

Speech Emotion Recognition using Data Augmentation Method by Cycle-Generative Adversarial Networks

Arash Shilandari · Hossein Marvi · Hossein Khosravi · Wenwu Wang

Received: date / Accepted: date

Abstract One of the obstacles in developing speech emotion recognition (SER) systems is the data scarcity problem, i.e., the lack of labeled data for training these systems. Data augmentation is an effective method for increasing the amount of training data. In this paper, we propose a cycle generative adversarial network (Cycle-GAN) for data augmentation in the SER systems. For each of the five emotions considered, an adversarial network is designed to generate data that has a similar distribution to the main data in that class but has a different distribution to those of other classes. These networks are trained in an adversarial way to produce feature vectors similar to those in the training set, which are then added to the original training sets. Instead of using the common cross-entropy loss to train Cycle-GANs, we use the Wasserstein divergence to mitigate the gradient vanishing problem and to generate high-quality samples. The proposed network has been applied to SER using the EMO-DB dataset. The quality of the generated data is evaluated using two classifiers based on Support Vector Machine (SVM)

and Deep Neural Network (DNN). The results showed that the recognition accuracy in un-weighted average recall (UAR) was about 83.33%, which is better than the baseline methods compared.

Keywords Speech processing · data augmentation · speech emotion recognition · generative adversarial networks

1 Introduction

The data scarcity problem is one of the critical challenges in developing speech emotion recognition (SER) systems. This problem can be examined from three aspects. The first aspect is related to the dramatized emotions (generated by actors), used to avoid legal and moral issues [1]. The second aspect is the mislabelling of the emotions of the speakers, and the third issue is related to the lack of balance in the number of samples available for each class. To train an emotion classifier, a balanced data set (equal number of emotional samples in each class) is often required.

Some standard data augmentation techniques used for images such as transfer and rotation [2] may not be applicable for text or speech. Synonymous substitution [3], which is mainly used to process text, is not appropriate for emotion classification and recognition from speech. Similarly, traditional data augmentation methods for speech, such as changes in voice and sound velocity [4], are also inappropriate for images or texts. In contrast, the data augmentation method based on generative adversarial networks (GANs) [12] is focused on learning and simulating real data distribution and is independent of the tasks. Therefore, GANs-based data augmentation method is our focus in this paper.

A. Shilandari
Faculty of Electrical Engineering, Shahrood University of Technology, Shahrood, Iran
E-mail: Shilandari@shahroodut.ac.ir

H. Marvi
Faculty of Electrical Engineering, Shahrood University of Technology, Shahrood, Iran
E-mail: h.marvi@shahroodut.ac.ir

H. Khosravi
Faculty of Electrical Engineering, Shahrood University of Technology, Shahrood, Iran
E-mail: hosseinkhosravi@shahroodut.ac.ir

W. Wang
Department of Electrical and Electronic Engineering, University of Surrey, Guildford, United Kingdom
E-mail: w.wang@surrey.ac.uk

Recently, end-to-end methods are used for speech emotion recognition [5, 6], where the input to the systems are feature vectors and the output is class labels. In [7], the features are extracted by convolution filters. With the development of DNNs in SER, various data augmentation methods have been proposed [8, 9]. Transfer learning can be used to address the data sparsity problem [10], e.g. in image and speech processing [11]. Deng et al. proposed a transfer learning method by transferring knowledge learned from source domain data to the target domain data [8].

One of the effective methods to augment data is the GANs introduced by Goodfellow et al. [12], which was shown to improve image recognition performance [13]. Zhang et al. introduced GAN to produce high-dimensional data and showed that data augmentation by GANs performs better than the typical data augmentation techniques [14], such as time warping, frequency masking, and time masking. Hybrid methods include four different combinations: LibriSpeech basic (LB), LibriSpeech doubles (LD), Switchboard mild (SM), and Switchboard strong (SS) [38].

Cycle adversarial data augmentation networks use Jensen-Shannon divergence as a divergence criterion. According to [15], if two data distributions are less overlapped or not overlapped, Jensen-Shannon divergence tends to be constant, which can lead to a gradient vanishing problem. The method proposed in this study can address this problem. In training, source and target data distribution are significantly overlapped, which makes it difficult for the discriminator to distinguish between these two vector groups. As a result, the discriminator network leads to increased cross-entropy errors, and the generator network then receives a gradient error. Moreover, with the adversarial data augmentation networks, other divergence methods such as the Wasserstein divergence can be easily used for gradient descent. As compared with the Jensen-Shannon divergence, the Wasserstein divergence can measure the distance between two data distributions even if they are not overlapped. The hidden space generated by adversarial data augmentation networks also makes it easy to learn emotional information due to the low dimensions of the vectors in the training data. In addition, practical programs [16, 17] have shown that models with the Wasserstein divergence are better than those with other divergences, such as Jensen-Shannon divergence and maximum mean discrepancy. Therefore, the Wasserstein divergence based adversarial data augmentation may offer improved performance in emotion recognition.

In this paper, we present a Cycle-GAN for data augmentation and then test it on SER with two classifier networks. The Cycle-GAN generates samples similar to actual data thereby augmenting the dataset with additional samples for emotion classification. In addition, we study the effectiveness of the GANs, and replace the standard cross-entropy error by the Wasserstein divergence to train the GAN to improve the classification performance. We evaluate the method using the EMO-DB database. The results show that the proposed data augmentation technique improves the SER performance on the EMO-DB dataset, and the Cycle-GAN with Wasserstein divergence outperforms the Cycle-GAN with the conventional cross-entropy loss. We show that the synthetic samples generated from an ordinary Cycle-GAN cover part of the actual data while the clusters created by the Cycle-GAN using Wasserstein distance (artificial samples generated from our method) completely cover the feature space for each five emotion classes.

Section 2 reviews existing methods for the data scarcity problem. Section 3 proposes the suggested network design and provides theoretical analysis. Section 4 presents experiment details, including dataset, features, experimental setup, and evaluation protocols. Section 5 analyzes experimental results. Finally, Section 6 concludes the paper.

2 Background

2.1 Related work

To address the data scarcity problem, we can use data augmentation methods to expand the data set by generating new samples using techniques, such as adding noise to the data [32], pitch shifting and time-stretching of the audio signal, varying the loudness of the speech signal, applying random frequency filters, and interpolating between samples in input space. However, these methods usually change the data, and may cause problems or introduce artefacts into the data, such as rotation, adding noise, speech echoing, and signal clipping [18]. Advanced data augmentation methods are based on GANs and their variants, such as Conditional-GANs and Cycle-GANs. Hu et al. used a deep convolutional neural network to produce extra features to train acoustic models and showed that data augmentation can improve the performance of speech recognition systems [19]. Sahu et al. synthesized feature vectors with automatic adversarial encoders using Gaussian mixed noise in the generator network [20]. Sahu et al. also developed a model based on a Conditional-GANs to generate artificial feature vectors [9]. Several methods were used to train the Conditional-

GANs, including generator initialization with detector weights, as well as, using an automatic adversarial encoder.

One fundamental issue in training GANs is that the generator and the discriminator are trained in parallel. Dynamic alternating training [14] can be used so that the number of training epochs in the generator network and the discriminator network do not have to match. This is because the ultimate goal is not about the number of training epochs, but the amount of training in each network.

2.2 Generative Adversarial Networks

As mentioned before, GANs consist of two deep neural networks. The generator network produces synthetic data, and the discriminator network distinguishes the real data from the synthetic data. The loss function of GANs can be expressed as follows [21]:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - DG(z))] \quad (1)$$

where D is a discriminator, G is a generator, Z is noise, $p_{\text{data}}(x)$ is the original data distribution, and $p_z(z)$ is the input noise distribution. In practice, according to [22], we train G to maximize $\log D(x)$, instead of training G to minimize $\log(1 - DG(z))$. This objective function can mitigate the vanishing gradient problem without compromising the equilibrium point of G and D .

$$J^{(D)}(D, G) = - \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] - \mathbb{E}_{z \sim p_z(z)} [\log(1 - DG(z))] \quad (2)$$

$$J^{(G)}(G) = - \mathbb{E}_{z \sim p_z(z)} [\log D(G(z))] \quad (3)$$

Figure 1 shows the network architecture designed. The entire process of training a GAN is shown in Algorithm 1.




Fig. 1 Diagram of the proposed SER system with Cycle-GAN data augmentation.

Algorithm 1 Training a GAN in the vanishing gradient method

Repeat for the number of training epochs:
While the stopping criterion is not met **do**:
 for each k **do**:
 Sample m data points with distribution $p_z(z)$.
 $z = \{z^{(1)}, z^{(2)}, z^{(3)}, \dots, z^{(m)}\}$
 Sample m data points with initial distribution p .
 $x = \{x^{(1)}, x^{(2)}, x^{(3)}, \dots, x^{(m)}\}$
 Calculate the loss of the discriminator network:
 $\nabla_{\theta_d} \frac{1}{m} \sum_i [\log D(x^{(i)}) + \log(1 - D(G(z^{(i)})))]$
 end for
 for each k **do**:
 Sample m data points from the initial noise space $pg(z)$.
 $z = \{z^{(1)}, z^{(2)}, z^{(3)}, \dots, z^{(m)}\}$
 Update the generator weights with the gradient descent method:
 $\nabla_{\theta_g} \frac{1}{m} \sum_i [\log(1 - D(G(z^{(i)})))]$
 end for
 end while

2.3 Cycle Generative Adversarial Networks

Cycle-GANs have been used for image generation for non-paired data [23]. Figure 2 shows the architecture of Cycle-GAN for data augmentation [25]. This network includes two transfer functions, F and G , where G learns how to transfer samples from the source domain S to target domain T , and F is the inverse of G . Both F and G may be considered as generators to produce the target and source data, respectively. Moreover, there are two adversarial discriminator networks, D^T and D^S , where D^T discriminates real targets from the synthetic targets, while D^S discriminates real sources from synthetic sources. This network sets its target so that $F(G(S)) \approx S$ and $G(F(T)) \approx T$. Therefore, it is called Cycle-GAN [24].

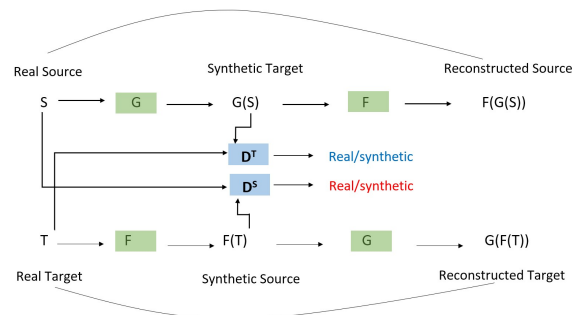


Fig. 2 The structure of Cycle-GANs.

The loss used in Cycle-GANs includes the adversarial loss and the cycle consistency loss. Removing adversity may be transformed into a part of target data generation and a part of source data generation. The loss function for target data generation is as follows [24]:

$$L^{GAN}(G, D^T, S, T) = E_{t \sim p_t} [\log D^T(t)] + E_{s \sim p_s} [\log(1 - D^T(G(s)))] \quad (4)$$

Losses are expressed as value functions. In the generation process, the objective is $\min_G \max_{D^T} L^{GAN}(G, D^T, S, T)$, and to reproduce real data, the objective is $\min_F \max_{D^S} L^{GAN}(F, D^S, T, S)$. Zou *et al.* have defined the cycle loss as follows [24]:

$$L^{cyc}(G, F) = E_{t \sim p_t} [\|G(F(t)) - t\|_1] + E_{s \sim p_s} [\|F(G(s)) - s\|_1] \quad (5)$$

where the L_1 norm may be substituted with other criteria in these losses. The total losses for Cycle-GANs are as follows:

$$L(G, F, D^T, D^S) = L^{GAN}(G, D^T, S, T) + L^{GAN}(F, D^S, T, S) + \lambda L^{cyc}(G, F) \quad (6)$$

where λ controls the relative importance of both losses [24]. We conducted an ablation study to analyze the impact of the proposed regularization term L^{cyc} by varying the corresponding weight λ using the EMO-DB dataset, and observed that increasing λ improves both the quality and diversity of the generated samples. Nevertheless, as the weighting parameter λ becomes larger than a threshold value e.g. 1.0, the training becomes unstable, which results in low quality, and even low diversity of synthesized samples. As a result, we empirically set the weighting parameter $\lambda = 1.0$ for all the experiments.

3 Methodology

3.1 The Proposed Method

For a labelled dataset X with N emotional classes, artificial samples for each emotion are generated using a separate Cycle-GAN. According to Figure 3, Cycle-GAN transfers between source domain S and target domain T_i , where S is a dataset without labels and T_i is emotional sample i in the labeled dataset X . Discriminator networks D_i^T and D_i^S are used to identify the artificial target which cannot be distinguished from real samples. The generator loss and discriminator loss are introduced by $L_i^{GAN}(G_i, D_i^T, S, T_i)$ and $L_i^{GAN}(F_i, D_i^S, S, T_i)$, respectively. We have

$$L_i^{GAN}(G_i, F_i, D_i^T, D_i^S, S, T_i) = L_i^{GAN}(G_i, D_i^T, S, T_i) + L_i^{GAN}(F_i, D_i^S, S, T_i) \quad (7)$$

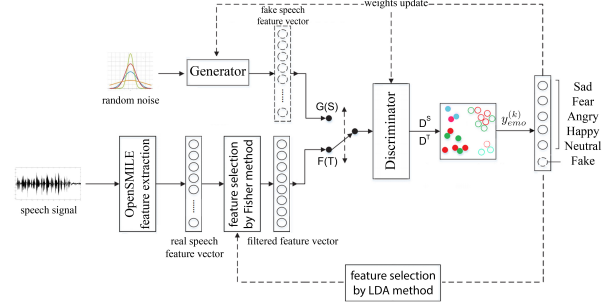


Fig. 3 Detailed architecture of the proposed SER system.

The cycle loss function can have an impact on the number training epochs in the Cycle-GANs. Therefore, we translate the synthetic target $G_i(S)$ back to the source domain and compute the mean squared error (MSE) between the real source S and reconstructed data $F_i G_i(S)$. This is similarly done for T_i and reconstructed target data $G_i(S)$. As a result, the total loss function in each cycle will be as follows:

$$L_i^{cyc}(G_i, F_i, S, T_i) = E_{s \sim p_s} [\|F_i(G_i(s)) - s\|_2^2] + E_{t \sim p_t} [\|G_i(F_i(t)) - t\|_2^2] \quad (8)$$

3.2 Overcoming Gradient Vanishing Problem in Training Cycle-GANs

To overcome the gradient vanishing problem in Cycle-GANs, we suggest using the Wasserstein distance. In extreme case, the gradient descent may be stopped during the process of weight modification and the training of generators and discriminators. Considering two probability distributions P_r and P_g , the Wasserstein distance is defined as follows:

$$W_1(P_r, P_g) = \sup_{\|f\|_L \leq 1} E_{x \sim P_r} \{f(x)\} - E_{\tilde{x} \sim P_g} \{f(\tilde{x})\} \quad (9)$$

where $\|f\|_L \leq 1$ shows that f satisfies the 1-Lipschitz limitation [26]. It is worth mentioning that W_1 is invariant up to a positive scalar K if the Lipschitz constraint is modified to be K . W_1 is believed to be more suitable for data distributed on low-dimensional manifolds. If the weights are greater or smaller than the expected limit, they will be changed into minimum or maximum predefined. In the gradient penalty method, the gradient penalty is based on Lipschitz, which is derived from the fact that if gradients are at most 1 everywhere, they are 1-Lipschitz functions. Their squared difference from one is used as a gradient penalty. According to [27], weight clipping may lead to a non-optimal solution. Gradient penalty was also applied to overcome weight clipping limitations [16]. However,

if there is a data sparsity problem, satisfying the K -Lipschitz condition is difficult for the whole data set. Accordingly, Wu *et al.* [27] suggested a new Wasserstein divergence, where the Wasserstein distance is calculated without applying Lipschitz condition:

$$L_D = \mathbb{E}_{x \sim P_r} \{f(x)\} - \mathbb{E}_{\tilde{x} \sim P_g} \{f(\tilde{x})\} + \lambda \mathbb{E}_{\tilde{x} \sim P_u} [\|\nabla f(\tilde{x})\|^p] \quad (10)$$

where λ controls the effect of gradient modification on the target function and $\lambda > 0$. P_u is a Radon distribution, and p is related to L_p space for function f and $p > 1$, [27]. Finally, the loss function in the generator and the discriminator is written as follows:

$$L_G^{(WC-GAN)} = \mathbb{E}_{p(x,y,z)} \{D(G(z,y))\} \left\{ -a \sum_{k=1}^K y_{emo}^{(k)} \log C((G(z,y))_k) \right\} \quad (11)$$

$$L_D^{(WC-GAN)} = \mathbb{E}_{p(x,z,\tilde{x},y)} \{D(E(x) - D(G(z,y)))\} \{+\lambda \|\nabla \tilde{x}\|^p\} \quad (12)$$

where $()_k$ denotes the k th element of a vector, C stands for the auxiliary classifier, \tilde{x} is a reconstructed sample of the source, y_{emo} is the output of emotion classifier, and a determines the contributions of the classification error to the loss in the generator. The structure of the Cycle-GAN with the Wasserstein distance is shown in Figure 2.

3.3 Recognizing samples generated by Cycle-GAN Augmentation Network

Figure 4 shows that transferring data by Cycle-GAN results in similarity between real and artificial data distribution. A classification loss function is used to ensure that the synthetic data can be correctly allocated to their target emotion class, which is defined here as the cross-entropy error:

$$L^{cls} = - \sum_i y_i \log(C(G_i(S))) \quad (13)$$

where y_i is the label of the target emotions. The total loss is defined as follows:

$$L = \sum_i L_i^{GAN} + \lambda^{cyc} \sum_i L_i^{cyc} + \lambda^{cls} L^{cls} \quad (14)$$

where λ^{cyc} and λ^{cls} are the weights corresponding to the Cycle-GAN loss and the classification loss.




Fig. 4 The difference between two mapping samples without the classification loss and with classification loss.

4 Experiments

4.1 Dataset

We performed experiments on the EMO-DB dataset [28], which is a small dataset of 535 training clips with seven emotional classes. All speech signals were recorded by ten professional actors in German. This database includes seven emotions. We used five emotions to perform the experiments: anger (127 samples), fear (69 samples), happiness (71 samples), sadness (62 samples), and neutral (79 samples). We did not use disgust (81 samples) and surprise (46 samples). The data was recorded at a 48 kHz sampling rate and then down-sampled to 16 kHz. The average length of each audio clip is 2.8 seconds.

The other datasets that are often used in speech emotion recognition include the Danish Emotional Speech Database (DES) with 200 samples, the Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) with 2496 samples, the Interactive Emotional Dyadic Motion Capture Database (IEMOCAP) with 1150 samples, and the Vera am Mittag Database (VAM) with 1018 samples. There are also audio-visual datasets such as SEWA [37] and MuSe-CAR [36] that are not discussed in this article because we only focus on emotion recognition from speech data. As it turns out, all of these databases suffer from data shortages due to the lack of data samples and are not suitable for deep neural network training. As a suitable solution, we suggest creating a synthetic dataset using GANs trained by available datasets. We chose EMO-DB for the reason that this dataset contains less training samples as compared to the remaining datasets. Another reason for our choice was to compare the results of our proposed data augmentation method with other data augmentation methods.

4.2 Experimental Setup

It is challenging to train the generators with high-dimensional feature vectors. To address this issue, we pre-trained both G_i and F_i generators based on the reconstruction error between S and $F_i G_i(S)$ and also the reconstruction error between T_i and $G_i(F_i(T_i))$.

We used the OpenSMILE software to extract the features, and then used the proposed method to generate new feature vectors to increase the number of training samples and to balance the number of samples in the dataset. The dimension of the feature vector is 2185 for each training sample.

DNN with two hidden layers and 800 hidden neurons were used in the proposed Cycle-GANs. We used ResNet for the generator network and PatchGAN for the discriminator network. In addition, DNN and SVM networks were used as classifiers, and Leaky ReLU was applied to all the layers. The linear kernel is used in the SVM classifier. We also used the Xavier Algorithm [29] and the Adam Optimizer [30] with a learning rate of 0.0002 which was reduced every 50 epochs. DNNs were implemented using Tensorflow (V2.1) in Python, while SVMs were implemented using Scikit-Learn Package.

To balance the training of G and D , the generator weights were updated two times per epoch, and the discriminator weights were updated one time per epoch. Moreover, unilateral label smoothing [31] was used to reduce the vulnerability of neural networks to adversarial examples, i.e. by replacing the binary output values 0 and 1 of the classifier with smoothed values, e.g. 0.1 and 0.9.

5 Results

The augmented data were gradually and randomly added to the original data, and two DNN and SVM classifiers were used for SER. The L_2 regulation was used to train deep neural networks, and each experiment was repeated three times, and the mean absolute accuracy was reported as the performance measure. Figure 5 shows the UAR results of the SVM and DNN on the EMO-DB dataset.

We compared the performance of the proposed method with those of the standard data augmentation techniques, such as sample reproduction, adding random noise to feature vectors and artificial sampling (SMOTE) [33]. The performance of data augmentation via adding noise depends on the amount of noise, and the results may not be stable, as shown in Figure 6. Generating synthetic data similar to the primary samples helps deep neural networks learn data distribution better, however, repetitive samples will not lead to better network training. The SMOTE method




Fig. 5 Comparing classification results with real samples.




Fig. 6 Classification results with data augmented by Gaussian noise.

is designed to augment samples in one class and cannot be used to augment samples in all classes but it has a relatively stable performance [33]. Figure 7 shows the results of this method.




Fig. 7 Classification results with data augmented by the SMOTE method.

The Cycle-GANs based data augmentation method could also lead to the improvement of SVM performance. Figure 8 shows the performance of two classifiers by combing real and augmented data based on a Cycle-GAN. The results show that augmenting data helps SVM recognize metadata in feature space and classify them with better performance. According to Figure 9, it is possible to improve the performance of the data augmentation approach based on the Wasserstein distance introduced in Section 3. The unweighted average recall is gradually augmented by




Fig. 8 Classification results with data augmented by Cycle-GAN.

adding artificial samples to the training set. These results show that data augmentation based on Cycle-GANs may generate new and meaningful emotional vectors which help improve the performance of the emotion classifier. Figure 10 shows the clusters




Fig. 9 Classification results with data augmented by Cycle-GANs with the Wasserstein distance.

created by the Cycle-GAN using the Wasserstein distance for the five emotional classes. In Figure 10(b), artificial samples generated from the proposed method completely cover the feature space for each emotion class, while the samples generated from an ordinary Cycle-GAN in Figure 10(a) cover only part of feature space of the actual data.




Fig. 10 Data distribution of each class: (a) samples generated by Cycle-GAN, (b) samples generated by Cycle-GAN with the Wasserstein distance.

We compared our method with the methods in [32–34] in Table 1. This table shows that with the proposed

method, the classifier can be better trained and our method outperforms [35] with the handcrafted features. Our method also outperforms Chen *et al.* [34], who used 3D-ACRNNs to extract features.

Table 1 Different data augmentation and SER techniques

Method	Classifier	WA%	UAR%
Add noise [32]	DNN	82.06	82.73
Add noise [32]	SVM	81.12	80.25
SMOTE [33]	DNN	82.43	81.51
SMOTE [33]	SVM	80.93	79.51
Cycle-GAN	DNN	83.55	82.50
Cycle-GAN	SVM	81.50	80.30
Cycle-GAN + Wasserstein	DNN	84.49	83.33
Cycle-GAN + Wasserstein	SVM	81.07	80.08
2D-ACRNN [34]	DNN	–	79.38
3D-ACRNN [34]	DNN	–	82.82

6 Conclusion

We have presented a method for generating synthetic samples based on Cycle-GAN to mitigate data scarcity and to improve speech emotion classification. We generated artificial data in the space of each emotional class that completely covers the leading data space. We showed that using the Wasserstein divergence can overcome the vanishing gradient problem during the training process. The results show that including synthetic samples in the real samples can improve the emotion recognition performance to as high as 83.33% in terms of UAR on the EMO-DB dataset. As explained, we only dealt with the case where the features were extracted by OpenSMILE software and this can be extended by providing raw data to the network.

References

1. M. El Ayadi, M. S. Kamel, F. Karray.: Survey on speech emotion recognition: Features, classification schemes, and databases. *Pattern Recognition*. 44.3, pp. 572–587 (2011).
2. J. Wang, L. Perez.: The effectiveness of data augmentation in image classification using deep learning. *Computer Vision and Pattern Recognition* (2017).
3. X. Zhang, Y. LeCun.: Text understanding from scratch. (2015) arXiv preprint arXiv:1502.01710 .
4. T. Ko, V. Peddinti, D. Povey, S. Khudanpur.: Audio augmentation for speech recognition. The Sixteenth annual conference of the international speech communication association, September 6–10, (2015), Dresden, Germany.
5. G. Trigeorgis, F. Ringeval, R. Brueckner, E. Marchi, M. A. Nicolaou, B. Schuller, S. Zafeiriou.: Adieu features? End-to-end speech emotion recognition using a deep

- convolutional recurrent network. Proc. IEEE Int. Conf. Acoust., Speech Signal Process (ICASSP), Mar. (2016), pp. 5200–5204.
6. X. Ma, Z. Wu, J. Jia, M. Xu, H. Meng, and L. Cai.: Emotion recognition from variable-length speech segments using deep learning on spectrograms. Proc. Interspeech, pp. 3683–3687, Sep. (2018).
 7. P. Li, Y. Song, I. McLoughlin, W. Guo, and L. Dai.: An attention pooling-based representation learning method for speech emotion recognition. Proc. Interspeech, pp. 3087–3091, Sep. (2018).
 8. J. Deng, Z. Zhang, E. Marchi, and B. Schuller.: Sparse autoencoder based feature transfer learning for speech emotion recognition. Proc. Humaine Assoc. Conf. Affect. Comput. Intell. Interact., Sep. (2013), pp. 511–516.
 9. S. Sahu, R. Gupta, and C. Espy-Wilson.: On enhancing speech emotion recognition using generative adversarial networks. (2018) arXiv:1806.06626. [Online]. Available: <http://arxiv.org/abs/1806.06626>
 10. S. J. Pan and Q. Yang.: A survey on transfer learning. IEEE Trans. Knowl. Data Eng., vol. 22, no. 10, pp. 1345–1359, Oct. (2010).
 11. M. Wang and W. Deng.: Deep visual domain adaptation: A survey. Neurocomputing, vol. 312, pp. 135–153, Oct. (2018).
 12. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio.: Generative adversarial nets. Advances in Neural Information Processing Systems, Jun. (2014).
 13. A. Antoniou, A. Storkey, and H. Edwards.: Data augmentation generative adversarial networks. (2017) arXiv:1711.04340. [Online]. Available: <http://arxiv.org/abs/1711.04340>
 14. Z. Zhang, J. Han, K. Qian, C. Jannett, Y. Guo, and B. Schuller.: Snore- GANs: Improving automatic snore sound classification with synthesized data. IEEE J. Biomed. Health Informat., vol. 24, no. 1, pp. 300–310, Jan. (2020).
 15. M. Arjovsky, S. Chintala, and L. Bottou.: Wasserstein generative adversarial networks. Proc. Int. Conf. Mach. Learn., Aug. (2017), pp. 214–223.
 16. I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville.: Improved training of Wasserstein GANs. Proc. Adv. Neural Inf. Process. Syst. pp. 5767–5777 (2017).
 17. J. Shen, Y. Qu, W. Zhang, and Y. Yu.: Wasserstein distance guided representation learning for domain adaptation. Proc. AAAI Conf. Artif. Intell. (2018), pp. 4058–4065.
 18. T. DeVries and G. W. Taylor.: Dataset augmentation in feature space. (2017) arXiv:1702.05538. [Online]. Available: <https://arxiv.org/abs/1702.05538>
 19. H. Hu, T. Tan, and Y. Qian.: Generative adversarial network-based data augmentation for noise-robust speech recognition. Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP), pp. 5044–5048 Apr. (2018).
 20. S. Sahu, R. Gupta, G. Sivaraman, W. Abdalmageed, and C. Espy-Wilson.: Adversarial auto-encoders for speech-based emotion recognition. Proc. Interspeech, pp. 1243–1247 Aug. (2017).
 21. N. Hajarolasvadi, E. Bashirov, and H. Demirel.: Video-based person-dependent and person-independent facial emotion recognition. Signal, Image and Video Processing, Volume 15 Issue 5, pp. 1049–1056 (2021).
 22. A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey.: Adversarial autoencoders. 4th International Conference on Learning Representations (ICLR), May 2–4, (2016), Puerto Rico.
 23. M. El Ayadi, M. S. Kamel, F. Karray.: Survey on speech emotion recognition: Features, classification schemes, and data-bases. Pattern Recognition 44.3 , pp. 572–587 (2011).
 24. J.-Y. Zhu, T. Park, P. Isola, A. A. Efros.: Unpaired image-to-image translation using cycle-consistent adversarial networks. Proceedings of the IEEE International Conference on Computer Vision (2017), pp. 2223–2232.
 25. F. Bao, M. Neumann, N. T. Vu.: CycleGAN-based Emotion Style Transfer as Data Augmentation for Speech Emotion Recognition. Interspeech, (2019).
 26. M. Arjovsky, S. Chintala, L. Bottou.: Wasserstein generative adversarial networks. ICML, (2017).
 27. J. Wu, Z. Huang, J. Thoma, D. Acharya, and L. Van Gool.: Wasserstein divergence for GANs. Proc. Eur. Conf. Comput. Vis. (ECCV), Sep. (2018), pp. 653–668.
 28. F. Burkhardt, A. Paeschke, M. Rolfes, W. F. Sendlmeier, and B. Weiss.: A database of German emotional speech. Proc. 9th Eur. Conf. Speech Commun. Technol (2005), pp. 1–4.
 29. X. Glorot and Y. Bengio.: Understanding the difficulty of training deep feedforward neural networks. Proc. 13th Int. Conf. Artif. Intell. Statist. (2010), pp. 249–256.
 30. D. P. Kingma and J. L. Ba.: Adam: A method for stochastic optimization. Proc. 3rd Int. Conf. Learn. Represent. (ICLR), (2015), pp.1–15.
 31. T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen.: Improved techniques for training GANs. 13th Conference on Neural Information Processing Systems, December 5–10, (2016), Barcelona, Spain.
 32. U. Tiwari, M. Soni, A. Panda, R. Chakraborty, and S. Kumar Kopparapu.: Multi-Conditioning and Data Augmentation using Generative Noise Model for Speech Emotion Recognition in Noisy Conditions. Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 4–8 May (2020).
 33. N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer.: SMOTE: Synthetic minority over-sampling technique. J. Artif. Intell. Res., vol. 16, pp. 321–357, Jun. (2002).
 34. M. Chen, X. He, J. Yang, and H. Zhang.: 3-D convolutional recurrent neural networks with attention model for speech emotion recognition. IEEE Signal Process. Lett., vol. 25, no. 10, pp. 1440–1444, Oct. (2018).
 35. I. Luengo, E. Navas, and I. Hernaez.: Feature analysis and evaluation for automatic emotion identification in speech. IEEE Trans. Multimedia, vol. 12, no. 6, pp. 490–501, Oct. (2010).
 36. L. Stappen, A. Baird, L. Schumann, and B. Schuller.: The multimodal sentiment analysis in car reviews (MuSe-CaR) dataset: collection, insights and improvements. IEEE Trans. on Affective Computing (EARLY ACCESS), Oct. (2021).
 37. J. Kossaifi, R. Walecki, Y. Panagakis, J. Shen, M. Schmitt, F. Ringeval, J. Han, V. Pandit, A. Toisoul, B. Schuller, K. Star, E. Hajjiev, and M. Pantic.: SEWA DB: a rich database for audio-visual emotion and sentiment research in the wild. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 43, no. 3, pp. 1022–1040, Mar. (2021).
 38. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, D. Cubuk, and Q. V. Le.: SpecAugment: A simple data augmentation method for automatic speech recognition. Proc. Interspeech, pp. 2613–2617, (2019).