



Accurate Emotion Strength Assessment for Seen and Unseen Speech Based on Data-Driven Deep Learning

Rui Liu^{1,3,4}, Berrak Sisman⁴, Björn W. Schuller⁵, Guanglai Gao¹, Haizhou Li^{2,3}

¹ Inner Mongolia University, China ² The Chinese University of Hong Kong, Shenzhen, China

³ National University of Singapore ⁴ Singapore University of Technology and Design

⁵ Imperial College London, United Kingdom

liurui_imu@163.com, berraksisman@u.nus.edu, bjoern.schuller@imperial.ac.uk,
csggl@imu.edu.cn, haizhouli@cuhk.edu.cn

Abstract

Emotion classification of speech and assessment of the emotion strength are required in applications such as emotional text-to-speech and voice conversion. The emotion attribute ranking function based on Support Vector Machine (SVM) was proposed to predict emotion strength for emotional speech corpus. However, the trained ranking function doesn't generalize to new domains, which limits the scope of applications, especially for out-of-domain or unseen speech. In this paper, we propose a data-driven deep learning model, i.e. StrengthNet, to improve the generalization of emotion strength assessment for seen and unseen speech. This is achieved by the fusion of emotional data from various domains. We follow a multi-task learning network architecture that includes an acoustic encoder, a strength predictor, and an auxiliary emotion predictor. Experiments show that the predicted emotion strength of the proposed StrengthNet is highly correlated with ground truth scores for both seen and unseen speech. We release the source codes at: <https://github.com/ttslr/StrengthNet>.

Index Terms: Emotion strength, deep learning, data-driven

1. Introduction

Accurate emotion classification of speech and assessment of its strength are essential to profile human behaviors, which has many potential applications, such as human-robot interface, human-machine dialogue, and social media. In recent years, there is an increasing interest in emotion control in expressive speech synthesis, such as emotional text-to-speech, emotional voice conversion, where accurate control of emotional strength in speech becomes critically important.

The research by Rui Liu is funded by the High-level Talents Introduction Project of Inner Mongolia University (No. 10000-22311201/002), with Rui Liu as the Principal Investigator. The research by Rui Liu was funded by the High-level Talents Introduction Project of Inner Mongolia University (No. 10000-22311201/002), with Rui Liu as the Principal Investigator. The research by Haizhou Li is supported by the Guangdong Provincial Key Laboratory of Big Data Computing under the Grant No. B10120210117-KP02, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen) and the Science and Engineering Research Council, Agency of Science, Technology and Research (A*STAR), Singapore, through the National Robotics Program under Human Robot Interaction Phase 1 (Grant No. 192 25 00054); Human Robot Collaborative AI under its AME Programmatic Funding Scheme (Project No. A18A2b0046); The research by Berrak Sisman is supported by the Ministry of Education, Singapore, under its MOE Tier 2 funding programme, award no: MOE-T2EP50220-0021, SUTD Startup Grant Artificial Intelligence for Human Voice Conversion (SRG ISTD 2020 158) and SUTD AI Grant – Thrust 2 Discovery by AI (SGPAIRS1821).

The simplest emotion strength control method is to linearly scale the emotion representation vector [1]. The effect of such a linear scale is hardly interpretable. To obtain a meaningful strength descriptor, some followed the idea of “relative attributes” [2–4] and quantify the emotion strength by learning from the <neutral, emotional> speech pairs. Support-Vector-Machine (SVM) based attribute ranking [4] learns the difference between two samples that are significantly different in a particular attribute, that has been widely studied in computer vision [5, 6].

In speech processing, Zhu et al. [2] proposed to learn an emotion attribute ranking function $R(\cdot)$ from the <neutral, emotional> paired speech features, then weight the emotional feature with a learnable weighting vector and return a weighted sum as the indicator of the emotion strength of one specific emotional speech. Lei et al. [3] further extended the utterance level emotion attribute ranking function to phoneme level and obtain a fine-grained ranking function. In this way, we can obtain a meaningful strength score for emotional speech on a specific dataset, which correlates with human perception.

We note that a trained ranking function $R(\cdot)$ on specific data is not easily generalized to new domains. In other words, $R(\cdot)$ is not able to calculate an accurate or appropriate strength score for unseen or out-of-domain speech. To extend to a new data, we need to retrain a new ranking function on the <neutral, emotional> paired samples from the new data. Furthermore, the learning of ranking functions for new data requires parallel samples. All these limit the scope of applications. Recently, it was shown that deep learning has the ability to learn a mapping function effectively [7, 8]. The neural solution learns complex non-linear mapping relationships, and exhibits good generalization ability with the support of a large number of model parameters [9]. Most importantly, the data-driven training strategy appears to be more powerful and have the potential to achieve good performance for out-of-domain data [10].

In this paper, we propose a novel neural solution to emotion strength assessment, termed StrengthNet. To improve the model generalization, we employ a data-driven strategy, named “domain fusion”, to mix emotional data from various domains for model training. StrengthNet is a multi-task framework that includes a convolutional neural network (CNN) based acoustic encoder, a bidirectional long short-term memory (BiLSTM) based strength predictor and an auxiliary BiLSTM based emotion predictor. The acoustic encoder extracts the high-level features from the input mel-spectrum. The strength predictor aims to predict the strength score for the input mel-spectrum. The emotion predictor is used to predict the emotion category,

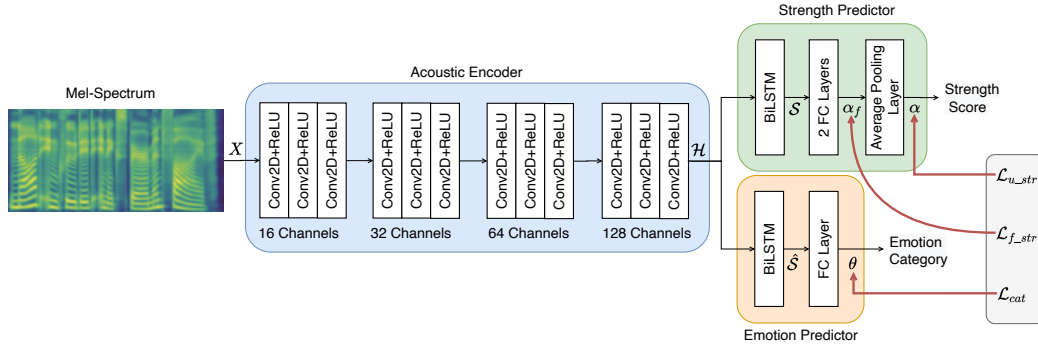


Figure 1: The proposed StrengthNet that consists of Acoustic Encoder, Strength Predictor and Emotion Predictor.

which serves as an auxiliary task.

The main contributions of this paper include, 1) We propose a novel data-driven deep learning based speech emotion strength assessment model, i.e., StrengthNet; 2) We show that the predicted emotion strength of seen and unseen speech is highly correlated with the ground truth; To our best knowledge, this is the first deep learning model for accurate emotion strength assessment for seen and unseen speech.

This paper is organized as follows: In Section 2, we formulate the proposed StrengthNet and the domain fusion training strategy. Section 3 report the experimental results. Finally, Section 4 concludes the study.

2. StrengthNet

We first describe the overall StrengthNet architecture, then explain the domain fusion training details. Lastly, we describe the model inference.

2.1. Model Overview

We formulate StrengthNet under the multi-task framework, which consists of an acoustic encoder, a strength predictor and an emotion predictor as in Fig. 1.

To improve the model generalization ability, we adopt a domain fusion training strategy. During training, the Mel-Spectrum features come from multiple emotional domains. The ground-truth strength scores, which is normalized to (0,1), are derived from a ranking function learned separately. The details will be described in Section 2.2. We employ a multi-task framework with an auxiliary emotion category prediction task. Next we will describe the model architecture first.

2.1.1. Acoustic Encoder

The acoustic encoder takes the acoustic feature sequence, that is mel-spectrum in this work, as input to extract a high-level feature representation. The acoustic encoder consists of 12 convolution layers. The strategy of stacking more convolutional layers to expand the receptive field of a CNN has been widely used to model time series data and yield satisfactory performance [7]. Given an input mel-spectrum sequence X , the CNN based acoustic encoder aims to extract a high-level feature \mathcal{H} . The high-level feature \mathcal{H} is then fed to two predictors to predict the emotion strength score and emotion category, respectively.

2.1.2. Strength Predictor

The strength predictor then reads the high-level feature representation to predict the emotion strength. Recent studies have confirmed the effectiveness of combining CNN and BiLSTM for classification [7], and recognition [11] tasks.

The strength predictor consists of a BiLSTM layer, two FC layers, followed by an average pooling layer. The BiLSTM layer takes the high-level feature \mathcal{H} as input to output the hidden states \mathcal{S} for each time step. We then use two FC layers to convert the frame-wise hidden states \mathcal{S} into the frame level scalar α_f to indicate the strength score of each frame. Finally, an average pooling layer is applied to the frame-level scores to obtain the utterance-level strength score α .

To supervise the training of the strength predictor, we define a mean absolute error (MAE) loss \mathcal{L}_{u_str} behind the average pooling layer to force the predicted utterance-level strength score α close to the ground-truth value.

To improve the convergence of StrengthNet, we further define another MAE loss [7] \mathcal{L}_{f_str} , denoted as frame-wise constraint, behind the last FC layer to minimize the difference between the predicted frame-level strength score α_f and the ground-truth strength.

2.1.3. Emotion Predictor

The auxiliary emotion predictor is devised to predict the emotion category. Similar to the strength predictor, the emotion predictor also consists of a BiLSTM layer and an additional softmax layer. The BiLSTM summarizes the temporal information of high-level feature \mathcal{H} into another latent states $\hat{\mathcal{S}}$. Finally, the softmax layer converts the latent states $\hat{\mathcal{S}}$ to the output probability for all emotion categories. Accordingly, we can obtain the predicted emotion category θ . We define a ‘‘categorical_cross_entropy’’ loss \mathcal{L}_{cat} to restrain the emotion predictor. We formulate the total objective function \mathcal{L}_{total} for training the StrengthNet as:

$$\mathcal{L}_{total} = \mathcal{L}_{f_str} + \mathcal{L}_{u_str} + \mathcal{L}_{cat}. \quad (1)$$

2.2. Domain Fusion

As shown in Fig. 2, to improve the model generalization, we employ a domain fusion strategy [12, 13], which mix multiple emotional speech datasets from various domains to train our StrengthNet. The fused data mixed from various domains will represent a more comprehensive set, thus minimizing the distance between the training and validation set, as well as any future testing sets.

Let us denote $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K$ as the K emotional speech datasets. We train the emotion attribute ranking function $R(\cdot)$ for all datasets, since they all provide <neutral, emotional> pairwise speech samples.

To achieve this, we first build two sets, that are O and S , which contain ordered and similar paired samples, respectively. Specifically, for each dataset \mathcal{D}_k ($k \in [1, K]$), we pick up one sample from neutral speech and another sample from emotional speech to build the ordered set O . We expect that

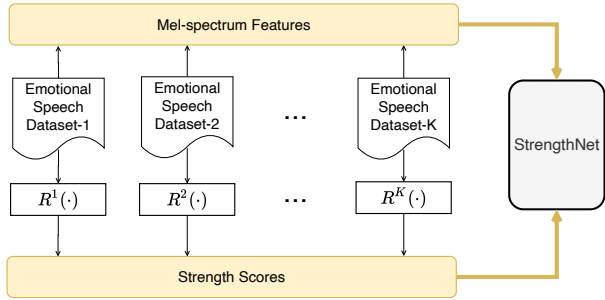


Figure 2: The shown domain fusion strategy is utilized to improve the model generalization.

the emotion strength of the emotional sample is higher than that of the neutral sample. For the similar set S , we pick up two samples from ‘neutral’ speech or emotional speech (such as ‘happy’). We assume that two samples from the same emotion category have more similar emotion strength than two others from different categories.

We follow [2] and build a support vector machine (SVM) [14] to learn the ranking function $R(\cdot)$ for the emotion strength attribute. Finally, we derive the strength scores across all datasets to serve as the ground-truth in the training objective of StrengthNet. Note that the strength score is normalized to (0,1) with 1 having the strongest emotion. We prepare mel-spectrum features for all datasets as the input of StrengthNet. The StrengthNet can be seen as the aggregation of multiple domain experts.

2.3. Run-time Inference

During inference, the StrengthNet takes a mel-spectrum extracted by any emotional speech as the input feature to predict its emotion strength score, as well as its emotion category. Furthermore, our StrengthNet can be used directly to predict emotion strength for a new emotional speech dataset without retraining, that is a clear advantage.

3. Experiments

3.1. Datasets

We use the ESD dataset [15] to validate the performance of StrengthNet in terms of strength prediction. The English corpus with a total of nearly 13 hours of speech is used in our experiments. In addition, we use two additional English SER datasets: the Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) [16] and the Surrey Audio-Visual Expressed Emotion (SAVEE) database [17] to achieve domain fusion and test the model generalization.

The datasets are summarized as follows:

- **ESD [15]:** ESD is a multilingual emotional speech dataset and has 350 parallel utterances spoken by 10 native English and 10 native Mandarin speakers. We use the English corpora with a total of nearly 13 hours of speech by 5 male and 5 female speakers in five emotions, namely happy, angry, neutral, sad and surprise. All speech samples are sampled at 16 kHz and coded in 16 bits.
- **RAVDESS [16]:** The RAVDESS dataset contains two emotional parts: a speech part and a song part. We mainly focus on its speech part, which has 1440 utterances acted by 24 professional actors (12 female, 12 male). There are 8 classes of emotion contained in

the speech part (calm, happy, sad, angry, fear, surprise, disgust, and neutral state). In the process of collecting data, actors are asked to speak two fixed sentences with different classes of emotion.

- **SAVEE [17]:** The SAVEE database consists of recordings from four male actors in seven different emotions, 480 British English utterances in total with seven emotional states (i.e., surprise, happy, sad, angry, fear, disgust, and neutral). And the sentences are selected from the standard TIMIT corpus and phonetically-balanced for each emotion.

Please note that the emotional speech databases can be divided into two types: acted and improvised [18]. All above three mentioned datasets belongs to acted family which widely used in emotional TTS [19]. We believe that it is appropriate to employ acted emotional speech dataset for our experiments, due to the improvised speech is hard to induce strong and well-differentiated emotions [20].

3.2. Experimental Setup

For all datasets, we take 5 emotion classes, that are happy, sad, angry, surprise, and neutral, to build the <neutral, emotional> paired speech to train the ranking function. During training of StrengthNet, we take 4 emotion classes, that are happy, sad, angry and surprise, into account.

We extract 80-channel mel-spectrum features with a frame size of 50ms and 12.5ms frame shift, that are further normalized to zero-mean and unit-variance, to serve as the model input. We employ openSMILE [21] to extract 384-d features for each utterance and follow the relative attributes algorithm¹ to train the ranking function. The utterance level strength scores for all datasets are obtained using their ranking function and serve as the ground truth score, or training target, for the strength predictor in our StrengthNet. To calculate the frame level MAE \mathcal{L}_{f_str} , the ground-truth strength score is used for all the frames in the speech utterance. The emotion predictor of StrengthNet aims to predict 4 emotion categories, including happy, sad, angry and surprise.

The acoustic encoder consists of 4 Conv2D blocks with filters size [16, 32, 64, 128], respectively. Each block includes 3 Conv2D layers with strides shape {[1,1], [1,1], [1,3]}, respectively. All layers share the same kernel size [3 × 3] and ReLU activation function. For BiLSTM layers in two predictors, each direction contains 128 cells.

All speech samples are resampled to 16 kHz. We train the models using the Adam optimizer with a learning rate=0.0001 and $\beta_1 = 0.9$, $\beta_2 = 0.98$. We set the batch size to 64. The dropout rate is set to 0.3. For each dataset, we partition the speech data into training, validation, and test set at a ratio of 8:1:1. We apply early stopping based on the MAE of the validation set with 30 epochs patience.

3.3. Experimental Results

3.3.1. Architecture Comparison Results

First, we intend to validate the effectiveness of the multi-task framework and frame constraint in our StrengthNet, in terms of strength prediction performance on the ESD dataset. We develop StrengthNet and two variants for an ablation study, reported as follows: 1) StrengthNet (*proposed*), which is our proposed model that consists of an acoustic encoder, strength predictor and auxiliary emotion predictor; 2) StrengthNet w/o \mathcal{L}_{cat} , which is the proposed model without an auxiliary emotion

¹<https://github.com/chaitanya100100/Relative-Attributes-Zero-Shot-Learning>




Figure 3: Histogram of the utterance level strength predictions for (a) StrengthNet; (b) StrengthNet w/o \mathcal{L}_{cat} ; (c) StrengthNet w/o \mathcal{L}_{f-str} . The X-axis and Y-axis of subfigures represent the strength scores and the utterance number, respectively.

Table 1: MAE results for the RAVDESS and SAVEE datasets in the comparison.

Method	MAE	
	RAVDESS	SAVEE
$R_{RAVDESS}(\cdot)$	NA	0.304
$R_{SAVEE}(\cdot)$	0.283	NA
$R_{ESD}(\cdot)$	0.266	0.272
StrengthNet _{ESD}	0.238	0.243
StrengthNet _{ESD+RAVDESS}	NA	0.173
StrengthNet _{ESD+SAVEE}	0.102	NA

predictor; and 3) StrengthNet w/o \mathcal{L}_{f-str} , which is the proposed model without a frame strength constraint term \mathcal{L}_{f-str} within the strength predictor.

We report the performance in terms of MAE values (range from 0 to ∞). It is negatively-oriented scores and lower values are better. Fig. 3 shows the overall performance of these systems at the utterance level. We report the MAE value of emotion strength prediction and the accuracy of emotion category prediction on test data. We observe in Fig. 3 that our proposed StrengthNet outperforms Strength w/o \mathcal{L}_{cat} and Strength w/o \mathcal{L}_{f-str} and achieves the best performance, that is attributed to the multi-task and frame constraint strategy. Specifically, we find that StrengthNet achieves the lowest MAE score of 0.072 and highest emotion recognition accuracy (denoted as *SER Acc* as shown in the axis titles of Fig. 3) of 0.874. To sum up, the above results indicate that the combination of multi-task learning and frame constraint can effectively learn emotion strength cues from the input mel-spectrum to perform emotion strength prediction well, along with category prediction.

3.3.2. Domain Fusion Results

We further verify the effectiveness of the proposed domain fusion for model generalization in terms of emotion strength prediction.

We train StrengthNet on three dataset settings, that are ESD, ESD+RAVDESS, ESD+SAVEE. StrengthNet_{ESD}, StrengthNet_{ESD+RAVDESS} and StrengthNet_{ESD+SAVEE} are used to represent three trained models, respectively. $R_{RAVDESS}(\cdot)$, $R_{SAVEE}(\cdot)$ and $R_{ESD}(\cdot)$ refer to the three trained ranking functions for three datasets, respectively. We adopt all trained ranking functions and models to predict the emotion strength score on unseen test data from RAVDESS or SAVEE, and report the MAE results in Table 1.

As shown in Table 1, we observe that our StrengthNet performs a lower MAE than ranking functions on both, the RAVDESS and SVAEE datasets. More importantly, in the case of StrengthNet_{ESD+RAVDESS} and StrengthNet_{ESD+SAVEE},

normal	0.88	0.12	normal	0.81	0.19	normal	0.78	0.22
	0.10	0.90		0.16	0.84		0.20	0.80
strong			strong			strong		
	normal	strong		normal	strong		normal	strong

Figure 4: Confusion matrices between predicted and perceived emotion strength categories for unseen speech. The X-axis and Y-axis of the figures represent the predicted and perceived category, namely “normal” and “strong”.

the MAE achieves the lowest values of 0.173 and 0.102 on SVAEE and RAVDESS respectively. As can be seen from the results, our proposed strength method can reduce the overall MAE on unseen data with the domain fusion strategy, which performs better model generalization than the attribute ranking function.

3.3.3. Human Perception Evaluation

We further evaluate how close the results of StrengthNet for unseen speech can get to the emotional strength as perceived by humans.

Since the RAVDESS dataset includes the manual emotion intensity label, which each emotion is produced at two levels of emotion intensity: normal and strong. Therefore, we adopt $R_{SAVEE}(\cdot)$, $R_{ESD}(\cdot)$ and StrengthNet_{ESD+SAVEE} to predict the emotion intensity for the test set of RAVDESS. We consider the predicted intensity scalars from 0 to 0.5 as ‘normal’ and 0.5 to 1.0 as ‘strong’ in two categories. We select 20 utterances from the test set.

Fig. 4 presents the intensity confusion matrices. It is observed that the StrengthNet shows a higher correlation between the predicted and perceived emotion intensity categories, with a correlation of over 85%, that is considered a competitive result against the trained ranking function. The experiments confirm the superiority of the proposed domain fusion based StrengthNet for unseen speech in terms of human perception.

4. Conclusion

This paper presents a data-driven deep learning-based speech emotion strength assessment model for the emotional speech synthesis task, referred to as StrengthNet. Experimental results demonstrate that our StrengthNet can achieve accurate emotion strength prediction for both seen and unseen speech with the help of domain fusion strategy. As per our knowledge, the proposed StrengthNet is the first end-to-end speech emotion strength assessment model. In future work, we intend to integrate our StrengthNet as a front-end or back-end for emotional speech synthesis models to enhance the emotion expressiveness of output emotional speech.

5. References

- [1] T. Li, S. Yang, L. Xue, and L. Xie, "Controllable emotion transfer for end-to-end speech synthesis," in *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.
- [2] X. Zhu, S. Yang, G. Yang, and L. Xie, "Controlling emotion strength with relative attribute for end-to-end speech synthesis," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 192–199.
- [3] Y. Lei, S. Yang, and L. Xie, "Fine-grained emotion strength transfer, control and prediction for emotional speech synthesis," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 423–430.
- [4] D. Parikh and K. Grauman, "Relative attributes," in *2011 International Conference on Computer Vision*. IEEE, 2011, pp. 503–510.
- [5] B. Siddiquie, R. S. Feris, and L. S. Davis, "Image ranking and retrieval based on multi-attribute queries," in *CVPR 2011*. IEEE, 2011, pp. 801–808.
- [6] Z. Meng, N. Adluru, H. J. Kim, G. Fung, and V. Singh, "Efficient relative attribute learning using graph neural networks," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 552–567.
- [7] C. Lo, S. Fu, W. Huang, X. Wang, J. Yamagishi, Y. Tsao, and H. Wang, "Mosnet: Deep learning-based objective assessment for voice conversion," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 1541–1545. [Online]. Available: <https://doi.org/10.21437/Interspeech.2019-2003>
- [8] S. Fu, Y. Tsao, H. Hwang, and H. Wang, "Quality-net: An end-to-end non-intrusive speech quality assessment model based on BLSTM," in *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, B. Yegnanarayana, Ed. ISCA, 2018, pp. 1873–1877. [Online]. Available: <https://doi.org/10.21437/Interspeech.2018-1802>
- [9] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [10] H. Dinkel, S. Wang, X. Xu, M. Wu, and K. Yu, "Voice activity detection in the wild: A data-driven approach using teacher-student training," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1542–1555, 2021.
- [11] R. Liu, B. Sisman, G. Gao, and H. Li, "Expressive tts training with frame and style reconstruction loss," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1806–1818, 2021.
- [12] R. Volpi, H. Namkoong, O. Sener, J. Duchi, V. Murino, and S. Savarese, "Generalizing to unseen domains via adversarial data augmentation," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 5339–5349.
- [13] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, pp. 1–48, 2019.
- [14] T. Joachims, "Optimizing search engines using clickthrough data," in *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2002, pp. 133–142.
- [15] K. Zhou, B. Sisman, R. Liu, and H. Li, "Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 920–924.
- [16] S. R. Livingstone and F. A. Russo, "The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english," *PLoS one*, vol. 13, no. 5, p. e0196391, 2018.
- [17] P. Jackson and S. Haq, "Surrey audio-visual expressed emotion (savee) database," *University of Surrey: Guildford, UK*, 2014.
- [18] P. N. Juslin and P. Laukka, "Impact of intended emotion intensity on cue utilization and decoding accuracy in vocal expression of emotion," *Emotion*, vol. 1, no. 4, p. 381, 2001.
- [19] C. Cui, Y. Ren, J. Liu, F. Chen, R. Huang, M. Lei, and Z. Zhao, "EMOVIE: A Mandarin Emotion Speech Dataset with a Simple Emotional Text-to-Speech Model," in *Proc. Interspeech 2021*, 2021, pp. 2766–2770.
- [20] R. Banse and K. R. Scherer, "Acoustic profiles in vocal emotion expression," *Journal of personality and social psychology*, vol. 70, no. 3, p. 614, 1996.
- [21] F. Eyben, F. Wenyinger, F. Gross, and B. Schuller, "Recent developments in opensmile, the munich open-source multimedia feature extractor," in *Proceedings of the 21st ACM international conference on Multimedia*, 2013, pp. 835–838.