More Than a Feeling: Emotion in Text

Candidacy Examination

Elsbeth Turcan

April 23, 2019

Columbia University, Department of Computer Science

Introduction

Incorporating Psychological Models

- Scientists make modeling assumptions to study real-world phenomena
- Natural language processing - author and audience
- Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this
- A domain-specific understanding of these phenomena can help incorporate them

Outline

Focus for this talk: emotion in text

Emotion and Influence

In Argumentation Theory

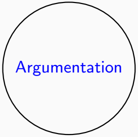

- **Studying emotion and influence is not a new idea**

In Argumentation Theory

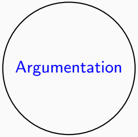

- Studying emotion and influence is **not** a new idea
- Western roots in Aristotle (**Alan Brinton**)
    - Recognition by some of the earliest, most influential Western scholars
    - Pathos - one of Aristotle's essential means of persuasion
    - Pathos related to morality and virtue for Aristotle
    - Persuading to feel emotion vs. emotion as basis for action

In Argumentation Theory

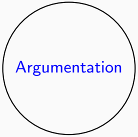

- Western roots in Aristotle (**Alan Brinton**)
- Argumentation and logic (**Michael Gilbert**)
  - Proposition: emotional and factual argument are *equally* fuzzy and ambiguous
  - Treats emotion under the **acceptability, relevance, sufficiency framework**

In Argumentation Theory

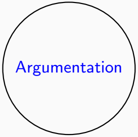

- Western roots in Aristotle (**Alan Brinton**)
- Argumentation and logic (**Michael Gilbert**)
  - Proposition: emotional and factual argument are *equally* fuzzy and ambiguous
  - Treats emotion under the **acceptability, relevance, sufficiency framework**
  - **In order to study the real world, we must study the things that actually happen, not just their idealized models**

In Psychology

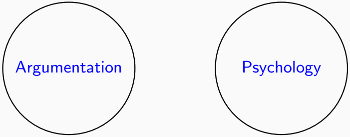

- **Emotion and influence highly studied in psychology**

In Psychology

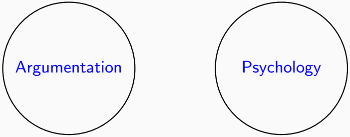

- Emotion and influence highly studied in psychology
- Processing arguments (**Schwarz et al.**)
  - Critical review of candidate mechanisms for how emotion affects persuasion
  - Central vs. peripheral processing

In Psychology

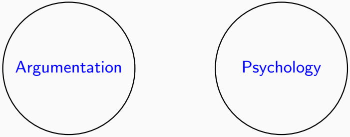

- Processing arguments (**Schwarz et al.**)
  - Mood as peripheral cue hypothesis
  - Mood congruency hypothesis
  - Change in criteria hypothesis
  - Motivational hypothesis
  - Cognitive capacity hypothesis

In Psychology

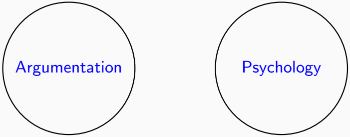

- Processing arguments (**Schwarz et al.**)
  - Mood as peripheral cue hypothesis
  - Mood congruency hypothesis
  - Change in criteria hypothesis
  - **Motivational hypothesis**
  - Cognitive capacity hypothesis

In Psychology

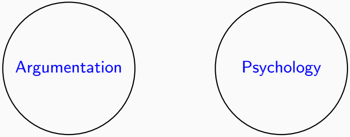

- Processing arguments (**Schwarz et al.**)
    - Critical review of candidate mechanisms for how emotion affects persuasion

In Psychology

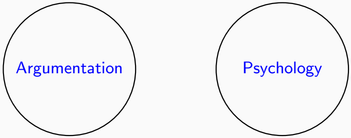

- Processing arguments (**Schwarz et al.**)
  - Critical review of candidate mechanisms for how emotion affects persuasion
  - **Conclusion: negative mood more conducive to central processing**

In Psychology

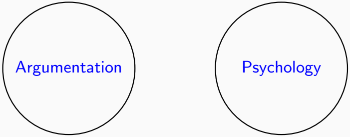

- Processing arguments (**Schwarz et al.**)
- Making arguments (**Villata et al.**)
  - Used facial expression to examine emotions felt while arguing
  - Some significant correlations observed (e.g., sadness $\propto$ withdrawal)

In Psychology

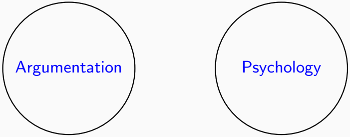

- Processing arguments (**Schwarz et al.**)
- Making arguments (**Villata et al.**)
  - Used facial expression to examine emotions felt while arguing
  - Some significant correlations observed (e.g., sadness $\propto$ withdrawal)
  - **Emotion and the process of making an argument do interact**

In Computer Science

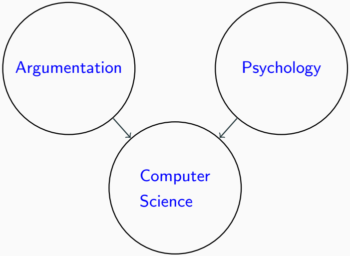

- Effects of argument type by audience (**Lukin et al.**)
  - Interaction of argument type and audience personality
  - No measures of long-lasting belief change, but definite short-term effects
- Characterizing emotional vs. logical arguments (**Oraby et al.**)
  - Syntactic patterns extracted from emotional and logical arguments
  - Logical arguments more structural, emotional more vivid and immediate

In Computer Science

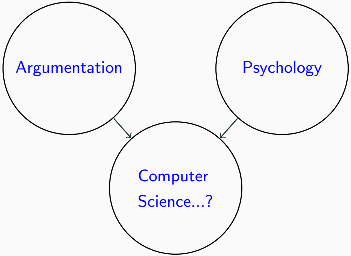

- Effects of argument type by audience (**Lukin et al.**)
  - Interaction of argument type and audience personality
  - No measures of long-lasting belief change, but definite short-term effects
- Characterizing emotional vs. logical arguments (**Oraby et al.**)
  - Syntactic patterns extracted from emotional and logical arguments
  - Logical arguments more structural, emotional more vivid and immediate
- **...and...?**

Summary

- Emotion plays a significant role in influence and has been studied extensively in multiple fields

Summary

- Emotion plays a significant role in influence and has been studied extensively in multiple fields
- Emotion interacts with cognition–in the author and audience

Summary

- Emotion plays a significant role in influence and has been studied extensively in multiple fields
- Emotion interacts with cognition–in the author and audience
- However, there is a dearth of computational work in this area

Emotion Classification

Emotion Classification

- Problem: given a piece of text, assign it one (or more) emotion label(s)

[1]

- Typically a supervised machine learning problem with discrete emotion labels
- Greatly accelerated with popularization of social media data

[1] https://twitter.com/YLKATDelta/status/1117913821513838597

A Very Abridged History of Emotion Classification

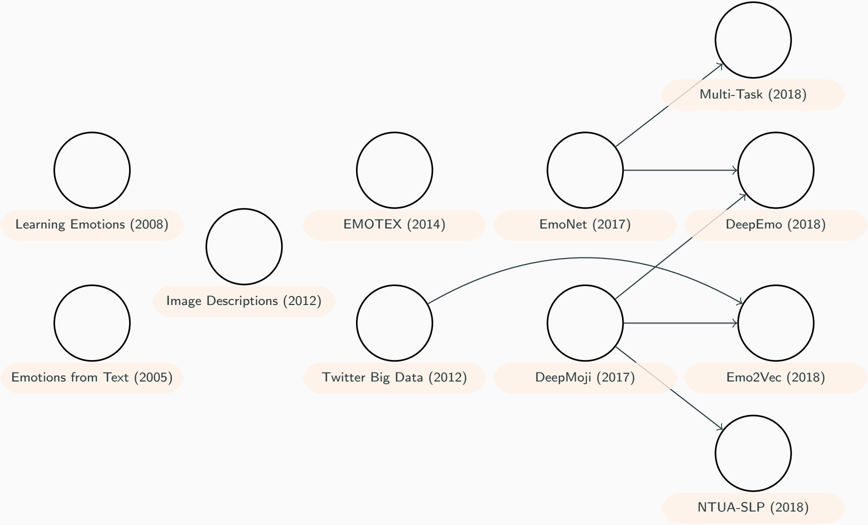

Emotion Categories

Ekman's Six Basic Emotions (universal facial expressions)

Plutchik's Wheel of Emotions (evolutionarily adaptive behaviors)

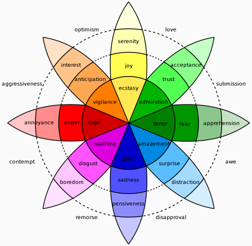

Other Models

- Circumplex model
- Geneva emotion wheel
- Valence/sentiment
- Still many more in psychology literature
- etc.....

A Very Abridged History of Emotion Classification

Ekman's Six

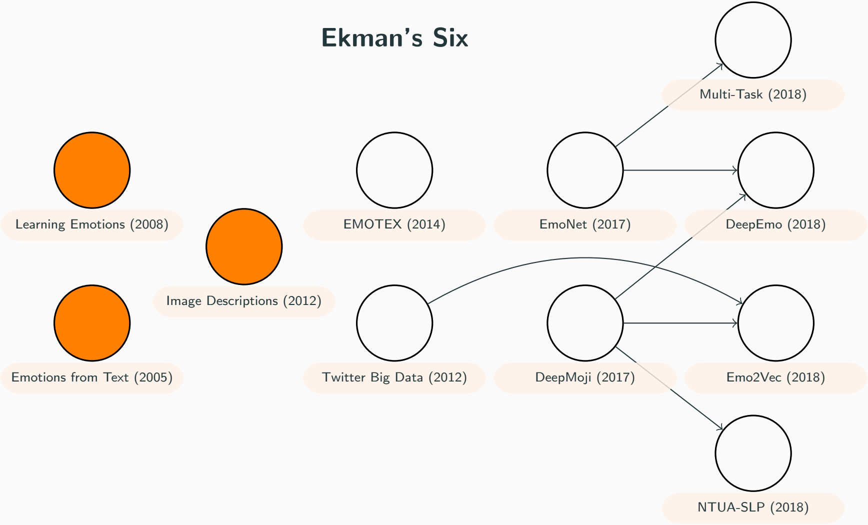

A Very Abridged History of Emotion Classification

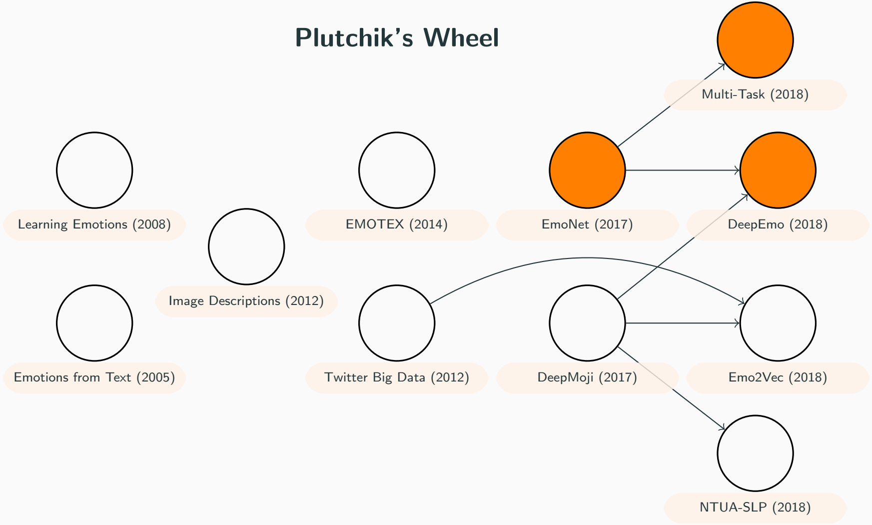

A Very Abridged History of Emotion Classification

Other Emotion Schemes

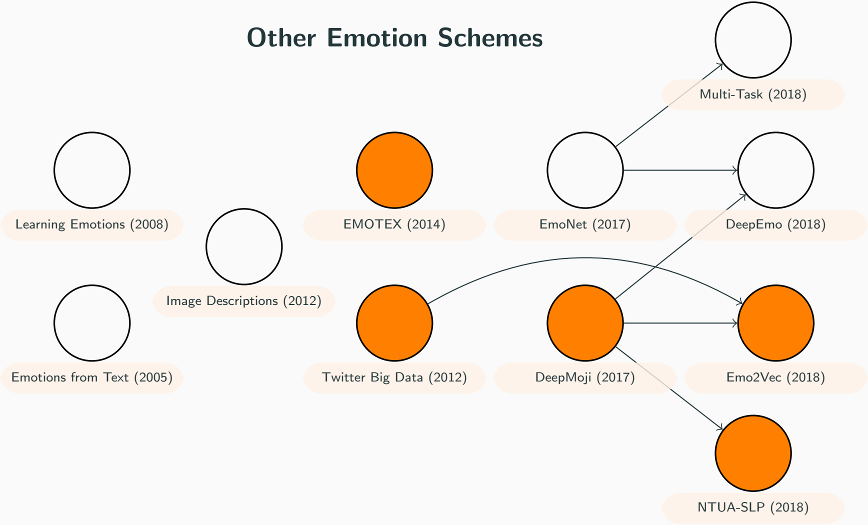

A Very Abridged History of Emotion Classification

Traditional ML

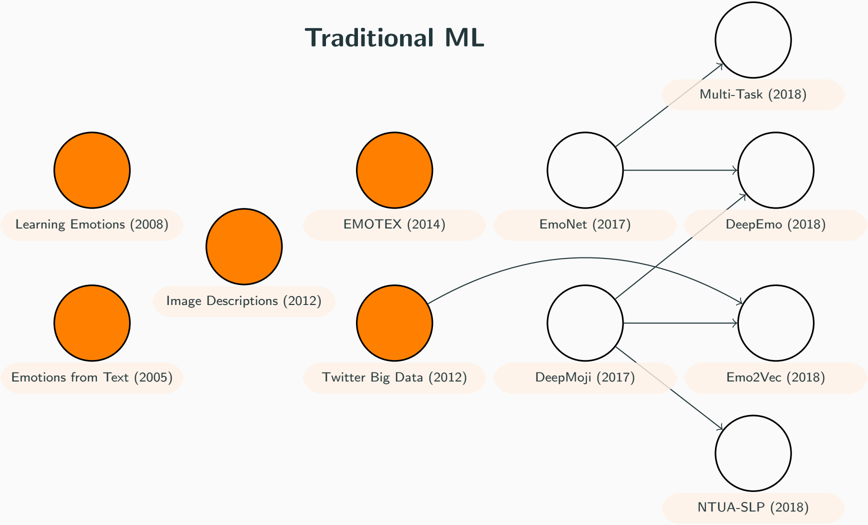

Non-Neural Approaches

- Linear classifiers (SVM, Naïve Bayes)
- Decision and distance algorithms (decision trees, k-nearest neighbor, latent semantic analysis)

Feature Represenation

- Bag-of-ngram features
- Lexical features (punctuation, emoticons, ALL CAPS)
- Semantic features (POS tags, negations)
- Curated lexicons (LIWC, WordNet, DAL, MPQA Subjectivity Lexicon)

A Very Abridged History of Emotion Classification

Deep Learning

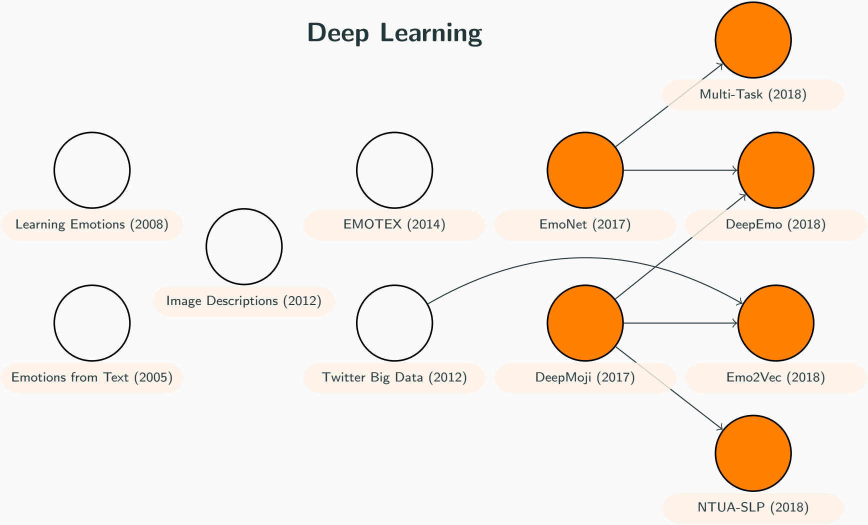

Deep Learning (Supervised)

Recurrent Networks

- Long Short-Term Memory Networks (**NTUA-SLP**, **DeepMoji**)
- Gated Recurrent Neural Networks (**EmoNet**, **Multi-Task**)
- "Tricks": bidirectional, attention

Convolutional Networks

- A sort of neural n-gram aproach (**DeepEmo**, **Emo2Vec**)

Emotional Embeddings

- Model-specific word embedding layers
- **NTUA-SLP** - add 10 affective dimensions and spread scores from hand-annotated words
- **DeepEmo** - collect syntactic patterns indicative of different emotions

A Very Abridged History of Emotion Classification

Hand-Annotated Corpora

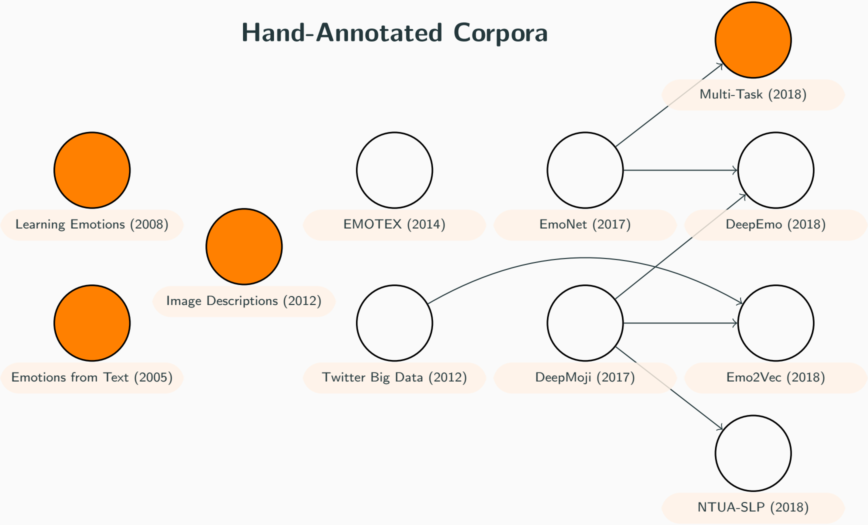

Hand-Annotated Corpora

- A small number of manual annotators (**Emotions from Text**, **Learning Emotions**)
- Crowdsourced annotations (**Image Descriptions**, **Multi-Task**)

A Very Abridged History of Emotion Classification

Distant Labeling

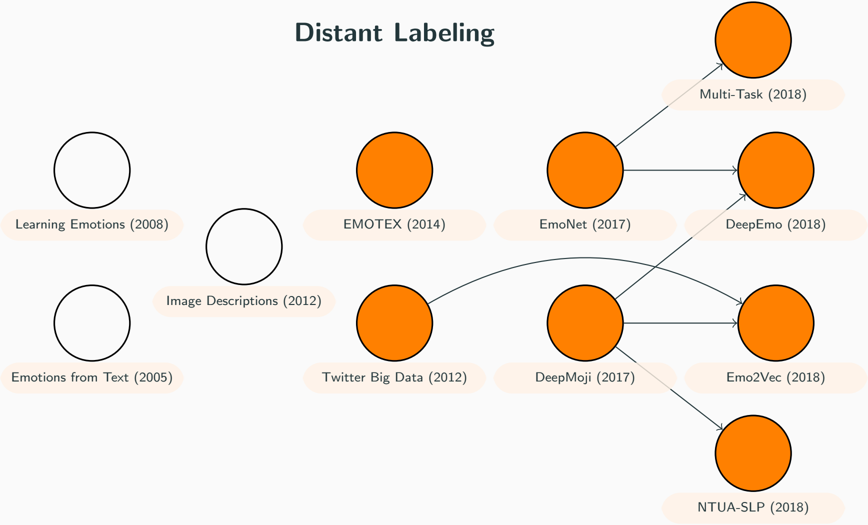

Distant Labeling - Social Media

- **Twitter hashtags** - **EmoNet**, **EMOTEX**, **DeepEmo**, **Twitter Big Data**, **Emo2Vec**, **Multi-Task**
- **DeepMoji** - Tweets with emojis
- **Image Descriptions** - LiveJournal posts (author provides a mood)

Validation studies - Crowd annotations match distant labels fairly well and inform preprocessing (e.g., hashtags at end of Tweet only)

A Very Abridged History of Emotion Classification

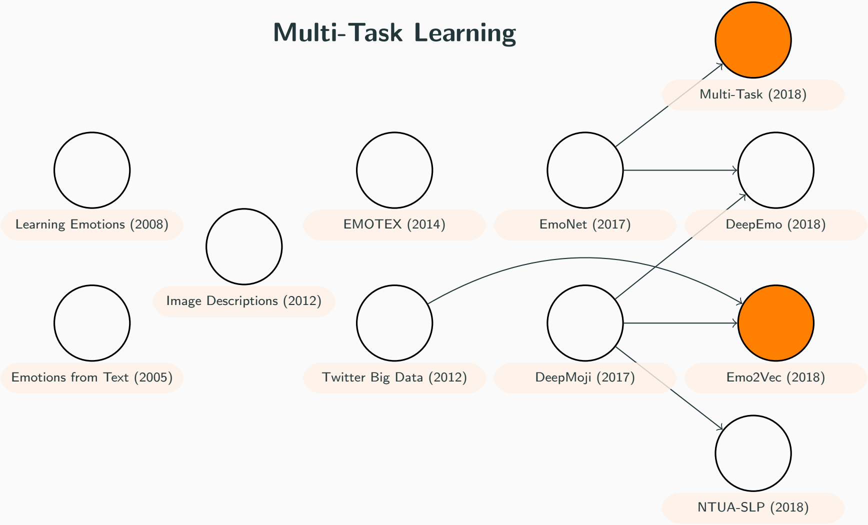

Multi-task Learning

- **Multi-Task** - same task, two different datasets (distantly labeled vs. hand-annotated)
- **Emo2Vec** - seven tasks (emotion classification/intensity, sentiment, sarcasm, stress, abusive language, personality, insults)

A Very Abridged History of Emotion Classification

Transfer Learning

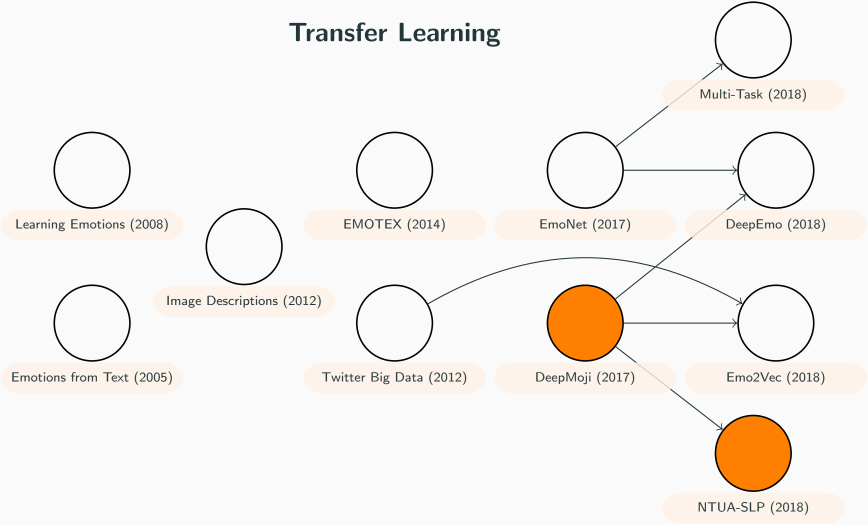

Transfer Learning

- **NTUA-SLP** - pretraining on much larger sentiment dataset (SemEval 2017); fine-tuned whole model simultaneously
- **DeepMoji** - pretraining on emoji prediction task; fine-tuned using chain-thaw approach

A Very Abridged History of Emotion Classification

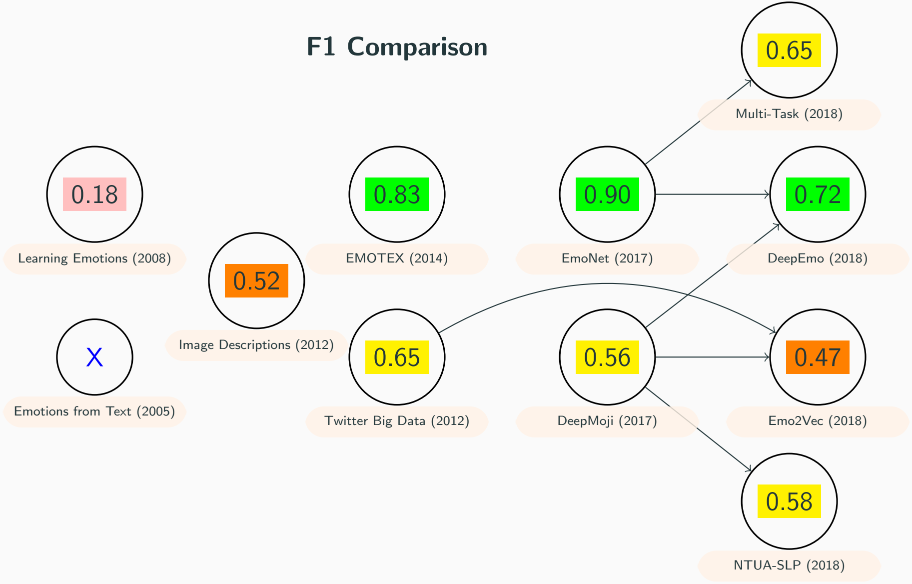

A Very Abridged History of Emotion Classification

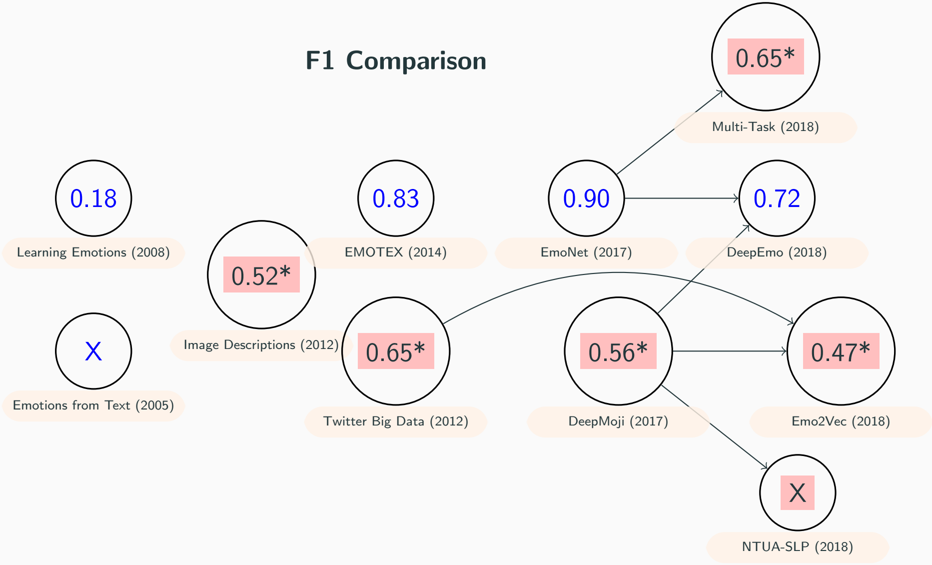

Summary

- Computational research in affectual text processing is expanding
  - Distant labeling and deep learning have had a huge impact
  - Adjacent problems: emotion intensity, causes of emotion, etc.

Summary

- Computational research in affectual text processing is expanding
  - **Distant labeling and deep learning** have had a huge impact
  - Adjacent problems: emotion intensity, causes of emotion, etc.

Summary

- Computational research in affectual text processing is expanding
  - Distant labeling and deep learning have had a huge impact
  - Adjacent problems: emotion intensity, causes of emotion, etc.
- Little work done to validate our choice or use of psychological theory
  - e.g., choice of emotion labels; domain-informed features

Summary

- Computational research in affectual text processing is expanding
  - Distant labeling and deep learning have had a huge impact
  - Adjacent problems: emotion intensity, causes of emotion, etc.
- Little work done to **validate our choice or use of psychological theory**
  - e.g., choice of emotion labels; domain-informed features

Summary

- Computational research in affectual text processing is expanding
  - Distant labeling and deep learning have had a huge impact
  - Adjacent problems: emotion intensity, causes of emotion, etc.
- Little work done to validate our choice or use of psychological theory
  - e.g., choice of emotion labels; domain-informed features
- Current models are not very interpretable

Summary

- Computational research in affectual text processing is expanding
  - Distant labeling and deep learning have had a huge impact
  - Adjacent problems: emotion intensity, causes of emotion, etc.
- Little work done to validate our choice or use of psychological theory
  - e.g., choice of emotion labels; domain-informed features
- Current models are **not very interpretable**

Uniting NLP Methods with Psychological Theories

NLP Methods and Psychological Theories

Useful Theories from Psychology

- Psychology theory can influence how we build models and collect our data

Useful Theories from Psychology

- Psychology theory can influence how we build models and collect our data
- Interaction of **language** and emotion perception (**Lindquist et al.**)
  - Satiating relevant emotion words made emotion-matching tasks more difficult

Useful Theories from Psychology

- Psychology theory can influence how we build models and collect our data
- Interaction of **language** and emotion perception (**Lindquist et al.**)
  - Satiating relevant emotion words made emotion-matching tasks more difficult
  - **Suggests humans use emotion words to categorize emotional stimuli**

Useful Theories from Psychology

- Psychology theory can influence how we build models and collect our data
- Interaction of **language** and emotion perception (**Lindquist et al.**)
    - Satiating relevant emotion words made emotion-matching tasks more difficult
    - Suggests humans use emotion words to categorize emotional stimuli
- Influence of **presentation** on emotion perception (**Sapute et al.**)
    - Expressing a choice as a binary or continuum changed the tipping point for choosing one label or the other

Useful Theories from Psychology

- Psychology theory can influence how we build models and collect our data
- Interaction of **language** and emotion perception (**Lindquist et al.**)
  - Satiating relevant emotion words made emotion-matching tasks more difficult
  - Suggests humans use emotion words to categorize emotional stimuli
- Influence of **presentation** on emotion perception (**Sapute et al.**)
  - Expressing a choice as a binary or continuum changed the tipping point for choosing one label or the other
  - **The change correlated with different patterns of brain activity, suggesting a difference in perception**

Useful Theories from Psychology

- Psychology theory can influence how we build models and collect our data
- Interaction of **language** and emotion perception (**Lindquist et al.**)
  - Satiating relevant emotion words made emotion-matching tasks more difficult
  - Suggests humans use emotion words to categorize emotional stimuli
- Influence of **presentation** on emotion perception (**Sapute et al.**)
  - Expressing a choice as a binary or continuum changed the tipping point for choosing one label or the other
  - The change correlated with different patterns of brain activity, suggesting a difference in perception
- **How we talk about and present emotions can change how we and our annotators perceive them**

NLP Methods and Psychological Theories

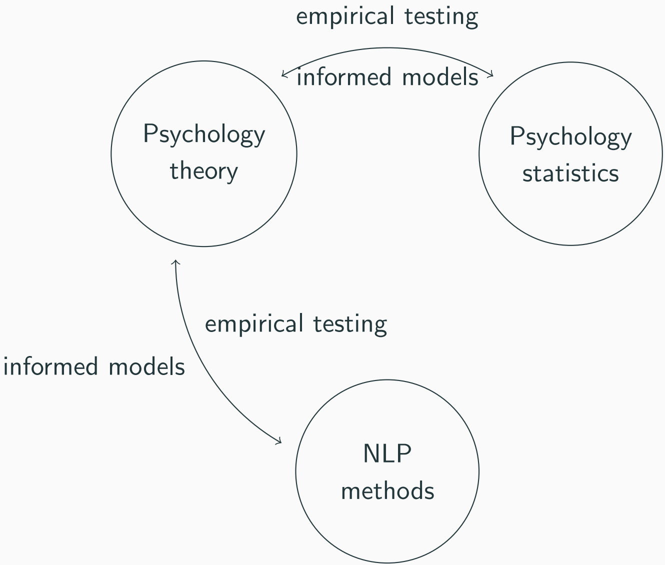

Benefits of NLP Methods

- Psychology and health research tends to rely on human experts and shy away from textual data

Benefits of NLP Methods

- Psychology and health research tends to rely on human experts and shy away from textual data
- Analyzing motives for self-injury (**Snir et al.**)
  - Participants kept numerical diaries about acts and urges of self-harm

Benefits of NLP Methods

- Psychology and health research tends to rely on human experts and shy away from textual data
- Analyzing motives for self-injury (**Snir et al.**)
  - Participants kept numerical diaries about acts and urges of self-harm
  - **Textual analysis can support or challege survey answers and provide deeper insight into the variables of interest**

Benefits of NLP Methods

- Psychology and health research tends to rely on human experts and shy away from textual data
- Analyzing motives for self-injury (**Snir et al.**)
  - Participants kept numerical diaries about acts and urges of self-harm
  - Textual analysis can support or challege survey answers and provide deeper insight into the variables of interest
- Analyzing the effects of the implicit power motive (**Ditlmann et al.**)
  - Implicit power motive: individual's tendency to seek influence

Benefits of NLP Methods

- Psychology and health research tends to rely on human experts and shy away from textual data
- Analyzing motives for self-injury (**Snir et al.**)
  - Participants kept numerical diaries about acts and urges of self-harm
  - Textual analysis can support or challege survey answers and provide deeper insight into the variables of interest
- Analyzing the effects of the implicit power motive (**Ditlmann et al.**)
  - Implicit power motive: individual's tendency to seek influence
  - Examined how this motive was expressed and received

Benefits of NLP Methods

- Psychology and health research tends to rely on human experts and shy away from textual data
- Analyzing motives for self-injury (**Snir et al.**)
    - Participants kept numerical diaries about acts and urges of self-harm
    - Textual analysis can support or challege survey answers and provide deeper insight into the variables of interest
- Analyzing the effects of the implicit power motive (**Ditlmann et al.**)
    - Implicit power motive: individual's tendency to seek influence
    - Examined how this motive was expressed and received
    - **Text data was coded and manipulated by hand, so sample size was small**

NLP Methods and Psychological Theories

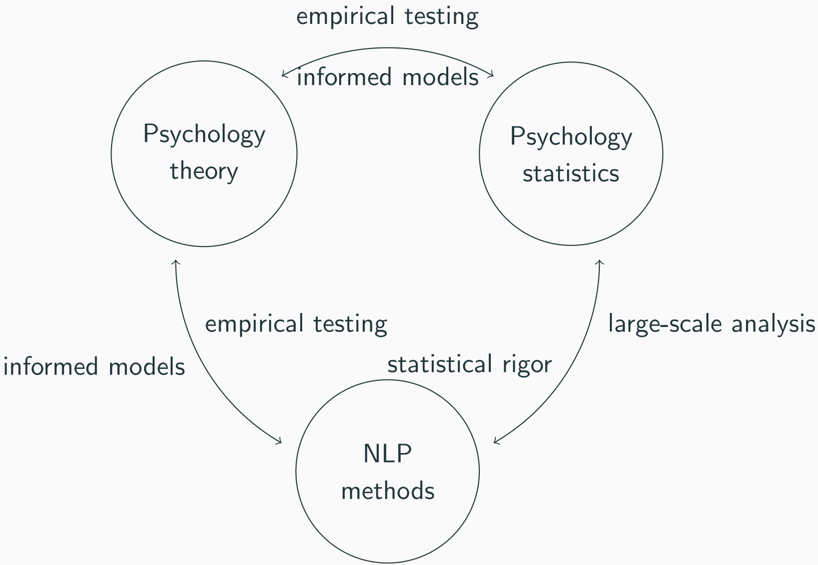

Integration: NLP Tasks

Paper	Li et al.	Rosenthal & McKeown
Task	targeted sentiment	influence detection
Theory	social cognitive theories	Cialdini's weapons of influence
Implementation	postprocessing predictions to conform to theory (e.g., attitudes are globally consistent)	features fed into linear classifier (predicted traits of author and of the majority in the discussion)

Integration: NLP Tasks

Paper	Li et al.	Rosenthal & McKeown
Task	targeted sentiment	influence detection
Theory	social cognitive theories	Cialdini's weapons of influence
Implementation	postprocessing predictions to conform to theory (e.g., attitudes are globally consistent)	features fed into linear classifier (predicted traits of author and of the majority in the discussion)

- Incorporating theory allows us to test that theory on a large scale

Integration: NLP Tasks

Paper	Li et al.	Rosenthal & McKeown
Task	targeted sentiment	influence detection
Theory	social cognitive theories	Cialdini's weapons of influence
Implementation	postprocessing predictions to conform to theory (e.g., attitudes are globally consistent)	features fed into linear classifier (predicted traits of author and of the majority in the discussion)

- Incorporating theory allows us to test that theory on a large scale
- **We must consider what conclusions we can make and what our models measure**

Integration: Mental Health and Psychology

Paper	**De Choudhury et al.**	**Doré et al.**	**Bedi et al.**
Task	predict future suicidal ideation	predict emotional features of text given time and location	predict future incidence of psychosis
Empirical work	mental health	emotion, trauma	schizophrenia
Implementation	handcrafted semantic and syntactic features	LIWC	LSA, syntactic features

Integration: Mental Health and Psychology

Paper	De Choudhury et al.	Doré et al.	Bedi et al.
Task	predict future suicidal ideation	predict emotional features of text given time and location	predict future incidence of psychosis
Empirical work	mental health	emotion, trauma	schizophrenia
Implementation	handcrafted semantic and syntactic features	LIWC	LSA, syntactic features

- Large-scale NLP methods complement domain knowledge
- Careful analysis of textual features bolsters their trustworthiness

Integration: Mental Health and Psychology

Paper	**De Choudhury et al.**	**Doré et al.**	**Bedi et al.**
Task	predict future suicidal ideation	predict emotional features of text given time and location	predict future incidence of psychosis
Empirical work	mental health	emotion, trauma	schizophrenia
Implementation	handcrafted semantic and syntactic features	LIWC	LSA, syntactic features

- Large-scale NLP methods complement domain knowledge
- Careful analysis of textual features bolsters their trustworthiness
- **We must consider what conclusions we can make and what our models measure**

Conclusion

Recap

- **Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this**

Recap

- Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this
- **Emotion has been recognized as a valid component of influence, but is understudied in the computational literature**

Recap

- Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this
- Emotion has been recognized as a valid component of influence, but is understudied in the computational literature
- **Affectual text processing is a growing field, but lacking in psychological grounding**

Recap

- Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this
- Emotion has been recognized as a valid component of influence, but is understudied in the computational literature
- Affectual text processing is a growing field, but lacking in psychological grounding
- **External theory and NLP methods complement one another to draw conclusions supported by domain experts and the power of big data**

Recap

- Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this
- Emotion has been recognized as a valid component of influence, but is understudied in the computational literature
- Affectual text processing is a growing field, but lacking in psychological grounding
- External theory and NLP methods complement one another to draw conclusions supported by domain experts and the power of big data
- **If we apply it carefully, external theory can help us build more informed and interpretable models that better represent the real world**

Recap

- Humans produce and consume the text we study (so far), and our models lack crucial information if we do not recognize this
- Emotion has been recognized as a valid component of influence, but is understudied in the computational literature
- Affectual text processing is a growing field, but lacking in psychological grounding
- External theory and NLP methods complement one another to draw conclusions supported by domain experts and the power of big data
- If we apply it carefully, external theory can help us build more informed and interpretable models that better represent the real world
- **Thank you for listening!**