

Automatically augmenting an emotion dataset improves classification using audio

Lakomkin Egor Cornelius Weber Stefan Wermter

Department of Informatics, Knowledge Technology Group

University of Hamburg

Vogt-Koelln Str. 30, 22527 Hamburg, Germany

{lakomkin, weber, wermter}@informatik.uni-hamburg.de

Abstract

In this work, we tackle a problem of speech emotion classification. One of the issues in the area of affective computation is that the amount of annotated data is very limited. On the other hand, the number of ways that the same emotion can be expressed verbally is enormous due to variability between speakers. This is one of the factors that limits performance and generalization. We propose a simple method that extracts audio samples from movies using textual sentiment analysis. As a result, it is possible to automatically construct a larger dataset of audio samples with positive, negative emotional and neutral speech. We show that pretraining recurrent neural network on such a dataset yields better results on the challenging EmotiW corpus. This experiment shows a potential benefit of combining textual sentiment analysis with vocal information.

1 Introduction

Emotion recognition recently gained a lot of attention in the literature. The evaluation of the human emotional state and its dynamics can be very useful for many areas such as safe human-robot interaction and health care. While recently deep neural networks achieved significant performance breakthroughs on tasks such as image classification (Simonyan and Zisserman, 2014), speech recognition (Hannun et al.,) and natural language understanding (Sutskever et al., 2014), the performance on emotion recognition benchmarks is still low. A limited amount of annotated emotional samples is one of the factors that negatively impacts the performance. While obtaining such data is a cumbersome and expensive process, there are plenty of

unlabelled audio samples that could be useful in the classifier learning (Ghosh et al., 2016).

The majority of recent works use neural networks combining facial expressions and auditory signals for emotion classification (Barros et al., 2015; Yao et al., 2015; Chao et al., 2016). There is a clear benefit of merging visual and auditory modalities, but only in those situations when the speaker’s face can be observed. In (Hines et al., 2015) it was shown that incorporating linguistic information along with acoustic representations can improve performance. Semantic representations of spoken text can help in emotional class disambiguation, but in this case, the model will rely on the accuracy of the speech-to-text recognition system. Pretraining convolutional neural network (Ebrahimi Kahou et al., 2015) on an external dataset of faces improves the performance of the emotion classification model. However, the problem of augmenting emotional datasets with audio samples to improve the performance of solely audio processing models remained unsolved.

Our motivation for this paper was to fill this gap and conduct experiments on automatically generating a larger and potentially richer dataset of emotional audio samples to make the classification model more robust and accurate. In this work, we describe a method of emotional corpus augmentation by extracting audio samples from the movies using sentiment analysis over subtitles. Our intuition is that there is a significant correlation between the sentiment of spoken text and an actually expressed emotion by the person. Following this intuition we collect positive, neutral and negative audio samples and test the hypothesis that such an additional dataset can be useful in learning more accurate classifiers for the emotional state prediction. Our contribution is two-fold: a) we introduce a simple method to extract automatically positive and negative audio training samples from

full-length movies b) we demonstrate that using an augmented dataset improves the results of the emotion classification.

2 Models and experimental setup

2.1 Dataset

For our experiment, we have used the EmotiW 2015 dataset (Dhall et al., 2015), which is a well-known corpus for emotional speech recognition composed of short video clips annotated with categorical labels such as *Happy*, *Sad*, *Angry*, *Fear*, *Neutral*, *Disgust* and *Surprise*. Each utterance is approximately 1-5 seconds in duration. The EmotiW dataset is considered as one of the most challenging datasets as it contains samples from very different actors and the lighting conditions, background noise and other overlapping sounds make the task even more difficult. The training set and the validation set contains 580 and 383 video clips respectively. We have used the official EmotiW validation set to report the performance as the test set labels were not released and 10% of the official training set as validation set for neural network early stopping.

2.2 Generating emotional audio samples

As a source for emotional speech utterance candidates, we use full-length movies taking the list of titles from the EmotiW corpus. For each of the films, there are subtitles available, which can be treated as a good approximation of a spoken text, even though sometimes there can be inaccuracies as producing subtitles is a manual process. Our intuition is that the movies contain a large variety of auditory emotional expressions by many different speakers and is a potentially valuable source of emotional speech utterances. For each of the movies, sentiment score was calculated for each of the subtitle phrases at the time of the utterance with the NLTK (Bird et al., 2009) toolkit. Sentiment score represents how positive or negative the text segment is. The NLTK sentiment analyzer was used for simplicity and effectiveness. Phrases longer than 100 characters and shorter than four words were filtered out to avoid having very long or very short utterances. Subtitle phrases with polarity score higher than 0.7 were treated as positive samples and the ones with sentiment score lower than -0.6 as negative samples. The thresholds were selected empirically to make the number of the positive and negative samples balanced.

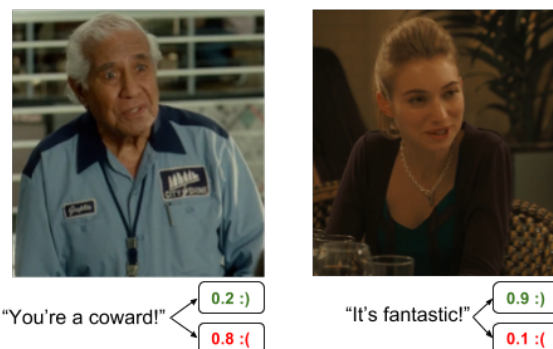


Figure 1: Visualization of the process of extraction of positive and negative speech utterances, based on sentiment analysis of the subtitles.

As the majority of the phrases were assigned a value of sentiment close to 0, we treated them as neutral and used only a random subsample of it. Corresponding audio samples were cut from the movie with respect to the timings of the subtitle phrase. Overall 2100 positive, negative and neutral speech utterances were automatically selected from 59 movies and used as the additional dataset for emotion classification for binary tasks and as a dataset for model pre-training in multi-class setup.

2.3 Features extracted

We extracted FFT (Fast Fourier Transform) spectrograms from the utterances with a window length of 1024 points and 512 points overlap. Frequencies above 8kHz and below 60Hz were discarded as higher frequencies usually contain more noise and a log-scale in the frequency domain was used as emphasizing lower frequencies appears to be more significant for the emotional state prediction (Ghosh et al., 2016). Maximum length of the utterance in the dataset is 515 frames.

2.4 GRU model

The Gated-recurrent unit (GRU) (Bahdanau et al., 2014) is a recurrent neural network (Elman, 1991) model trained to classify a sequence of input vectors. One of the main reasons for its success is that the GRU is less sensitive to the vanishing gradient problem during training, which is especially crucial for acoustic processing as the length of the sequences can easily reach hundreds or even thousands of time steps, as opposed to NLP tasks.

As a first stage, a single layer bi-directional GRU model has been used in our experiments with a 32 dimension cell size. Temporal mean pooling over all intermediate hidden memory representa-

tions was used to construct the final memory vector.

$$z = \sigma(x_t U^z + s_{t-1} W^z) \quad (1)$$

$$r = \sigma(x_t U^r + s_{t-1} W^r) \quad (2)$$

$$h^{fw} = \tanh(x_t U^h + (s_{t-1}^{fw} \circ r) W^h) \quad (3)$$

$$h^{bw} = \tanh(x_t U^h + (s_{t+1}^{bw} \circ r) W^h) \quad (4)$$

$$s_t^{fw} = (1 - z) \circ h_t^{fw} + z \circ s_{t-1}^{fw} \quad (5)$$

$$s_t^{bw} = (1 - z) \circ h_t^{bw} + z \circ s_{t+1}^{bw} \quad (6)$$

$$s_t = \text{concat}([s_t^{fw}, s_t^{bw}]) \quad (7)$$

$$c = \frac{\sum_{t=1}^T s_t}{T} \quad (8)$$

In these equations, the c vector is used to represent the whole speech utterance as an average of intermediate memory vectors s_t^{fw} and s_t^{bw} , where fw index corresponds to forward GRU execution and bw for backward. x_t is a spectrogram frame, r and z represent reset and update gates and s_t is a GRU memory representation at timestamp t , following notation in (Bahdanau et al., 2014). We have used Keras (Chollet, 2015) and Theano (Bastien et al., 2012) frameworks for our implementation.

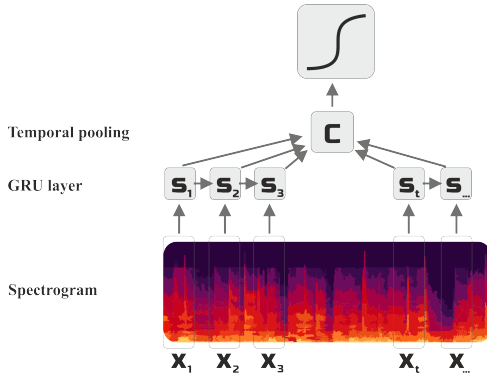


Figure 2: Recurrent neural network model for emotion speech utterance classification with temporal pooling.

2.5 Transfer learning

In multi-class setup, firstly we trained the neural network on the augmented corpus predicting labels generated by sentiment analyzer. We refer to it as a pre-trained network. As our goal is to predict emotional categories like happy or anger, we afterward replaced the softmax layer of the pre-trained network comprised of positive, negative and neutral classes with the new softmax

layer for emotion prediction with angry, happy, sad and neutral classes. By using such a procedure, GRU layer hopefully can grasp meaningful representation of positive, neutral and negative speech which, as a result, will be helpful for emotion classification by means of transfer learning. Fine-tuning was done on the training data of the EmotiW corpus.

2.6 Results

We compare our results in three binary emotion classification tasks: *happy-vs-fear*, *happy-vs-disgust* and *happy-vs-anger* and multi-class setup, where we considered *Happy*, *Angry*, *Sad* and *Neutral* samples. For each of the tasks we treated generated negative samples as either *fear*, *disgust* or *anger* samples respectively and positive samples as *happy*. For the multi-class setup, we follow the transfer learning routine by adapting neural network trained on the augmented data to the 4-way emotional classification. Accuracy is reported for binary tasks and F-score for multi-class setup. Results are presented in Table 1. By using automatically generated emotional samples there is a slight decrease in the accuracy for *happy-vs-anger* task and an improvement in the accuracy for *happy-vs-fear* and *happy-vs-disgust* tasks. Also, in our experiments, temporal pooling worked significantly better than using the memory vector at the last time step.

Table 1: Utterance level emotion classification performance (accuracy) in 3 binary tasks: happy vs fear, happy vs angry and happy vs disgust. Also, multi-class performance (F-measure) is reported with 4 basic emotions: Angry, Happy, Sad and Neutral. BM - baseline method without augmentation, PM - proposed method with augmentation.

Experiment	BM	PM
Binary classification:		
Happy vs Fear	58.7	66.1
Happy vs Angry	70.7	68.9
Happy vs Disgust	61.1	64.1
Multi-class:		
Angry, Happy, Sad, Neutral	36	38

3 Conclusion

In this paper, we proposed a novel method for automatically generating positive, neutral and negative audio samples for emotion classification from full-length movies. We experimented with three different binary classification problems: happy vs anger, happy vs fear and happy vs disgust and found that for the latter two there is an improvement in the accuracy on the official EmotiW validation set. Also, we observed the improvements of the results in multi-class setup. We found that the augmented larger dataset even though contains noisy and weak labels, contribute positively to the accuracy of the classifier.

For future work, we want to explore jointly learning sentiment and acoustic representations of the spoken text, which appears to be beneficial for accurate speech emotion classification, as it allows to deal with the ambiguity of the spoken text sentiment.

Acknowledgments

The authors gratefully acknowledge partial support from the German Research Foundation DFG under project CML (TRR 169), the European Union under project SECURE (No 642667), and the Hamburg Landesforschungsförderungsprojekt.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. In *ICLR 2015*.
- Pablo Barros, Cornelius Weber, and Stefan Wermter. 2015. Emotional expression recognition with a cross-channel convolutional neural network for human-robot interaction. In *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th International Conference on*. IEEE.
- Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian J. Goodfellow, Arnaud Bergeron, Nicolas Bouchard, and Yoshua Bengio. 2012. Theano: new features and speed improvements. Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. O'Reilly Media.
- Linlin Chao, Jianhua Tao, Minghao Yang, Ya Li, and Zhengqi Wen. 2016. Audio visual emotion recognition with temporal alignment and perception attention. *CoRR*, abs/1603.08321.
- François Chollet. 2015. Keras. <https://github.com/fchollet/keras>.
- Abhinav Dhall, OV Ramana Murthy, Roland Goecke, Jyoti Joshi, and Tom Gedeon. 2015. Video and image based emotion recognition challenges in the wild: EmotiW 2015. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. ACM.
- Samira Ebrahimi Kahou, Vincent Michalski, Kishore Konda, Roland Memisevic, and Christopher Pal. 2015. Recurrent neural networks for emotion recognition in video. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. ACM.
- Jeffrey L Elman. 1991. Distributed representations, simple recurrent networks, and grammatical structure. *Machine learning*, 7(2-3).
- Sayan Ghosh, Eugene Laksana, Louis-Philippe Morency, and Stefan Scherer. 2016. Representation Learning for Speech Emotion Recognition. *Interspeech 2016*, pages 3603–3607, September.
- Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y Ng. Deep Speech: Scaling up end-to-end speech recognition.
- Christopher Hines, Vidhyasaharan Sethu, and Julien Epps. 2015. Twitter: A New Online Source of Automatically Tagged Data for Conversational Speech Emotion Recognition. *Proceedings of the 1st International Workshop on Affect & Sentiment in Multimedia*.
- Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR*, 1409.1556.
- I Sutskever, O Vinyals, and QV Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information*.
- Anbang Yao, Junchao Shao, Ningning Ma, and Yurong Chen. 2015. Capturing au-aware facial features and their latent relations for emotion recognition in the wild. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. ACM.