



Generative Data Augmentation Guided by Triplet Loss for Speech Emotion Recognition

Shijun Wang¹, Hamed Hemati¹, Jón Guðnason², Damian Borth¹

¹University of St. Gallen, Switzerland

²Reykjavik University, Iceland

shijun.wang@unisg.ch, hamed.hemati@unisg.ch, jg@ru.is, damian.borth@unisg.ch

Abstract

Speech Emotion Recognition (SER) is crucial for human-computer interaction but still remains a challenging problem because of two major obstacles: data scarcity and imbalance. Many datasets for SER are substantially imbalanced, where data utterances of one class (most often Neutral) are much more frequent than those of other classes. Furthermore, only a few data resources are available for many existing spoken languages. To address these problems, we exploit a GAN-based augmentation model guided by a triplet network, to improve SER performance given imbalanced and insufficient training data. We conduct experiments and demonstrate: 1) With a highly imbalanced dataset, our augmentation strategy significantly improves the SER performance (+8% recall score compared with the baseline). 2) Moreover, in a cross-lingual benchmark, where we train a model with enough source language utterances but very few target language utterances (around 50 in our experiments), our augmentation strategy brings benefits for the SER performance of all three target languages.

Index Terms: speech emotion recognition, speech augmentation, cross lingual

1. Introduction

Speech Emotion Recognition (SER) is a task aiming to understand the underlying emotional information in speech. In recent years, there has been a surge of interest in SER, because emotion conveys crucial information during human-computer interaction. However, two obstacles usually hinder the application of SER: the imbalance and scarcity issues of the SER dataset. We can further describe the issues from two perspectives.

Firstly, SER datasets have a common issue that the speech data distributions are non-uniform or highly skewed among different emotion classes. Speech utterances labeled as "Neutral" are much more frequent than those labeled as "Happy, Angry, etc.". The data scarcity of various emotion classes results in highly imbalanced datasets. To alleviate the data imbalance issue, a common way is to generate synthetic data through data augmentation techniques. Authors in [1] use vanilla Generative Adversarial Networks (GANs) [2] to generate synthetic openSMILE acoustic features [3] for the training of the SER task. In [4], the authors employ a CycleGAN [5] to transfer openSMILE feature vectors from one emotion class to another. Instead of openSMILE features, some approaches do augmentation on low-level acoustic features (e.g. Mel-Spectrogram), which is an advantage for other tasks. Approaches in [6, 7] apply a starGAN [8] to convert speech utterances of Neutral emotion to other emotions. Authors in [9] modify a GAN to perform augmentation for imbalanced data and they empirically demonstrate that this GAN can improve the SER performance under imbalanced datasets.

Another important aspect to consider the scarcity issue is the availability of data for various existing spoken languages. There are very few data resources available for many languages, which is a significant barrier to the research and application of SER for these languages. One way to address this problem is by training an SER model on one or more languages with enough data (source), combined with very few training utterances of other languages (target). The idea is to use the knowledge gained from source languages to improve the performance on low-resource target languages [10]. Such SER training strategy that spans different languages, i.e. cross-lingual SER, has been widely studied [11, 12, 13]. Although adding source languages improves the SER performance on target languages, there is still room to improve the cross-lingual SER task.

In this paper, We propose an augmentation strategy, which can be used to address the data imbalance and scarcity problem. One component of our augmentation approach is a GAN, whose generator is employed for augmentation. The generator is inspired by [14], where authors successfully perform automatic augmentation for multiple downstream tasks. To stabilize the GAN or improve its performance, additional auxiliary objectives are usually introduced during training [15, 16, 17]. Thus, we apply a representation learner based on triplet loss to learn emotion representations from the original data. And the learned knowledge of the representation learner is used to increase and stabilize the augmentation performance for the SER task.

We summarize our contributions as follows: 1) We propose a GAN-based augmentation for the SER task. 2) We demonstrate that our augmentation approach can largely improve the SER performance with a highly imbalanced dataset (+8% recall with our augmentation). 3) Moreover, for the cross-lingual SER task, our augmentation improves the performance for three different target languages with only around 50 utterances.

2. Method

The proposed SER data augmentation method is designed with two augmentation requirements: preserving the emotion class and providing rich variance. This means that an augmented utterance should have the same emotion label as the original one while also differing sufficiently from it. Our system setup and corresponding loss functions reflect these requirements.

2.1. Model Components

Our approach consists of a GAN and a representation learner. The generator in our GAN augment the Mel-Spectrograms. The representation learner enables the GAN to present high augmentation variance and stabilize the augmentation.

The GAN is shown on the top of Fig. 1, it contains one generator and one discriminator, while we denote the inputted Mel-

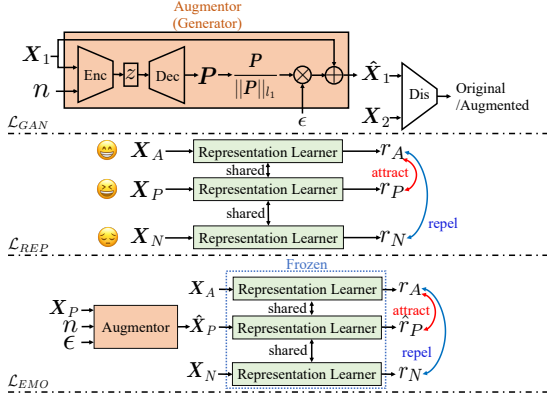


Figure 1: **TOP:** The training pipeline of our GAN. **MIDDLE:** The training pipeline of our representation learner. **BOTTOM:** The loss to guide the augmentor to preserve the emotion information.

Spectrograms as \mathbf{X} . In the rest of the paper, we refer to the generator as the augmentor since it is used to perform augmentation on Mel-Spectrograms. The augmentor employs an encoder to map a Mel-Spectrogram \mathbf{X}_1 into a latent vector z , given a Gaussian noise \mathbf{n} . Then, z is fed to the decoder to generate a tensor P with the same shape as \mathbf{X}_1 . To avoid excessive removal of features from the original Mel-Spectrogram, we constrain the decoder output P with an l_1 -norm operation as used in [14]. We then use a hyper-parameter ϵ to control the augmentation intensity and obtain the augmented Mel-Spectrogram $\hat{\mathbf{X}}_1$. For brevity, in this paper, we denote our augmentor as $AUG(\mathbf{X})$. Finally, we feed the discriminator $D(\mathbf{X})$ with $\mathbf{X} = \hat{\mathbf{X}}_1$ or \mathbf{X}_2 , where \mathbf{X}_2 is another sample from the original dataset. $D(\mathbf{X})$ aims to determine whether the input is augmented or not. Overall, we apply the non-saturating loss [2] to train our GAN,

$$\mathcal{L}_{GAN} = \min_{AUG} \max_D [\log D(\mathbf{X}) + \log(-D(AUG(\mathbf{X})))] \quad (1)$$

To achieve better augmentation, we further implement a representation learner. Its target is to learn meaningful emotion representations that can later be used to improve our augmentor. In contrast to a classification loss, a triplet loss is applied to train more discriminative emotion representations [18, 19]. As illustrated in the middle of Fig. 1, our representation learner maps the Mel-Spectrogram \mathbf{X} into an emotion representation vector r . The objective is to attract the representations of \mathbf{X}_A (Anchor) and \mathbf{X}_P (Positive), whose emotion classes are the same, while repelling representations of \mathbf{X}_A and \mathbf{X}_N (Negative), since they are from different emotion classes. We define the triplet loss of our representation learner as,

$$\mathcal{L}_{REP} = \max(\text{dist}(r_A, r_P) - \text{dist}(r_A, r_N) + \beta, 0), \quad (2)$$

where $\text{dist}(\cdot, \cdot)$ is l_1 distance, and β is the margin of the triplet loss. The margin is a hyper-parameter that defines how far away the dissimilarities should be.

The representation learner is then employed to guide the augmentor to preserve the emotion class after the augmentation. Thus, we introduce the loss \mathcal{L}_{EMO} to penalize the augmentor for generating an output that is of different emotion class to the input. The setup is shown at the bottom of Fig. 1. As we can see, we freeze the representation learner, and an augmented Mel-Spectrogram $\hat{\mathbf{X}}_P$ is used to supersede \mathbf{X}_P as one of the inputs for the triplet loss. The motivation is to use the knowledge

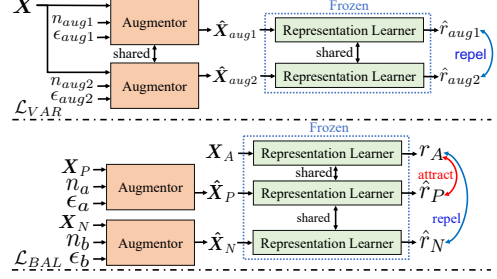


Figure 2: **TOP:** The auxiliary loss to guide the augmentor to provide augmentation variance. **BOTTOM:** The auxiliary loss to balance emotion preservation and augmentation variance.

of the representation learner to guide the augmentor to preserve emotional information. We define the loss as,

$$\mathcal{L}_{EMO} = \max(\text{dist}(r_A, \hat{r}_P) - \text{dist}(r_A, r_N) + \beta, 0). \quad (3)$$

Then, we define the model loss as,

$$\mathcal{L}_{Model} = w_g \mathcal{L}_{GAN} + w_r \mathcal{L}_{REP} + w_e \mathcal{L}_{EMO}, \quad (4)$$

where \mathcal{L}_{Model} are weighted by w_g , w_r and w_e .

2.2. Auxiliary Objectives for Augmentation

To achieve better augmentation, we apply the representation learner to provide auxiliary objectives during training. We first expect the augmentor can provide augmentation variance. In other words, for the same Mel-Spectrogram, the augmentor can produce various augmented versions. Otherwise, identical mapping or barely perceivable augmentation of the augmentor might not be helpful for the SER task. Thus, as shown on the top of Fig. 2, for the same input \mathbf{X} , we first obtain two augmented Mel-Spectrograms $\hat{\mathbf{X}}_{aug1}$ and $\hat{\mathbf{X}}_{aug2}$ based on different noise n and augmentation intensity ϵ . We then send these two augmented Mel-Spectrograms into the representation learner to acquire representations r_{aug1} and r_{aug2} . We freeze the representation learner as well in order to utilize its knowledge of the original data. Our target is to repel these two representations, which can force the augmentor to generate various augmented versions whose representations are dissimilar, even from the same input. We define the auxiliary variance loss as,

$$\mathcal{L}_{VAR} = \text{dot}(\hat{r}_{aug1}, \hat{r}_{aug2}). \quad (5)$$

We enforce the dissimilarity by minimizing the dot product ($\text{dot}(\cdot, \cdot)$) between the representations from two augmented Mel-Spectrograms.

However, our initial experiments showed that \mathcal{L}_{VAR} leads to unstable augmentation, we observe this issue by t-SNE visualisation [20]. As shown on the left in Fig. 3, each point depicts the emotion representation r derives from Mel-Spectrogram in the dataset IEMOCAP [21]. For each emotion class, we randomly pick 30 speech utterances and augment each sample 3 times with different noise and augmentation intensities. In the figure, the augmented representations from different emotions are not distinct and do not form clusters. The reason is that to meet the augmentation variance loss \mathcal{L}_{VAR} , the augmentor might perform arbitrarily to generate augmented versions whose representations are excessively far from original data, which means the emotion class information is damaged. Such behavior is against the loss \mathcal{L}_{EMO} (Eq. 3, aiming to preserve




Figure 3: *T-SNE visualization of the emotion representations from the representation learner. LEFT: the augmented utterances are from the augmentor trained without \mathcal{L}_{BAL} ; RIGHT: the augmented utterances are from the augmentor trained with \mathcal{L}_{BAL} , which leads to a more clear cluster boundary.*

the emotion information), however, since the dataset is small, the augmentor is less penalized by the original data.

To enable our augmentor to preserve the emotion information while providing variant augmentations, we need to balance \mathcal{L}_{EMO} (Eq. 3) and \mathcal{L}_{VAR} (Eq. 5). Rather than doing tedious fine-tuning of the weights of these two losses, we achieve the balance by introducing another auxiliary loss as shown at the bottom of Fig. 2. As we can see, another triplet loss is applied, but the positive and the negative inputs are both augmented. The motivation is to prevent the representations of the augmented Mel-Spectrograms from tangling together, by constraining the augmentor not only with original data, but also with augmented versions themselves. In other words, the augmented versions should be close to original versions if they belong to the same emotion, while be far from other augmented version when they from different emotion classes. Therefore, the triplet loss is defined as,

$$\mathcal{L}_{BAL} = \max(\text{dist}(r_A, \hat{r}_P) - \text{dist}(r_A, \hat{r}_N) + \beta, 0). \quad (6)$$

As we can observe from the t-SNE visualisation on the right of Fig. 3, with the addition of the balancing loss, the augmented representations have a clear boundary between different classes. We conduct ablation study to demonstrate the contributions of our proposed losses further. Finally, we define the total loss as,

$$\mathcal{L}_{Total} = \mathcal{L}_{Model} + w_v \mathcal{L}_{VAR} + w_b \mathcal{L}_{BAL}, \quad (7)$$

where w_v and w_b are loss weights.

2.3. Training Phases

We repeat the following steps until our models converge:

- 1) **Representation Learner:** We update our representation learner by \mathcal{L}_{REP} (Eq. 2) from original data \mathbf{X} .
- 2) **Discriminator:** We freeze the augmentor, and obtain augmented Mel-Spectrograms $\hat{\mathbf{X}}$ from original data \mathbf{X} . The discriminator D is updated by \mathbf{X} and $\hat{\mathbf{X}}$.
- 3) **Augmentation Variance:** We freeze the representation learner, and utilize it to enhance the augmentation variance of the augmentor, by the loss \mathcal{L}_{VAR} (Eq. 5).
- 4) **Emotion Preservation and Balance:** We update the augmentor. We freeze the representation learner, the discriminator D , then generate augmented Mel-Spectrograms $\hat{\mathbf{X}}$ from the augmentor. The augmentor is finally updated from the discriminator loss, \mathcal{L}_{EMO} (Eq. 3) and \mathcal{L}_{BAL} (Eq. 6).

2.4. Implementation Details

In our augmentor, the encoder consists of 3 convolutional layers with LeakyReLU activation function and maps the input \mathbf{X} into

a latent 128D vector z . The decoder employs two deconvolutional layers with ReLU activation function to convert z into a tensor \mathbf{P} whose shape is the same as the input. The discriminator first uses four convolutional layers with the LeakyReLU activation function to process the input, then uses an LSTM and an attention layer [22] to discern whether the input is original or augmented. Our representation learner employs 5 convolutional layers with LeakyReLU and an LSTM to process the Mel-Spectrogram, then the same attention layer in the discriminator is used to produce the 128 dimensional emotion representation.

For augmentation intensity ϵ , we randomly sample ϵ from a uniform distribution $U(0.05, 0.3)$, because we found a random ϵ performs better than a fixed one as in [14]. We set the β in \mathcal{L}_{REP} (Eq. 2), \mathcal{L}_{EMO} (Eq. 3) and \mathcal{L}_{BAL} (Eq. 6) as 7. For loss weights, we set w_g , w_r and w_v to 1, w_e to 10 and w_b to 8. We train all of our models with Adam optimizer for 30k iterations, and set the learning rate to 1e-6.

3. Experiment and Results

The proposed augmentation method is demonstrated with SER tasks. In short, the results show that: 1) Our augmentor can improve the SER with a highly imbalanced dataset. 2) Our augmentor can improve the cross-lingual SER task with very few training utterances from the target language. 3) the importance of the contribution of \mathcal{L}_{VAR} (Eq. 5, for adding augmentation variance), and \mathcal{L}_{BAL} (Eq. 6, for balancing the emotion preservation and providing variance).

3.1. Experimental Setup

3.1.1. Data Preprocessing and SER Classifier

In all our experiments, we use the same acoustic features and same emotion classifier as in [9]. To extract Mel-Spectrograms from waveform files, we use a 50-millisecond window and a 50 percent overlap ratio, with 128 Mel Coefficients. A VGG19 architecture [23] is employed as our emotion classifier, which takes a fix-sized 128×128 Mel-Spectrogram (128 frames) as the input and predicts the emotion class. During the evaluation, the predicted class is determined by the majority voting of segments for each Mel-Spectrogram. To train our GAN and representation learner, we use 512×128 Mel-Spectrograms.

3.1.2. Datasets

Imbalanced SER and Ablation Study: For the experiments of the Imbalanced SER (Sec. 3.2.1) and Ablation Study (Sec. 3.2.3) We train and test our models on IEMOCAP [24]. Like many other SER works [1, 9, 25], we only focus on 4 classes (Angry, Sad, Neutral and Happy), resulting in 5531 speech utterances of about 7 hours total duration. Furthermore, in order to simulate the data imbalance issue, we follow [9] and randomly remove 80% of each class except Neutral. Since IEMOCAP is originally split into 5 sessions, we conduct 5 fold cross-validation, by using 4 sessions for training and 1 for testing. We train our SER classifier with original and augmented utterances. We augment each original training sample four times, resulting in a hybrid original-augmented training dataset that is roughly the same size as the original non-reduced dataset.

Cross-Lingual SER: For the experiment of the Cross-Lingual SER (Sec. 3.2.2), our datasets are listed in Tab. 1. These datasets span four languages, with IEMOCAP and ESD serving as source languages and the remaining datasets serving as target languages. Although ESD originally has both Chinese and En-

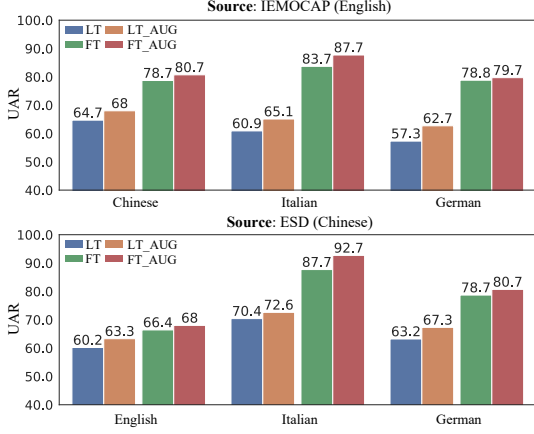


Figure 4: UAR results of cross-lingual SER. **TOP**: English (source) to 3 target languages. **BOTTOM**: Chinese (source) to 3 target languages.

English speech utterances, we only use the Chinese portion in this experiment. To create IEMOCAP.SUB and ESD.SUB, we randomly selected utterances from IEMOCAP and ESD, in order to make all of the target language datasets the same size. Since each dataset has different class numbers, one of the consistent ways to investigate cross-lingual SER task is by considering the binary positive/negative valence classification problem, like in many other related works [11, 12]. We follow the same method to map emotions into valence, and split 25% of the target language datasets for our evaluation. During the training of the SER classifier, we augment each original training sample from the target language 20 times, since the target language’s data is severely lacking (only 50 utterances in some experiment cases).

3.2. Results

3.2.1. SER on Imbalanced Dataset

Models: We follow [9] and train our SER classifier with: NoAUG: the 80% reduced dataset; AUG: the 80% reduced dataset and the augmented utterances.

Results: We list the Unweighted Average Recall (UAR) results reported by [9] and ours in Tab. 2. One thing to note is that, although using the identical setup in [9] to train our NoAUG model, the results (the 2nd and 4th rows) differ because the 80% removed component is randomly chosen. The average UAR results, on the other hand, are close, implying that our comparison is fair. By comparing the two AUG models, we can see that our augmentation strategy is roughly 5% higher than the augmentation approach in [9], which means a better SER performance can be achieved with our augmentation approach.

3.2.2. Cross-Lingual SER

Models: We train our SER classifier with: Low Target (LT), 100% of the source language data but only 10% of the target language dataset; LT_AUG, same as LT, but the augmented data is also fed during training; Full Target (FT), 100% of the source language data and 75% of the target; FT_AUG, same as FT, but we add the augmented data from our augmentor.

Results: The UAR results on 25% of the target language datasets are shown in Fig. 4. By comparing with LT and LT_AUG, our augmentation approach increases UAR performance by around 5%, given only 10% of the target language

Table 1: Cross-lingual datasets. IEMOCAP and ESD are used as source languages, the others serve as target languages.

Dataset	Language	Valence	
		Negative	Positive
IEMOCAP [24]	English	2187	3344
ESD [26]	Chinese	7000	7000
EMO-DB [27]	German	385	150
EMOVO [28]	Italian	336	252
IEMOCAP.SUB	English	250	250
ESD.SUB	Chinese	250	250

Table 2: UAR results on the 80%-reduced imbalanced dataset.

	Angry	Sad	Happy	Average
NoAUG [9]	47.8	46.9	52.2	49.0
AUG [9]	53.5	52.1	55.2	53.6
NoAUG (Ours)	54.23	49.03	47.02	50.09
AUG (Ours)	62.75	58.11	53.71	58.19

Table 3: Ablation study results (UAR) of the contribution of \mathcal{L}_{VAR} and \mathcal{L}_{BAL} .

Model	Angry	Sad	Happy	Average
NoAUG	54.23	49.03	47.02	50.09
\mathcal{L}_{Model}	56.03	48.43	50.71	51.52
$\mathcal{L}_{Model} + \mathcal{L}_{VAR}$	36.47	46.61	45.31	42.80
\mathcal{L}_{Total}	62.75	58.11	53.71	58.19

data. The performance can be further improved with the augmented data, when the target language data is additionally given, based on the results between FT and FT_AUG.

3.2.3. Ablation Study

Models: For ablation study, we use the 80% reduced dataset to train an SER classifier with and without augmentation. To train the augmentor, we consider three cases: 1) with only \mathcal{L}_{Model} ; 2) with \mathcal{L}_{Model} and auxiliary variance loss \mathcal{L}_{VAR} ; 3) with the total loss that consists of both auxiliary terms \mathcal{L}_{VAR} and \mathcal{L}_{BAL} .

Results: The UAR results are listed in Tab. 3. We can observe that without \mathcal{L}_{VAR} and \mathcal{L}_{BAL} , performance can only improve little because the assistance of an augmentor with low augmentation variance is limited. After we apply the loss \mathcal{L}_{VAR} , as we mentioned in Sec. 2.2, The augmentor arbitrarily performs augmentation and destroys the emotion information, resulting in a significant fall in performance. Lastly, by introducing \mathcal{L}_{BAL} , we can obtain the best result by balancing augmentation variance and emotion preservation.

4. Conclusion

This work presents a GAN-based augmentation approach to alleviate the data imbalance and scarcity issue for the SER task. Specifically, we conduct experiments and demonstrate: 1) Even with a severely imbalanced dataset, our augmentation approach can significantly increase SER performance. 2) With only about 50 training utterances of the target languages provided, our augmentation approach considerably enhances SER performance for these low-resource languages.

5. References

- [1] S. Sahu, R. Gupta, and C. Y. Espy-Wilson, "On enhancing speech emotion recognition using generative adversarial networks," *ArXiv*, vol. abs/1806.06626, 2018.
- [2] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, "Generative adversarial nets," in *NIPS*, 2014.
- [3] F. Eyben, F. Weninger, F. Groß, and B. Schuller, "Recent developments in opensmile, the munich open-source multimedia feature extractor," in *MM '13*, 2013.
- [4] F. Bao, M. Neumann, and N. T. Vu, "CycleGAN-based emotion style transfer as data augmentation for speech emotion recognition," in *INTERSPEECH*, 2019.
- [5] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2242–2251, 2017.
- [6] G. Rizos, A. Baird, M. Elliott, and B. Schuller, "Stargan for emotional speech conversion: Validated by data augmentation of end-to-end emotion recognition," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3502–3506, 2020.
- [7] X. He, J. Chen, G. Rizos, and B. Schuller, "An improved stargan for emotional voice conversion: Enhancing voice quality and data augmentation," *ArXiv*, vol. abs/2107.08361, 2021.
- [8] Y. Choi, M.-J. Choi, M. S. Kim, J.-W. Ha, S. Kim, and J. Choo, "Stargan: Unified generative adversarial networks for multi-domain image-to-image translation," *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8789–8797, 2018.
- [9] A. Chatziagapi, G. Paraskevopoulos, D. Sgouropoulos, G. Pantazopoulos, M. Nikandrou, T. Giannakopoulos, A. Katsamanis, A. Potamianos, and S. S. Narayanan, "Data augmentation using gans for speech emotion recognition," in *INTERSPEECH*, 2019.
- [10] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, pp. 1345–1359, 2010.
- [11] H. Sagha, J. Deng, M. Gavryukova, J. Han, and B. Schuller, "Cross lingual speech emotion recognition using canonical correlation analysis on principal component subspace," *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5800–5804, 2016.
- [12] S. Latif, R. K. Rana, S. Younis, J. Qadir, and J. Epps, "Transfer learning for improving speech emotion classification accuracy," in *INTERSPEECH*, 2018.
- [13] S. Goel and H. S. M. Beigi, "Cross lingual cross corpus speech emotion recognition," *ArXiv*, vol. abs/2003.07996, 2020.
- [14] A. Tamkin, M. Wu, and N. D. Goodman, "Viewmaker networks: Learning views for unsupervised representation learning," *ArXiv*, vol. abs/2010.07432, 2021.
- [15] T. Chen, X. Zhai, M. Ritter, M. Lucic, and N. Houlsby, "Self-supervised gans via auxiliary rotation loss," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12 146–12 155, 2019.
- [16] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6199–6203, 2020.
- [17] J. Jeong and J. Shin, "Training gans with stronger augmentations via contrastive discriminator," *ArXiv*, vol. abs/2103.09742, 2021.
- [18] J. Harvill, M. Abdelwahab, R. Lotfian, and C. Busso, "Retrieving speech samples with similar emotional content using a triplet loss function," *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7400–7404, 2019.
- [19] J. Huang, Y. Li, J. Tao, and Z. Lian, "Speech emotion recognition from variable-length inputs with triplet loss function," in *INTER-SPEECH*, 2018.
- [20] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [21] 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.
- [22] Z. Zhang, B. Wu, and B. Schuller, "Attention-augmented end-to-end multi-task learning for emotion prediction from speech," *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6705–6709, 2019.
- [23] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2015.
- [24] C. Busso, M. Bulut, C.-C. Lee, E. A. Kazemzadeh, E. M. Provost, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: interactive emotional dyadic motion capture database," *Language Resources and Evaluation*, vol. 42, pp. 335–359, 2008.
- [25] S. Latif, M. Asim, R. K. Rana, S. Khalifa, R. Jurdak, and B. W. Schuller, "Augmenting generative adversarial networks for speech emotion recognition," *ArXiv*, vol. abs/2005.08447, 2020.
- [26] K. Zhou, B. Sisman, R. Liu, and H. Li, "Emotional voice conversion: Theory, databases and esd," *ArXiv*, vol. abs/2105.14762, 2021.
- [27] F. Burkhardt, A. Paeschke, M. Rolfes, W. F. Sendlmeier, and B. Weiss, "A database of german emotional speech," in *INTER-SPEECH*, 2005.
- [28] G. Costantini, I. Iaderola, A. Paoloni, and M. Todisco, "Emovo corpus: an italian emotional speech database," in *LREC*, 2014.