Emotion Dynamics of Fictional Characters
using large emotion lexicons

Saif M. Mohammad
Senior Research Scientist, National Research Council Canada
Saif.Mohammad@nrc-cnrc.gc.ca @SaifMMohammad

Basic Emotions Theory

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

Many tenets of BET discredited

- See Theory of Constructed Emotion (Barrett, 2017)
- Still useful work on categorical emotions such as joy, sadness, fear, etc.

Plutchik's Emotion Wheel
Image credit: Julia Belyanevych

Core Dimensions of Connotative Meaning

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

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

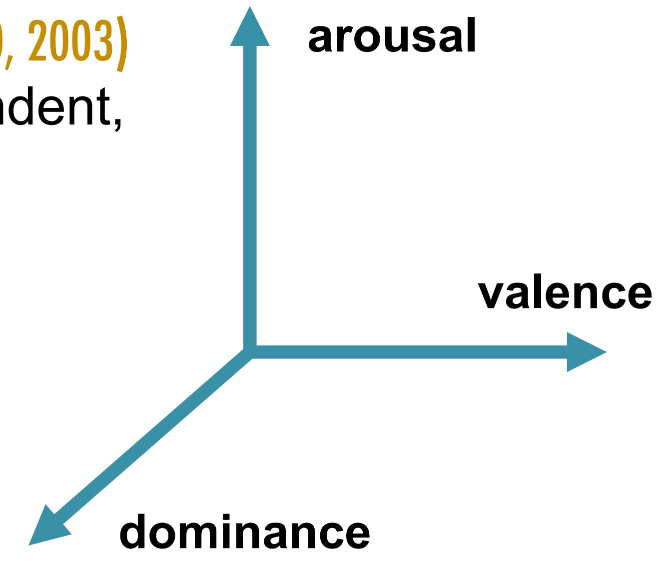

Thus, when comparing the meanings of two words, we can compare their V, A, D scores. For example:

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

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

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

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

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

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

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

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

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

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

The NRC Word–Colour Association Lexicon (2010)

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

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

Crowdsourcing and Quality Control

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

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

Comparative annotations (not Likert scales)

- Avoids various biases

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

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

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

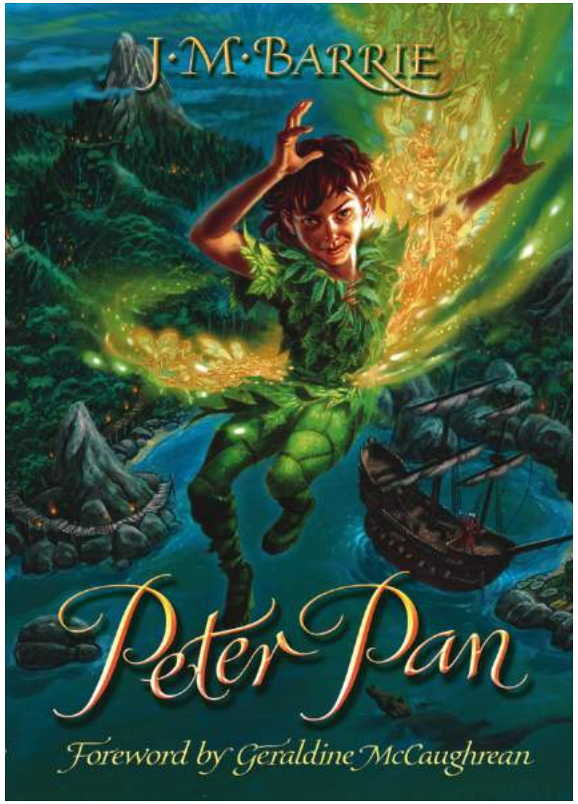

Detecting Emotions in Stories

Tracking Emotions in Stories (Kurt Vonnegut inspired)

- Can we automatically track the emotions of characters?
- Are there some canonical shapes common to most stories?
- Can we track the change in distribution of emotion words?

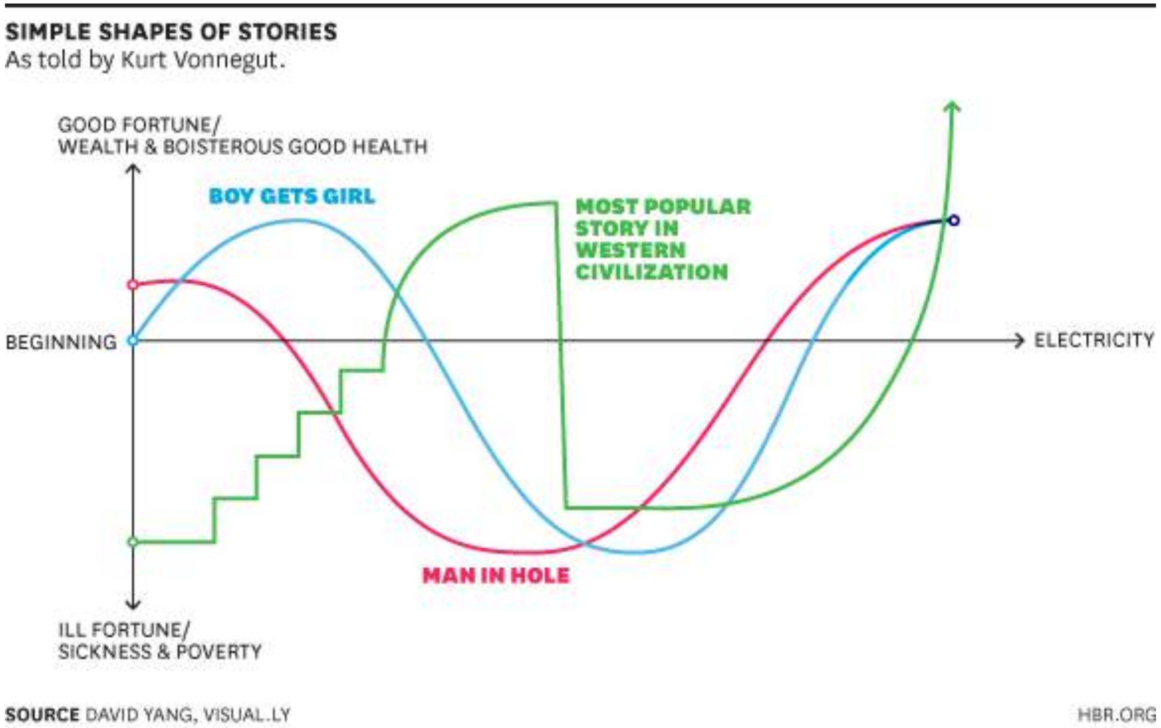

Back in 2011:

Tracking emotion word distribution in novels and fairy tales.

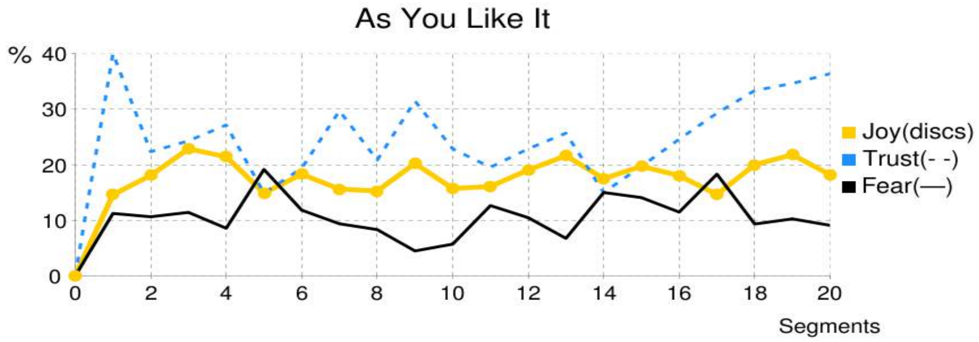

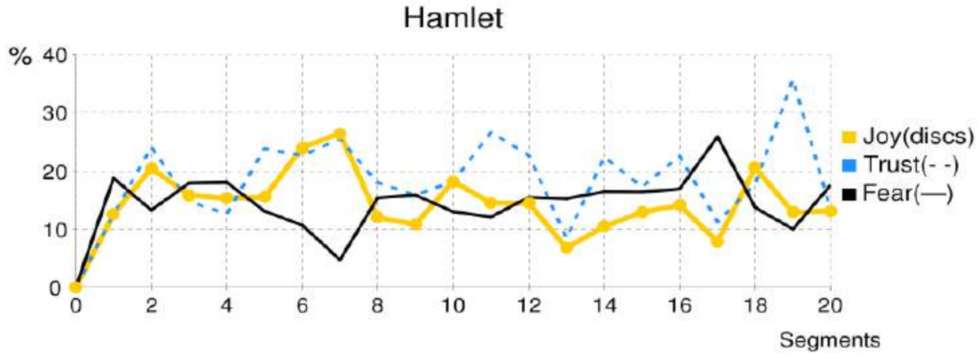

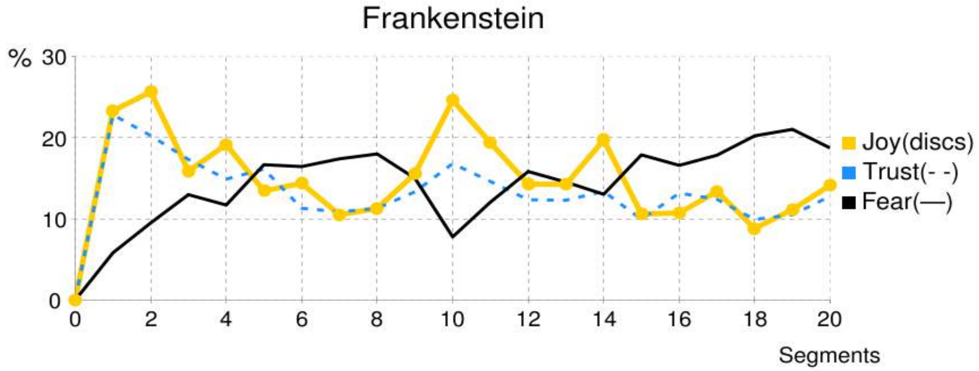

From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales, Saif Mohammad, In Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), June 2011, Portland, OR.

Work on shapes of stories

- From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales, Saif Mohammad, In Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), June 2011, Portland, OR.
- Character-based kernels for novelistic plot structure. Elsner, M., 2012, April. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 634-644). Association for Computational Linguistics.
- A novel method for detecting plot. M. Jockers http://www.matthewjockers.net/2014/06/05/a-novel-method-for-detecting-plot/, June 2014.
- The emotional arcs of stories are dominated by six basic shapes. Reagan, A.J., Mitchell, L., Kiley, D., Danforth, C.M. and Dodds, P.S., 2016. EPJ Data Science, 5(1), p.31.

Hannah Davis
Artist/Programmer

Generating music from text

Paper:

- Generating Music from Literature. Hannah Davis and Saif M. Mohammad, In Proceedings of the EACL Workshop on Computational Linguistics for Literature, April 2014, Gothenburg, Sweden.

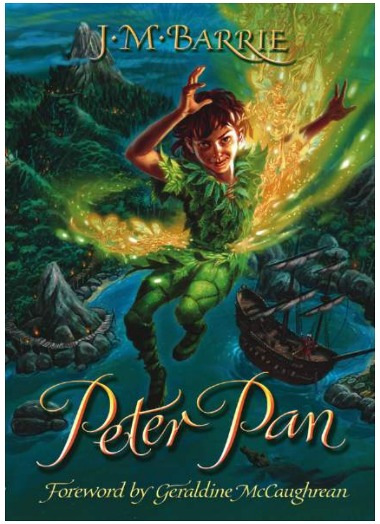

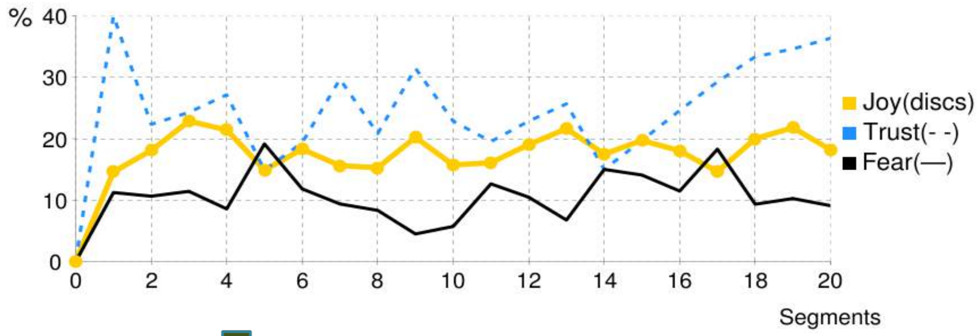

A method to generate music from literature

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

TransProse

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

TransProse

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

Examples

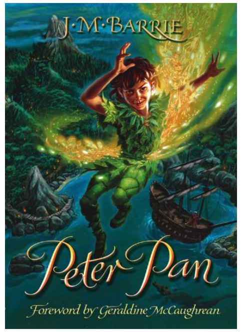

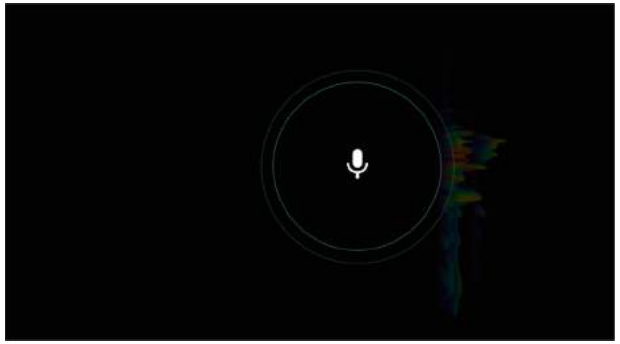

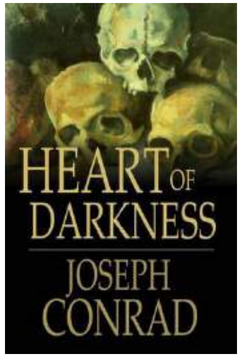

TransProse Music Played by an Orchestra, at the Louvre Museum, Paris

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

Emotion Dynamics (from Psychology)

Study of change in emotional state with time

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

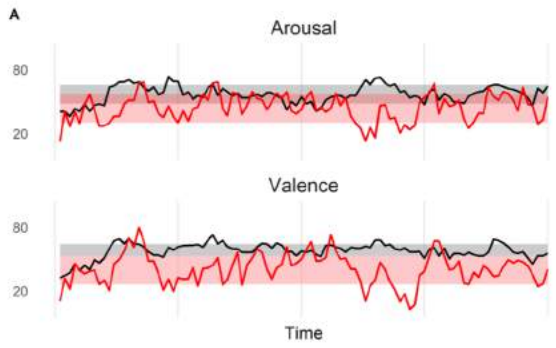

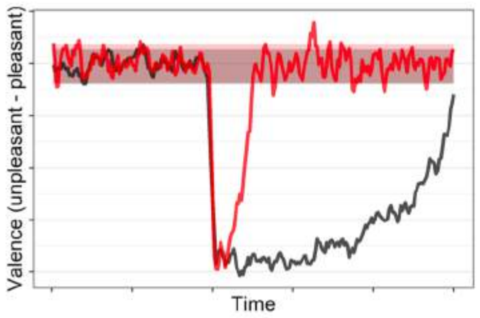

Another window into emotions is through our words:

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

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

Utterance Emotion Dynamics: Metrics

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

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

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

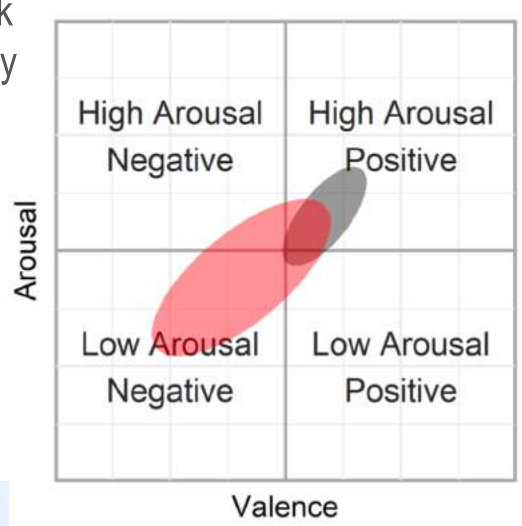

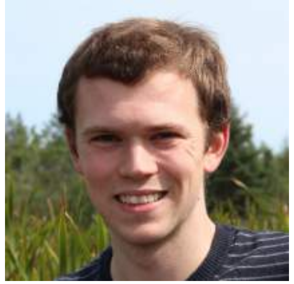

Will Hipson

Emotion Dynamics in Movie Dialogues

Will E. Hipson, Saif M. Mohammad, Under Review.

Character Dialogue from Literature and Film

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

Data we used:

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

Analyzing Characters

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

Affect Dimensions: Valence, Arousal, Valence – Arousal

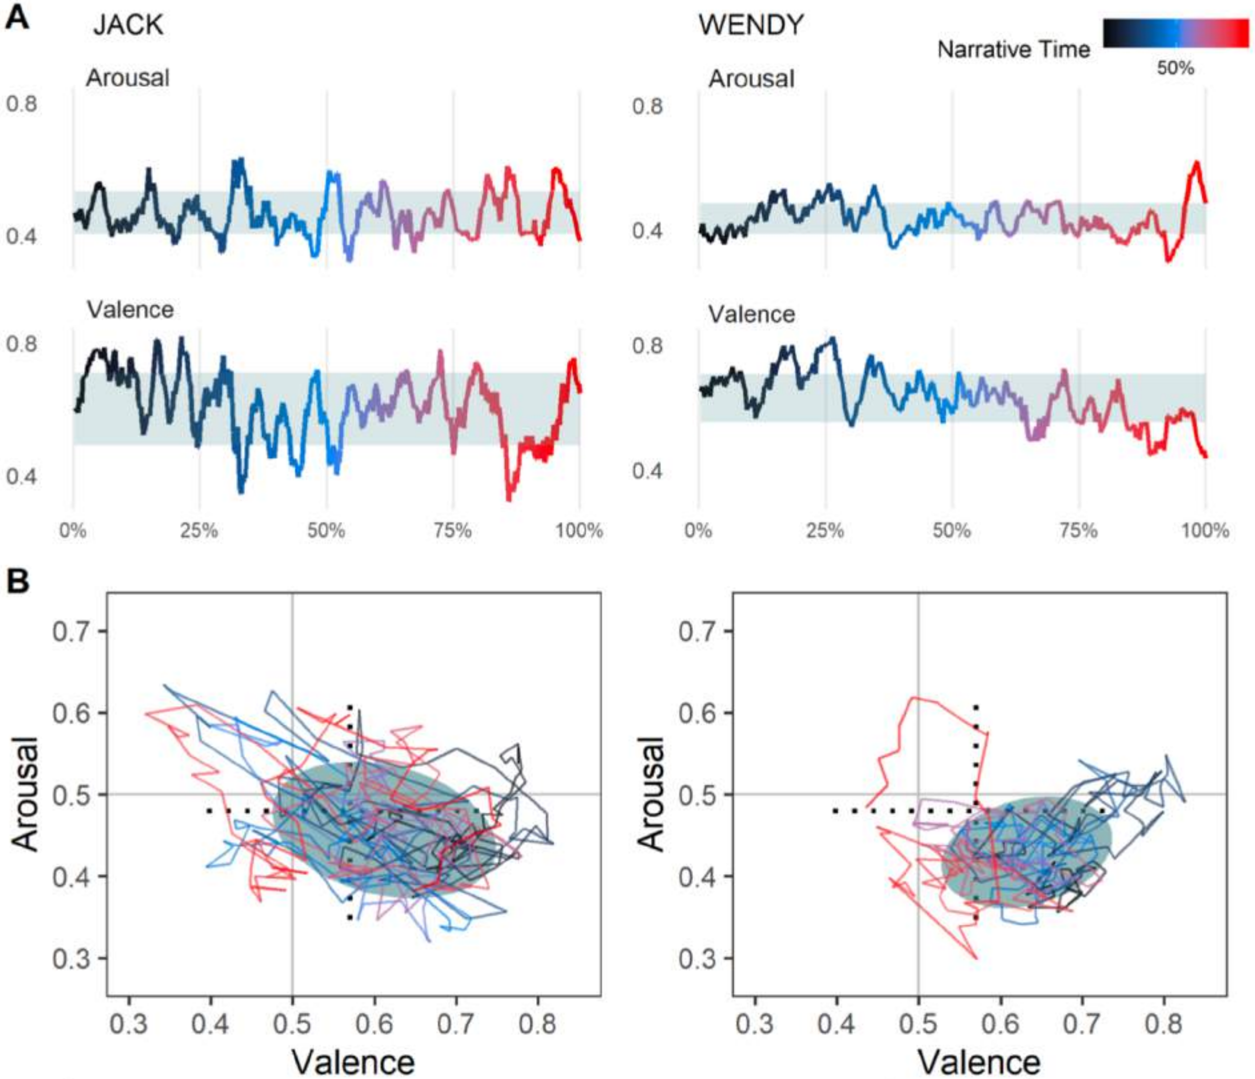

Fig 3. One dimensional and two dimensional state spaces for Jack ($n = 389$ words) and Wendy ($n = 279$ words), two main characters from The Shining (1980). Color of line corresponds to narrative time, with dark blue meaning earlier in the movie and red meaning later. The black dotted lines show the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

UED Metric: Emotion word density

Affect Dimensions: Positive, Negative, Eight Plutchik emotions

Table 1. Average emotion word density (Av. EWD) and standard deviation (SD) of main characters in IMSDb ($N = 2{,}687$).

Emotion	**Av. EWD**	**SD**
Negative	16.5	3.8
Positive	20.3	4.5
Anger	7.6	2.6
Anticipation	12.0	2.8
Disgust	5.6	2.3
Fear	9.9	3.0
Joy	9.8	3.5
Sadness	8.3	2.5
Surprise	6.8	2.0
Trust	13.5	3.4

UED Metric: Various metrics

Affect Dimensions: Valence – Arousal

Table 2. Average UED metrics (2–6) and standard deviation (SD) for main characters in IMSDb (N = 2,687).

Metric	**Av. UED**	**SD**
Home Base-Major Width	0.13	0.02
Home Base-Minor Width	0.09	0.01
Emotion Variability	0.15	0.02
Displacement Length	9.13	1.90
Displacement Count	34.46	18.01
Peak Distance	0.17	0.03
Rise Rate	0.05	0.01
Recovery Rate	0.05	0.01

UED Metric: Emotional variability

Affect Dimensions: Valence – Arousal

Table 3. Characters with the highest/lowest emotional variability (Var.). Note that the bottom rank number is less than the total number of characters in the data because some characters had insufficient number of displacements to obtain reliable averages.

Rank	Character	Movie Title	Var.
1	Jessica	Little Athens	0.228
2	TJ	Hesher	0.220
3	Ginger	Casino	0.215
4	Dennis	Hostage	0.208
5	Wes	Three Kings	0.204
2610	Lynn	L.A. Confidential	0.107
2611	Diane	Horse Whisperer	0.104
2612	Dolores	Sweet Hereafter	0.103
2613	Riker	Star Trek	0.103
2614	Data	Star Trek	0.100

UED Metric: Recovery rate

Affect Dimensions: Valence – Arousal

Table 4. Characters with highest/lowest recovery rate (Rec.). Note that the bottom rank number is less than the total number of characters in the data because some characters had insufficient number of displacements to obtain reliable averages.


Rank	Character	Movie Title	Rec.
1	Jacob	Nightmare on Elm Street	0.107
2	Chad	Burn After Reading	0.106
3	Andrew	The Breakfast Club	0.105
4	Rennie	Friday the 13th	0.101
5	Jimmy	Magnolia	0.101
2610	Jonson	Anonymous	0.019
2611	Agnis	Shipping News, The	0.018
2612	Paul	Manhattan Murder Mystery	0.017
2613	Jack	Burlesque	0.017
2614	Chigurh	No Country for Old Men	0.015

Analyzing Characters

Across emotion space and narrative time

Plot: Topological map showing where peak displacement tends to occurs

Affect Dimensions: Valence – Arousal

Analyzing Characters Across Emotion Space

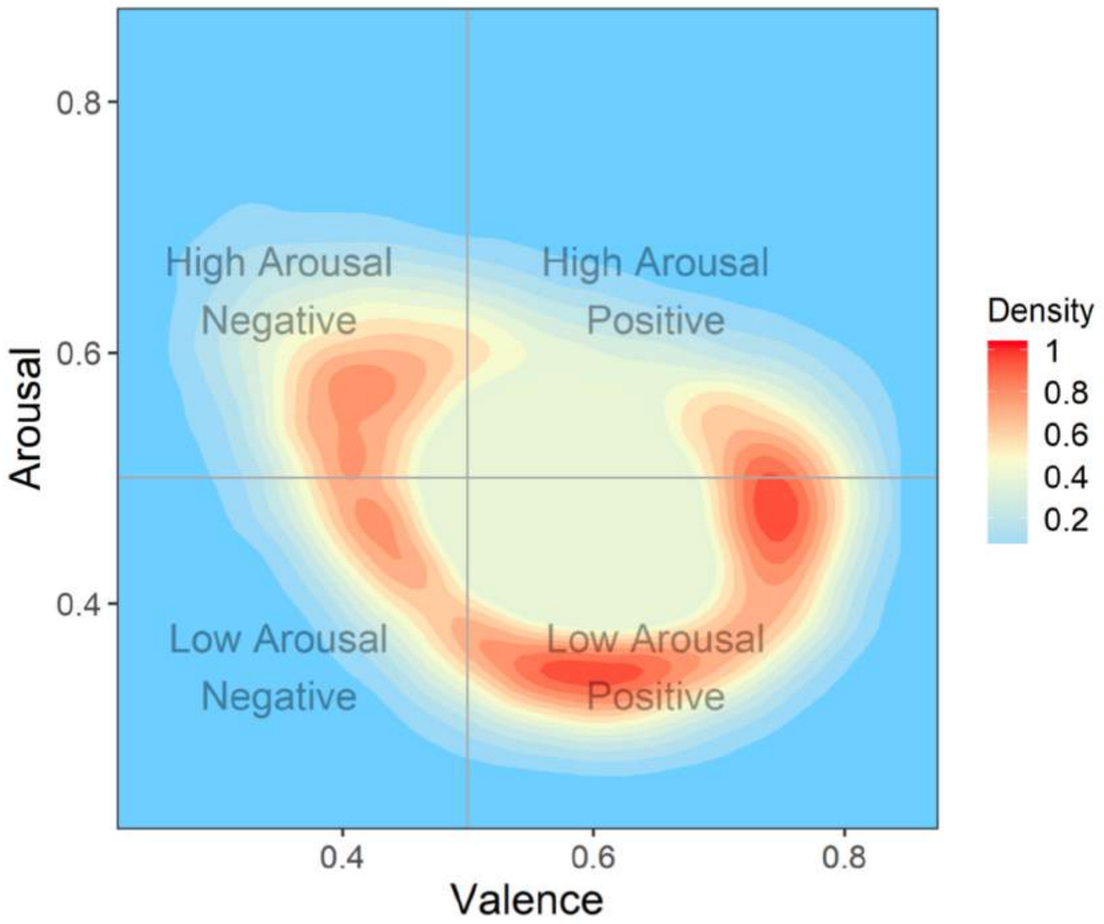

Fig 5. Density map showing where peak displacements tend to occur. Red corresponds to more peaks. Density is normalized to go from 0–1.

Plot: Average emotion word densities across narrative time

Affect Dimensions: Positive, Negative

Analyzing Characters Across Narrative Time

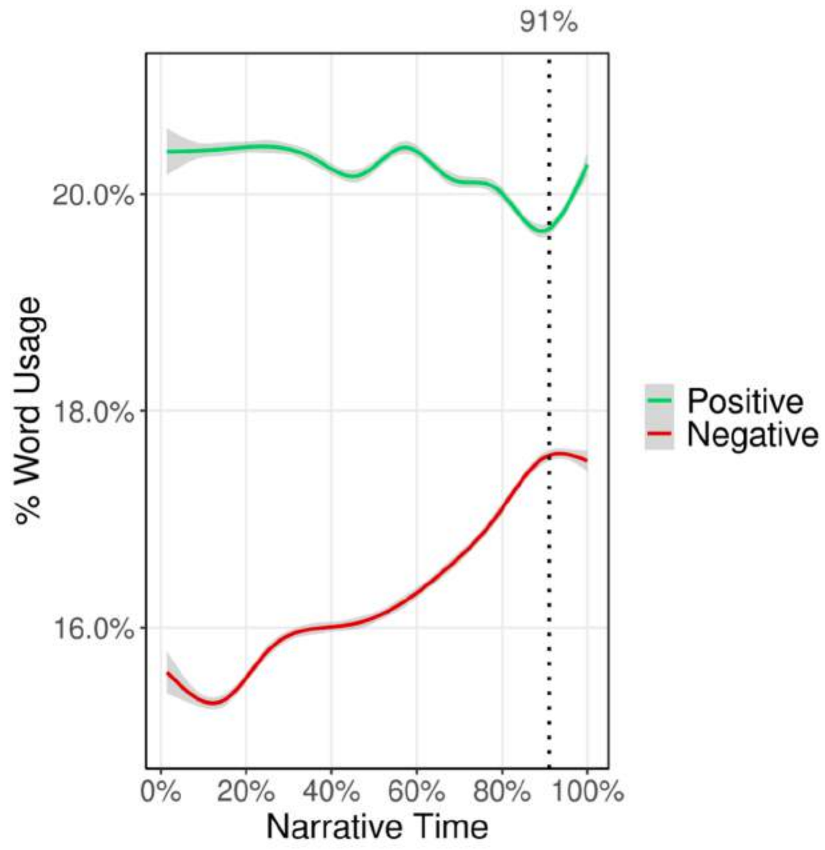

Fig 6. Average trends in proportion of positive and negative word usage over time. Vertical dotted line shows location of peak negative density and lowest positive density. Grey band is the 95% confidence interval around the estimated mean. n = 965,147 words.

Analyzing Character–Character Interactions

At any given point, pairs of characters may be:

- in-sync: use emotion words similarly
  (e.g., both use lots of low valence words)
- discordant: use emotion words dissimilarly
  (e.g., one uses high-arousal words, other uses low-arousal words)

To what extent do character–character discordances vary throughout the movie plot?

Plot: Average Character–Character Discordance

Affect Dimensions: Valence – Arousal

Analyzing Character–Character Discordance across Narrative Time

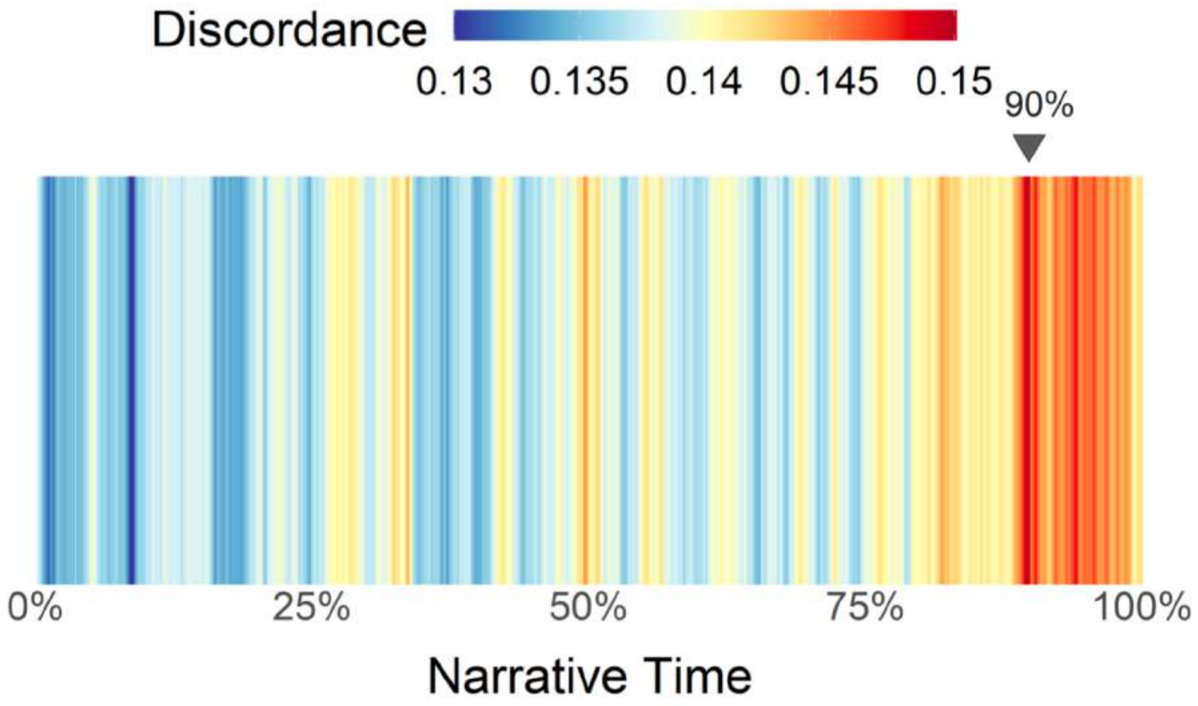

Fig 7. Character discordance over narrative time (n = 1,246,990 words). Red indicates more discordance, blue indicates less discordance. Discordance is lowest during first quarter of a movie and peaks at 90%. Score is measured in the same scale as the v-a space (i.e., 0.15 implies a Euclidean distance of 0.15 in the state space).

Summary

Showed the use of large word–emotion lexicons in literary analysis

- Simple, yet powerful when applied to text streams
- Easily interpretable

Resources Available at: www.saifmohammad.com

- Emotion lexicons
- Practical and Ethical considerations in using emotion lexicons
- Tutorials and book chapters on emotion recognition

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