

# DISENTANGLEMENT FOR AUDIO-VISUAL EMOTION RECOGNITION USING MULTITASK SETUP

Raghuveer Peri<sup>†</sup>\*      Srinivas Parthasarathy\*      Charles Bradshaw\*      Shiva Sundaram\*

<sup>†</sup> University of Southern California, Los Angeles, CA, USA

\*Amazon Inc., Sunnyvale, CA, USA

## ABSTRACT

Deep learning models trained on audio-visual data have been successfully used to achieve state-of-the-art performance for emotion recognition. In particular, models trained with multitask learning have shown additional performance improvements. However, such multitask models entangle information between the tasks, encoding the mutual dependencies present in label distributions in the real world data used for training. This work explores the disentanglement of multimodal signal representations for the primary task of emotion recognition and a secondary person identification task. In particular, we developed a multitask framework to extract low-dimensional embeddings that aim to capture emotion specific information, while containing minimal information related to person identity. We evaluate three different techniques for disentanglement and report results of up to 13% disentanglement while maintaining emotion recognition performance.

*Index Terms*— Keywords: Emotion recognition, multimodal learning, disentanglement, multitask learning

## 1. INTRODUCTION

Emotions play an important role in human communication. Humans externalize their reactions to surrounding stimuli through a change in the tone of their voice, facial expressions, hand and body gestures. Therefore, automatic emotion recognition is of interest for building natural interfaces and effective human-machine interaction. [1]. With regards to human communication, emotion is primarily manifested through speech and facial expressions, each providing complementary information [2]. Therefore, multimodal techniques have been widely used for reliable emotion prediction [3, 4, 5].

Several studies have shown that emotion recognition benefits from training with secondary related tasks through multitask learning (MTL). In Parthasarathy and Busso [6], predicting the continuous affective attributes of valence, arousal and dominance are treated as the multiple tasks and trained jointly. In Li et al. [7] and Kim et al. [8], gender prediction as a secondary task improves emotion recognition performance by upto 7.7% as measured by weighted accuracy on a standard corpus. A more comprehensive study involving domain, gender and corpus differences was performed in Zhang et al. [9], where cross-corpus evaluations showed that, in general, information sharing across tasks yields improvements in performance of emotion recognition across corpora. These studies indicate that several paralinguistic tasks help generalize shared representations that improve overall performance of the primary task. This motivates us to use person identification as a secondary task to help improve performance on the primary emotion task.

With MTL the shared representations among tasks retain information pertaining to all the tasks. While this generalizes the overall architecture, it does so by entangling information between multiple tasks [10, 11, 12]. Since most machine learning models are trained on human-annotated, unconstrained real-world data, several factors that should theoretically be independent end up being dependent. For e.g. in the case of emotions, studies have shown the correlation with demographical information [13]. Therefore, MTL inherently captures the joint dependencies between different factors in the data. This is problematic as the gains through generalization across tasks may lead to bias and subsequently poor performance on unseen data.

To address the entanglement of information in MTL, this paper develops a multimodal emotion recognition model, improves its performance using person identification as a secondary task and subsequently disentangles the learned person identity information, while still maintaining the improved emotion recognition performance. As an additional contribution, we analyze how much emotion information is present in the identity representations when models are trained in a MTL setup. For disentanglement, we experiment with three distinct disentanglement techniques to minimize the information transfer between speaker embeddings and emotional labels and vice-versa. We present experiments that make use of alternate adversarial training strategy, gradient reversal based technique adapted from Domain Adversarial Training (DAT) literature and a confusion loss based technique inspired from [14]. We evaluate the models pre and post disentanglement, showing that disentanglement retains or improves performance on primary tasks upto 2% absolute, while reducing the leakage of information between the tasks with disentanglement upto 13% as measured by F-score.

## 2. RELATED WORK

In the context of representation learning for emotion recognition, the goal is to extract low dimensional embeddings that are invariant to factors such as domain and speaker. Abdelwahab and Busso [15] used gradient reversal (GR) to extract emotion representations that are invariant to domain. Mao et al. [16] imposed an explicit orthogonality criterion to encourage the learning of domain invariant and emotion discriminative features. Similarly, to extract speaker-invariant emotion representations, adversarial learning approach was explored in addition to an online data augmentation technique by Tu et al. [17]. They showed improvements in the emotion recognition performance while testing on speakers unseen during training. More recently Li et al. [18] proposed an entropy-based loss function along with GR and showed improved performance compared to [17]. Kang et al. [19] propose channel and emotion invariant speaker embeddings. However, most of these works consider emotion recognition using speech modality alone. Jaiswal and Provost [20] explored

\*Work was completed while interning at Amazon

privacy-preserving multimodal emotion representations, where audio and text modalities were utilized. Our study differs from previous studies by using a secondary task to improve primary emotion recognition performance while being invariant to the auxiliary factors.

With regards to identity embeddings, Williams and King [12] have shown that speaker embeddings capture significant amount of affect information. It has been found that differences in the affective states of a person between training and testing conditions can degrade the performance on the task of identity verification from speech [21, 22]. Techniques have been proposed to compensate this by transforming features from expressive speech to neutral speech domain [23, 24]. While most of the existing works learn identity representations separately and then try to make them invariant to emotional states, we co-learn identity representations with an emotion recognition task while simultaneously removing emotion information from them.

### 3. METHODOLOGY

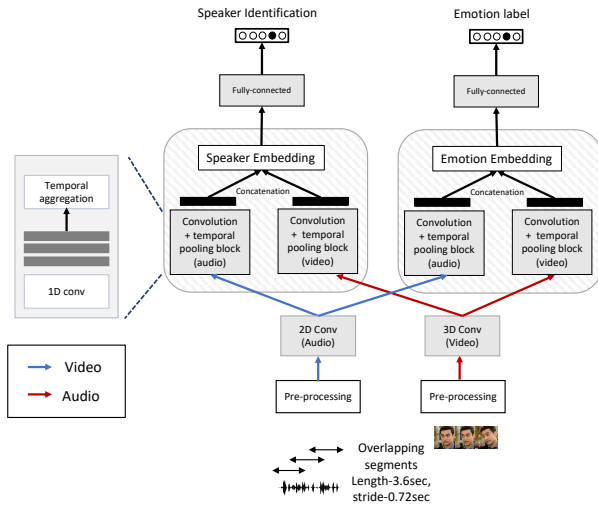


Fig. 1. Block diagram depicting the baseline multimodal, multitask training

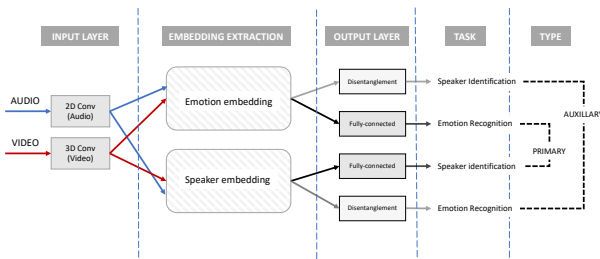


Fig. 2. Block diagram depicting the baseline model with an auxiliary disentanglement task

Fig. 1 illustrates the multitask architecture for emotion recognition and person identification. The inputs to the model are both audio and video frames that are time-synchronized. The first step is a shared convolutional feature extraction stage where a data-driven

representation is extracted for both audio and video independently. The architectures for this first stage blocks are adopted from [25]. A second level temporal aggregation block pools the feature representation for audio and video separately over entire clips to fixed dimensional representation. The outputs of the audio and video pooling blocks are concatenated; resulting in independent embedding layers *emotion embedding* and *speaker embedding*. The final output layers for task-specific outputs are fully connected layers with a softmax activation function to predict the emotion and person identity labels respectively. Please note that we have used the terms *speaker identity* and *person identity* interchangeably throughout the paper.

Fig. 2 illustrates the addition of auxiliary branches to the baseline multitask architecture. The auxiliary branches are used to assess the amount of emotion information in the speaker embeddings and vice versa. These auxiliary branches are also used for disentanglement as explained in Section 3.2.

#### 3.1. Pre-processing

The input audio and face crop streams from a video clip are first fed into corresponding pre-processing blocks. On the audio stream, pre-processing includes extracting log Mel frequency spectrogram features on overlapping segments of fixed length and stride. This results in one feature vector per segment, with varying number of segments per video clip, depending on the length of the clip. In order to perform efficient batch processing, we pad the features with a constant value to ensure that each video clip contains the same number of segments,  $N$ . The resulting features have the dimensions  $B * N * D_a$  where  $B$  is the minibatch size and  $D_a$  is the dimension of the Mel spectrogram features. On the face crops, pre-processing includes resizing them into a fixed size of  $D_v * D_v$  pixels and rescaling the values to between  $-1$  and  $1$ . The resulting face crops have the dimensions  $B * N * D_v * D_v$ .

#### 3.2. Auxiliary branch for disentanglement

The multitask outputs are built on top of the common embedding layers for the emotion and person identification tasks respectively. As a result, when training the model, it tends to train an entangled embedding that is optimized for both tasks. This form of entanglement could lead to learning needless dependencies in the train set that may affect the overall generalization. In this work, both for person identification and emotion recognition, the second output or auxiliary task is used to disentangle the emotion information from the speaker embeddings and vice-versa (Fig. 2). Disentanglement is achieved using the auxiliary branch. The basic intuition here is similar to domain adversarial training, where the goal is to learn representations that are optimized for the primary task, while simultaneously training it to perform poorly on the auxiliary task. To this end, we experiment with three techniques for disentanglement: (1) gradient reversal, (2) alternate primary-auxiliary training and (2) and confusion loss (CONF).

**Gradient reversal** was originally developed in Ganin and Lempitsky [26] to make digit recognition task invariant to domain through adversarial training. As discussed in Section 2, it was adapted to extract speaker-invariant speech emotion representations in Tu et al. [17]. Gradient reversal is achieved by introducing it in the stages of a network where the auxiliary branch separates from the primary branch. This layer has no effect in the forward pass of training, while in the backward pass the gradients from the auxiliary branch are multiplied by a negative value before backpropagating it to the embedding layer.

**Alternate training** strategy for disentanglement was inspired from adversarial training literature [27], where two models are trained with competing objectives. In our setup, for emotion embeddings, the primary task is to predict the emotion labels, while the auxiliary task is to predict person identity labels. Equations 1 and 2 show the loss functions of the primary and auxiliary branch respectively, which are modeled as cross-entropy loss.  $\hat{e}_{prim}$  and  $\hat{s}_{prim}$  denote the primary predictions from the emotion and speaker identification branches respectively. Similarly,  $\hat{e}_{aux}$  and  $\hat{s}_{aux}$  denote the auxiliary predictions from the speaker identification and emotion recognition branches respectively.  $e_{target}$  and  $s_{target}$  denote the groundtruth emotion and speaker identity labels.

$$L_{primary} = w_{em.prim} * L(\hat{e}_{prim}, e_{target}) + w_{spk.prim} * L(\hat{s}_{prim}, s_{target}) \quad (1)$$

$$L_{auxiliary} = w_{spk.aux} * L(\hat{e}_{aux}, e_{target}) + w_{em.aux} * L(\hat{s}_{aux}, s_{target}) \quad (2)$$

Alternate training proceeds in a minimax fashion. The auxiliary branch is trained to minimize  $L_{auxiliary}$ , while the primary branch is trained to minimize  $L_{primary}$  and simultaneously maximize  $L_{auxiliary}$ .

**Confusion loss** for disentanglement has been introduced in Tzeng et al. [28] and adapted for disentangling person identity and spoken content representations in Nagrani et al. [25]. We apply a similar strategy to disentangle the emotion and person identity representations. On a high level, the loss forces the embeddings such that, for the auxiliary task, each class is predicted with the same probability. Similar to [25], we implement the confusion loss as the cross-entropy between the predictions and a uniform distribution.

## 4. EXPERIMENTAL FRAMEWORK

### 4.1. Dataset

For the primary task and disentanglement experiments for multimodal emotion recognition, we use the EmoVox dataset [29]. The EmoVox dataset comprises of emotional labels on the VoxCeleb dataset obtained by predictions using a strong teacher network over eight emotional states: neutral, happiness, surprise, sadness, anger, disgust, fear and contempt. Note that the teacher model was trained only using facial features (visual only). Overall, the dataset consists of interview videos from 1251 celebrities spanning a wide range of ages and nationalities. For each video clip, we find the most dominant emotion based on the distribution and use that as our ground-truth label similar to [29]. The label distribution is heavily skewed towards a few emotion classes because emotions such as disgust, fear, contempt and surprise are rarely exhibited in interviews. Following previous approaches that deal with such imbalanced datasets [30], we combine these labels into a single class ‘other’, resulting in 5 emotion classes. Further, we discard videos corresponding to speakers belonging to the bottom 5 percentile w.r.t the number of segments to reduce the imbalance in the number of speech segments per speaker. We create three splits from the database: *EmoVox-Train* to train models, *EmoVox-Validation* for hyperparameter tuning, *EmoVox-Test* to evaluate models on held out speech segments from speakers present in the train set. The subset *EmoVox-Train* corresponds to the *Train* partition in [29], whereas the *EmoVox-Validation* and *EmoVox-Test* were created from the *Heard-Val* partition in [29].

### 4.2. Experimental Settings

The model architecture for the shared 2D Convolutional layers and the fully connected layers was adapted from [25] and modified to suit the dimensions of our inputs and outputs. We use uniform duration videos of 12 seconds each as input to our system. For the audio features we use  $D_a = 40$ , and for the visual features we use  $D_v = 224$ . We fix the emotion embedding dimension to 2048, while varying the speaker embedding dimension 2048, 256 and 64. We use Adam optimizer with an initial learning rate of  $1e - 4$  and  $1e - 3$  for the primary branch and auxiliary branch updates respectively, decaying exponentially with a factor of  $\gamma = 0.9$ . For alternate training (Eqs. 1 and 2), we chose  $w_{em.prim}$  and  $w_{spk.prim}$  to be 0.5 each and  $w_{em.aux}$  and  $w_{spk.aux}$  to 0.3 each. All parameters were chosen based on preliminary experiments on a subset of *EmoVox-Train*. The emotion recognition performance was evaluated using unweighted F-score averaged across the 5 emotion classes and for person identity with identification accuracy scores. Disentanglement is measured by combining both the F-score on emotion recognition using speaker embeddings and accuracy on person identification using emotion embeddings. Optimal models were chosen to give the best disentanglement (lowest score) on the *EmoVox-Validation* set. All results are presented on the *EmoVox-Test* set.

## 5. RESULTS



### 5.1. Baseline models without disentanglement

**Emotion Recognition:** Figure 3(a) illustrates the primary emotion recognition results. The blue bars show the performance of all models trained using MTL and the dashed line shows the performance of Single-task learning (STL) setup where the models are not trained on person identification. It is evident that MTL gives substantial gains in performance compared to STL setup. It is also observed that emotion recognition performance improves as the person identification embedding dimension is reduced, which may indicate better regularization with fewer embedding dimensions.

**Person identification:** Table 1 shows the person identification accuracy, trained with varying speaker embedding dimensions. It is worth noting that, despite the reduction in speaker embedding dimension, the models retain performance, pointing to the fact that the task of learning identity representations when both audio and visual modalities are available does not require many degrees of freedom.

**Identity information in emotion embeddings:** Our preliminary experiments showed that the amount of person identity information entangled in emotion embeddings was minimal. Evaluating the person identification task using emotion embeddings produced an accuracy of 0.1%, which was close to random chance performance. Therefore we focus on disentangling emotion information in identity embeddings.

**Emotion information in identity embeddings** To baseline the amount of emotion information entangled in the speaker embeddings, we separately train single hidden layer neural network classifiers that predict the emotion labels from speaker embeddings. Figure 3(b) illustrates the performance. First, it is worth noting that speaker embeddings from models trained for the single task of person identification retain substantial amount of emotion information, as shown by the red dashed line, compared to a random chance F-score of 17.40% if all samples were predicted as ‘neutral’ class (shown by the green dashed line). Further the blue bars illustrate the performance in the MTL setup where the F-scores are well above random chance as there is more information entanglement. This motivates the need for disentanglement to minimize the emotion



**Fig. 3.** Unweighted Average F-scores for AER on *EmoVox-Test* by varying speaker embedding dimension using (a) Emotion embeddings (Higher is better) (b) Speaker embeddings (Lower is better). A Stuart-Maxwell marginal homogeneity test comparing the results found, with statistical significance, that all the models with disentanglement were different compared to the baseline model

**Table 1.** Person identification accuracy (%) comparing models with varying speaker embedding dimensions without and with disentanglement on *EmoVox-Test*

Emb Dim	Baseline	ALT	GR	CONF
2048	90.98	92.40	93.19	93.12
256	94.75	95.04	95.86	95.42
64	90.62	92.83	91.17	90.75

information present in speaker embeddings without compromising performance on the emotion recognition, person identification tasks.

## 5.2. Proposed models with disentanglement

Next we report the results of the proposed disentanglement techniques and compare them to the baseline models. We trained each disentanglement technique for all three configurations of speaker embedding dimension, 2048, 256 and 64 to investigate their effect on disentanglement performance

**Emotion Recognition** From Fig. 3(a), we observe that models trained with all three disentanglement strategies outperform the baseline models trained without disentanglement in all but one case. In particular, ALT and CONF methods provide gains consistently across the various embedding dimensions. We performed a Stuart-Maxwell marginal homogeneity test comparing the results and found, with statistical significance, that all the models with disentanglement were different compared to the baseline models<sup>1</sup>. We also observe that, similar to the baseline models, models trained with disentanglement tend to perform better for reduced speaker embedding dimensions, though with smaller gains.

**Person identification** Table 1 shows the person identification accuracy for the models with disentanglement compared to the baseline without disentanglement. We observe that, in general, all models perform better after disentanglement when compared to the baseline without disentanglement. There is no clear evidence of one technique performing better than the other, though GR and ALT seem to

<sup>1</sup> $H_0$ : The predictions from the compared models are the same. Reject  $H_0$  if  $p < \alpha$  with  $\alpha = 0.01$

perform marginally better compared to CONF.

**Emotion information in identity embeddings** Fig. 3(b) illustrates the amount of emotion information in the person identity embeddings after explicit disentanglement. The drop in unweighted average F-score for emotion recognition shows the measure of obtained disentanglement. Compared to the models trained without disentanglement, we observe that the models trained with explicit disentanglement show reduction in F-score of predicting emotions from speaker embeddings. This is noticeable in all the three disentanglement techniques. ALT, CONF training show better disentanglement than GR. Overall, these results show the efficacy of using a separate auxiliary branch to disentangle the emotion information from speaker embeddings. Furthermore, it can be observed that the models trained using the smallest speaker embedding dimension of 64 shows the least amount of emotion information. This is expected because a reduced person identity embedding dimension creates a bottleneck to capture the primary identity information, and thus retains lesser amount of entangled emotion information. Considering the person identity dimension of 64, we see absolute gains of 2% for emotion recognition while ALT training gives 13.5% disentanglement.

## 6. CONCLUSIONS

This study analyses disentanglement techniques for emotion recognition in a multitask learning setup, where person identification is the secondary task. We showed with an audio-visual architecture that person identification helps emotion recognition performance. This comes at a cost, as there is significant information transfer between the tasks, which lets us predict emotional categories from speaker embeddings well above chance percentage. To combat this we studied three disentanglement techniques, each reducing the amount of information that is entangled while maintaining or improving performance on the primary task. For our next steps we will explore and validate these methods on other databases which have stronger emotion labels. Furthermore, it is of interest to dig deeper into the reasons for differences in performance across the various disentanglement methods. Finally, this paper shows that there is significant emotional information in the speaker embeddings and the contrary is not necessarily true. Therefore we will explore a hierarchical structure where emotion recognition is more downstream than the person identification task.

## 7. REFERENCES

- [1] M. Pantic and J. M. L. Rothkrantz, "Toward an affect-sensitive multimodal human-computer interaction," *Proceedings of the IEEE*, vol. 91, no. 9, pp. 1370–1390, 2003.
- [2] A. Mehrabian, "Communication without words," *Communication theory*, vol. 6, pp. 193–200, 2008.
- [3] Y. Kim, H. Lee, and E. M. Provost, "Deep learning for robust feature generation in audiovisual emotion recognition," in *2013 IEEE international conference on acoustics, speech and signal processing*, Vancouver, BC, Canada, 2013, IEEE, pp. 3687–3691.
- [4] Y. Wang and L. Guan, "Recognizing human emotional state from audiovisual signals," *IEEE transactions on multimedia*, vol. 10, no. 5, pp. 936–946, 2008.
- [5] M. Song, J. Bu, C. Chen, and N. Li, "Audio-visual based emotion recognition—a new approach," in *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, Washington, DC, USA, 2004, IEEE, vol. 2, pp. II–II.
- [6] S. Parthasarathy and C. Busso, "Jointly predicting arousal, valence and dominance with multi-task learning.," in *Interspeech*, Stockholm, Sweden, 2017, pp. 1103–1107.
- [7] Y. Li, T. Zhao, and T. Kawahara, "Improved end-to-end speech emotion recognition using self attention mechanism and multi-task learning.," in *Interspeech*, Graz, Austria, 2019, pp. 2803–2807.
- [8] J. Kim, G. Englebienne, P. K. Truong, and V. Evers, "Towards speech emotion recognition" in the wild" using aggregated corpora and deep multi-task learning," *arXiv preprint arXiv:1708.03920*, 2017.
- [9] B. Zhang, E. M. Provost, and G. Essl, "Cross-corpus acoustic emotion recognition with multi-task learning: Seeking common ground while preserving differences," *IEEE Transactions on Affective Computing*, vol. 10, no. 1, pp. 85–99, 2017.
- [10] J. Liang, Z. Liu, J. Zhou, X. Jiang, C. Zhang, and F. Wang, "Model-protected multi-task learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2020.
- [11] L. Xiao, H. Zhang, W. Chen, Y. Wang, and Y. Jin, "Learning what to share: Leaky multi-task network for text classification," in *Proceedings of the 27th International Conference on Computational Linguistics*, Santa Fe, NM, USA, 2018, pp. 2055–2065.
- [12] J. Williams and S. King, "Disentangling style factors from speaker representations.," in *INTERSPEECH*, Graz, Austria, 2019, pp. 3945–3949.
- [13] T. M. Chaplin, "Gender and emotion expression: A developmental contextual perspective," *Emotion Review*, vol. 7, no. 1, pp. 14–21, 2015.
- [14] M. Alvi, A. Zisserman, and C. Nellåker, "Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings," in *Proceedings of the European Conference on Computer Vision (ECCV)*, Munich, Germany, 2018, pp. 556–572.
- [15] M. Abdelwahab and C. Busso, "Domain adversarial for acoustic emotion recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 12, pp. 2423–2435, 2018.
- [16] Q. Mao, G. Xu, W. Xue, J. Gou, and Y. Zhan, "Learning emotion-discriminative and domain-invariant features for domain adaptation in speech emotion recognition," *Speech Communication*, vol. 93, pp. 1–10, 2017.
- [17] M. Tu, Y. Tang, J. Huang, X. He, and B. Zhou, "Towards adversarial learning of speaker-invariant representation for speech emotion recognition," *arXiv preprint arXiv:1903.09606*, 2019.
- [18] H. Li, M. Tu, J. Huang, S. Narayanan, and P. Georgiou, "Speaker-invariant affective representation learning via adversarial training," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Barcelona, Spain, 2020, IEEE, pp. 7144–7148.
- [19] W. H. Kang, S. H. Mun, M. H. Han, and N. S. Kim, "Disentangled speaker and nuisance attribute embedding for robust speaker verification," *IEEE Access*, vol. 8, pp. 141838–141849, 2020.
- [20] M. Jaiswal and E. M. Provost, "Privacy enhanced multimodal neural representations for emotion recognition.," in *AAAI*, 2020, pp. 7985–7993.
- [21] S. Parthasarathy, C. Zhang, J. H. L. Hansen, and C. Busso, "A study of speaker verification performance with expressive speech," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, USA, 2017, IEEE, pp. 5540–5544.
- [22] W. Wu, T. F. Zheng, M. X. Xu, and H. J. Bao, "Study on speaker verification on emotional speech," in *Ninth International Conference on Spoken Language Processing*, Pittsburgh, PA, USA, 2006, pp. 2102–2105.
- [23] H. Bao, M. X. Xu, and T. F. Zheng, "Emotion attribute projection for speaker recognition on emotional speech," in *Eighth Annual Conference of the International Speech Communication Association*, Antwerp, Belgium, 2007, pp. 758–761.
- [24] S. R. Krothapalli, J. Yadav, S. Sarkar, G. S. Koolagudi, and A. K. Vuppala, "Neural network based feature transformation for emotion independent speaker identification," *International Journal of Speech Technology*, vol. 15, no. 3, pp. 335–349, 2012.
- [25] A. Nagrani, J. S. Chung, S. Albanie, and A. Zisserman, "Disentangled speech embeddings using cross-modal self-supervision," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Barcelona, Spain, 2020, IEEE, pp. 6829–6833.
- [26] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *International conference on machine learning*, Lille, France, 2015, PMLR, pp. 1180–1189.
- [27] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, Montreal, Canada, 2014, pp. 2672–2680.
- [28] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *Proceedings of the IEEE International Conference on Computer Vision*, Santiago, Chile, pp. 4068–4076.
- [29] S. Albanie, A. Nagrani, A. Vedaldi, and A. Zisserman, "Emotion recognition in speech using cross-modal transfer in the wild," in *Proceedings of the 26th ACM international conference on Multimedia*, Seoul, Republic of Korea, pp. 292–301.
- [30] C. Busso, S. Parthasarathy, A. Burmania, M. AbdelWahab, N. Sadoughi, and E. M. Provost, "Msp-improv: An acted corpus of dyadic interactions to study emotion perception," *IEEE Transactions on Affective Computing*, vol. 8, no. 1, pp. 67–80, 2016.