



Leveraging Pre-trained Language Model for Speech Sentiment Analysis

Suwon Shon^{1*}, Pablo Brusco^{1*}, Jing Pan¹, Kyu J. Han¹, Shinji Watanabe²

¹ASAPP, USA

²Carnegie Mellon University, USA

{sshon, pbrusco, jpan, khan}@asapp.com, shinjiw@ieee.org

Abstract


In this paper, we explore the use of pre-trained language models to learn sentiment information of written texts for speech sentiment analysis. First, we investigate how useful a pre-trained language model would be in a 2-step pipeline approach employing Automatic Speech Recognition (ASR) and transcripts-based sentiment analysis separately. Second, we propose a pseudo label-based semi-supervised training strategy using a language model on an end-to-end speech sentiment approach to take advantage of a large, but unlabeled speech dataset for training. Although spoken and written texts have different linguistic characteristics, they can complement each other in understanding sentiment. Therefore, the proposed system can not only model acoustic characteristics to bear sentiment-specific information in speech signals, but learn latent information to carry sentiments in the text representation. In these experiments, we demonstrate the proposed approaches improve F1 scores consistently compared to systems without a language model. Moreover, we also show that the proposed framework can reduce 65% of human supervision by leveraging a large amount of data without human sentiment annotation and boost performance in a low-resource condition where the human sentiment annotation is not available enough.

Index Terms: speech sentiment analysis, pre-trained language model, end-to-end speech recognition


1. Introduction

Speech sentiment analysis is the task of classifying positive/neutral/negative sentiments of a given speech. Compared to emotion recognition, it is a more abstract level of a recognition task. For example, negative sentiment not only contains anger emotion, but it also includes disparagement, sarcasm, doubt, suspicion, frustration, etc [1]. These negative sentiments may be related to the acoustic/prosodic features of speech as well as relevant to the context of the speech.

The conventional approach for speech sentiment analysis is using ASR on speech then employing sentiment analysis on the ASR transcripts so that it becomes a text classification task in a 2-step pipeline or cascade pipeline. However, this 2-step approach has two major disadvantages. First, it loses rich acoustic/prosodic information which is critical to understand spoken language. Second, there is a lack of sentiment-annotated datasets available when it comes to the spoken conversations domain, therefore, systems trained on a different communication channel from the conversational speech would tend to under-perform in production environments in the wild [2]. To address the first drawback, recently end-to-end (E2E) type speech sentiment analysis systems were proposed [2, 3, 4, 5, 6, 7, 8]. For the second problem, one op-



(a) 2-step pipeline



(b) E2E system with semi-supervised training

Figure 1: Proposed speech sentiment analysis

tion can be collecting more sentiment-labeled spoken data such as [9], but such options would require a costly effort.

In this paper, to tackle these two problems all together, we investigate how we can minimize human supervision and, at the same time, still efficiently train a model for speech sentiment analysis. In the Natural Language Processing (NLP) field, we can easily obtain robust sentiment analysis models by fine-tuning powerful pre-trained LMs such as BERT [10]. We thus leverage the pre-trained LMs to provide the additional information obtained through large training datasets for speech sentiment analysis models. We evaluate this approach in a 2-step pipeline setup first and compare it against an E2E framework in learning speech sentiments with limited training data. For the 2-step setup, we use a pre-trained LM as is, then process ASR transcripts as input for text-based sentiment analysis as shown in Fig.1(a). For the E2E model, we use a semi-supervised training approach using LM as a pseudo labeler as shown in Fig. 1 (b). In this system, we use an ASR encoder and pseudo labeler to train a sentiment classification model semi-supervisedly with pseudo labeled data then further tune the sentiment classification model with smaller, annotated data in the fine-tuning stage. We compare our system with the state-of-the-art system proposed in [2]. The evaluation was performed on a large sentiment dataset published by [9] that mitigates drawbacks in conventional datasets such as being scripted [11] or single speaker monologues [12].

2. Related work

Learning good representation from speech signals is the key to a speech sentiment/emotion analysis task. A recent study suggests to use a pre-trained ASR encoder [2] to prevent overfitting, and it showed promising results by surpassing the traditional audio + text multimodal systems [13, 14, 15]. Without the pre-trained ASR encoder, the model tends to overfit to the training data and the same model working on emotion recognition gives mediocre results on the sentiment analysis task [2].

*Equal contribution

Similarly, in the study of Spoken Language Understanding (SLU), pre-training approaches were proposed in combination with ASR [16, 17, 18] or acoustic classification modules [19], using ground truth (GT) text or ASR transcripts to improve SLU performance under limited resources.

The aforementioned pre-training approaches are based on the assumption that if a model is pre-trained to recognize words or phonemes, the fine-tuning result of downstream tasks will be improved. Our approach is also based on the same assumption, but we propose the use of powerful pre-trained LMs to transfer more abstract knowledge from the written text-domain to speech sentiment analysis. Specifically, we leverage pre-trained BERT models to extract robust embedding from text tokens for the 2-step pipeline, and to generate pseudo labels for semi-supervised training a model for the E2E speech sentiment analysis system.

3. Approaches

In the field of NLP, great advances have been made through pre-training task-agnostic LMs without any supervision. These pre-trained models can be fine-tuned using downstream task-specific data and showed state-of-the-art performance in many problems such as text classification, question answering, and summarization [10, 20, 21, 22].

In this section, we describe how we use LMs, such as BERT, for speech sentiment analysis. First, we place the pre-trained LM as embedding layers for the sentiment classification model in the 2-step pipeline framework. Second, we explore how the speech sentiment classifier in the E2E system can be enhanced through the use of a pseudo label-based semi-supervised training approach applied to large audio corpora without human annotations. By leveraging the LM in both the 2-step pipeline and E2E framework, we expect models to generalize better by being able to integrate the text-domain sentiment-related knowledge into the speech sentiment analysis space.

3.1. 2-step pipeline

Here we describe how we build two systems that use ASR transcripts as input: a baseline system in which embeddings are trained from scratch, and a BERT-based model that uses the pre-trained BERT as an embedding layer.

Suppose an input acoustic feature sequence is $\mathbf{x}_{1:T}$ and an ASR transcript or GT transcript token sequence is $\mathbf{o}_{1:L}$. Then, for the baseline system, the token sequence $\mathbf{o}_{1:L}$ is followed by BLSTM layers and an output layer to predict the three possible sentiments, e.g. Negative, Neutral, and Positive. Therefore, the model is trained by maximizing $P(y|\theta_s, \mathbf{o}_{1:L})$ where θ represents the model parameters and y the GT sentiment label. The BERT-based system uses a pre-trained BERT model (bert-base-uncased [23]) to encode the token sequence $\mathbf{o}_{1:L}$ into a BERT output sequence $\mathbf{z}_{1:L}$, after which the same sentiment classification layers follow as the baseline system has. The function to maximize in this case is $P(y|\theta_b, \mathbf{z}_{1:L})$. In Section 4.2 we describe more details about these two architectures.

3.2. Semi-supervised E2E speech sentiment analysis

As for the E2E systems, we start by creating a baseline system that uses an ASR encoder output as features for speech sentiment analysis [2]. Based on this framework, we propose a pseudo label-based semi-supervised training approach, shown in Figure 1.(b), that we describe in the next subsections.

3.2.1. Sentiment classifier

Let $\mathbf{h}_{1:T}$ be the ASR encoder output given $\mathbf{x}_{1:T}$. The sentiment classifier block takes $\mathbf{h}_{1:T}$ as input and predicts a sentiment class. We follow a similar architecture to the one proposed in [8]. This architecture has two BLSTM layers after a fully connected (FC) layer. Then, an attention-based weighted pooling takes the output sequence of the BLSTM, so that the average is performed to summarize the frame-level embedding into an utterance level embedding. Finally, the output layer maps the utterance-level embedding into a sentiment class. As the cost function, we used a cross entropy loss.

3.2.2. Semi-supervised training with pseudo label

To transfer the knowledge from the text domain, we generated sentiment pseudo labels \hat{y} using a pre-trained LM, a *pseudo labeler* from the given token sequence $\mathbf{o}_{1:L}$, which can be generated either from the GT or the ASR transcripts. Then, we use the pseudo labels to train the sentiment classifier.

For building this pseudo labeler, we first chose a few state-of-the-art pre-trained LMs, i.e. BERT [10], DistilBERT [20], RoBERTa [22], XLNet [21], that we fine-tuned with the Stanford Sentiment Treebank (SST) data [24]. At the end of this process, we obtained four different text-based sentiment analysis models that we will use as pseudo labelers in our experiments. The semi-supervised training of the sentiment classifier using pseudo labels can be done by maximizing $P(\hat{y}|\theta_p, \mathbf{h}_{1:T})$ to pre-train $\hat{\theta}_p$, and fine-tuning by maximizing $P(y|\theta_f, \mathbf{h}_{1:T}; \hat{\theta}_p)$.

4. Experiments

4.1. Datasets and Metric

For our experiments, we used the SWBD-Sentiment dataset [9] labeled with 3 sentiments (negative, neutral, and positive) by 3 different human annotators for every segment. From each segment, we computed the majority vote and discarded utterances in which there was a 3-way disagreement. We split the resulting data into a 86h training set (SWBD-train), a 5h test set (SWBD-test), and a 5h holdout set (SWBD-holdout)¹. We use SWBD-test as our validation set during training for choosing the best hyperparameters. We used SWBD-holdout as our evaluation dataset.

During evaluation, we computed weighted and unweighted averages of recall (REC), precision (PRE), and F1 scores (F1). Note that the conventional weighted/unweighted accuracy is equivalent to the weighted/unweighted REC since speech sentiment analysis is a closed-set multi-class classification task.

4.2. 2-step pipeline experiment setup

In some of our experiments, instead of using the GT texts, we used the ASR transcripts of the SWBD-Sentiment and Fisher datasets [25]. These were generated by using an HMM-DNN hybrid ASR model with a multistream CNN architecture [26] for acoustic modeling, and 4- and 5-gram LMs for 1st pass decoding and rescoring, respectively. This ASR model was trained on approximately 1,900h speech data consisting of in-house phone call data and the Switchboard Cellular Part 1 dataset.

As for the transcript-based sentiment classification component, (either GT or ASR) transcripts were tokenized in a subword unit with a max length of 500 tokens. Regarding the classification model, using the SWBD-train and validation sets,

¹The split information was provided by the authors of [9].

Table 1: 2-step pipeline evaluation on pre-trained BERT systems vs. the baseline systems. All the models were evaluated on ASR transcripts (SWBD-test/holdout-ASR means ASR transcripts). If not specified, the models use the full SWBD-train set (86h).

Architecture	SWBD-train transcript type	Validation Set (SWBD-test-ASR)						Evaluation Set (SWBD-holdout-ASR)					
		Unweighted			Weighted			Unweighted			Weighted		
		REC	PRE	F1	REC	PRE	F1	REC	PRE	F1	REC	PRE	F1
Baseline	GT	59.27	55.55	55.06	56.39	62.88	57.50	59.21	55.77	55.55	56.93	62.10	57.91
	ASR	52.57	50.07	47.38	47.60	58.44	48.30	52.43	49.84	47.66	47.92	57.12	48.48
	ASR (5h)	33.33	18.31	23.64	54.93	30.18	38.96	33.33	17.82	23.22	53.46	28.58	37.24
BERT	GT	63.87	64.64	64.12	68.16	68.01	67.96	64.53	65.05	64.56	67.87	67.98	67.73
	ASR	63.75	65.21	63.87	68.08	68.46	67.78	63.63	65.13	63.64	67.29	67.94	66.99
	ASR (5h)	50.18	55.01	50.99	61.08	58.88	58.82	50.09	56.06	51.03	60.85	58.91	58.34

Table 2: E2E speech sentiment analysis baseline evaluation. SWBD-train set was used for training.

Input feature	Sentiment Classifier Architecture	Validation Set (SWBD-test)						Evaluation Set (SWBD-holdout)					
		Unweighted			Weighted			Unweighted			Weighted		
		REC	PRE	F1	REC	PRE	F1	REC	PRE	F1	REC	PRE	F1
FBank	CNN	41.94	47.88	41.87	56.21	52.94	51.73	40.00	45.62	38.90	51.68	49.68	46.70
RNN-T encoder [2]	BLSTM	62.39	-	-	70.10	-	-	-	-	-	-	-	-
CTC-Attention encoder	BLSTM	64.59	68.89	66.24	71.41	70.86	70.72	61.21	65.92	62.74	67.73	67.89	66.99

we performed a hyperparameter search for finding the optimal weights for the trainable layers and also explored different design decisions, like how to utilize the pre-trained BERT embedding (i.e., in the BERT-based system, using the last layer or the last four layers or the sentence embedding provided by BERT – similar techniques to the ones proposed in [10]). We also tested pre-trained embeddings specifically tuned for sentiment classification (i.e. DistilBERT-SST2), but the plain BERT model was shown to provide better results. The selected baseline model has an embedding dimension of 200, two BLSTM layers with a hidden dimension of 128, trained with a weighted cross-entropy loss as the optimization objective. The selected BERT model uses three BLSTM layers applied on top of the sum of the last four BERT layers and an unweighted cross-entropy loss.

In our experiments, we trained and evaluated systems using both GT and ASR transcripts from the SWBD-sentiments corpus. In this way, we can not only compare our system results with other systems (that usually use GT transcripts), but also measure a performance gap in the case of a real production system that would run on top of ASR results.

4.3. E2E system experiment setup

For the E2E speech sentiment analysis systems, we utilized the encoder part of an E2E ASR system that we trained using our in-house data. The ASR model has an encoder-decoder architecture where each component is based on Transformer jointly optimized with the CTC loss [27]. This system included a byte-pair encoding tokenizer (token size of 2,000), and the encoder consisted of 12 transformer blocks that generated a 512 dimension embedding vector. All parameters in the encoder were fixed in the experiments.

For choosing the pseudo labeler, we evaluated several pre-trained LMs fine-tuned with the Stanford Sentiment Treebank (SST) dataset [24], as shown in Table 3. In the SST corpus, there are two types of labels, fine-grained (5-classes, SST5) and binary (negative/positive, SST2), and we used the SST2 portion to fine-tune the models. The table shows the REC score on the GT transcripts of the evaluation set. Since the family of SST2 models produces the binary classes, we ran this evaluation only on the negative and positive utterances in the evaluation set (thus these numbers are incomparable to the 3-way classification task with the SWBD-Sentiment dataset in our other experiments). Based on the results in the table, we chose BERT-SST2 and XLNet-SST2 as our pseudo labelers.

The sentiment classifier in the E2E system was trained us-

Table 3: Pseudo labeler performance on SWBD-holdout for negative and positive classes

	Unweighted REC	Weighted REC
BERT-SST2	70.99	71.05
DistilBERT-SST2	70.95	71.16
RoBERTa-SST2	70.16	70.25
XLNet-SST2	72.15	72.19
BERT-SST2 (ASR transcript)	69.08	69.20

ing the SWBD-train portion of the data. In the semi-supervised training phase with pseudo labels (Figure 1(b)), we discarded the output layer of the classifier since the pseudo labeler is binary as $\hat{y} = \{Neg, Pos\}$. Then, we replaced it with a randomly initialized output layer which has a 3 class output in the fine-tuning stage. When fine-tuning the sentiment classifier, we updated the whole parameters in the model.

4.4. Experiment result

4.4.1. 2-step pipeline

Table 1 shows 2-step pipeline experiment results. First, we observe that the BERT-based models outperform the baseline systems trained with the simple neural net architecture described in Section 3.1 across all the metrics. Second, the baseline model trained with GT transcripts displayed a remarkable performance drop compared to the baseline model trained with ASR transcripts (55.55 vs. 47.66 Unweighted F1). On the other hand, the BERT-based model trained on ASR transcripts exhibited similar performance to the one trained on GT transcripts (64.56 vs. 63.64 Unweighted F1).

This table also contrasts performance drop in the baseline and BERT-based systems when using a subset of our training data (5h). The baseline system showed a notable drop of 51% Unweighted F1 score. Although, the BERT-based model showed only a 19% decrease under the same condition. We also observe that the 5h BERT-based model outperforms the baseline model trained on ASR transcripts corresponding to 86 hours by approximately 7% Unweighted F1 on both the validation and evaluation set (47.38 vs. 50.99, and 51.03 vs. 47.66, respectively). This shows the knowledge embedded in the pre-trained LM can be distilled to the sentiment classifier in the 2-step pipeline even with a small fine-tuning data.

4.4.2. E2E system

Table 2 shows the performance of E2E systems. Compared to the model presented in [2] (which used RNN-T to train an

Table 4: Semi-supervised approach on E2E speech sentiment analysis system evaluation. \mathcal{S} :SWBD-train with GT transcripts, \mathcal{F} :Fisher with GT transcripts, \mathcal{S}_{asr} :SWBD-train with ASR transcripts, \mathcal{F}_{asr} :Fisher with ASR transcripts.

Fine-tuning dataset	Pseudo labeler	Semi-supervised training dataset	Validation Set (SWBD-test)						Evaluation Set (SWBD-holdout)					
			Unweighted			Weighted			Unweighted			Weighted		
			REC	PRE	F1	REC	PRE	F1	REC	PRE	F1	REC	PRE	F1
SWBD-train (86h)	-	-	64.59	68.89	66.24	71.41	70.86	70.72	61.21	65.92	62.74	67.73	67.89	66.99
	BERT-SST2	\mathcal{S}	63.68	67.65	65.23	70.37	69.79	69.71	62.37	66.68	63.85	68.47	68.58	67.85
	BERT-SST2	\mathcal{S}, \mathcal{F}	64.87	68.05	66.15	70.82	70.28	70.31	63.23	66.82	64.55	69.05	68.77	68.46
	XLNet-SST2	\mathcal{S}, \mathcal{F}	63.64	67.56	65.17	70.45	69.86	69.78	61.61	65.48	62.94	67.73	67.58	67.06
	BERT-SST2	$\mathcal{S}_{asr}, \mathcal{F}_{asr}$	65.74	66.51	66.11	70.23	70.01	70.10	64.18	65.28	64.57	68.27	68.35	68.14
SWBD-train (5h)	-	-	51.33	53.82	51.98	60.66	58.73	59.24	47.76	49.86	48.16	56.62	54.88	55.12
	BERT-SST2	\mathcal{S}	54.16	58.08	54.96	62.74	61.40	61.33	52.12	56.61	53.06	60.40	59.11	58.84
	BERT-SST2	\mathcal{S}, \mathcal{F}	58.72	58.67	58.54	63.92	63.74	63.72	57.45	57.92	57.63	61.98	61.67	61.79
	XLNet-SST2	\mathcal{S}, \mathcal{F}	58.19	57.89	58.00	62.63	63.07	62.82	56.86	57.39	56.75	60.59	61.52	60.74
	BERT-SST2	$\mathcal{S}_{asr}, \mathcal{F}_{asr}$	54.78	55.51	55.02	61.10	60.38	60.67	52.23	53.16	52.60	57.39	57.00	57.10

ASR encoder with a word-level token), our joint CTC-Attention model (that uses sub-word unit tokens) showed slightly better performance. We used this CTC-Attention encoder for the rest of the experiments.

The upper section of Table 4 shows the performances of the proposed semi-supervised training approaches in the E2E system using SWBD-train (86h) and Fisher (2,000h) for pseudo labeling and the full SWBD-train data alone as a fine-tuning set. We observed a similar performance for all the systems on the validation set, however, the semi-supervised training approaches with the pseudo labeled data generally outperformed the baseline system without the semi-supervised training phase on the evaluation set. It is also shown that the marginal gain can be obtained when we added the Fisher data on top of SWBD-train as the semi-supervised training dataset. Comparing the results obtained using XLNet-SST2 and BERT-SST2, we did not see a meaningful difference. Besides, using ASR transcripts as input to the pseudo labelers presented similar performance to using GT transcripts, showing the proposed system is robust to ASR errors from a sentiment classification perspective.

In the bottom section of Table 4, we show the results of the semi-supervised training approaches with only a 5h subset of SWBD-train for fine-tuning. In this setting, the semi-supervised training approaches showed significant improvements as opposed to the baseline without taking an advantage of pseudo labeling. Also, we observe that the best system showed about 20% improvement on unweighted F1 score on the evaluation set (57.63%) compared to the baseline (48.16%). Finally, using ASR transcripts for pseudo labeling did not give a similar improvement compared to using GT transcripts, but it shows still a better performance than the baseline.

If we compare to the upper section of the table, that is comparing a fine-tuning set of 86h vs. 5h, we observed that the baseline dropped about 20% (62.74% to 48.16%) while the best system decreased only by 10% (64.55% to 57.63%).

Additionally, we checked the performance by increasing fine-tuning data by 2.5h to predict how much we can reduce the human sentiment annotation job in Figure 2. From this experiment, we verified the best performing system with a semi-supervised training approach reaches the baseline performance when we use at least a 30h subset of SWBD-train. This means that we can save about 65% of human sentiment annotation jobs if we use the semi-supervised training approach.

4.5. Discussion

The 2-step pipeline experiments showed that the use of pre-trained LMs can achieve good performance even with a low amount of labeled data. Also, the pre-trained LM enables the model to use ASR transcript for training without performance degradation. We attribute this robustness to errors in words,




Figure 2: Semi-supervised training approach efficiency on evaluation set. Note that baseline used all of SWBD-train set (86h)

to the way BERT is trained – using token masking techniques. This result suggests that when new data for speech sentiment analysis is needed, we can skip the expensive human transcripts by using any off-the-shelf ASR.

Given that the most of the SST-2 benchmark results² are above 90% recall using BERT, the text sentiment classifier is not performing well on spoken speech as shown in Table 3, displaying once more the difficulties of using models trained on written data to adapt to transcribed conversational data. However, our noisy pseudo label-based semi-supervised training approach still showed encouraging results, and generally outperformed trained-from-scratch models in various conditions. These results suggest that sentiment pseudo labels carry the text-domain sentiment knowledge that could transfer some knowledge to speech sentiment classifiers on a semi-supervised training stage.

A limitation of this study is that we did not consider another text corpus for building the pseudo labeler. There are more fine-grained sentiment datasets such as SST-5 (5 classes), IMDb (10 classes), Yelp (5 classes). We believe these fine-grained data could benefit in different ways to what a binary sentiment system does. Another limitation is that we did not update the ASR encoder for the speech sentiment analysis system. We expect that updating the ASR encoder on both semi-supervised training and fine-tuning steps could considerably affect the results.

5. Conclusion

In this paper, we investigated an approach to transfer knowledge from the written text to spoken text or speech domain using an LM to reduce and use efficiently the human annotation on speech dataset. The experiments explored two scenarios, a 2-step pipeline, and an E2E speech sentiment analysis system, to verify the effectiveness of leveraging BERT. From the experiments, we observed that the proposed approach is able to encode the information robustly and generalize better with less supervision. While the proposed approaches show improvement in all conditions, we verified that it has a greater advantage in the case where a large amount of audio is available whether it is transcribed or not.

²<https://gluebenchmark.com/>

6. References

- [1] S. Mohammad, "A practical guide to sentiment annotation: Challenges and solutions," in *Proceedings of the 7th workshop on computational approaches to subjectivity, sentiment and social media analysis*, 2016, pp. 174–179.
- [2] Z. Lu, L. Cao, Y. Zhang, C.-C. Chiu, and J. Fan, "Speech sentiment analysis via pre-trained features from end-to-end ASR models," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7149–7153.
- [3] R. Li, Z. Wu, J. Jia, S. Zhao, and H. Meng, "Dilated residual network with multi-head self-attention for speech emotion recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6675–6679.
- [4] P. Li, Y. Song, I. McLoughlin, W. Guo, and L. Dai, "An attention pooling based representation learning method for speech emotion recognition," in *Interspeech*, 2018, pp. 3087–3091.
- [5] X. Wu, S. Liu, Y. Cao, X. Li, J. Yu, D. Dai, X. Ma, S. Hu, Z. Wu, X. Liu *et al.*, "Speech emotion recognition using capsule networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6695–6699.
- [6] Y. Xie, R. Liang, Z. Liang, C. Huang, C. Zou, and B. Schuller, "Speech emotion classification using attention-based LSTM," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1675–1685, 2019.
- [7] P. Tzirakis, J. Zhang, and B. W. Schuller, "End-to-end speech emotion recognition using deep neural networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5089–5093.
- [8] S. Mirsamadi, E. Barsoum, and C. Zhang, "Automatic speech emotion recognition using recurrent neural networks with local attention," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 2227–2231.
- [9] E. Chen, Z. Lu, H. Xu, L. Cao, Y. Zhang, and J. Fan, "A large scale speech sentiment corpus," in *Proceedings of The 12th Language Resources and Evaluation Conference*, 2020, pp. 6549–6555.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [11] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMOCAP: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [12] A. B. Zadeh, P. P. Liang, S. Poria, E. Cambria, and L.-P. Morency, "Multimodal language analysis in the wild: CMU-MOSEI dataset and interpretable dynamic fusion graph," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 2236–2246.
- [13] E. Kim and J. W. Shin, "DNN-based emotion recognition based on bottleneck acoustic features and lexical features," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6720–6724.
- [14] S. Siriwardhana, A. Reis, R. Weerasekera, and S. Nanayakkara, "Jointly fine-tuning "BERT-like" self supervised models to improve multimodal speech emotion recognition," in *Interspeech*, 2020, pp. 3755–3759.
- [15] J. Cho, R. Pappagari, P. Kulkarni, J. Villalba, Y. Carmiel, and N. Dehak, "Deep neural networks for emotion recognition combining audio and transcripts," in *Interspeech*, 2018, pp. 247–251.
- [16] H.-K. J. Kuo, Z. Tüske, S. Thomas, Y. Huang, K. Audhkhasi, B. Kingsbury, G. Kurata, Z. Kons, R. Hoory, and L. Lastras, "End-to-end spoken language understanding without full transcripts," *arXiv preprint arXiv:2009.14386*, 2020.
- [17] Y.-P. Chen, R. Price, and S. Bangalore, "Spoken language understanding without speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 6189–6193.
- [18] P. Haghani, A. Narayanan, M. Bacchiani, G. Chuang, N. Gaur, P. Moreno, R. Prabhavalkar, Z. Qu, and A. Waters, "From audio to semantics: Approaches to end-to-end spoken language understanding," in *IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 720–726.
- [19] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, "Speech model pre-training for end-to-end spoken language understanding," in *Interspeech*, 2019, pp. 814–818.
- [20] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.
- [21] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. V. Le, "XLNet: Generalized autoregressive pretraining for language understanding," *arXiv preprint arXiv:1906.08237*, 2019.
- [22] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A robustly optimized BERT pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [23] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Huggingface's transformers: State-of-the-art natural language processing," *arXiv preprint arXiv:1910.03771*, 2019.
- [24] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts, "Recursive deep models for semantic compositionality over a sentiment treebank," in *EMNLP*, 2013, pp. 1631–1642.
- [25] C. Cieri, D. Miller, and K. Walker, "The Fisher corpus: a resource for the next generations of speech-to-text," in *LREC*, vol. 4, 2004, pp. 69–71.
- [26] K. J. Han, J. Pan, V. K. N. Tadala, T. Ma, and D. Povey, "Multistream CNN for robust acoustic modeling," *arXiv preprint arXiv:2005.10470*, 2020.
- [27] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 4835–4839.