Hybrid approaches to fine-grained emotion detection in social media data

Annika Marie Schoene

The University of Hull Department of Computer Science amschoene@googlemail.com

Abstract

This paper states the challenges in fine-grained targetdependent Sentiment Analysis for social media data using recurrent neural networks. Firstly, we outline the problem statement and give a brief overview of related work in the area. Then we outline progress and results achieved to date, a brief research plan and future directions of this work.

Problem Statement

Whilst recurrent neural networks (RNN) are well suited towards natural language processing tasks such as Sentiment Analysis (SA) due to their ability to handle sequential data, there are still shortcomings which ultimately affect the accurate classification of longer sequences. The focus of this research is in SA of social media data, utilising RNNs. More specifically, we look at applying RNNs to classify finegrained emotions, based on Ekman's six basic emotions (Ekman et al. 1987) in Twitter data. Traditionally, tweets have been categorised as short sequences or sentence-level SA (Kouloumpis, Wilson, and Moore 2011), however, it could be argued that this should no longer be the case especially since Twitter increased its allowed character limit from 140 to 280 (Twitter 2018). Subsequently, many tweets could also be seen as longer sequences, that face problems associated to other long sequence classification tasks in natural language processing tasks (Hochreiter and Schmidhuber 1997). As a result of this it is even more important to classify long sequences accurately, because often important information that indicates a person's emotions, attitude or opinions are expressed through the use of emojis and often appear towards the end of a tweet (Novak et al. 2015). Therefore we investigate, the use of multiscale RNNs to overcome the issue of learning more frequently occurring short-term information whilst retaining long-term information. Current SA methodologies predominantly focus on the classification of polarities or pre-determined moods (Cambria et al. 2017). SA has been further extended by introducing the task of identifying the aspects or targets of polarities (Ma, Peng, and Cambria 2018) in reviews. Whilst aspect-based SA has made

tremendous progress using RNNs and attention, in comparison target-dependent SA for Twitter data has been limited to attempts utilising dependency parsing, but using some RNNs in more recent work. Nonetheless, data produced on social media platforms such as Twitter are still important for generating insights into political events, product marketing or mental health. Only utilising polarities and single targets may lead to not generating the right insights into the data and may make people draw misleading conclusions from it (Wang et al. 2017). The future of this work will look at investigating the use of external knowledge to identify target entities and fine-grained emotions using multi-scale RNNs that are capable of learning long-term dependencies. The main contributions of this thesis can be summarised as follows: (1) a hybridised multi-scale recurrent neural network based learning model that can model long sequences of textual data. (2) Such a model will be able to utilise external knowledge to classify textual data into fine-grained emotions and identify the target-entities related to these emotions. It is hypothesised that a hybrid approach combining the strength of both statistical and knowledge-based approaches will enable more fine-grained sentiment detection. This hybrid approach would be able to capture the overarching emotions in large unstructured textual datasets and identify the relations and concepts on a lower level of the data.

Related work

SA has often been summarized as the 'computational study of opinions, sentiments and attitudes concerning different topics, as expressed in textual input' (Ravi and Ravi 2015). There are a range of different tasks that are associated to the broader field of SA and it has been described as a suitcase of different tasks, where an essential function of SA is polarity detection and emotion recognition (Cambria et al. 2017). Ekman's six basic emotions have often been used in fine-grained emotion detection within the field of Sentiment Analysis. This includes, but is not limited to emotion annotation for blog posts and chat message annotation of online conversations to determine if the content is emotional or not (Meiselman 2016). Other research has looked into the use of affect lexicons for classifying Ekman's six emotions in sentences (Mohammad 2012). Most recently, a shared

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

task (Klinger et al. 2018) has taken place using Tweets that were collected based on Ekman's six emotions and the main task was to classify tweets accurately when the 'emotion keyword' was removed, where the most successful model was RNN-based. Several different approaches have looked at utilising RNNs for sequence learning. Most notably work by Hochreiter and Schmidhuber has proposed the LSTM cell to circumvent problems associated with RNNs, such as vanishing and exploding gradient and learning long-term dependencies. Other successful approaches to this are the use of multiscale RNNs (Chung, Ahn, and Bengio 2016) which group the hidden units of the network into multiple modules that operate on different timescales. Another research direction has looked at the optimisation of the recurrence matrix in RNNs by proposing approaches such as orthogonal matrices for initialisation (Le, Jaitly, and Hinton 2015) or gradient clipping (Pascanu, Mikolov, and Bengio 2013). Targetdependent SA has focused on both identifying single and multiple target entities (Wang et al. 2017), where early work has used traditional machine learning algorithms as well as using external knowledge (i.e. lexicons). More recent work has looked at using dependency parsing (Wang et al. 2017) as well as RNNs (Tang et al. 2015). More recent approaches using deep learning techniques (Song et al. 2019) have also noted the limitations when using RNNs for such tasks due to their lacking ability to learn long-term dependencies.

Progress and Research Plan

The first phase of this work has focused on an in-depth literature review and data collection from Twitter. This included applying for ethical approval and cleaning, anonymising and storing collected data. The second phase concentrated on establishing benchmarks on the collected dataset by implementing various RNN architectures and conducting experiments to compare model performance. This include the implementation and extension of a bi-directional dilated RNN with attention for long sequence classification of emotions in tweets. At the point of this submission and the start of the third phase of this work we will extend current work to target-dependent SA, where there are three main objectives. First, we will focus on identifying target-entities in my data. This will be done by implementing, comparing and evaluating existing methodologies and potentially using human annotators. Then we will focus on building a knowledge-base containing target-entity and emotion pairs that will later be used to add 'internal' knowledge into the RNN's cell. Furthermore, we have worked on applying multi-scale RNNs to document classification tasks, where we classify suicide note from other types of discourse using both linguistic, cognitive and sentiment features (Schoene et al. 2019).

References

Cambria, E.; Poria, S.; Gelbukh, A.; and Thelwall, M. 2017. Sentiment analysis is a big suitcase. *IEEE Intelligent Systems* 32(6):74–80.

Chung, J.; Ahn, S.; and Bengio, Y. 2016. Hierarchical multiscale recurrent neural networks. *arXiv preprint arXiv:1609.01704*. Ekman, P.; Friesen, W. V.; O'sullivan, M.; Chan, A.; Diacoyanni-Tarlatzis, I.; Heider, K.; Krause, R.; LeCompte, W. A.; Pitcairn, T.; Ricci-Bitti, P. E.; et al. 1987. Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of personality and social psychology* 53(4):712.

Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Comput.* 9(8):1735–1780.

Klinger, R.; De Clercq, O.; Mohammad, S. M.; and Balahur, A. 2018. Iest: Wassa-2018 implicit emotions shared task. *arXiv preprint arXiv:1809.01083*.

Kouloumpis, E.; Wilson, T.; and Moore, J. D. 2011. Twitter sentiment analysis: The good the bad and the omg! *Icwsm* 11(538-541):164.

Le, Q. V.; Jaitly, N.; and Hinton, G. E. 2015. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*.

Ma, Y.; Peng, H.; and Cambria, E. 2018. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive lstm. In *Proceedings of AAAI*.

Meiselman, H. L. 2016. *Emotion Measurement*. Woodhead Publishing.

Mohammad, S. M. 2012. # emotional tweets. In *Proceedings of 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, 246–255. Association for Computational Linguistics.*

Novak, P. K.; Smailović, J.; Sluban, B.; and Mozetič, I. 2015. Sentiment of emojis. *PloS one* 10(12):e0144296.

Pascanu, R.; Mikolov, T.; and Bengio, Y. 2013. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, 1310–1318.

Ravi, K., and Ravi, V. 2015. A survey on opinion mining and sentiment analysis: tasks, approaches and applications. *Knowledge-Based Systems* 89:14–46.

Schoene, A. M.; Lacey, G.; Turner, A. P.; and Dethlefs, N. 2019. Dilated lstm with attention for classification of suicide notes. In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, 136–145.

Song, Y.; Wang, J.; Jiang, T.; Liu, Z.; and Rao, Y. 2019. Attentional encoder network for targeted sentiment classification. *arXiv preprint arXiv:1902.09314*.

Tang, D.; Qin, B.; Feng, X.; and Liu, T. 2015. Effective lstms for target-dependent sentiment classification. *arXiv preprint arXiv:1512.01100*.

Twitter. 2018. Counting characters. https://developer.twitter. com/en/docs/basics/counting-characters.html. Accessed on 2018-11-11.

Wang, B.; Liakata, M.; Zubiaga, A.; and Procter, R. 2017. Tdparse: Multi-target-specific sentiment recognition on twitter. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, 483–493.