



HAL
open science

Attentive Conversation Modeling for Emotion Detection and Classification

Waleed Ragheb, Jérôme Azé, Sandra Bringay, Maximilien Servajean

► **To cite this version:**

Waleed Ragheb, Jérôme Azé, Sandra Bringay, Maximilien Servajean. Attentive Conversation Modeling for Emotion Detection and Classification. SemEval 2019 - 13th International Workshop on Semantic Evaluation in NAACL-HLT, Jun 2019, Minneapolis, MN, United States. pp.251-255. lirmm-02145395

HAL Id: lirmm-02145395

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-02145395v1>

Submitted on 2 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LIRMM-Advanse at SemEval-2019 Task 3: Attentive Conversation Modeling for Emotion Detection and Classification

Waleed Ragheb^{1,2}, Jérôme Azé^{1,2}, Sandra Bringay^{1,3} and Maximilien Servajean^{1,3}

¹LIRMM UMR 5506, CNRS, University of Montpellier, Montpellier, France

²IUT de Béziers, University of Montpellier, Béziers, France

³AMIS, Paul Valéry University - Montpellier 3, Montpellier, France
{First.Last}@lirmm.fr

Abstract

This paper addresses the problem of modeling textual conversations and detecting emotions. Our proposed model makes use of 1) deep transfer learning rather than the classical shallow methods of word embedding; 2) self-attention mechanisms to focus on the most important parts of the texts and 3) turn-based conversational modeling for classifying the emotions. Our model was evaluated on the data provided by the SemEval-2019 shared task on contextual emotion detection in text. The model shows very competitive results.

1 Introduction

Emotional intelligence has played a significant role in many applications in recent years (Krakovsky, 2018). It is one of the essential abilities to move from narrow to general human-like intelligence. Being able to recognize expressions of human emotion such as interest, distress, and pleasure in communication is vital for helping machines choose more helpful and less aggravating behavior. Human emotions are a mental state that can be sensed and hence recognized in many sources such as visual features in images or videos (Boubenna and Lee, 2018), as textual semantics and sentiments in texts (Calefato et al., 2017) or even patterns in EEG brain signals (Jenke et al., 2014). With the increasing number of messaging platforms and with the growing demand of customer chat bot applications, detecting the emotional state in conversations becomes highly important for more personalized and human-like conversations (Zhou et al., 2018).

This paper addresses the problem of modeling a conversation that comes with multiple turns for detecting and classifying emotions. The proposed model makes use of transfer learning through the universal language modeling that is composed of

consecutive layers of Bi-directional Long Term Short Term Memory (Bi-LSTM) units. These layers are learned first in sequence-to-sequence fashion on a general text and then fine-tuned to a specific target task. The model also makes use of an attention mechanism in order to focus on the most important parts of each text turn. Finally, the proposed classifier models the changing of the emotional state of a specific user across turns.

The rest of the paper is organized as follows. In Section 2, the related work is introduced. Then, we present a quick overview of the task and the datasets in Section 3. Section 4 describes the proposed model architecture, some variants and hyperparameters settings. The experiments and results are presented in Section 5. Section 6 concludes the study.

2 Related Work

Transfer learning or domain adaptation has been widely used in machine learning especially in the era of deep neural networks (Goodfellow et al., 2016). In natural language processing (NLP), this is done through Language Modeling (LM). Through this step, the model aims to predict a word given some context. This is considered as a vital and important basics in most of NLP applications. Not only because it tries to understand the long-term dependencies and hierarchical structure of the text but also for its open and free resources. LM is considered as unsupervised learning process which needs only corpus of unlabeled text. The problem is that LMs get overfitted to small datasets and suffer catastrophic forgetting when fine-tuned with a classifier. Compared to Computer Vision (CV), NLP models are typically more shallow and thus require different fine-tuning methods. The development of the Universal Language Model Fine-tuning

(ULMFiT) (Howard and Ruder, 2018) is considered like moving from shallow to deep pre-training word representation. This idea has been proved to achieve CV-like transfer learning for many NLP tasks. ULMFiT makes use of the state-of-the-art AWD-LSTM (Average stochastic gradient descent - Weighted Dropout) language model (Merity et al., 2018). Weight-dropped LSTM is a strategy that uses a DropConnect (Wan et al., 2013) mask on the hidden-to-hidden weight matrices, as a means to prevent overfitting.

On the other hand, one of the recent trend in deep learning models is the attention Mechanism (Young et al., 2018). Attention in neural networks are inspired from the visual attention mechanism found in humans. The main principle is being able to focus on a certain region of an image with “high resolution” while perceiving the surrounding image in “low resolution”, and then adjusting the focal point over time. This is why the early applications for attention were in the field of image recognition and computer vision (Larochelle and Hinton, 2010). In NLP, most competitive neural sequence transduction models have an encoder-decoder structure (Vaswani et al., 2017). A limitation of these architectures is that it encodes the input sequence to a fixed length internal representation. This cause the results going worse performance for very long input sequences. Simply, attention tries to overcome this limitation by guiding the network to learn where to pay close attention in the input sequence. Neural Machine Translation (NMT) is one of the early birds that make use of attention mechanism (Bahdanau et al., 2014). It has recently been applied to other problems like sentiment analysis (Ma et al., 2018) and emotional classification (Majumder et al., 2018).

3 Data

The datasets are collections of labeled conversations (Chatterjee et al., 2019b). Each conversation is a three turn talk between two persons. The conversation labels correspond to the emotional state of the last turn. Conversations are manually classified into three emotional states for *happy*, *sad*, *angry* and one additional class for *others*. In general, released datasets are highly imbalanced and contains about 4% for each emotion in the validation (development) set and final test set. Table 1 shows the number of conversations examples and emotions provided in the official released datasets.

Dataset	Data size	Happy	Sad	Angry
Training	30160	5191	6357	6027
Validation (Dev)	2755	180	151	182
Testing	5509	369	308	324

Table 1: Used datasets.

4 Proposed Models

4.1 Model Architecture

In figure 1, we present our proposed model architecture. The model consists of two main steps: encoder and classifier. We used a linear decoder to learn the language model encoder as we will discuss later. This decoder is replaced by the classifier layers. The input conversations come in turns of three. After tokenization, we concatenate the conversation text but keep track of each turn boundaries. The overall conversation is inputted to the encoder. The encoder is a normal embedding layer followed by AWD-LSTM block. This uses three stacked different size Bi-LSTM units trained by ASGD (Average Stochastic Gradient Descent) and managed dropout between LSTM units to prevent overfitting. The conversation encoded output has the form of $C_{Enc} = [T_{Enc}^1 \oplus T_{Enc}^2 \oplus T_{Enc}^3]$ where T^i is the i^{th} turn in the conversation and \oplus denotes a concatenation operation and $T_{Enc}^i = \{T_1^i, T_2^i, \dots, T_{N_i}^i\}$. The sequence length of turn i is denoted by N_i . The size of T_j^i is the final encoding of the j 's sequence item of turn i .

For classification, the proposed model pays close attention to the first and last turns. The reasons behind this are that the problem is to classify the emotion of the last turn. Also, the effect of the middle turn appear implicitly on the encoding of the last turn as we used Bi-LSTM encoding on the concatenated conversation. In addition to these, tracking the difference between the first and the last turn of the same person may be beneficial in modeling the semantic and emotional changes. So, we apply self-attention mechanism followed by an average pooling to get turn-based representation of the conversation. The attention scores for the i^{th} turn S^i is given by:

$$S^i = \text{Softmax}\{W_i \cdot T_{Enc}^i\} \quad (1)$$

Where W_i is the weight of the attention layer of the i^{th} turn and S^i has the form of $S^i = \{S_1^i, S_2^i, \dots, S_{N_i}^i\}$. The output of the attention layer is the scoring of the encoded turn sequence $O^i =$




Figure 1: Proposed model architecture (*Model-A*).

$\{o_1^i, o_2^i, \dots, o_{N_i}^i\}$ which has the same length as the turn sequence and is given by $O^i = S^i \odot T_{Enc}^i$ where \odot is the element-wise multiplication. The difference of the pooled scored output of O^1 and O^3 is computed as O_{diff} . The Input of the linear block is X_{in} is formed by:

$$X_{in} = [O_{diff} \oplus O_{pool}^3] \quad (2)$$

The fully connected linear block consist of two different sized dense layers followed by a Softmax to determine the target emotion of the conversation.

4.2 Training Procedures

Training the overall models comes into three main steps: 1) The LM is randomly initialized and then trained by stacking a linear decoder in top of the encoder. The LM is trained on a general-domain corpus. This helps the model to get the general features of the language. 2) The same full LM after training is used as an initialization to be fine-tuned using the data of the target task (conversation text). In this step we limit the vocabulary of the LM to the frequent words (repeated more than twice) of target task. 3) We keep the encoder and replace the decoder with the classifier and both are fine-tuned on the target task.

For training the language model, we used the Wikitext-103 dataset (Merity et al., 2016). We train the model on the forward and backward LMs for both the general-domain and task specific datasets. Both LMs -backward and forward- are used to build two versions of the same proposed architecture. The final decision

is the ensemble of both. Our code is released at <https://github.com/WaleedRagheb/Attentive-Emocontext>.

4.3 Model Variations

In addition to the model -*Model-A*- described by Figure 1, we tried five different variants.

The first variant -(*Model-B*)- is formed by bypassing the self attention layer. This will pass the output of the encoder directly to the average pooling layer such that $X_{in}^B = [T_{diff} \oplus T_{pool}^3]$ where T_{diff} is the difference between the first and third pooled encoded turns of the conversations.

-(*Model-C*)- is to input a pooled condensed representation to the whole conversation C_{pool} rather than the last turn to the linear layer block. In this case: $X_{in}^C = [O_{diff} \oplus C_{pool}]$. We also studied two versions of the basic model where only one input is used $X_{in}^D = O_{diff}$ -(*Model-D*)- and $X_{in}^E = O_{pool}^3$ -(*Model-E*). In these two variants, we just change the size of the first linear layer.

Also, we apply the forward direction LM and classifier only without ensemble them with the backward direction and keep the same basic architecture -(*Model-F*).

4.4 Hyperparameters

We use the same set of hyperparameters across all model variants. For training and fine-tuning the LM, we use the same set of hyperparameter of AWD-LSTM proposed by (Merity et al., 2018) replacing the LSTM with Bi-LSTM. For classifier, we used masked self-attention layers and average pooling. For the linear block, we used hidden linear layer of size 100 and apply dropout of 0.4. We

Models	Results									
	Happy			Sad			Angry			Micro
	P	R	F1	P	R	F1	P	R	F1	F1
A	0.7256	0.7077	0.7166	0.8291	0.776	0.8017	0.7229	0.8054	0.7619	0.7582
B	0.7341	0.6514	0.6903	0.7401	0.82	0.778	0.7049	0.8255	0.7604	0.7439
C	0.7279	0.6972	0.7122	0.7765	0.792	0.7842	0.6941	0.8221	0.7527	0.7488
D	0.7214	0.7113	0.7163	0.8128	0.764	0.7876	0.6965	0.8087	0.7484	0.749
E	0.7204	0.7077	0.714	0.8205	0.768	0.7934	0.7026	0.8087	0.752	0.7512
F	0.7336	0.669	0.6998	0.8377	0.764	0.7992	0.738	0.7752	0.7561	0.75

Table 2: Test set results of the basic proposed model and it’s variants.

used Adam optimizer (Dozat and Manning, 2017) with $\beta_1 = 0.8$ and $\beta_2 = 0.99$. The base learning rate is 0.01. We used the same batch size used in training LMs but we create each batch using weight random sampling. We used the same weights (0.4 for each emotion). We train the classifier on training set for 30 epochs and select the best model on validation set to get the final model.

5 Results & Discussions

The results of the test set for different variants of the model for each emotion is shown in table 2. The table shows the value of precision (P), recall (R) and F1 measure for each emotion and the micro-F1 for all three emotional classes. The micro-F1 scores are the official metrics used in this task. *Model-A* gives the best performance F1 for each emotion and the overall micro-F1 score. However some variants of this model give better recall or precision values for different emotions, *Model-A* compromise between these values to give the best F1 for each emotion. Removing the self-attention layer in the classifier -*Model-B*- degraded the results. Also, inputting a condensed representation of the all conversation rather than the last turn -*Model-C*- did not improve the results. Even modeling the turns difference only -*Model-D*- gives better results over *Model-C*. These proves empirically the importance of the last turn in the classification performance. This is clear for *Model-E* where the classifier is learned only by inputting the last turn of the conversation. Ensemble the forward and backward models was more useful than using the forward model only -*Model-F*.

Comparing the results for different emotions and different models, we notice the low performance in detecting happy emotion. This validate

the same conclusion of Chatterjee et.al in (2019a). The model shows a significant improvement over the EmoContext organizer baseline (F1: 0.5868). Also, comparing to other participants in the same task with the same datasets, the proposed model gives competitive performance and ranked 11th out of more than 150 participants. The proposed model can be used to model multi-turn and multi-parties conversations. It can be used also to track the emotional changes in long conversations.

6 Conclusions

In this paper, we present a new model used for Semeval-2019 Task-3 (Chatterjee et al., 2019b). The proposed model makes use of deep transfer learning rather than the shallow models for language modeling. The model pays close attention to the first and the last turns written by the same person in 3-turn conversations. The classifier uses self-attention layers and the overall model does not use any special emotional lexicons or feature engineering steps. The results of the model and it’s variants show a competitive results compared to the organizers baseline and other participants. Our best model gives micro-F1 score of 0.7582. The model can be applied to other emotional and sentiment classification problems and can be modified to accept external attention signals and emotional specific word embedding.

Acknowledgement

We would like to acknowledge La Région Occitanie and Communauté d’Agglomération Béziers Méditerranée which finance the thesis of Waleed Ragheb.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, volume abs/1409.0473.
- Hadjer Boubenna and Dohoon Lee. 2018. [Image-based emotion recognition using evolutionary algorithms](#). *Biologically Inspired Cognitive Architectures*, 24:70–76.
- Fabio Calefato, Filippo Lanubile, and Nicole Novielli. 2017. Emotxt: A toolkit for emotion recognition from text. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pages 79–80. IEEE.
- Ankush Chatterjee, Umang Gupta, Manoj Kumar Chinnakotla, Radhakrishnan Srikanth, Michel Galley, and Puneet Agrawal. 2019a. Understanding emotions in text using deep learning and big data. *Computers in Human Behavior*, 93:309–317.
- Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019b. Semeval-2019 task 3: Emocontext: Contextual emotion detection in text. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval-2019)*, Minneapolis, Minnesota.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. volume abs/1611.01734.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339.
- R. Jenke, A. Peer, and M. Buss. 2014. [Feature extraction and selection for emotion recognition from eeg](#). *IEEE Transactions on Affective Computing*, 5(3):327–339.
- Marina Krakovsky. 2018. [Artificial \(emotional\) intelligence](#). *Commun. ACM*, 61(4):18–19.
- Hugo Larochelle and Geoffrey E Hinton. 2010. Learning to combine foveal glimpses with a third-order boltzmann machine. In *Advances in Neural Information Processing Systems 23*, pages 1243–1251.
- Yukun Ma, Haiyun Peng, and Erik Cambria. 2018. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive lstm. In *Association for the Advancement of Artificial Intelligence (AAAI 2018)*.
- Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander F. Gelbukh, and Erik Cambria. 2018. Dialoguernn: An attentive rnn for emotion detection in conversations. *CoRR*, Association for the Advancement of Artificial Intelligence (AAAI 2019).
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. Regularizing and optimizing LSTM language models. In *International Conference on Learning Representations (ICLR)*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. [Pointer sentinel mixture models](#). *CoRR*, abs/1609.07843.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. 2013. [Regularization of neural networks using dropconnect](#). In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1058–1066, Atlanta, Georgia, USA. PMLR.
- Tom Young, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria. 2018. Recent trends in deep learning based natural language processing [review article]. In *IEEE Computational Intelligence Magazine*, volume 13, pages 55–75.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 730–739.