Twitter Stance Detection with Bidirectional Conditional Encoding

Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, Kalina Bontcheva
{i.augenstein | t.rocktaschel}@cs.ucl.ac.uk, {a.vlachos | k.bontcheva}@sheffield.ac.uk

Stance Detection

Classify attitude of tweet towards target as “favor”, “against”, “none”

Tweet: “No more Hillary Clinton” **Target**: Donald Trump **Stance**: FAVOR

Training targets: Climate Change is a Real Concern, Feminist Movement, Atheism, Legalization of Abortion, Hillary Clinton

Testing target: Donald Trump

Challenges

- Tweet interpretation depends on target
  - **Solution**: bidirectional conditional model
- Labelled data not available for the test target
  - **Solution**: 1) domain adaptation; 2) weak labelling

Data

- SemEval 2016 Task 6 Twitter Stance Detection corpus
- 5 628 **labelled training** tweets **(**1 278 about Hillary Clinton, used for **dev)**
- 278 013 **unlabelled** Donald Trump tweets
- 395 212 additional **unlabelled** tweets about all targets
- 707 Donald Trump **testing** tweets

Conclusions

- Learning conditional sequence representation of targets and tweets better than representations without conditioning
- State of the art performance with weakly supervised model
- Good target representations with unsupervised pre-training
- Successful sequence representation learning with small data

References

SemEval 2016 Stance Detection: http://alt.qcri.org/semeval2016/task6/
Code: https://github.com/sheffieldnlp/stance-conditional

Bidirectional Conditional Encoding

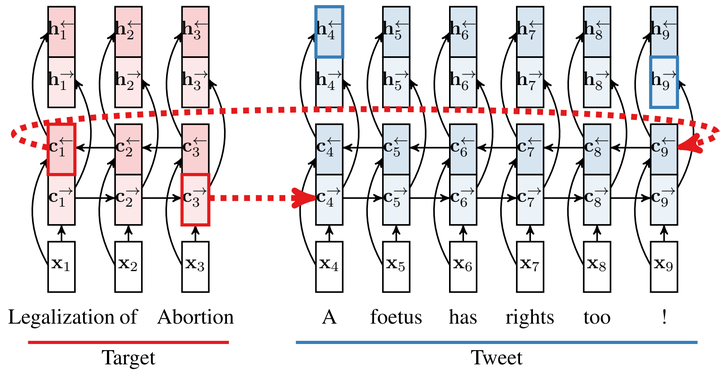

Bidirectional encoding of tweet conditioned on bidirectional encoding of target $[\vec{\mathbf{c}_3}\ \overleftarrow{\mathbf{c}_1}]$
The stance is predicted using the last forward and reversed output representations $[\vec{\mathbf{h}_9}\overleftarrow{\mathbf{h}_4}]$

Experiments

Stance Detection Models

- Sequence-agnostic: *BoWV*
- Target-agnostic: Tweet-only LSTM encoding (*TweetOnly*)
- Target+Tweet, no conditioning: Concatenated target and tweet LSTM encodings (*Concat*)
- Target+Tweet, conditioning: Target conditioned on tweet (*TarCondTweet*); tweet conditioned on target (*TweetCondTar*); bidirectional conditioning (*BiCond*)

Word Embedding Training

- Random initialisation
- Trained fixed word embeddings on tweets
- Pre-trained word embeddings on tweets, continue training with supervised objective

Acknowledgements

This work was partially supported by the European Union, grant No. 611233 PHEME (http://www.pheme.eu) and by the Microsoft Research PhD Scholarship Programme.

Setting

Unseen Target Setting

- Train on Climate Change Is A Real Concern, Feminist Movement, Atheism, Legalization of Abortion, Hillary Clinton tweets
- Test on Donald Trump tweets

Weakly Supervised Setting

- Weakly label Donald tweets using hashtags for training
- Test on Donald Trump tweets

Results

Unseen Target Models

Weakly Supervised Models

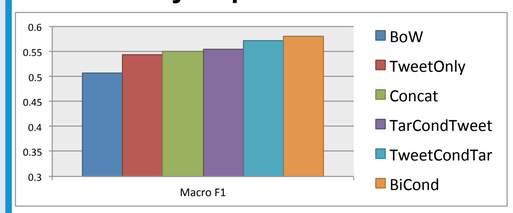

Comparison Against SOA

Model	Macro F1
SVM-ngrams-comb (official)	0.2843
BiCond (Unseen Target)	0.4901
pkudblab (Weakly Supervised*)	0.5628
BiCond (Weakly Supervised)	**0.5803**