CARER: Contextualized Affect Representations for Emotion Recognition

Elvis Saravia, Hsien-Chi Toby Liu, Yen-Hao Huang, Junlin Wu, Yi-Shin Chen





Latinx in AI Coalition @ NeurIPS, Montreal, Canada, 2018
National Tsing Hua University, Taiwan

Emotion Analysis

Task: To detect the *fine-grained emotions* expressed in textual information.

Challenge: Emotions are communicated using a ***variety of linguistic phenomena*** due to social and cultural differences: slang, emoticons, abbreviations, etc.

Thanks God for everything #safe

Tnx mom for waaaaking me two hours early.
Can’t fall asleep now #angry

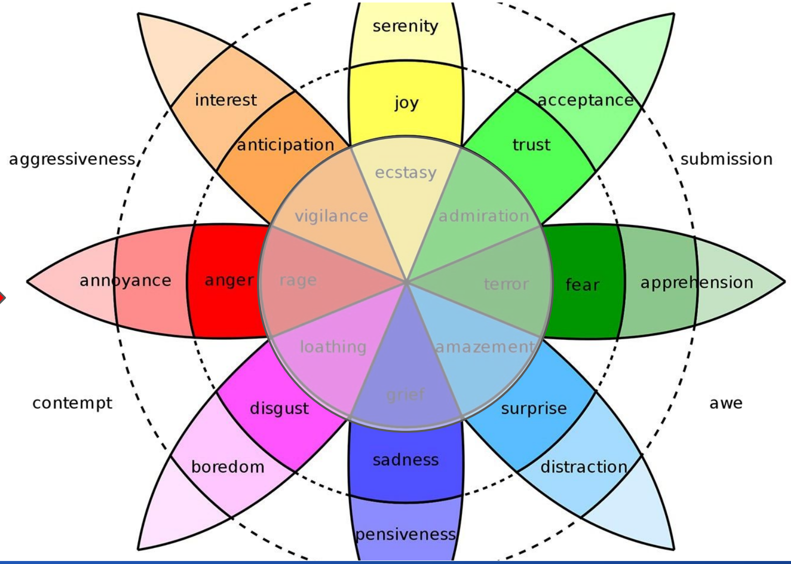

Discovering Emotional Language

Due to linguistic variability, we need a robust method that properly models and captures both ***contextual*** and ***implicit*** emotional information

"Thanks God for everything"

"thanks goodness for a great team .. "

"thanks **fro** all the continued support and prayers gotta keep working hard!!"

"Thanks for all the tweets… onto the next path now."

"Thanks mom for waaaaking me two hours early. Can't get asleep now."

"thanks dad i can always count on u to mess up my day"

Graph-Based Pattern Extraction

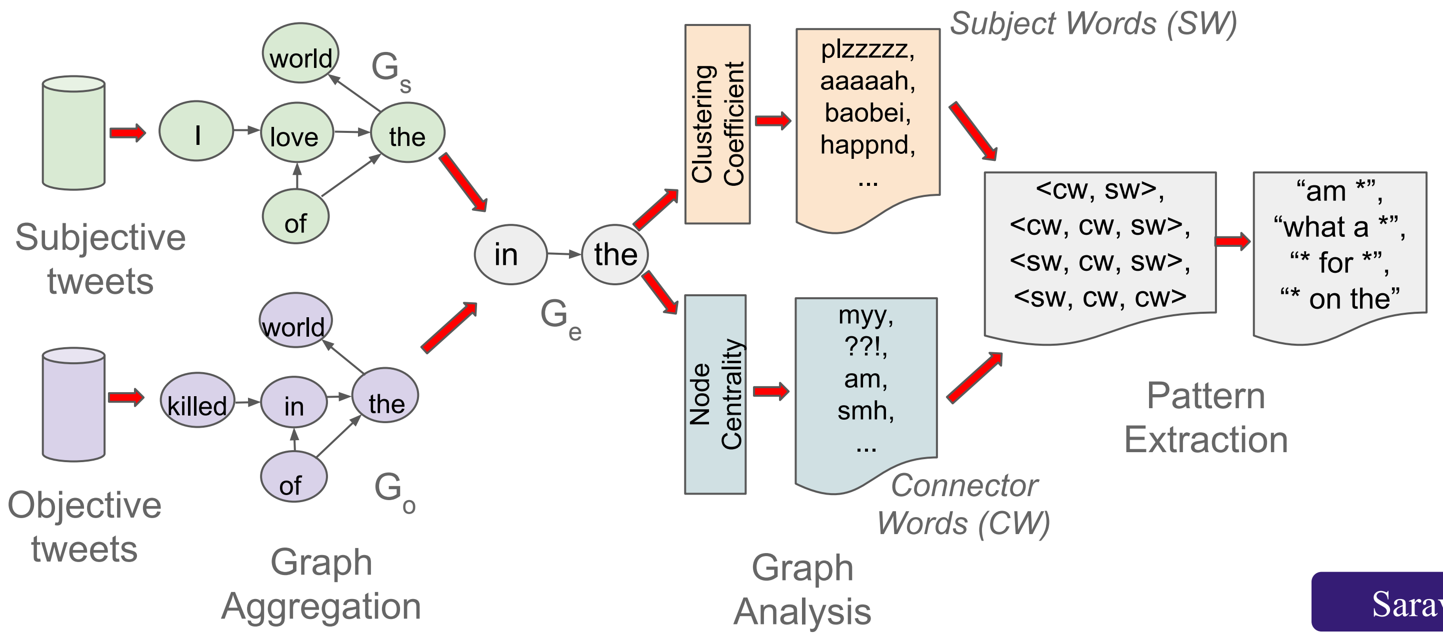

Saravia et al., 2016

Are Graph-Based Emotion Patterns Enough?

Joy Pattern: "thanks *"

"Thanks God for everything"
"thanks goodness for a great team .. "

"thanks **fro** all the continued support and prayers gotta keep working hard!!"
"Thanks for all the tweets… onto the next path now."

"Thanks mom for waaaaking me two hours early. Can't get asleep now."
"thanks dad i can always count on u to mess up my day"

Wildcards (i.e., *) are helpful for preserving structure and generalizing **but cannot preserve semantic relationships**

Objectives

- Build an automatic graph-based algorithm for **emotion-relevant feature extraction**
- Construct **contextualized representations that preserve semantic relationships**
- Analyse model for **robustness and explainability** given the proposed representations

Affective

Contextual

Explainable

Methodology

- → Building Syntactic Patterns
- → Contextualizing Patterns
- → Representation Learning

Pattern Refinement & Enrichment

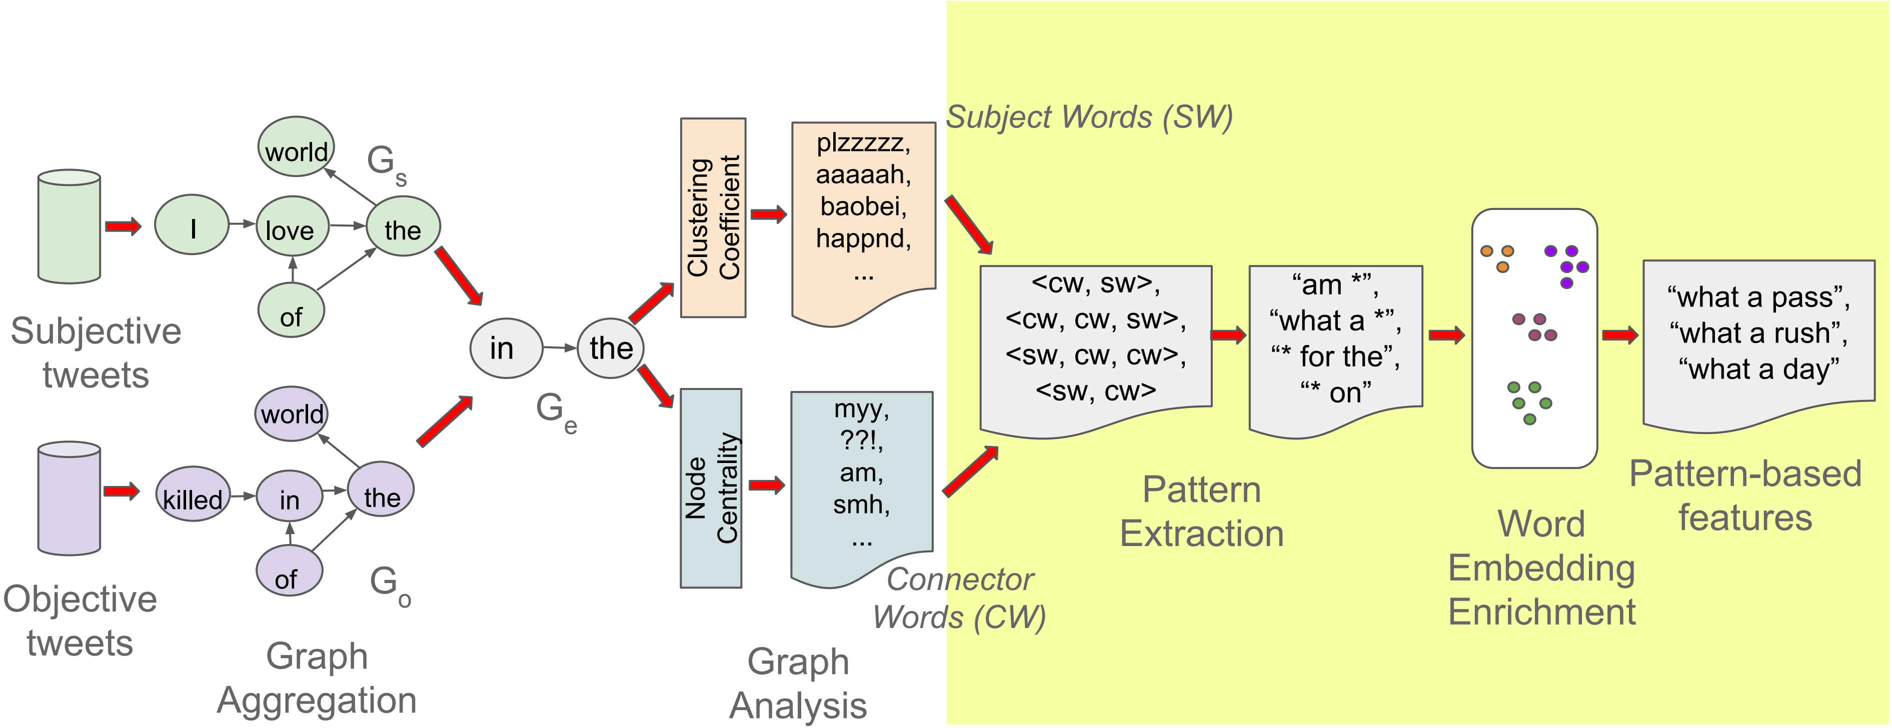

Step 1: Cluster Word Embeddings

Word2Vec

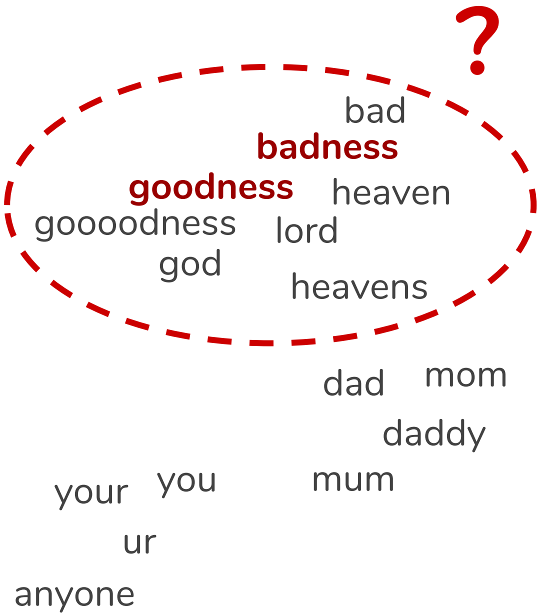

Mikolov et al., 2013

Goal:

- Model words using Word2Vec
- Cluster words based on similarity measure
- Antonyms are close due to similar context

So its **badness** would be …
would its **goodness** be revealed…
… about the **badness** of human.
Goodness of a human ...

Step 2: Update Word Embeddings

Pretrained Word2Vec Embeddings

Sentiment Updated Word2Vec

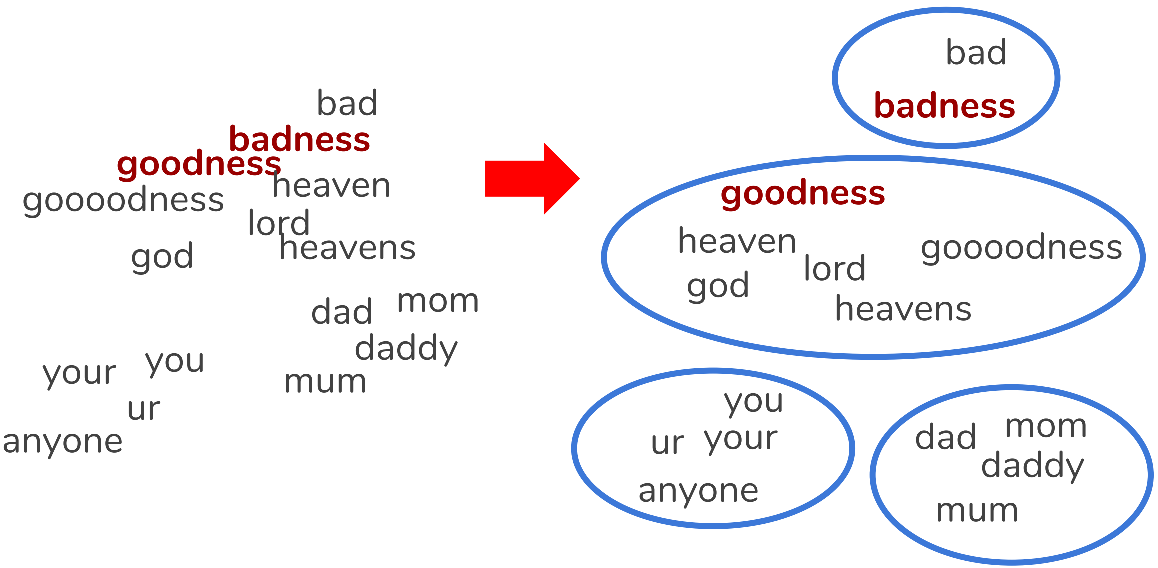

Deriu et al., 2016

Sentiment Task

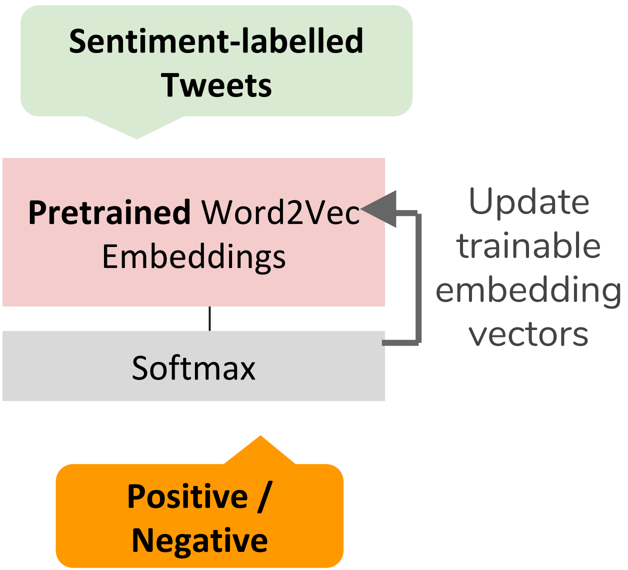

Step 3: Preserving semantic relationship

Pattern	Text	Contextualized Pattern
thanks *	"thanks god" "Thanks goodness" "Thanks gooodness"	thanks C#58
	"thanks your"	thanks C#90
	"thanks mom" "thanks mum"	thanks C#28

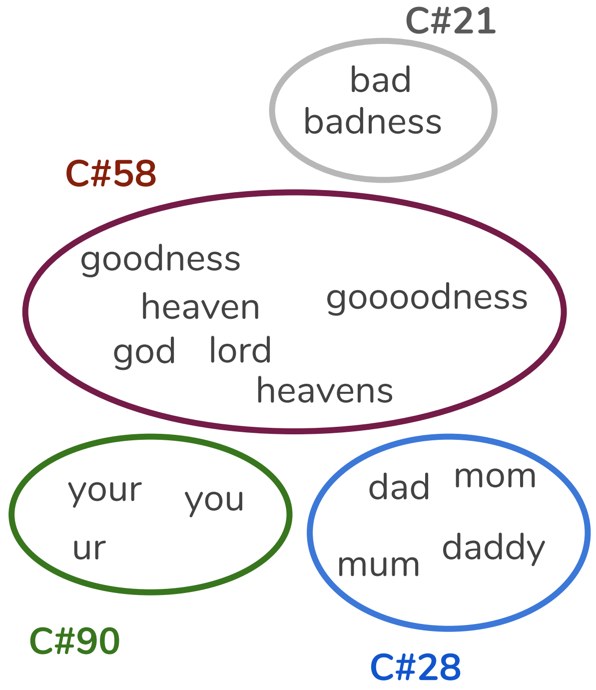

Input: Contextualized Patterns

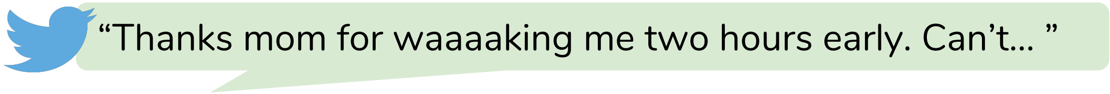

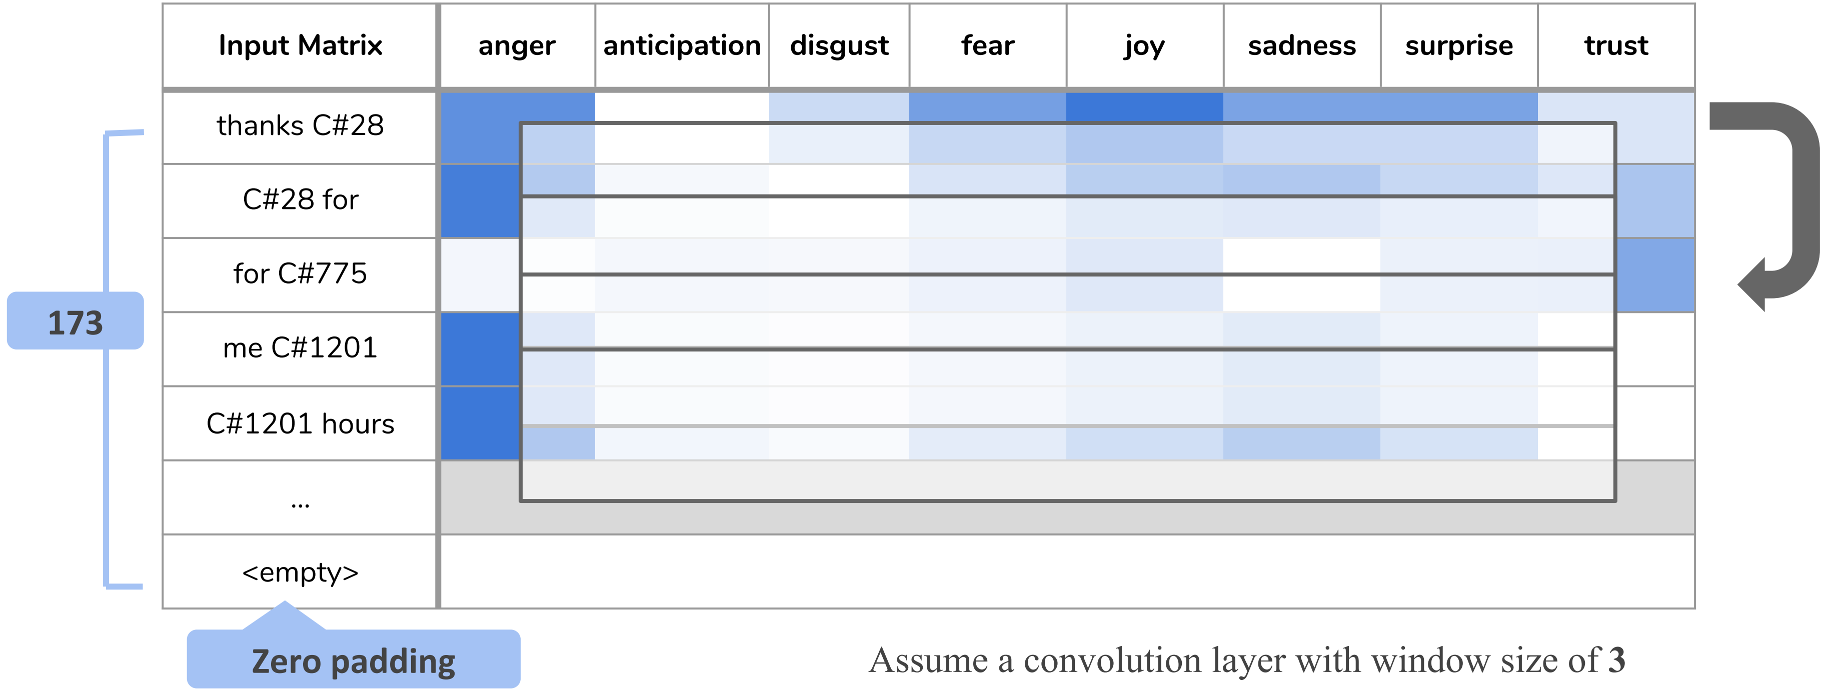

Assume a convolution layer with window size of **3**

CNN-Based Emotion Recognition Model

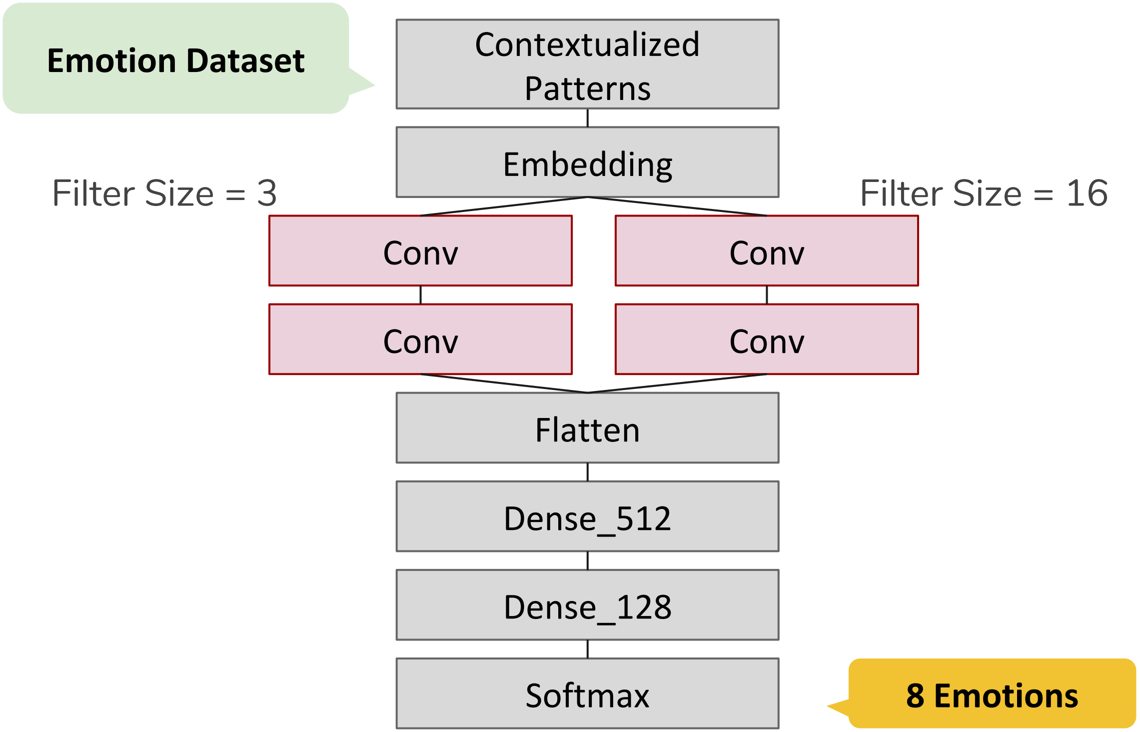

- CNN-based (multilayer)
- Zero padding
- **Pattern scores** as embedding vectors
- Two filter sizes for features with different length

Experiments

Emotion Recognition for English Short Texts

Experimental Setup: Dataset

- Crawl English tweets
  - Annotated via distant supervision
  - Total 339 hashtags corresponding to the eight emotions
- Refining process (Abdul et al., 2017)
  - 0.66M tweets in total

Emotion	Amount	Hashtags
sadness	214,454	#depressed, #grief
joy	167,027	#fun, #joy
fear	102,460	#fear, #worried
anger	102,289	#mad, #pissed
surprise	46,101	#strange, #surprise
trust	19,222	#hope, #secure
disgust	8,934	#awful, #eww
anticipation	3,975	#pumped, #ready

Experimental Results: 8 Emotions Task

	Model	Feature	anger	antici-pation	dis-gust	fear	joy	sad-ness	sur-prise	trust	F1-avg
Traditional Methods	BoW	word freq.	0.53	0.08	0.17	0.53	0.71	0.60	0.36	0.33	0.57
	N-gram	word freq.	0.56	0.09	0.17	0.57	0.73	0.64	0.42	0.39	**0.61**
	char	char. freq.	0.35	0.03	0.04	0.20	0.51	0.46	0.10	0.12	0.37
Lexica-based	LIWC	affect lexicons	0.35	0.03	0.11	0.30	0.49	0.35	0.18	0.19	0.35
State-of-the-Art Methods	CNN_{sw2v}	s-word embed.	0.57	0.10	0.15	0.63	**0.75**	0.64	0.61	**0.70**	**0.65**
	EmoNet	word embed.	0.36	0.00	0.00	0.46	0.69	0.61	0.13	0.25	0.52
	FastText	word embed.	0.57	0.01	0.01	0.65	0.77	0.71	0.50	0.54	0.66
	DeepMoji	word embed.	0.60	0.00	0.03	0.49	**0.75**	0.67	0.20	0.27	0.59
Baseline and our work	CNN_{EP}	EmoPattern	0.65	0.10	0.22	0.64	0.73	0.56	0.15	0.08	0.52
	CARER_{β}	cont. patt.‡	0.61	0.31	0.34	0.67	**0.75**	0.68	0.60	0.55	**0.67**
	CARER	cont. patt.	**0.74**	**0.41**	**0.43**	**0.79**	**0.83**	**0.82**	**0.76**	**0.75**	**0.79**

Note: CARER uses a recent dataset and fewer pattern templates (details in paper)

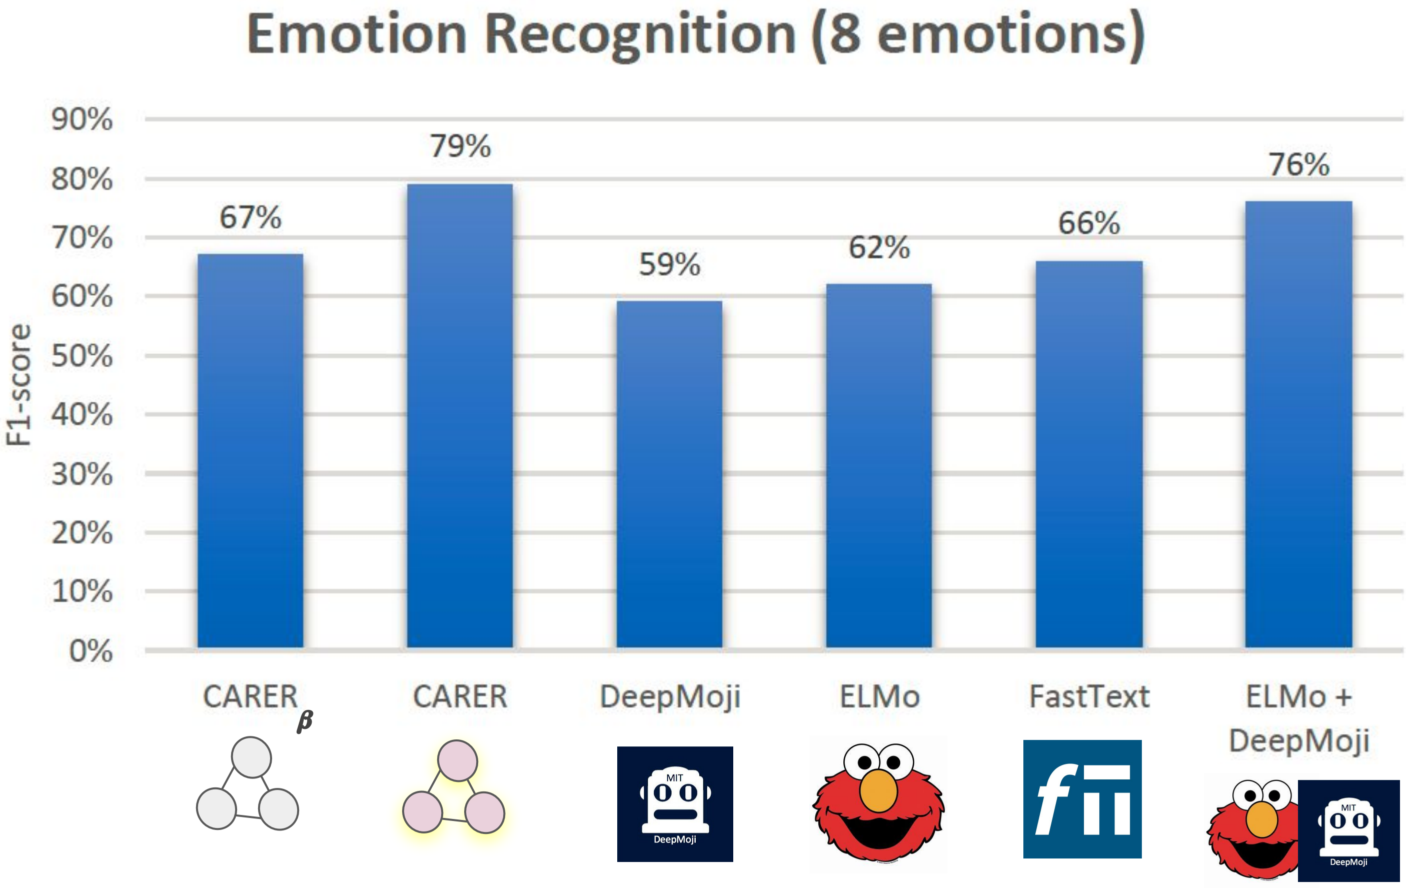
Emotion Recognition (8 emotions)
F1-score
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
67%
79%
59%
62%
66%
76%
CARER
β
CARER
DeepMoji
ELMo
FastText
ELMo +
DeepMoji
MIT
DeepMoji
MIT
DeepMoji

What's Captured by CARER?

Our proposed method can grasp emotional cues in cases of ***short text***, ***rare words*** and ***mixed emotions***

Case	Document	Label	DeepMoji	EmoNet	CARER	Contextualized Pattern
Short text	damn what a night	joy	surprise	sadness	joy	what a {**night, day, rush, pass**}
Rare words	got thee worst sleep ever	anger	sadness	sadness	anger	got {**thee, madd, thatt, bacc**}
Mixed emotions	what the h**k is going on !?	fear	anger	sadness	fear	is {**going, ends, finishes**}

Conclusion

- We proposed **contextualized affect representations** for improving emotion recognition
- In the future, we anticipate a comprehensive study of how contextualized patterns can be adapted to other **emotion-related tasks**

References

- Elvis Saravia, Carlos Argueta, and Yi-Shin Chen. *"Unsupervised graph-based pattern extraction for multilingual emotion classification."* Social Network Analysis and Mining, 2016.
- J. Deriu, et al. *"Swisscheese at semeval-2016 task 4: Sentiment classification using an ensemble of convolutional neural networks with distant supervision."* Procs. of SemEval, 2016.
- Abdul-Mageed, et al., *"Emonet: Fine-grained emotion detection with gated recurrent neural networks."* Procs. of ACL, 2007.
- Felbo, Bjarke, et al. *"Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm."* arXiv preprint arXiv:1708.00524, 2017.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. *Enriching word vectors with subword information.* Transactions of the Association for Computational Linguistics, 2017.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. *"Efficient estimation of word representations in vector space."* arXiv preprint arXiv:1301.3781, 2013