

BIMODAL SPEECH EMOTION RECOGNITION USING PRE-TRAINED LANGUAGE MODELS

Verena Heusser*, Niklas Freymuth*, Stefan Constantin, Alex Waibel

Karlsruhe Institute of Technology
Institute for Anthropomatics and Robotics
{verena.heusser, niklas.freymuth}@student.kit.edu, {stefan.constantin, waibel}@kit.edu

ABSTRACT

Speech emotion recognition is a challenging task and an important step towards more natural human-machine interaction. We show that pre-trained language models can be fine-tuned for text emotion recognition, achieving an accuracy of 69.5 % on Task 4A of SemEval 2017, improving upon the previous state of the art by over 3 % absolute. We combine these language models with speech emotion recognition, achieving results of 73.5 % accuracy when using provided transcriptions and speech data on a subset of four classes of the IEMOCAP dataset. The use of noise-induced transcriptions and speech data results in an accuracy of 71.4 %. For our experiments, we created IEMoNet, a modular and adaptable bimodal framework for speech emotion recognition based on pre-trained language models. Lastly, we discuss the idea of using an emotional classifier as a reward for reinforcement learning as a step towards more successful and convenient human-machine interaction.

Index Terms— Speech Emotion Recognition, Text Emotion Recognition, Bimodal Emotion Recognition, IEMOCAP, Self Attention, Pre-trained Language Models

1. INTRODUCTION

Emotions are an important aspect of human behavior. They do not only influence the reaction to our environment [1, 2], but also actively change our perception of it [3] and sometimes even contribute to how well we remember specific events [4]. As such, they influence both human-human and human-machine interaction. However, in human-machine interaction, emotions are often not at all or only scarcely considered. We propose to use automatically generated emotional feedback as a reward in Reinforcement Learning (RL). This could be used to improve human-machine interaction, for example when used with dialogue systems. In this paper, we will focus on the tasks of automatic Text Emotion Recognition (TER) and automatic Speech Emotion Recognition (SER).

It has been shown that combining textual features and a SER model into a multimodal system improves the overall classification performance [5, 6, 7, 8]. Additionally, the tasks of Automatic Speech Recognition (ASR) and language modelling have seen great improvements over the last year. This leads to a steady improvement in the quality of both transcripts and language models. Since most language models have been shown to be able to solve a plethora of different tasks [9, 10], the combination of ASR and TER can be used to automatically transcribe a given utterance and use the language model for an emotion classification. To show the effectiveness of language models for TER, we fine-tune both BERT [11] and XLNet [12], two self-attentive language models, on Task 4A of SemEval 2017 [13].

Motivated by these two recent advancements, we present IEMoNet (Interactive Emotion Network) which is a modular and adaptable bimodal SER framework based on pre-trained language models. IEMoNet consists of independently trainable ASR, SER and TER sub-modules. They can be combined by fine-tuning only a few layers, omitting time-consuming re-training of the whole model. This allows us to utilize both textual and paralinguistic features for emotion classification. We test our model on the ‘Text+Speech’ subtask of the Interactive Emotional dyadic MOTion CAPture (IEMOCAP) dataset [14]. Using 10-fold cross validation, IEMoNet achieves results of 73.5 % when combining text and speech inputs. Using speech and (simulated) automatically transcribed texts instead of the provided transcriptions results in 71.4 % accuracy.

2. RELATED WORK

IEMoNet is an approach to bimodal SER, meaning that it uses different modalities to extract richer and more varied features. In general, bi- and multimodal models combine modalities such as video recordings [15], more specific facial features [16], movements [17] and even EEG signals [18]. In our case text, transcribed speech and speech are used, closely following [7, 8, 19]. Previous research focused on the Time Delay Neural Network (TDNN) [20] and recurrent architectures like the Long Short-Term Memory (LSTM) [21].

*equal contribution

Recently, ELMo [22], an LSTM-based pre-trained language model, has been successfully used to generate textual features for multimodal emotion recognition [5]. We use pre-trained language models based on Transformer encoder stacks [23] instead. These Transformer encoder stacks rely solely on the self-attention mechanism [24] for contextualization. They have shown great potential for ASR, SER and (emotional) language modelling individually, as will be shown in the following sections.

2.1. Speech Emotion Recognition

SER has been studied for several decades [25], with early work consisting of the classification of a given utterance to one of usually four distinct emotions based on prosodic features such as pitch, speech and intensity [26, 27]. While these relatively simple features are still used today, deep neural architectures have led researchers to shift towards either more features such as the IS09 emotional feature set [28], or more or less unprocessed audio/signal data as input [29, 30].

Training and evaluating SER models is challenging because most datasets only include a few hours of spoken data, making it difficult to train large neural networks without severe overfitting. Research on SER utilizes Mel Frequency Cepstral Coefficients (MFCCs), spectrograms and raw waveforms as inputs, while the architectures vary from gated recurrent units [31] to a combination of LSTMs and Time Delay Neural Networks (TDNNs) [32, 33]. Recently, the Transformer has also been used in SER, either directly [34] or as a means of richer representations via predictive coding [35], achieving competitive results in both cases.

2.2. Automatic Speech Recognition

While ASR faces similar challenges as SER, research has shown that transfer between ASR and SER models is only possible in early layers of the respective models, meaning that both tasks require their respective models to learn vastly different features [36].

As of today, a few architectures have shown to be well suited for ASR, all of which make either direct or indirect use of neural networks. Looking at the LibriSpeech corpus [37] as an example, Hybrid Models [38], Convolutional networks [39] and LSTMs [40, 41, 42] all achieve Word Error Rates (WER) between 2.5 % and 3 % on the 'test-clean' sub-task when combined with a language model. Additionally, researchers found that models based on the Transformer are also competitive and very promising, especially in regards to more complicated tasks [43, 44].

2.3. Text Emotion Recognition

TER is similar to sentiment analysis. Given an emotional model, usually Ekman's Base Emotions [45], the task is to find the dominant emotion in a text. Common approaches are

either keyword based or end-to-end. Keyword based methods operate on strong emotional words [46] or are based around meta expressions like hashtags [47], emoticons [48], and emojis [49]. End-to-end architectures are mostly based on LSTMs and the attention mechanism [50, 51], allowing the models to correctly model long-term dependencies.

However, with the introduction of attention-based models like the Transformer [23] model and later BERT [11], general language models are currently achieving great results on most Natural Language Understanding tasks, including sentiment analysis. An example is the GLUE Benchmark [52], a collection of various Natural Language Understanding tasks. Since GLUE's release in 2018, BERT [11] and similar models [12, 53, 54, 55] gave way to over a dozen improvements over the state of the art and eventually led to surpassing human performance on this benchmark¹.

Additionally, some of these language models are provided with pre-trained models and weights, allowing for efficient use of resources via transfer learning. These models can be fine-tuned on a given task, which means that some or all of the existing weights are trained on the new task for relatively few iterations. This drastically reduces both training time and the necessary amount of training samples, making it ideal for the data-sparse field of Emotion Recognition. Utilizing this, researchers have achieved good results in both combining language models with LSTM architectures [56] and fine-tuning them on sentiment analysis tasks [57], showing that they are suitable for the task at hand.

2.4. Emotion in Reinforcement Learning

In the context of dialogue systems (and human-machine interaction in general) it seems rather important to integrate the emotional state of the speaker into the overall decision making process of an agent. As of now, emotions have been used in both task-oriented [58] and non task-oriented dialogue [59], in both cases increasing the performance of the given system. Additionally, there has been a lot of research regarding emotionally intelligent systems which implement and operate based on their own emotional state [60].

3. REINFORCEMENT LEARNING

Using emotion recognition for reinforcement learning in dialogue systems has the potential to greatly improve both the systems usefulness and the users overall experience. For example, a call center could track a caller's anger [61] and annoyance [62] levels to see how satisfied they are with their current call. In the case of a unsatisfactory calling experience, the emotional reaction can then be interpreted as a negative reward and used to adjust the dialogue system, causing it to

¹See <https://gluebenchmark.com/leaderboard> (accessed 2019-10-14)

slightly change its behaviour and react more appropriately in future calls.

However, there is an important distinction between emotion and mood [63, 64], the main difference being the timescale of both concepts [65]. While emotion itself is short-term, its effects on a person’s mood might persist for a relatively long time. In the context of dialogue systems, this means that a particularly good dialogue turn might be enough to cheer a person up and make them sound happier for the rest of the interaction. This again generates more ‘happy’ features for the next turns, causing the SER system to classify what would normally be ‘neutral’ as ‘happy’. To counteract this effect of a ‘lingering’ mood, we propose to use an approximated derivative of the emotional state instead of the state itself. For an emotion e and a dialogue history h_t at dialogue turn t , the output of the emotion recognition system can be seen as the posterior probability $P(e|h)$. The delta $\Delta P(e|h_t) = P(e|h_t) - P(e|h_{t-1})$ of each individual probability can then be seen as the resulting change in its corresponding emotion based on the last dialogue turn.

Additionally, certain messages are expected to have certain emotional responses. Being the proverbial bearer of bad news, for example, a dialogue system should not perceive an anticipated negative emotion as a negative reward. Mathematically, this can be seen as the difference between the probabilities of the estimated emotional response $P^{est}(e|h_t)$ and the actual response $P(e|h_t)$ of the dialogue turn in question. Considering mood, the unexpected change becomes $\Delta P(e|h_t) - \Delta P^{est}(e|h_t)$.

Since a reward for reinforcement learning is usually bounded by a fixed interval, a mapping $f : E \rightarrow [p, q]$ between the set of emotions E and their reward is necessary. In our case, a negative emotion like ‘anger’ would be mapped to a negative value, the simplest case being -1 . Taking this into account, the actual reward r_t used for reinforcement learning would be

$$r_t = \sum_{e \in E} f(e) \cdot (\Delta P(e|h_t) - \Delta P^{est}(e|h_t)). \quad (1)$$

Given that it is easier to give an estimated reward rather than a set of changes in emotions, this can be simplified to

$$r_t = \sum_{e \in E} f(e) \cdot \Delta P(e|h_t) - r_t^{est}, \quad (2)$$

where r_t^{est} is the estimated reward for dialogue turn t and could be inferred from e.g. another neural network via regression.

4. MODEL

The IEemoNet framework consists of five parts. A given utterance is first forwarded into a pre-processing block, where both ASR and SER features are extracted from the same raw

signal. These features are then fed into an ASR and a (neural) SER System respectively. The ASR System automatically generates a transcript of the utterance, which is then used as an input for the (also neural) TER System. Both the spoken and the TER system then provide their outputs as independent features for the classification block. Finally, the classification block outputs the detected emotion. An overview can be seen in Figure 1.

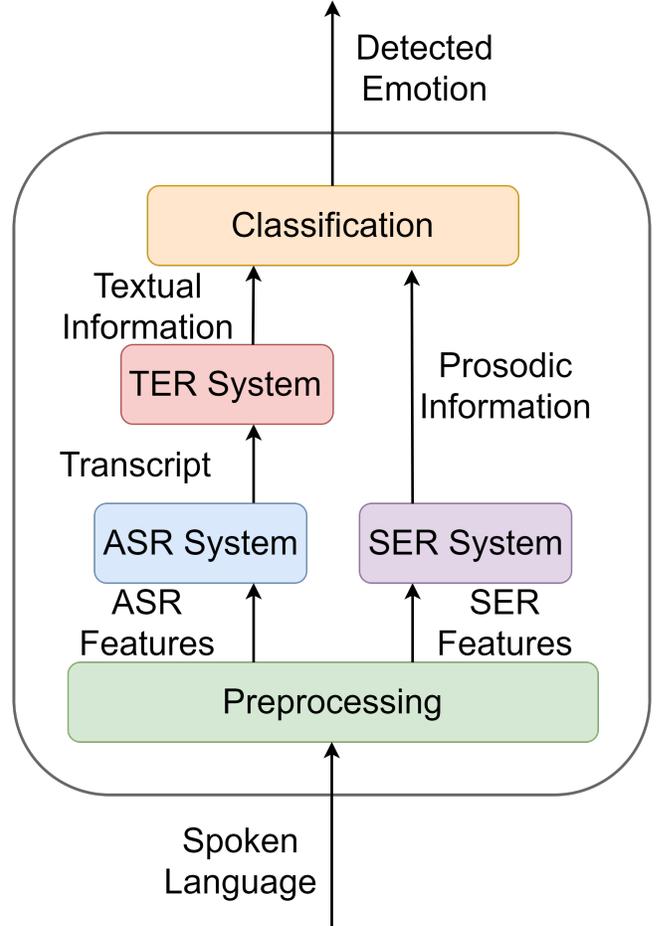


Fig. 1. The IEemoNet framework where each intermediate system (ASR, SER, TER) can be chosen and optimized individually.

This framework is very versatile and easily adaptable. The intermediate systems, i.e. the ASR, SER and TER systems, can be chosen and optimized independently, the only fixed variables being the input and the output size. Training IEemoNet is done in multiple steps. First the preprocessing block extracts ASR and SER features. Depending on the used sub-systems, the ASR features can range from classic MFCCs+Deltas to augmented waveforms as used in [42]. For the SER features either classical features such as pitch and prosody [26, 27], complex emotion recognition feature sets

(for example the IS09 set [28]) or augmented waveforms [32] are plausible.

After this, the utterances need to be transcribed to provide textual data. We found that modern out-of-the-box ASR systems generally provide high-quality transcriptions without the need for additional fine-tuning, allowing us to simply use one of these systems to transcribe the given utterances. Note that fine-tuning an ASR system towards a given dataset would probably provide slightly better transcriptions, leading to a slightly better performance overall. However, this comes at the cost of significantly more training time.

The SER and TER systems are trained individually on their parts of the given task, i.e. speech to emotion for SER and (possibly previously transcribed/generated) text to emotion for TER. For this, both models are extended by an auxiliary classification block followed by a softmax layer, effectively creating two independent recognition systems.

After this step is finished, we combine the SER and TER systems by removing the classification and softmax layers and replacing them with the classification block, leaving their second-to-last layers as the next block’s input features. The classification block usually consists of a simple feedforward network with a softmax classification layer. However, more sophisticated approaches employing the attention mechanism might be better suited at aligning the two sets of features. The classification block is then trained while keeping the intermediate systems (ASR, SER and TER) unchanged. This is done to prevent potential overfitting due to huge relative differences in model size and structure.

A big advantage of this modular setup is that any intermediate system can simply be replaced by one another, requiring at worst the re-training of the classification block and the extraction of new/other features in the pre-processing block. This proves to be especially useful considering the ongoing rapid development in both emotion recognition and deep learning in general. IEmoNet can adapt to these improvements by upgrading an outdated sub-system to a newer one without re-training all of the model.

5. EXPERIMENTAL DETAILS / RESULTS

5.1. Pre-trained Language Models

The performance of IEmoNet only relies on the ability of its intermediate systems to process sequence data. We found that using pre-trained language models for the TER part drastically improves performance while only taking very few training epochs to converge. For our experiments, we use BERT [11] and XLNet [12], two popular transformer-based general language model that are available with pre-trained weights. They both consist of either 24 (Large) or 12 (Base) Encoder Transformer blocks and are pre-trained on large resources of textual data. For a detailed description of the Encoder Transformer block, see [23]. We use the case-sensitive (cased)

Model	Accuracy
All Neutral	0.483
DataStories [50]	0.651
BB_twtr [51]	0.658
LIA [68]	0.661
NNEMBs [69]	0.664
XLNet-Large	0.663 ± 0.030
Bert-Base	0.681 ± 0.007
Bert-Large	0.687 ± 0.015
XLNet-Base	0.695 ± 0.007

Table 1. Test accuracy on Task 4A of SemEval 2017. The best models are in bold. We repeat each experiment five times and report the mean and standard deviation.

versions of both models. Training BERT and XLNet on all tasks was conducted using the Adam-optimizer [66] with a learning rate of $3e - 5$. We use the provided WordPiece embeddings for BERT and byte-pair encodings for XLNet. The sequence length is set to 48. Longer/Shorter inputs are truncated/padded. We add a fully-connected layer of 768 units with Dropout [67] of 0.2 and L2-regularization of 0.03 between the pre-trained models and the softmax layer. We use weighted cross-entropy as our loss function to counteract the imbalanced datasets. These hyperparameters were chosen experimentally and work equally well on both SemEval and IEMOCAP.

5.2. SemEval Experiments

We found that using pre-trained language models and fine-tuning them on emotional datasets achieves good results. In addition, fine-tuning instead of training from scratch significantly reduces training time and resource usage. To show the effectiveness of pre-trained language models for TER, we run some tests on Task 4A of the SemEval 2017 competition [13], which is a 3-class sentiment analysis task. The input is a tweet (a twitter post containing no more than 280 characters), the output is the predicted sentiment, in this case either ‘positive’, ‘neutral’, or ‘negative’. We choose this task because of its similarity to TER. The sentiments may be roughly mapped to the emotions ‘happy’, ‘neutral’ and ‘sad’/‘angry’.

The dataset is hand-annotated and consists of about 50.000 training and 12.000 test samples consisting of tweets from 2012 to 2017. In the training set, the data consists of roughly 39.5/44.8/15.6 % positive/neutral/negative tweets while the test set has a distribution of 19.3/48.3/32.3 % for the respective classes. In both cases, neutral tweets are the most common. However, the training data has a lot more positive tweets while the test set is more shifted towards negative ones. This causes the task to be more challenging because the models need to detect negative sentiments with relatively little training data.

We experiment with both BERT and XLNet and find XLNet to be slightly superior in performance. The tweets were not pre-processed.

The results can be seen in Table 1. We train each model five times with different initial weights and report the mean and standard deviation for each model. Our best model (XLNet-Base) achieves an accuracy of 69.5% after only one epoch of training. This is an improvement of over 3% absolute over previous models, showing the effectiveness of pre-trained models for this task.

5.3. IEMOCAP Experiments

The IEMOCAP dataset provides 12 hours of emotional audio dialogues, recorded with five pairs of professional actors in both scripted and spontaneous sessions. At least three annotators were asked to label each of the dialogues using Ekman’s Base Emotions [45] and a variation of Russell’s Circumplex Model [70]. The inter-annotator agreement for the task was 74.6% (66.0% for scripted and 83.1% for spontaneous utterances) [14]. Because of the scarcity of other emotions, we follow common research [8, 33, 7] and only classify the emotions ‘happiness’, ‘sadness’, ‘anger’ and ‘neutral’. Our final dataset contains 38.04% neutral, 24.57% angry, 24.14% sad and 13.24% happy samples. We split each dialogue into utterances from 3 to 15 seconds, with one utterance holding the information of one speaker (i.e. mono audio data and the respective transcription). Each utterance is then classified individually and without context information. We use both scripted and spontaneous dialogues for all experiments.

We use 10-fold cross-validation with randomly shuffled data with a fixed random seed. Note that some researchers (e.g. [33, 5]) test their models in a speaker-exclusive setting instead. They split the dataset by speakers, training on most speakers and testing exclusively on the remaining few. We denote the respective papers with a ‘*’. We train each model three times with different initial weights and report the mean of all models and the mean of the three standard deviations.

We evaluate the models with respect to weighted and unweighted accuracy. The Weighted Accuracy (WA) is given by the number of correctly classified samples divided by the total number of samples. The Unweighted Accuracy (UA) corresponds to first calculating the accuracy for each emotion class and then averaging by the number of samples per class.

5.3.1. IEMOCAP Text Experiments

As described in Section 4, we train IEMoNet’s modalities independently at first. Motivated by their performance on the SemEval task (see Section 5.2), we use both BERT and XLNet as TER models. We use the Base versions (i.e. 12 Transformer Encoder Blocks) for all experiments because they are roughly three times smaller than the large ones, making them more feasible for real-time applications. To compare

Model	WA	UA	T
TRE [8]	0.635 ± 0.018	–	P
<i>LSTM*</i> [7]	<i>0.648</i>	–	P
XLNet	0.694 ± 0.015	0.671 ± 0.019	P
BERT	0.709 ± 0.015	0.691 ± 0.018	P
TRE-ASR [8]	0.593 ± 0.022	–	A
XLNet-A	0.668 ± 0.020	0.645 ± 0.023	A
BERT-A	0.685 ± 0.017	0.665 ± 0.023	A

Table 2. Weighted Accuracy (WA) and Unweighted Accuracy (UA) for TER on the IEMOCAP dataset. Models without cross-validation are in *curly*. Models with ‘*’ use speaker-exclusive testing. The T row denotes the Transcription. P stands for Provided and A for Automatic.

the theoretical and the practical performance, we train on both provided and artificial ASR transcriptions [71, 72] with 6% WER. The latter are used as an alternative to an ASR system and are configured to simulate typical transcription errors as described in [71]. This WER was chosen to match the one reported in [8]. We use artificial noise instead of an ASR system due to time constraints to integrate an ASR of high quality.

For all text-related experiments, we omitted the use of a validation split, instead training on a fixed number of seven epochs for all folds. The number seven was chosen because the training loss usually stabilized after seven epochs. The results can be seen in Table 2. Here, BERT slightly outperforms XLNet, achieving an accuracy of 70.9% on the provided transcripts and 68.5% on the noisy ones. While our models’ performances still drops significantly when using noisy instead of provided transcriptions, they seem to be more robust to this noise than the one reported in [8]. We hypothesize that the extensive pre-training on large amounts of textual data leads to this increase in stability.

5.3.2. IEMOCAP Speech Experiments

For the SER part, we train our models on the IS09 feature set [28], following previous research [34, 73]. The features were extracted using the openSMILE feature extraction tool [74] with a frame size of 300 frames and a stride of 0.1, resulting in an overlap of 0.9 between two frames. We then pad/truncate the sequences to a fixed length of 384 frames.

The results in Table 3 are achieved using simple Attention-BiLSTM models. The models use a TDNN layer followed by two BiLSTM layers and an attention layer. The TDNN layer is used to map the inputs to the dimension of the BiLSTMs. The output of the second BiLSTM layer is then fed into the attention layer. Then, the outputs of the second BiLSTM layer and the attention layer are concatenated and used as the input for a softmax layer. Our best model, BiLSTM-128attDim, used a TDNN with filter size 16, 64 dimensional LSTMs

Model	WA	UA
IS09-Self-Att [34]	0.681	0.638
CNN-LSTM* [33]	0.688	0.594
<i>Att-TDNN/LSTM*</i> [32]	<i>0.701</i>	<i>0.607</i>
BiLSTM-64attDim	0.600 ± 0.020	0.531 ± 0.020
BiLSTM-128attdim	0.608 ± 0.017	0.526 ± 0.013

Table 3. Weighted Accuracy (WA) and Unweighted Accuracy (UA) on SER on the IEMOCAP dataset. Models without cross-validation are in *curive*. Models with ‘*’ use speaker-exclusive testing.

Model	WA	UA	T
H-MM-4* [5]	0.717	–	P
MDRE [8]	0.718 ± 0.019	–	P
<i>Att-align*</i> [19]	<i>0.725</i>	<i>0.709</i>	P
IEmoNet(XL)	0.715 ± 0.015	0.689 ± 0.016	P
IEmoNet(BE)	0.735 ± 0.016	0.710 ± 0.016	P
MDRE-ASR[8]	0.691 ± 0.019	–	A
<i>Att-align-A*</i> [19]	<i>0.704</i>	<i>0.695</i>	A
IEmoNet(XL)-A	0.687 ± 0.020	0.657 ± 0.020	A
IEmoNet(BE)-A	0.714 ± 0.017	0.686 ± 0.022	A

Table 4. Weighted Accuracy (WA) and Unweighted Accuracy (UA) on IEMOCAP using both speech and text data. Models without cross-validation are in *curive*. Models with ‘*’ use speaker-exclusive testing. The T row denotes the Transcription. P stands for Provided and A for Automatic.

(in both directions) and self-attention with dimension 128. It achieved an accuracy of 60.8%. The second best model, BiLSTM-64attDim, uses the same configuration, except for an attention dimension of 64, achieving 60.0% weighted accuracy. Note that more sophisticated SER models have been shown to yield far better performance and would thus be better suited for a combined bimodal architecture. However, we primarily want to showcase the capability of pre-trained language models for bimodal Emotion Recognition, making these basic models sufficient for this purpose.

5.3.3. IEMOCAP Bimodal Experiments

We combine our best individual models as described in Section 4. Again, each model is trained for a fixed number of seven epochs. The results can be found in Table 4. IEmoNet(XL) uses the best XLNet models from Table 2 as its text sub-system while IEmoNet(BE) uses the BERT models instead. Both combine them with the best SER model from Table 3. We also tried fine-tuning the whole model (i.e. the classification block, the TER and the SER parts together) after the seven epochs. However, the performance was significantly worse than the models that only trained the

classification block. This is likely caused by the size of BERT and XLNet allowing them to easily dominate the classification outcome. The best models achieve 73.5% and 71.4% accuracy for provided and automatic transcripts respectively.

6. CONCLUSION / FUTURE WORK

Detecting emotions in speech is a challenging task that marks an important step towards more natural and adaptive human-computer interaction. We show that self-attentive pre-trained language models are well suited for text emotion recognition and that the Transformer model in general works well for bimodal Emotion Recognition. Our experiments result in 69.5% accuracy on Task 4A of SemEval 2017. We achieve 73.5% and on the ‘Text+Speech’ subtask of the IEMOCAP dataset. Using speech and (simulated) automatic transcriptions results in 71.4% instead. The results on IEMOCAP are achieved by fine-tuning BERT and combining it with a traditional speech emotion recognition system. This is done using IEmoNet, a modular framework that classifies both the speech data and a textual transcription in independent sub-modules that can be switched out without pre-training the whole model.

We suggest using Speech Emotion Recognition as a reward for Reinforcement Learning. It could be used to improve the quality of dialogue systems by allowing them to adjust their behaviour based on the emotional reactions of their dialogue partner. We mention possible problems that arise due to different aspects of human emotion and propose a suitable reward function to circumvent them.

Future work will involve the implementation of our framework into a Reinforcement Learning agent as well as more sophisticated combinations of Speech Emotion Recognition sub-modules. We will also repeat our experiments with an ASR system instead of artificial noise. Another possible direction is to optimize pre-trained language models for emotional context, for example by adding an additional training step before fine-tuning them on a given task.

7. ACKNOWLEDGEMENTS

The project on which this report is based was funded by the Federal Ministry of Education and Research (BMBF) of Germany under the number 01IS18040A. The authors are responsible for the content of this publication.

8. REFERENCES

- [1] Benedetto De Martino, Dharshan Kumaran, Ben Seymour, and Raymond J. Dolan, “Frames, biases, and rational decision-making in the human brain,” *Science*, vol. 313, no. 5787, pp. 684–687, 2006.

- [2] Jeffrey Bouffard, “The influence of emotion on rational decision making in sexual aggression,” *Journal of Criminal Justice*, vol. 30, pp. 121–134, 02 2002.
- [3] Elizabeth A. Phelps, Sam Ling, and Marisa Carrasco, “Emotion facilitates perception and potentiates the perceptual benefits of attention,” *Psychological Science*, vol. 17, no. 4, pp. 292–299, 2006, PMID: 16623685.
- [4] Elizabeth A. Kensinger, “Negative emotion enhances memory accuracy: Behavioral and neuroimaging evidence,” *Current Directions in Psychological Science*, vol. 16, no. 4, pp. 213–218, 2007.
- [5] Gustavo Aguilar, Viktor Rozgic, Weiran Wang, and Chao Wang, “Multimodal and multi-view models for emotion recognition,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July 2019, pp. 991–1002, Association for Computational Linguistics.
- [6] Jaejin Cho, Raghavendra Pappagari, Purva Kulkarni, Jesús Villalba, Yishay Carmiel, and Najim Dehak, “Deep neural networks for emotion recognition combining audio and transcripts,” in *Proc. Interspeech 2018*, 2018, pp. 247–251.
- [7] Samarth Tripathi and Homayoon S. M. Beigi, “Multimodal emotion recognition on IEMOCAP dataset using deep learning,” *CoRR*, vol. abs/1804.05788, 2018.
- [8] Seunghyun Yoon, Seokhyun Byun, and Kyomin Jung, “Multimodal speech emotion recognition using audio and text,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 112–118.
- [9] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever, “Language models are unsupervised multitask learners,” *OpenAI Blog*, vol. 1, no. 8, 2019.
- [10] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao, “Multi-task deep neural networks for natural language understanding,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July 2019, pp. 4487–4496, Association for Computational Linguistics.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 4171–4186.
- [12] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le, “Xlnet: Generalized autoregressive pretraining for language understanding,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Álché-Buc, E. Fox, and R. Garnett, Eds., pp. 5754–5764. Curran Associates, Inc., 2019.
- [13] Sara Rosenthal, Noura Farra, and Preslav Nakov, “SemEval-2017 task 4: Sentiment analysis in twitter,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Vancouver, Canada, Aug. 2017, pp. 502–518, Association for Computational Linguistics.
- [14] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan, “Iemocap: interactive emotional dyadic motion capture database,” *Language Resources and Evaluation*, vol. 42, no. 4, pp. 335, Nov 2008.
- [15] S. Poria, I. Chaturvedi, E. Cambria, and A. Hussain, “Convolutional mkl based multimodal emotion recognition and sentiment analysis,” in *2016 IEEE 16th International Conference on Data Mining (ICDM)*, Dec 2016, pp. 439–448.
- [16] Carlos Busso, Zhigang Deng, Serdar Yildirim, Murtaza Bulut, Chul Min Lee, Abe Kazemzadeh, Sungbok Lee, Ulrich Neumann, and Shrikanth Narayanan, “Analysis of emotion recognition using facial expressions, speech and multimodal information,” in *Proceedings of the 6th International Conference on Multimodal Interfaces*, New York, NY, USA, 2004, ICMI ’04, pp. 205–211, ACM.
- [17] Fatemeh Noroozi, Ciprian Corneanu, Dorota Kamińska, Tomasz Sapiński, Sergio Escalera, and Gholamreza Anbarjafari, “Survey on emotional body gesture recognition,” *IEEE Transactions on Affective Computing*, vol. PP, 01 2018.
- [18] M. Sreeshakthy and J. Preethi, “Classification of human emotion from deap eeg signal using hybrid improved neural networks with cuckoo search,” *BRAIN. Broad Research in Artificial Intelligence and Neuroscience*, vol. 6, no. 3-4, pp. 60–73, 2016.
- [19] Haiyang Xu, Hui Zhang, Kun Han, Yun Wang, Yiping Peng, and Xiangang Li, “Learning Alignment for Multimodal Emotion Recognition from Speech,” in *Proc. Interspeech 2019*, 2019, pp. 3569–3573.
- [20] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, “Phoneme recognition using time-delay neural networks,” *IEEE Transactions on Acoustics, Speech,*

and *Signal Processing*, vol. 37, no. 3, pp. 328–339, March 1989.

- [21] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [22] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer, “Deep contextualized word representations,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, New Orleans, Louisiana, June 2018, pp. 2227–2237, Association for Computational Linguistics.
- [23] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [24] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [25] Purnima Chandrasekar, Santosh V. Chapaneri, and Deepak J. Jayaswal, “Automatic speech emotion recognition: A survey,” *2014 International Conference on Circuits, Systems, Communication and Information Technology Applications (CSCITA)*, pp. 341–346, 2014.
- [26] Frank Dellaert, Thomas Polzin, and Alex Waibel, “Recognizing emotion in speech,” in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP’96*. IEEE, 1996, vol. 3, pp. 1970–1973.
- [27] Thomas Polzin and Alex Waibel, “Detecting emotions in speech,” *Proceedings of the 2nd International Conference on Cooperative Multimodal Communication, CMC*, 1 1998.
- [28] Björn Schuller, Stefan Steidl, and Anton Batliner, “A.: The interspeech 2009 emotion challenge,” in *In ISCA, ed.: Proceedings of Interspeech*, 2009, pp. 312–315.
- [29] Moataz El Ayadi, Mohamed S. Kamel, and Fakhri Karrray, “Survey on speech emotion recognition: Features, classification schemes, and databases,” *Pattern Recognition*, vol. 44, no. 3, pp. 572 – 587, 2011.
- [30] Björn W Schuller, “Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends,” *Communications of the ACM*, vol. 61, no. 5, pp. 90–99, 2018.
- [31] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio, “On the properties of neural machine translation: Encoder–decoder approaches,” in *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, Doha, Qatar, Oct. 2014, pp. 103–111, Association for Computational Linguistics.
- [32] Mousmita Sarma, Pegah Ghahremani, Daniel Povey, Nagendra Kumar Goel, Kandarpa Kumar Sarma, and Najim Dehak, “Emotion identification from raw speech signals using dnns,” in *Proc. Interspeech 2018*, 2018, pp. 3097–3101.
- [33] Aharon Satt, Shai Rozenberg, and Ron Hoory, “Efficient emotion recognition from speech using deep learning on spectrograms,” in *INTERSPEECH*, 2017, pp. 1089–1093.
- [34] Lorenzo Tarantino, Philip N Garner, and Alexandros Lazaridis, “Self-attention for speech emotion recognition,” *Proc. Interspeech 2019*, pp. 2578–2582, 2019.
- [35] Zheng Lian, Ya Li, Jianhua Tao, and Jian Huang, “Improving speech emotion recognition via transformer-based predictive coding through transfer learning,” *arXiv preprint arXiv:1811.07691*, 2018.
- [36] Haytham M. Fayek, Margaret Lech, and Lawrence Cavedon, “On the correlation and transferability of features between automatic speech recognition and speech emotion recognition,” in *Interspeech 2016*, 2016, pp. 3618–3622.
- [37] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. apr 2015, IEEE.
- [38] Xuerui Yang, Jiwei Li, and Xi Zhou, “A novel pyramidal-fsmn architecture with lattice-free mmi for speech recognition,” *CoRR*, vol. abs/1810.11352, 2018.
- [39] Jason Li, Vitaly Lavrukhin, Boris Ginsburg, Ryan Leary, Oleksii Kuchaiev, Jonathan M. Cohen, Huyen Nguyen, and Ravi Teja Gadde, “Jasper: An End-to-End Convolutional Neural Acoustic Model,” in *Proc. Interspeech 2019*, 2019, pp. 71–75.
- [40] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 4960–4964.
- [41] Kazuki Irie, Rohit Prabhavalkar, Anjuli Kannan, Antoine Bruguier, David Rybach, and Patrick Nguyen, “On

- the Choice of Modeling Unit for Sequence-to-Sequence Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 3800–3804.
- [42] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 2613–2617.
- [43] Linhao Dong, Shuang Xu, and Bo Xu, “Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition,” *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5884–5888, 2018.
- [44] Ngoc-Quan Pham, Thai-Son Nguyen, Jan Niehues, Markus Müller, and Alex Waibel, “Very Deep Self-Attention Networks for End-to-End Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 66–70.
- [45] Paul Ekman, “Basic emotions,” *Handbook of cognition and emotion*, pp. 45–60, 1999.
- [46] Pollyanna Gonçalves, Matheus Araújo, Fabrício Benvenuto, and Meeyoung Cha, “Comparing and combining sentiment analysis methods,” in *Proceedings of the First ACM Conference on Online Social Networks*, New York, NY, USA, 2013, COSN ’13, pp. 27–38, ACM.
- [47] Efthymios Kouloumpis, Theresa Wilson, and Johanna D. Moore, “Twitter sentiment analysis: The good the bad and the omg!,” in *ICWSM*, 2011.
- [48] Alexander Hogenboom, Daniella Bal, Flavius Frascar, Malissa Bal, Franciska de Jong, and Uzay Kaymak, “Exploiting emoticons in sentiment analysis,” in *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, New York, NY, USA, 2013, SAC ’13, pp. 703–710, ACM.
- [49] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann, “Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm,” in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, Sept. 2017, pp. 1615–1625, Association for Computational Linguistics.
- [50] Christos Baziotis, Nikos Pelekis, and Christos Douk-eridis, “DataStories at SemEval-2017 task 4: Deep LSTM with attention for message-level and topic-based sentiment analysis,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Vancouver, Canada, Aug. 2017, pp. 747–754, Association for Computational Linguistics.
- [51] Mathieu Cliche, “BB_twtr at SemEval-2017 task 4: Twitter sentiment analysis with CNNs and LSTMs,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Vancouver, Canada, Aug. 2017, pp. 573–580, Association for Computational Linguistics.
- [52] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman, “GLUE: A multi-task benchmark and analysis platform for natural language understanding,” in *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, Brussels, Belgium, Nov. 2018, pp. 353–355, Association for Computational Linguistics.
- [53] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” 2019.
- [54] Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu, “Freelb: Enhanced adversarial training for language understanding,” 2019.
- [55] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut, “Albert: A lite bert for self-supervised learning of language representations,” 2019.
- [56] Chenyang Huang, Amine Trabelsi, and Osmar R. Zaiane, “ANA at semeval-2019 task 3: Contextual emotion detection in conversations through hierarchical lstms and BERT,” in *Proceedings of the 13th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2019, Minneapolis, MN, USA, June 6-7, 2019*, 2019, pp. 49–53.
- [57] Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu, “BERT post-training for review reading comprehension and aspect-based sentiment analysis,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, 2019, pp. 2324–2335.
- [58] Milan Gnjatović and Dietmar Rösner, “Adaptive dialogue management in the nimitex prototype system,” in *Perception in Multimodal Dialogue Systems*, Elisabeth André, Laila Dybkjær, Wolfgang Minker, Heiko Neumann, Roberto Pieraccini, and Michael Weber, Eds., Berlin, Heidelberg, 2008, pp. 14–25, Springer Berlin Heidelberg.
- [59] Yuya Chiba, Takashi Nose, Taketo Kase, Mai Yamanaka, and Akinori Ito, “An analysis of the effect

of emotional speech synthesis on non-task-oriented dialog system,” in *Proceedings of the 19th Annual SIG-Dial Meeting on Discourse and Dialogue*, Melbourne, Australia, July 2018, pp. 371–375, Association for Computational Linguistics.

- [60] Thomas M. Moerland, Joost Broekens, and Catholijn M. Jonker, “Emotion in reinforcement learning agents and robots: a survey,” *Machine Learning*, vol. 107, no. 2, pp. 443–480, Feb 2018.
- [61] F. Burkhardt, J. Ajmera, R. Englert, J. Stegmann, and W. Bursleson, “Detecting anger in automated voice portal dialogs,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2006, vol. 2, pp. 1053–1056.
- [62] Jon Irastorza and María Inés Torres, “Tracking the expression of annoyance in call centers,” in *Cognitive Infocommunications, Theory and Applications*, pp. 131–151. Springer, 2019.
- [63] Chantal Berna, Siri Leknes, Emily A Holmes, Robert R Edwards, Guy M Goodwin, and Irene Tracey, “Induction of depressed mood disrupts emotion regulation neurocircuitry and enhances pain unpleasantness,” *Biological psychiatry*, vol. 67, no. 11, pp. 1083–1090, 2010.
- [64] Sara C Sereno, Graham G Scott, Bo Yao, Elske J Thaden, and Patrick J O’Donnell, “Emotion word processing: does mood make a difference?,” *Frontiers in psychology*, vol. 6, pp. 1191, 2015.
- [65] Andrew Lane, Christopher Beedie, and Peter Terry, “Distinctions between emotion and mood,” *Cognition and Emotion*, vol. 19, 09 2005.
- [66] Diederik Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *International Conference on Learning Representations*, 12 2014.
- [67] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, Jan. 2014.
- [68] Mickael Rouvier, “LIA at SemEval-2017 task 4: An ensemble of neural networks for sentiment classification,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Vancouver, Canada, Aug. 2017, pp. 760–765, Association for Computational Linguistics.
- [69] Yichun Yin, Yangqiu Song, and Ming Zhang, “NNEMBs at SemEval-2017 task 4: Neural twitter sentiment classification: a simple ensemble method with different embeddings,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Vancouver, Canada, Aug. 2017, pp. 621–625, Association for Computational Linguistics.
- [70] James A Russell, “A circumplex model of affect.,” *Journal of personality and social psychology*, vol. 39, no. 6, pp. 1161, 1980.
- [71] Niehues-J. Waibel A. Sperber, M., “Toward robust neural machine translation for noisy input sequences,” in *In International Workshop on Spoken Language Translation (IWSLT)*., 2017.
- [72] Stefan Constantin, Jan Niehues, and Alex Waibel, “Incremental processing of noisy user utterances in the spoken language understanding task,” in *Proceedings of the Fifth Workshop on Noisy User-generated Text (W-NUT)*, 2019.
- [73] Gaetan Ramet, Philip N. Garner, Michael Baeriswyl, and Alexandros Lazaridis, “Context-aware attention mechanism for speech emotion recognition,” in *IEEE Workshop on Spoken Language Technology*, Dec. 2018, pp. 126–131.
- [74] Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller, “Recent developments in opensmile, the munich open-source multimedia feature extractor,” in *Proceedings of the 21st ACM International Conference on Multimedia*, New York, NY, USA, 2013, MM ’13, pp. 835–838, ACM.