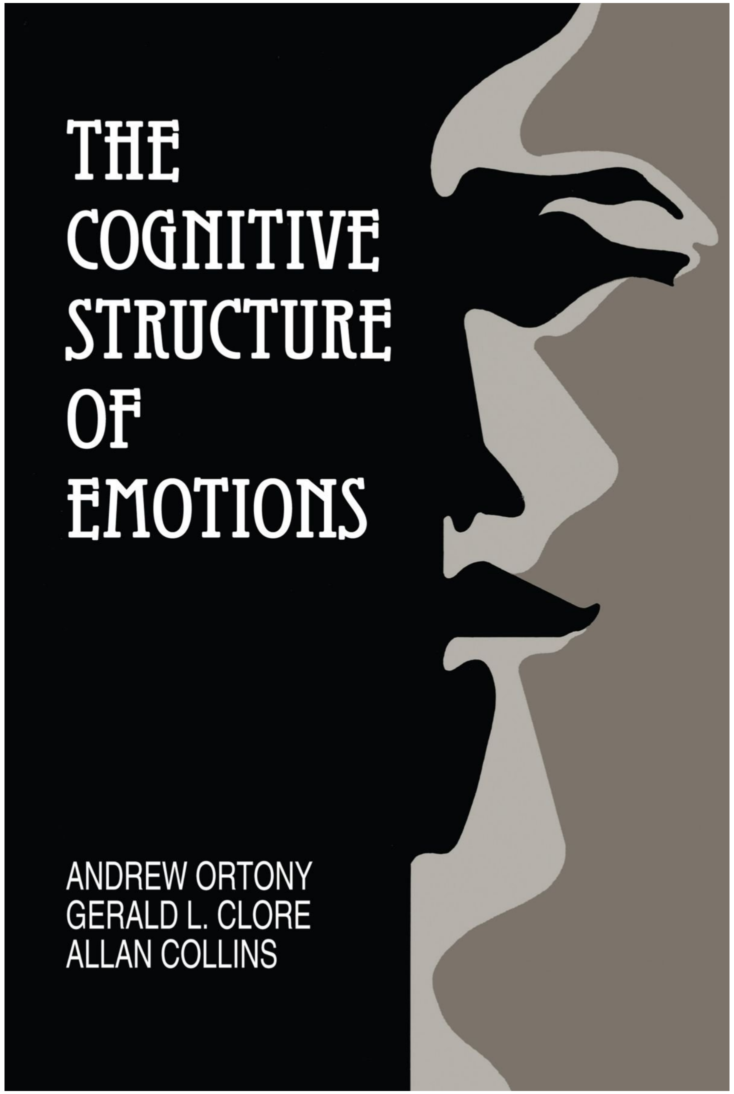

(1988)

Published in final edited form as:
Emot Rev. 2013 October ; 5(4): 335–343. doi:10.1177/1754073913489751.

Psychological Construction in the OCC Model of Emotion

Gerald L. Clore and
Department of Psychology, University of Virginia, USA

Andrew Ortony
Department of Psychology, Northwestern University, USA

Abstract

This article presents six ideas about the construction of emotion: (a) Emotions are more readily distinguished by the situations they signify than by patterns of bodily responses; (b) emotions emerge from, rather than cause, emotional thoughts, feelings, and expressions; (c) the impact of emotions is constrained by the nature of the situations they represent; (d) in the OCC account (the model proposed by Ortony, Clore, and Collins in 1988), appraisals are psychological aspects of situations that distinguish one emotion from another, rather than triggers that elicit emotions; (e) analyses of the affective lexicon indicate that emotion words refer to *internal mental states focused on affect*; (f) the modularity of emotion, long sought in biology and behavior, exists as mental schemas for interpreting human experience in story, song, drama, and conversation.

Warm-Up Example

How to interpret the emotion?

I passed the exam.

past: confirmed

about the self

generally positive event

potentially scary, when unconfirmed

⇒ Relief, Joy

OCC Model

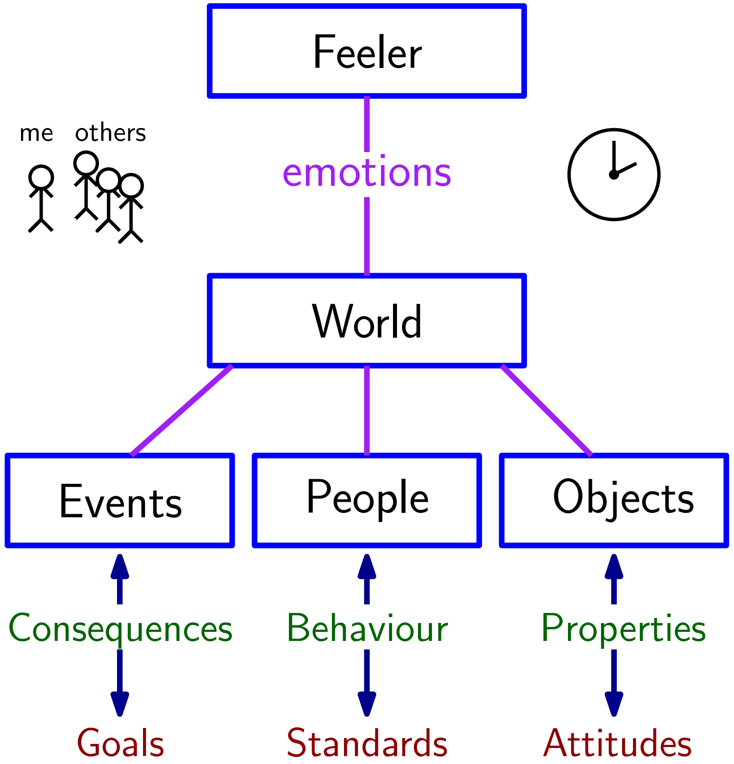

- The OCC Model explains how emotions happen in the interaction of a person and the world
  - The world consists of: Events, People, Objects
- Main components to evaluate the world:
  - Are events in line with goals?
  - Are people behaving in line with standards?
  - Does the person have a positive attitude towards objects?
- Further components
  - Point of view
  - Time

Exercise

- The employee thinks that he might be fired.
- Mary learns that her husband cheated to win in the lottery.

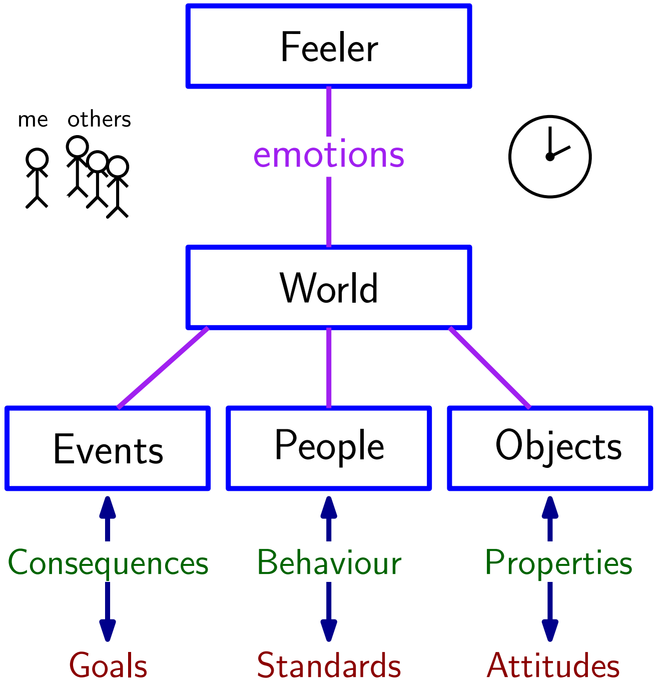

How can we interpret the different components in the OCC?

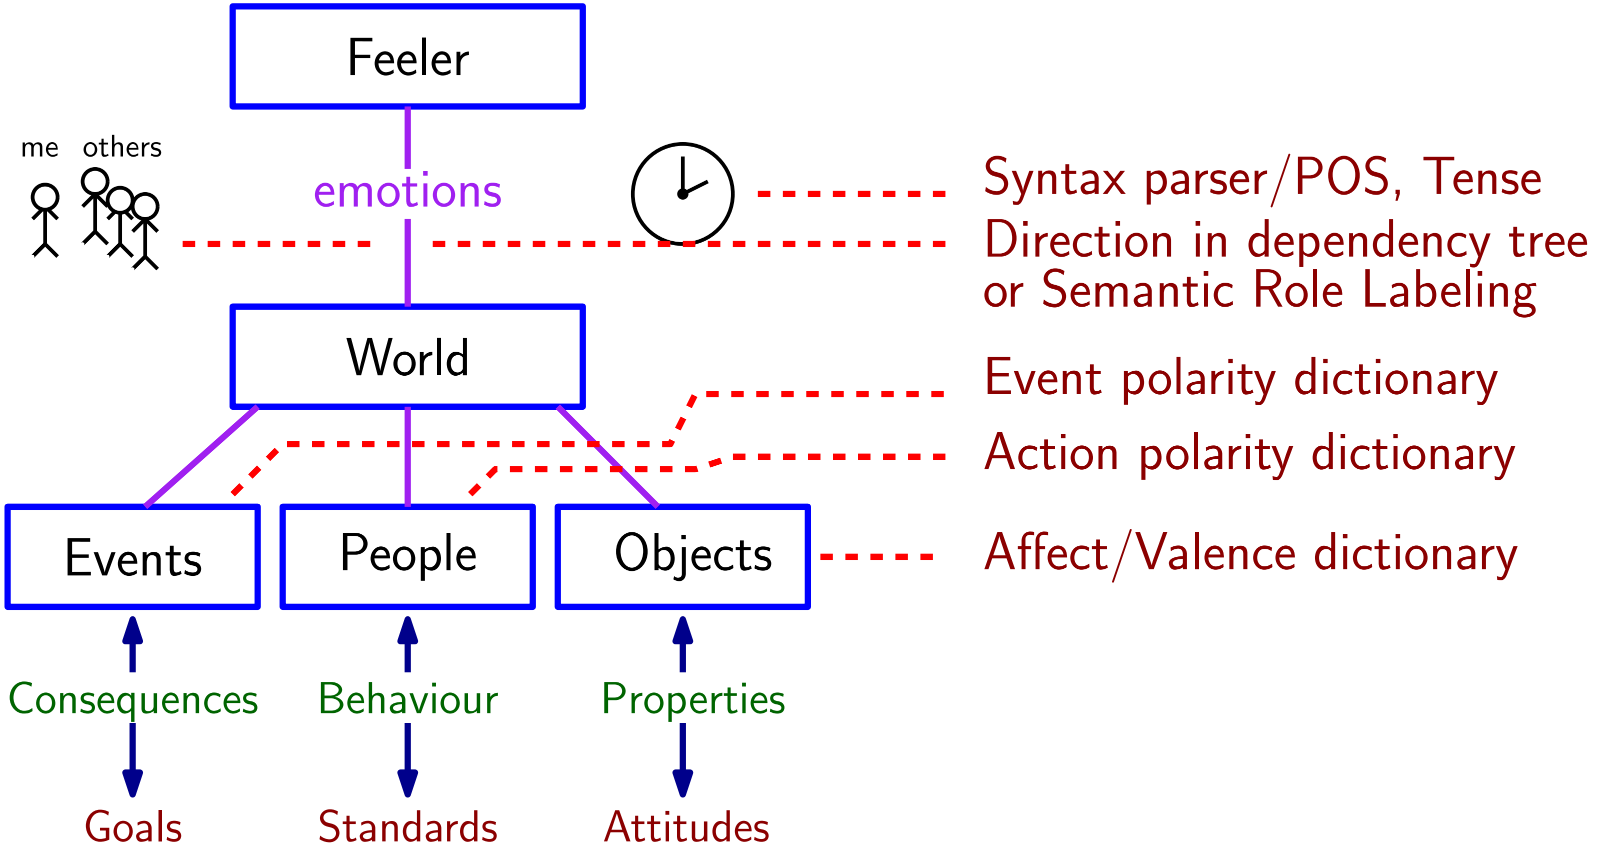

OCC Text Interpretation

Chapter 4
A Linguistic Interpretation of the OCC Emotion Model for Affect Sensing from Text

Mostafa Al Masum Shaikh, Helmut Prendinger, and Mitsuru Ishizuka

Abstract Numerous approaches have already been employed to 'sense' affective information from text; but none of those ever employed the OCC emotion model, an influential theory of the cognitive and appraisal structure of emotion. The OCC model derives 22 emotion types and two cognitive states as consequences of several cognitive variables. In this chapter, we propose to relate cognitive variables of the emotion model to linguistic components in text, in order to achieve emotion recognition for a much larger set of emotions than handled in comparable approaches. In particular, we provide tailored rules for textural emotion recognition, which are inspired by the rules of the OCC emotion model. Hereby, we clarify how text components can be mapped to specific values of the cognitive variables of the emotion model. The resulting linguistics-based rule set for the OCC emotion types and cognitive states allows us to determine a broad class of emotions conveyed by text.

A Rule-Based Approach to Implicit Emotion Detection in Text

Orizu Udochukwu(✉) and Yulan He

School of Engineering and Applied Science, Aston University, Birmingham, UK
{orizuus,y.he9}@aston.ac.uk

Abstract. Most research in the area of emotion detection in written text focused on detecting explicit expressions of emotions in text. In this paper, we present a rule-based pipeline approach for detecting implicit emotions in written text without emotion-bearing words based on the OCC Model. We have evaluated our approach on three different datasets with five emotion categories. Our results show that the proposed approach outperforms the lexicon matching method consistently across all the three datasets by a large margin of 17–30 % in F-measure and gives competitive performance compared to a supervised classifier. In particular, when dealing with formal text which follows grammatical rules strictly, our approach gives an average F-measure of 82.7 % on "Happy", "Angry-Disgust" and "Sad", even outperforming the supervised baseline by nearly 17 % in F-measure. Our preliminary results show the feasibility of the approach for the task of implicit emotion detection in written text.

Keywords: Implicit emotions · OCC model · Emotion detection · Rule-based approach

Example Rules (à la Shaikh)

"The employee thinks that he might be fired."

Variables:

- vr: valenced reaction
  as sentence valence
- sr: self reaction
  valence of event≈ desirability
- pros: prospect
  valence of verb
- sp: self presumption
  valence of event≈ desirability
- status
  tense of verb
- de: direction of emotion
  other if object is person/pronoun

- If (vr = true & sr = 'displeased' & pros = 'negative' & sp = 'undesirable' & status = 'unconfirmed' & de = 'self')
  ⇒ fear

The rules for the emotion are listed as follows.

- If (*vr* = true & *sr* = ‘displeased’ & *sp* = ‘undesirable’ & *de* = ‘self’), ‘distress’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *op* = ‘undesirable’ & *af* = ‘liked’ & *de* = ‘other’), ‘sorry-for’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *op* = ‘desirable’ & *af* = ‘not liked’ & *de* = ‘other’), ‘resentment’ is true.
- If (*vr* = true & *sr* = ‘pleased’ & *op* = ‘undesirable’ & *af* = ‘not liked’ & *de* = ‘other’), ‘gloating’ is true.
- If (*vr* = true & *sr* = ‘pleased’ & *pros* = ‘positive’ & *sp* = ‘desirable’ & *status* = ‘unconfirmed’ & *de* = ‘self’), ‘hope’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *pros* = ‘negative’ & *sp* = ‘undesirable’ & *status* = ‘unconfirmed’ & *de* = ‘self’), ‘fear’ is true.
- If (*vr* = true & *sr* = ‘pleased’ & *pros* = ‘positive’ & *sp* = ‘desirable’ & *status* = ‘confirmed’ & *de* = ‘self’), ‘satisfaction’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *pros* = ‘negative’ & *sp* = ‘undesirable’ & *status* = ‘confirmed’ & *de* = ‘self’), ‘fears-confirmed’ is true.
- If (*vr* = true & *sr* = ‘pleased’ & *pros* = ‘negative” & *sp* = ‘undesirable & *status* = ‘disconfirmed’ & *de* = ‘self’), ‘relief’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *pros* = ‘positive’ & *sp* = ‘desirable’ & *status* = ‘disconfirmed‘ & *de* = ‘self’), ‘disappointment’ is true.
- If (*vr* = true & *sr* = ‘pleased’ & *sa* = ‘praiseworthy’ & *sp* = ‘desirable’ & *de* = ‘self’), ‘pride’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *sa* = ‘blameworthy’ & *sp* = ‘undesirable’ & *de* = ‘self’), ‘shame’ is true.
- If (*vr* = true & *sr* = ‘pleased’ & *sa* = ‘praiseworthy’ & *op* = ‘desirable’ & *de* = ‘other’), ‘admiration’ is true.
- If (*vr* = true & *sr* = ‘displeased’ & *sa* = ‘blameworthy’ & *op* =
- ‘undesirable’ & *de* = ‘other’), ‘reproach’ is true.
- If (*vr* = true & *sp* = ‘desirable’ & *sr* = ‘pleased’ & *of* = ‘liked’ & *oa* = ‘attractive’ & event valence = ‘positive’ & *de* = ‘other’), ‘love’ is true.
- If (*vr* = true & *sp* = ‘undesirable’ & *sr* = ‘displeased’ & *of* = ‘not liked’ & *oa* = ‘not attractive’ & event valence= ‘negative’ & *de*= ‘other’), ‘hate’ is true.

The OCC model has four complex emotions, namely, ‘gratification,’ ‘remorse,’ ‘gratitude,’ and ‘anger.’ The rules for these emotions are as follows.

- If both ‘joy’ and ‘pride’ are true, ‘gratification’ is true.
- If both ‘distress’ and ‘shame’ are true, ‘remorse’ is true.
- If both ‘joy’ and ‘admiration’ are true, ‘gratitude’ is true.
- If both ‘distress’ and ‘reproach’ are true, ‘anger’ is true.

The cognitive states ‘shock’ and ‘surprise’ are ruled as follows.

- If both ‘distress’ and *unexp* are true, ‘shock’ is true (e.g., the bad news came unexpectedly).
- If both ‘joy’ and *unexp* are true, ‘surprise’ is true (e.g., I suddenly met my school friend in Tokyo University).

Results (Udochukwu/He 2015)

Emotion	ISEAR			SemEval			Alm's		
	Lexicon	NB	Rule	Lexicon	NB	Rule	Lexicon	NB	Rule
Joy/Happy	33.4	61.2	**69.6**	39.7	**71.7**	59.9	58.8	63.5	**81.8**
Fear/Fearful	0	**47.6**	18.3	0	**52.2**	31.8	0	**26.7**	14.0
Anger/Angry-Disgusted	23.0	47.1	**61.3**	55.8	16.2	**61.3**	48.9	58.6	**86.6**
Sadness/Sad	25.6	55.4	**68.0**	47.8	56.0	**71.5**	61.0	56.0	**79.6**
Disgust	25.6	**51.0**	39.2	38.5	34.5	**61.7**	-	-	-
Average	21.5	**52.5**	51.3	36.4	**58.2**	57.3	42.2	56.0	**65.5**
Average (− Fear)	27.0	53.7	**59.5**	45.5	44.6	**63.6**	56.12	65.8	**82.7**

Outline

Appraisal Prediction following Scherer's Model

Relevance	Implication	Coping	Normative Significance
Novelty (1) suddenness (2) familiarity (3) predictability (16) attention* (17) att. removal* Intrinsic Pleasantness (4) pleasant (5) unpleasant Goal Relevance (6) goal-related	Causality: agent (7) own responsibility (8) other's respons. (9) situational respons. Goal conduciveness (10) goal support Outcome probability (11) consequence anticipation Urgency (12) response urgency	Control (19) own control* (20) others' control* (21) chance control* Adjustment (13) anticipated acceptance (18) effort*	Internal standards compatibility (14) clash with own standards/ideals External standards compatibility (15) clash with laws/norms

Troiano, Oberlaender, Klinger, MIT CL 2023: Dimensional Modeling of Emotions in Text with Appraisal Theories: Corpus Creation, Annotation Reliability, and Prediction.

- Can appraisals be annotated reliably?
- Do appraisals help emotion categorization?

Approach

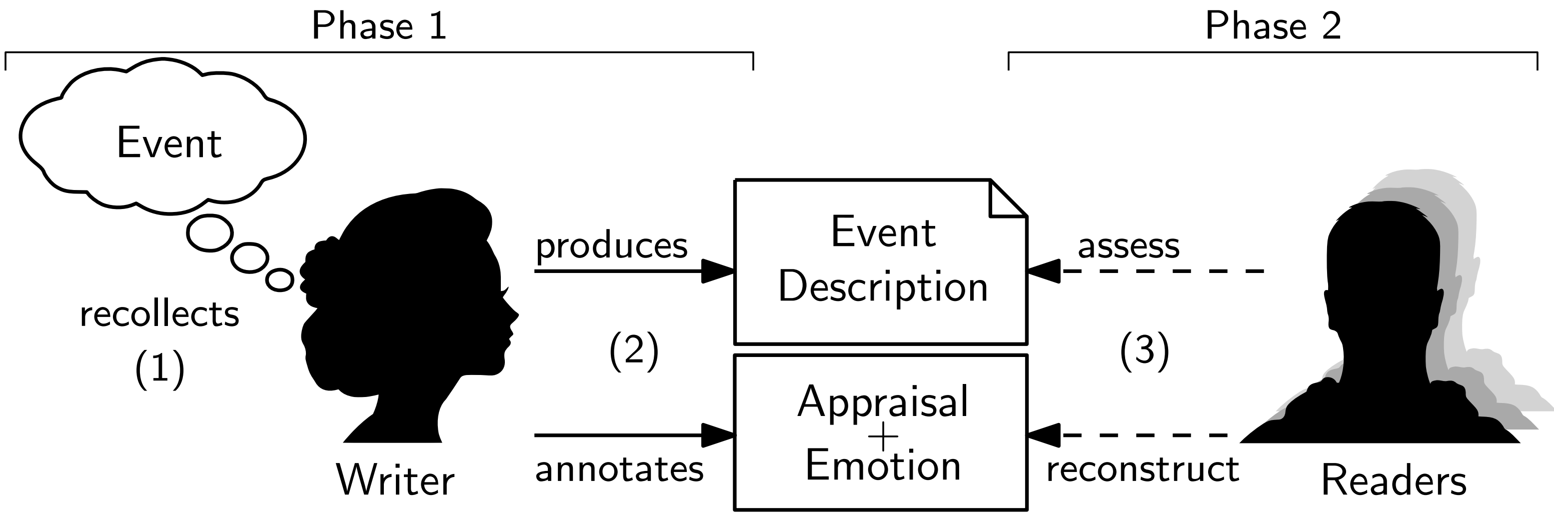

- Production: 550 event descriptions for anger, boredom, disgust, fear, guilt/shame, joy, pride, relief, sadness, surprise, trust, no emotion

Questions and Answers

- Do readers agree more with each other than with the writers?
  (does the writer make use of information that the readers do not have)
  - Yes, a bit for emotions; clearly for the appraisals.
- Does it matter if annotators share demographic properties?
  - Females agree more with each other, but men less.
  - People of similar age agree more.
- Does personality matter?
  - Extraverted, conscientious, agreeable annotators perform better.

Setup:

- Filter instances for attribute, compare with F_1/RMSE
- Significance test with bootstrap resampling for .95 confidence interval

Appraisals add additional information to emotion analysis

“That I put together a funeral service for my Aunt”

Dimension	Writer	Readers	△
Emotion	Pride	Sadness	
Suddenness	4	3.6	0.4
Familiarity	1	2.0	−1.0
Predictability	1	1.8	−0.8
Pleasantness	4	1.0	3.0
Unpleasantness	2	4.8	−2.8
Goal-Relevance	4	2.6	1.4
Chance-Resp.	4	4.4	−0.4
Self-Resp.	1	1.2	−0.2
Other-Resp.	1	1.4	−0.4
Conseq.-Predict.	2	1.8	0.2
Goal Support	1	1.2	−0.2
Urgency	2	3.8	−1.8
Self-Control	5	3.2	1.8
Other-Control	3	2.0	1.0
Chance-Control	1	4.6	−3.6
Accept-Conseq.	4	2.4	1.6
Standards	1	2.4	−1.4
Social Norms	1	1.2	−0.2
Attention	4	4.4	−0.4
Not-Consider	1	3.8	−2.8
Effort	4	4.6	−0.6

Modeling Results

- Classification with RoBERTa-based models
- Appraisal Classification: 75 F_1
- Emotion classification: 59 F_1
- $+$ Appraisals: $+$2pp F_1
  ($+$10 for guilt, $+$6 for sadness)

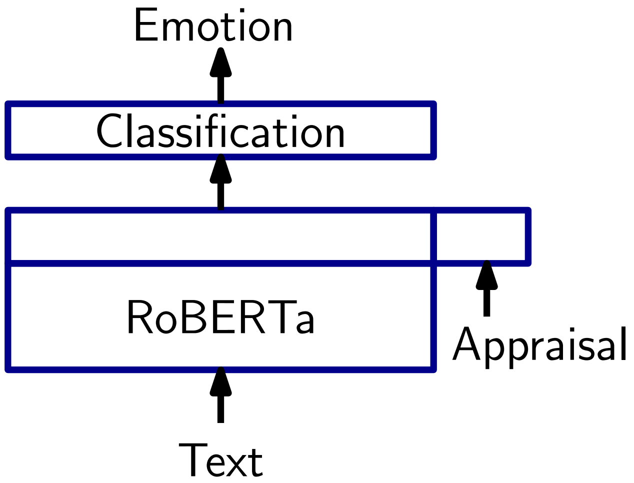

Examples where Appraisals correct the Emotion Classifier

- When my child settled well into school

  trust→relief

- broke an expensive item in a shop accidently

  guilt→shame

- my mother made me feel like a child

  shame→anger

- I passed my Irish language test

  pride→relief

- His toenails where massive

  pride→disgust

Outline

Other Approaches

- Balahur et al., 2011, EmotiNet:
  Knowledge base of events motivated by appraisal theories
- Stranisci et al., 2022, APPReddit:
  Reddit post corpus, focus on coping strategies
- Hofmann et al., 2020:
  Appraisal-based Emotion Analysis, annotated corpus for Smith/Ellsworth concepts

Take-Away

- Appraisal dimensions are an additional emotion model that serves as a fundamental for analysis in text
- It provides additional knowledge and supports the categorization into emotion concepts
- Could it support affect (valence/arousal) prediction? Not yet known.

Questions?

About this tutorial

Session 1 (09:00–10:30)

- Introduction
- Psychological Models
- Use Cases/Social Impact
- Resources
- Annotation Exercise

Break (10:30–11:15)

Session 2 (11:15–12:45)

- Non-Neural Methods
- Multi-task, transfer, zero-shot methods
- Open Challenges
- Appraisal Theories
- Role Labeling
- Ethical Considerations
- Closing

Universität Stuttgart
Institut für
Maschinelle Sprachverarbeitung

Emotion Analysis

Role Labeling and Stimulus Detection

EACL 2023 Tutorial

Sanja Štajner and Roman Klinger

Outline

Motivation (1)

What cannot be done with document/sentence-level emotion analysis?

Given a corpus, extract the information:

who is the target? who experiences the emotion?

Couple infuriated officials by landing helicopter in nature reserve

which words describe the emotion? what caused the emotion?

- Relevancy: Social media mining, literature analysis, network analysis, ...

Relation to Aspect-based sentiment analysis

Formulation 1:

- Closed set of aspects, classify polarity for each aspect.
- "The food was good, but the waiter was unfriendly.
  food → positive; staff → negative.
- e.g., Ganu et al. (2009). "Beyond the Stars: Improving Rating Predictions using Review Text Content."

Formulation 2:

- Given text, detect phrases that describe an aspect.
- Classify these aspects into sentiment polarities.
- "The food^{+} was good, but the waiter^{-} was unfriendly.
- e.g., Kessler et al. 2010. The 2010 ICWSM JDPA Sentiment Corpus for the Automotive Domain.

Task Definition: Relations, spans, or clauses?

- Relation detection:

cue stimulus

target experiencer

A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

- Sequence labeling:

A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

target cue experiencer stimulus

- Clause classification:

A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

emotion clause cause/stimulus clause

→ trade-off between task complexity and accurateness

Outline

Corpora: SRL4E

Resource	Original	SRL4E	%
Blogs	5,202	4,855	93.3
Elections	1,385	1,024	73.9
EmoTweet	15,553	15,553	100.0
GNE	5,000	5,000	100.0
NTCIR (ZH)	2,022	1,956	96.7
NTCIR (EN)	1,826	1,796	98.4
REMAN	1,720	1,705	99.1
All	32,708	31,889	97.5

Resource	cue	stim.	exp.	targ.
Blogs	✔	–	–	–
Elections	✔	✔	✔	✔
EmoTweet	✔	–	–	–
GNE	✔	✔	✔	✔
NTCIR	✔	✔	–	–
REMAN	✔	✔	✔	✔

- Campagnano et al., ACL 2022 aggregate a set of corpora into common format and conduct prediction experiments for the identification of all roles
- `https://github.com/sapienzanlp/srl4e`

Corpora

	Whole Instance		Stimulus	
Dataset	#	avg. len	#	avg. len
ES, Ghazi2015	2414	20.60	820	7.29
ET, Mohammad2014	4056	19.14	2427	6.25
GNE, Bostan2020	5000	13.00	4798	7.29
REMAN, Kim2018	1720	72.03	609	9.33
ECA, Gao2017	2558	62.24	2485	9.52

	Cue		Target		Exp.	
Dataset	#	avg. len	#	avg. len	#	avg. len
ET	2930	5.08	2824	1.71	29	1.76
GNE	4736	1.60	4474	4.86	3458	2.03
REMAN	1720	3.82	706	5.35	1050	2.04

Oberlaender et al. (2020), Experiencers, Stimuli, or Targets: Which Semantic Roles Enable Machine Learning to Infer the Emotions? PEOPLES

Corpus Examples (1)

- Mohammad et al. (2014). Semantic role labeling of emotions in tweets.
    - Crowdsourced span annotations in electoral Tweets
    - Modeling as stimulus classification task
- Ghazi et al. (2015). Detecting emotion stimuli in emotion-bearing sentences.
    - Expert-based span annotations in FrameNet sentences
    - Modeling span-based with feature-based CRF
- Kim/Klinger (2018). Who feels what and why? Annotation of a literature corpus with semantic roles of emotions.
    - Expert-annotated role graph in sentence triples of literature.
    - Modeling span-based with BiLSTM+CRF

Corpus Examples (2)

- Bostan et al. (2020). GoodNewsEveryone: A Corpus of News Headlines Annotated with Emotions, Semantic Roles, and Reader Perception.
  - Crowdsourced annotation of full graph.
  - Modeling span-based with ELMo+BiLSTM+CRF
- Gao et al. (2017). Overview of NTCIR-13 ECA task; Xia (2019). Emotion-Cause Pair Extraction: A New Task to Emotion Analysis in Texts.
  - Annotation of emotion and stimulus clauses
  - Modeling as clause classification

Examples: Emotion Stimulus

- happy: I suppose I am happy being so ‘ tiny’ ; it means I am able to surprise people with what is generally seen as my confident and outgoing personality .
- sad: Anne was sad at the death of the Misses Dolan but too much was happening for her to dwell on it .
- anger: I was very very angry to read Batty 's comments about Leeds .

Examples: REMAN

cause experiencer

When I mentioned the house, he seemed surprised.

event character surprise

experiencer target target experiencer

All laughed at the mistake, and none louder than the forth member of the parliament . . .

character disgust joy other strong joy character

Examples: Good News Everyone

Headline: *A couple infuriated officials by landing their helicopter in the middle of a nature reserve.*

phase 1

Emotion: Anger, Anger, Disgust
Reader Perception: Yes, No, Yes

phase 2

Emotion: Anger, Anger, Disgust
Intensity: Medium, High, High
Other emotions: None, None, None
Reader emotions: Annoyance, Negative Surprise, No Emotion

Experiencer: A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

Cue: A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

Cause: A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

Target: A couple infuriated officials by landing their helicopter in the middle of a nature reserve.

aggregated

Emotion: Anger
Intensity: High
Other emotions: None
Reader perception: Yes
Reader emotions: Annoyance, Negative Surprise, No Emotion

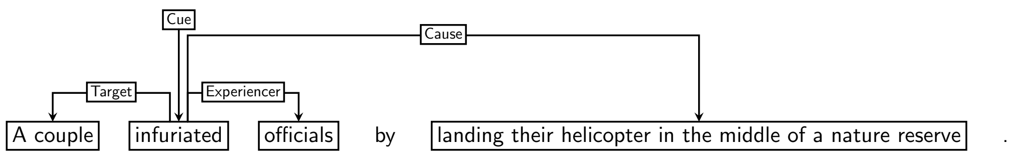

Examples: ECPE and ECA

Document

Yesterday morning, a policeman visited the old man with the lost money, and told him that the thief was caught. The old man was very happy, and deposited the money in the bank.

Emotion Cause Extraction (ECE)

happy → a policeman visited the old man with the lost money

happy → and told him that the thief was caught

Emotion-Cause Pair Extraction (ECPE)

(The old man was very happy, a policeman visited the old man with the lost money)

(The old man was very happy, and told him that the thief was caught)

ECPE – Modeling

- Attracted a lot of attention
- Often two steps:
  - (1) detect emotion (clauses) and cause clauses separately
  - (2) pair emotion and cause
- Example for one approach which does end-to-end modeling:
  Wei, Zhao, Mao. ACL 2020.
- Oberländer/Klinger *SEM 2020 compared clause classification and sequence labeling settings for English corpora: task formulation seems to be appropriate for Mandarin, but not for English.

Outline

Take Home

- Quite some work on clause classification and sequence labeling
- Nearly (?) no work on full graph reconstruction
- No work on linking stimulus detection with appraisal analysis

Questions?

About this tutorial

Session 1 (09:00–10:30)

- Introduction
- Psychological Models
- Use Cases/Social Impact
- Resources
- Annotation Exercise

Break (10:30–11:15)

Session 2 (11:15–12:45)

- Non-Neural Methods
- Multi-task, transfer, zero-shot methods
- Open Challenges
- Appraisal Theories
- Role Labeling
- Ethical Considerations
- Closing

ETHICAL CONSIDERATIONS

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

ETHICAL CONSIDERATIONS: DISCUSSION

- Privacy
- Failure modes and their consequences
- Who should be responsible?

ETHICAL CONSIDERATIONS: FURTHER READING

- Gremsl and Hödl. 2022. "Emotional AI: Legal and ethical challenges":

https://www.researchgate.net/publication/360210704_Emotional_AI_Legal_and_ethical_challenges

- Stark and Hoey. 2021. "The Ethics of Emotion in Artificial Intelligence Systems":

https://dl.acm.org/doi/10.1145/3442188.3445939

- Brian Green. 2016. "Social Robots, AI, and Ethics":

https://www.scu.edu/ethics/focus-areas/technology-ethics/resources/social-robots-ai-and-ethics/

Questions?

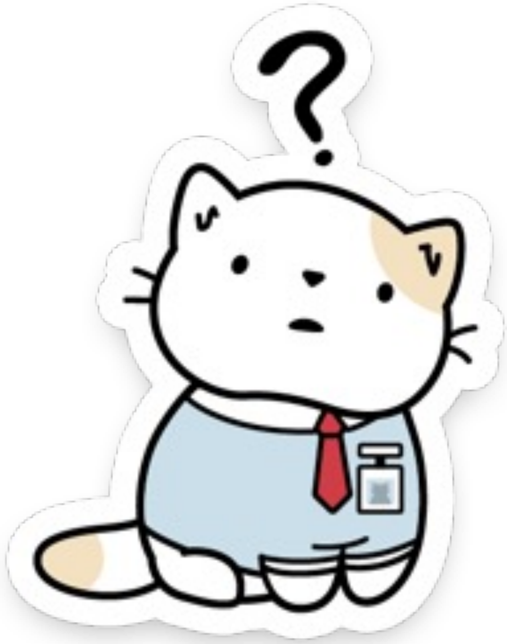

ETHICAL CONSIDERATIONS

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

CLOSING

- TOPICS COVERED:
  - Emotions in psychology
  - Use cases
  - Resources for emotion analysis in texts
  - Computational approaches to emotion analysis in texts
  - Challenges
  - Ethical considerations
- TOPICS NOT COVERED (only mentioned):
  - Emotion analysis from audio or video sequences
  - Multimodal emotion analysis
  - Resources for languages other than English
  - Universality of emotions

Contact:

Sanja Štajner:
stajner.sanja@gmail.com

Roman Klinger:
roman.klinger@ims.uni-stuttgart.de

Sanja Štajner and Roman Klinger

Emotion Analysis in Text

EACL 2023 Tutorial – 05.05.2023

Bibliography

Chochlakis, G., G. Mahajan, S. Baruah, K. Burghardt, K. Lerman, and S. Narayanan (2023). “Using Emotion Embeddings to Transfer Knowledge between Emotions, Languages, and Annotation Formats”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. DOI: 10.1109/ICASSP49357.2023.10095597.

Troiano, E., L. Oberländer, and R. Klinger (2023). “Dimensional Modeling of Emotions in Text with Appraisal Theories: Corpus Creation, Annotation Reliability, and Prediction”. In: *Computational Linguistics* 49.1. DOI: 10.1162/coli_a_00461.

Campagnano, C., S. Conia, and R. Navigli (2022). “SRL4E – Semantic Role Labeling for Emotions: A Unified Evaluation Framework”. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, pp. 4586–4601. DOI: 10.18653/v1/2022.acl-long.314.

Gremsl, T., E. Hödl, J. Čas, P. De Hert, M. G. Porcedda, and C. D. Raab (2022). “Emotional AI: Legal and Ethical Challenges1”. In: *Info. Pol.* 27.2, pp. 163–174. DOI: 10.3233/IP-211529.

Plaza-del-Arco, F. M., M.-T. Martin-Valdivia, and R. Klinger (2022). “Natural Language Inference Prompts for Zero-shot Emotion Classification in Text across Corpora”. In: *Proceedings of the 29th International Conference on Computational Linguistics*. Gyeongju, Republic of Korea: International Committee on Computational Linguistics, pp. 6805–6817. URL: https://aclanthology.org/2022.coling-1.592.

Shanthi, N., A. A. Stonier, A. Sherine, T. Devaraju, S. Abinash, R. Ajay, V. Arul Prasath, and V. Ganji (2022). *An integrated approach for mental health assessment using emotion analysis and scales*. Healthcare Technology Letters. DOI: 10.1049/htl2.12040.

Stranisci, M. A., S. Frenda, E. Ceccaldi, V. Basile, R. Damiano, and V. Patti (June 2022). “APPReddit: a Corpus of Reddit Posts Annotated for Appraisal”. In: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, pp. 3809–3818. URL: https://aclanthology.org/2022.lrec-1.406.

Banerjee, A., U. Bhattacharya, and A. Bera (2021). *Learning Unseen Emotions from Gestures via Semantically-Conditioned Zero-Shot Perception with Adversarial Autoencoders*. arXiv:2009.08906 [cs]. URL: http://arxiv.org/abs/2009.08906.

Casel, F., A. Heindl, and R. Klinger (2021). “Emotion Recognition under Consideration of the Emotion Component Process Model”. In: *Proceedings of the 17th Conference on Natural Language Processing (KONVENS 2021)*. Düsseldorf, Germany: KONVENS 2021 Organizers, pp. 49–61. URL: https://aclanthology.org/2021.konvens-1.5.

Doellinger, L., P. Laukka, L. B. Högman, T. Bänziger, I. Makower, H. Fischer, and S. Hau (2021). “Training Emotion Recognition Accuracy: Results for Multimodal Expressions and Facial Micro Expressions”. In: *Frontiers in Psychology* 12. DOI: 10.3389/fpsyg.2021.708867. URL: https://www.frontiersin.org/article/10.3389/fpsyg.2021.708867.

Stajner, S. (2021). “Exploring Reliability of Gold Labels for Emotion Detection in Twitter”. In: *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*. Held Online: INCOMA Ltd., pp. 1350–1359. URL: https://aclanthology.org/2021.ranlp-1.151.

Štajner, S. (2021). “Exploring Reliability of Gold Labels for Emotion Detection in Twitter”. In: *Proceedings of the 13th international conference on Recent Advances in Natural Language Processing (RANLP)*, pp. 1350–1359.

Stark, L. and J. Hoey (2021). “The Ethics of Emotion in Artificial Intelligence Systems”. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’21. Virtual Event, Canada: Association for Computing Machinery, pp. 782–793. DOI: 10.1145/3442188.3445939.

Bostan, L. A. M., E. Kim, and R. Klinger (2020). “GoodNewsEveryone: A Corpus of News Headlines Annotated with Emotions, Semantic Roles, and Reader Perception”. In: *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, pp. 1554–1566. URL: https://aclanthology.org/2020.lrec-1.194.

Chauhan, D. S., D. S R, A. Ekbal, and P. Bhattacharyya (2020). “Sentiment and Emotion help Sarcasm? A Multi-task Learning Framework for Multi-Modal Sarcasm, Sentiment and Emotion Analysis”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, pp. 4351–4360. DOI: 10.18653/v1/2020.acl-main.401.

Demszky, D., D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. Ravi (2020). “GoEmotions: A Dataset of Fine-Grained Emotions”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, pp. 4040–4054. DOI: 10.18653/v1/2020.acl-main.372.

Hofmann, J., E. Troiano, K. Sassenberg, and R. Klinger (2020). “Appraisal Theories for Emotion Classification in Text”. In: *Proceedings of the 28th International Conference on Computational Linguistics*. Barcelona, Spain (Online): International Committee on Computational Linguistics, pp. 125–138. DOI: 10.18653/v1/2020.coling-main.11.

Loureiro, M. L. and M. Alló (2020). “Sensing climate change and energy issues: Sentiment and emotion analysis with social media in the U.K. and Spain”. In: *Energy Policy* 143, p. 111490. DOI: 10.1016/j.enpol.2020.111490.

Oberländer, L. A. M. and R. Klinger (2020). “Token Sequence Labeling vs. Clause Classification for English Emotion Stimulus Detection”. In: *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*. Barcelona, Spain (Online): Association for Computational Linguistics, pp. 58–70. URL: https://aclanthology.org/2020.starsem-1.7.

Öhman, E., M. Pàmies, K. Kajava, and J. Tiedemann (2020). “XED: A Multilingual Dataset for Sentiment Analysis and Emotion Detection”. In: *Proceedings of the 28th International Conference on Computational Linguistics*. Barcelona, Spain (Online): International Committee on Computational Linguistics, pp. 6542–6552. DOI: 10.18653/v1/2020.coling-main.575.

Rajamanickam, S., P. Mishra, H. Yannakoudakis, and E. Shutova (2020). “Joint Modelling of Emotion and Abusive Language Detection”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, pp. 4270–4279. DOI: 10.18653/v1/2020.acl-main.394.

Saha, T., A. Patra, S. Saha, and P. Bhattacharyya (2020). “Towards Emotion-aided Multi-modal Dialogue Act Classification”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, pp. 4361–4372. DOI: 10.18653/v1/2020.acl-main.402.

Wei, P., J. Zhao, and W. Mao (2020). “Effective Inter-Clause Modeling for End-to-End Emotion-Cause Pair Extraction”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, pp. 3171–3181. DOI: 10.18653/v1/2020.acl-main.289.

Akhtar, M. S., D. Chauhan, D. Ghosal, S. Poria, A. Ekbal, and P. Bhattacharyya (2019). “Multi-task Learning for Multi-modal Emotion Recognition and Sentiment Analysis”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, pp. 370–379. DOI: 10.18653/v1/N19-1034.

Dankers, V., M. Rei, M. Lewis, and E. Shutova (2019). “Modelling the interplay of metaphor and emotion through multitask learning”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, pp. 2218–2229. DOI: 10.18653/v1/D19-1227.

Poria, S., D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea (2019). “MELD: A Multimodal Multi-Party Dataset for Emotion Recognition in Conversations”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, pp. 527–536. DOI: 10.18653/v1/P19-1050. URL: https://www.aclweb.org/anthology/P19-1050.

Rashkin, H., E. M. Smith, M. Li, and Y.-L. Boureau (2019). “Towards Empathetic Open-domain Conversation Models: A New Benchmark and Dataset”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, pp. 5370–5381. DOI: 10.18653/v1/P19-1534. URL: https://www.aclweb.org/anthology/P19-1534.

Rodríguez, A., C. Argueta, and Y.-L. Chen (2019). “Automatic Detection of Hate Speech on Facebook Using Sentiment and Emotion Analysis”. In: *2019 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*, pp. 169–174. DOI: 10.1109/ICAIIC.2019.8669073.

Srinivasan, S. M., R. S. Sangwan, C. J. Neill, and T. Zu (2019a). “Twitter Data for Predicting Election Results: Insights from Emotion Classification”. In: *IEEE Technology and Society Magazine* 38.1, pp. 58–63. DOI: 10.1016/j.enpol.2020.111490.

Srinivasan, S. M., R. S. Sangwan, C. J. Neill, and T. Zu (2019b). “Power of Predictive Analytics: Using Emotion Classification of Twitter Data for Predicting 2016 US Presidential Elections”. In: *Social media and society* 8, pp. 211–230.

Wang, S. and X. Chen (2019). “Recognizing CEO personality and its impact on business performance: Mining linguistic cues from social media”. In: *Information & Management*. DOI: 10.1016/j.im.2019.103173.

Yin, W., J. Hay, and D. Roth (2019). “Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, pp. 3914–3923. DOI: 10.18653/v1/D19-1404.

Fortuna, P. and S. Nunes (2018). “A Survey on Automatic Detection of Hate Speech in Text”. In: *ACM Comput. Surv.* 51.4. DOI: 10.1145/3232676.

Hsu, C.-C., S.-Y. Chen, C.-C. Kuo, T.-H. Huang, and L.-W. Ku (2018). “EmotionLines: An Emotion Corpus of Multi-Party Conversations”. In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki, Japan: European Language Resources Association (ELRA). URL: https://aclanthology.org/L18-1252.

Islam, M. R., A. Kabir, A. Ahmed, A. Kamal, H. Wang, and A. Ulhaq (2018a). “Depression detection from social network data using machine learning techniques”. In: *Health Information Science and Systems* 6, p. 8. DOI: 10.1007/s13755-018-0046-0.

Islam, M. R., M. A. Kabir, A. Ahmed, A. R. M. Kamal, H. Wang, and A. Ulhaq (2018b). “Depression detection from social network data using machine learning techniques”. In: *Health Inf. Sci. Syst.* 6.1, p. 8.

Kim, E. and R. Klinger (2018). “Who Feels What and Why? Annotation of a Literature Corpus with Semantic Roles of Emotions”. In: *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, pp. 1345–1359. URL: https://aclanthology.org/C18-1114.

Klinger, R., O. De Clercq, S. Mohammad, and A. Balahur (2018). “IEST: WASSA-2018 Implicit Emotions Shared Task”. In: *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Brussels, Belgium: Association for Computational Linguistics, pp. 31–42. DOI: 10.18653/v1/W18-6206.

Martins, R., M. Gomes, J. J. Almeida, P. Novais, and P. Henriques (2018). “Hate Speech Classification in Social Media Using Emotional Analysis”. In: *2018 7th Brazilian Conference on Intelligent Systems (BRACIS)*, pp. 61–66. DOI: 10.1109/BRACIS.2018.00019.

Mohammad, S., F. Bravo-Marquez, M. Salameh, and S. Kiritchenko (2018). “SemEval-2018 Task 1: Affect in Tweets”. In: *Proceedings of The 12th International Workshop on Semantic Evaluation*. New Orleans, Louisiana: Association for Computational Linguistics, pp. 1–17. DOI: 10.18653/v1/S18-1001.

Tafreshi, S. and M. Diab (2018). “Emotion Detection and Classification in a Multigenre Corpus with Joint Multi-Task Deep Learning”. In: *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, pp. 2905–2913. URL: https://aclanthology.org/C18-1246.

Feldman Barrett, L. (2017). “The theory of constructed emotion: an active inference account of interoception and categorization”. In: *Social Cognitive and Affective Neuroscience* 12.11, p. 1833.

Gao, Q., H. Jiannan, X. Ruifeng, G. Lin, Y. He, K.-F. Wong, and Q. Lu (2017). “Overview of NTCIR-13 ECA Task”. In: *Proceedings of the NTCIR-13 Conference*.

Kim, E., S. Padó, and R. Klinger (2017). “Investigating the Relationship between Literary Genres and Emotional Plot Development”. In: *Proceedings of the Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*. Vancouver, Canada: Association for Computational Linguistics, pp. 17–26. URL: http://www.aclanthology.org/W17-2203.

Schuff, H., J. Barnes, J. Mohme, S. Padó, and R. Klinger (2017). “Annotation, Modelling and Analysis of Fine-Grained Emotions on a Stance and Sentiment Detection Corpus”. In: *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Copenhagen, Denmark: Association for Computational Linguistics, pp. 13–23. DOI: 10.18653/v1/W17-5203.

Shahraki, A. G. and O. R. Zaïane (2017). “Lexical and learning-based emotion mining from text”. In: *Proceedings of the International Conference on Computational Linguistics and Intelligent Text Processing*.

Barrett, E. L. F., M. Lewis, and J. M. Haviland-Jones (2016). *Handbook of Emotions, Fourth Edition*.

Green, B. (2016). *Social Robots, AI, and Ethics*. URL: https://www.scu.edu/ethics/focus-areas/technology-ethics/resources/social-robots-ai-and-ethics/.

Reagan, A. J., L. Mitchell, D. Kiley, C. M. Danforth, and P. S. Dodds (2016). “The emotional arcs of stories are dominated by six basic shapes”. In: *EPJ Data Science* 31. DOI: 10.1140/epjds/s13688-016-0093-1.

Scarantino, A. (2016). “The philosophy of emotions and its impact on affective science”. In: *Handbook of emotions*. Guilford Press New York, NY. Chap. 4, pp. 3–48.
Ghazi, D., D. Inkpen, and S. Szpakowicz (2015). “Detecting emotion stimuli in emotion-bearing sentences”. In: *CICLing*.
Mohammad, S. M., X. Zhu, S. Kiritchenko, and J. Martin (2015). “Sentiment, emotion, purpose, and style in electoral tweets”. In: *Information Processing & Management* 51.4, pp. 480–499. DOI: 10.1016/j.ipm.2014.09.003.
Udochukwu, O. and Y. He (2015). “A Rule-Based Approach to Implicit Emotion Detection in Text”. In: *Natural Language Processing and Information Systems*. Ed. by C. Biemann, S. Handschuh, A. Freitas, F. Meziane, and E. Métais. Cham: Springer International Publishing, pp. 197–203.
Brynielsson, J., F. Johansson, C. Jonsson, and A. Westling (2014). “Emotion classification of social media posts for estimating people’s reactions to communicated alert messages during crises”. In: *Secur. Informatics* 3.1, p. 7. DOI: 10.1186/s13388-014-0007-3.
Mohammad, S., X. Zhu, and J. Martin (2014). “Semantic Role Labeling of Emotions in Tweets”. In: *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Baltimore, Maryland: Association for Computational Linguistics, pp. 32–41. DOI: 10.3115/v1/W14-2607. URL: https://aclanthology.org/W14-2607.
Clore, G. L. and A. Ortony (2013). “Psychological Construction in the OCC Model of Emotion”. In: *Emotion Review* 5.4, pp. 335–343.
Desmet, B. and V. Hoste (2013). “Emotion Detection in Suicide Notes”. In: *Expert Syst. Appl.* 40.16, pp. 6351–6358. DOI: 10.1016/j.eswa.2013.05.050.
Mohammad, S. and P. D. Turney (2013). “Crowdsourcing a Word-Emotion Association Lexicon.” In: *Computational Intelligence* 29.3, pp. 436–465. DOI: 10.1111/j.1467-8640.2012.00460.
Balahur, A., J. M. Hermida, and A. Montoyo (2012). “Building and Exploiting EmotiNet, a Knowledge Base for Emotion Detection Based on the Appraisal Theory Model”. In: *IEEE Transactions on Affective Computing* 3.1, pp. 88–101. DOI: 10.1109/T-AFFC.2011.33.
Liu, B. (2012). *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Cham: Springer International Publishing. DOI: 10.1007/978-3-031-02145-9.
Mohammad, S. (2012). “#Emotional Tweets”. In: **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*. Montréal, Canada: Association for Computational Linguistics, pp. 246–255. URL: https://aclanthology.org/S12-1033.
Pestian, J. P., P. Matykiewicz, M. Linn-Gust, B. South, O. Uzuner, J. Wiebe, K. B. Cohen, J. Hurdle, and C. Brew (2012a). “Sentiment analysis of suicide notes: A shared task”. In: *Biomed. Inform. Insights* 5.Suppl 1, pp. 3–16.
Pestian, J. P., P. Matykiewicz, M. Linngust, B. South, O. Uzuner, J. Wiebe, K. B. Cohen, J. Hurdle, and C. Brew (2012b). “Sentiment Analysis of Suicide Notes: A Shared Task”. In: *Biomedical Informatics Insights* 5, pp. 3–16.
Wang, W., L. Chen, K. Thirunarayan, and A. Sheth (2012a). “Harnessing Twitter ”Big Data” for Automatic Emotion Identification”. In: *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Confernece on Social Computing*, pp. 587–592.
Wang, W., L. Chen, K. Thirunarayan, and A. P. Sheth (2012b). “Harnessing Twitter "Big Data" for Automatic Emotion Identification”. In: IEEE, pp. 587–592. DOI: 10.1109/SocialCom-PASSAT.2012.119.
Dodds, P. S., K. D. Harris, I. M. Kloumann, C. A. Bliss, and C. M. Danforth (2011). “Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter”. In: *PLOS ONE* 6.12, pp. 1–1. DOI: 10.1371/journal.pone.0026752. URL: https://doi.org/10.1371/journal.pone.0026752.
Kessler, J. S., M. Eckert, L. Clark, and N. Nicolov (2010). “The ICWSM 2010 JDPA Sentiment Corpus for the Automotive Domain”. In: *ICWSM*.
Neviarouskaya, A., H. Prendinger, and M. Ishizuka (2010). “@AM: Textual Attitude Analysis Model”. In: *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*. Los Angeles, CA, pp. 80–88. URL: https://www.aclweb.org/anthology/W10-0210.
Ganu, G., N. Elhadad, and A. Marian (2009). “Beyond the Stars: Improving Rating Predictions using Review Text Content”. In: *International Workshop on the Web and Databases*.
Neviarouskaya, A., H. Prendinger, and M. Ishizuka (2009). “Compositionality Principle in Recognition of Fine-Grained Emotions from Text”. In: *Proceedings of the International AAAI Conference on Web and Social Media* 3, pp. 278–281. DOI: 10.1609/icwsm.v3i1.13987.

Shaikh, M. A. M., H. Prendinger, and M. Ishizuka (2009). “A Linguistic Interpretation of the OCC Emotion Model for Affect Sensing from Text”. In: *Affective Information Processing*. Ed. by J. Tao and T. Tan, pp. 45–73. DOI: 10.1007/978-1-84800-306-4_4.

Aman, S. and S. Szpakowicz (2007). “Identifying Expressions of Emotion in Text”. In: *Text, Speech and Dialogue*. Ed. by V. Matoušek and P. Mautner. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 196–205. DOI: 10.1007/978-3-540-74628-7_27.

Strapparava, C. and R. Mihalcea (2007). “SemEval-2007 Task 14: Affective Text”. In: *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*. Prague, Czech Republic: Association for Computational Linguistics, pp. 70–74. URL: https://aclanthology.org/S07-1013.

Alm, C. O., D. Roth, and R. Sproat (2005). “Emotions from Text: Machine Learning for Text-based Emotion Prediction”. In: *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*. Vancouver, British Columbia, Canada: Association for Computational Linguistics, pp. 579–586. URL: https://www.aclweb.org/anthology/H05-1073.

Alm, C. O. and R. Sproat (2005). “Emotional Sequencing and Development in Fairy Tales”. In: *Affective Computing and Intelligent Interaction*. Ed. by J. Tao, T. Tan, and R. W. Picard. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 668–674. DOI: 10.1007/11573548_86.

Posner, J., J. A. Russell, and B. S. Peterson (2005). “The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology.” In: *Development and Psychopathology* 17.3, pp. 715–734. DOI: 10.1017/S0954579405050340.

Plutchik, R. (2001). “The Nature of Emotions Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice”. In: *American Scientist* 89.4, pp. 344–350.

Scherer, K. R., A. Schorr, and T. Johnstone, eds. (2001). *Appraisal processes in emotion: theory, methods, research*. Series in affective science. Oxford, New York: Oxford University Press.

Picard, R. W. (1997). *Affective Computing*. MIT Press.

Ekman, P. (1992). “An argument for basic emotions”. In: *Cognition & emotion* 6.3-4, pp. 169–200.

Smith, C. A. and P. C. Ellsworth (1985). “Patterns of cognitive appraisal in emotion”. In: *Journal of Personality and Social Psychology* 48.4, pp. 813–838.

Darwin, C. (1872). *The Expression of the Emotions in Man and Animals*. John Murray.