

Acted vs. Improvised: Domain Adaptation for Elicitation Approaches in Audio-Visual Emotion Recognition

Haoqi Li^{1,3,*}, Yelin Kim², Cheng-Hao Kuo², Shrikanth Narayanan³

¹Amazon AWS AI ²Amazon Lab126

³Dept. of Electrical and Computer Engineering, University of Southern California, USA

{haoqili, kimyelin, chkuo}@amazon.com, shri@sipi.usc.edu

Abstract

Key challenges in developing generalized automatic emotion recognition systems include scarcity of labeled data and lack of gold-standard references. Even for the cues that are labeled as the same emotion category, the variability of associated expressions can be high depending on the elicitation context e.g., emotion elicited during improvised conversations vs. acted sessions with predefined scripts. In this work, we regard the emotion elicitation approach as domain knowledge, and explore domain transfer learning techniques on emotional utterances collected under different emotion elicitation approaches, particularly with limited labeled target samples. Our emotion recognition model combines the gradient reversal technique with an entropy loss function as well as the softlabel loss, and the experiment results show that domain transfer learning methods can be employed to alleviate the domain mismatch between different elicitation approaches. Our work provides new insights into emotion data collection, particularly the impact of its elicitation strategies, and the importance of domain adaptation in emotion recognition aiming for generalized systems.

Index Terms: emotion recognition, emotion elicitation, domain adaptation, domain transfer learning

1. Introduction

Human emotion expressions can provide clues about individual needs, preferences, and attitudes; emotion-related cues can be used to create more natural and human-centered interactive technology. Recent deep learning techniques have shown promising improvements in automatic emotion recognition [1, 2]; however, building an accurate emotion recognition system is still challenging due to the shortage of labeled data, and the lack of gold-standard references. Moreover, the performance of emotion recognition system is often sensitive to the variations across different domains. These domain variations can be potentially caused by many factors such as the speaker characteristics [3] and recording conditions [4], emotion elicitation approaches [5, 6, 7, 8] or even socio-cultural context [9, 10]. Hence, emotion recognition systems that are only evaluated within a single domain condition may not be robust to these variations.

In automated emotion recognition (ER) [11, 12, 13], transfer learning has been employed to alleviate the domain mismatch in the cross-corpora [14, 15, 16, 17, 18], cross-language [16, 19, 20], cross-speaker [21, 22] and cross-modality [23, 24] scenarios. Besides those above, emotion elicitation approaches can be another important mismatch factor in ER. Existing works [5, 6] already show utterances with the same emotion label can have different expressed behaviors in natural conversations,

scripted acted talks and improvised conversations. However, related work of quantitative inferential capacity across different emotion elicitation approaches is still very limited.

The elicitation approaches for most emotion recognition corpora can be divided into two categories: (1) read-speech with fixed lexical scripts and (2) spontaneous speech collected from improvised interaction sessions. In this work, we analyze the domain mismatch and inferential capacity across utterances generated from different *emotion elicitation approaches*. We explore two hypotheses: **H1**: Improvised utterances as a source domain can help the emotion recognition of the scripted utterances and **H2**: Using scripted utterances as a source domain can improve the emotion recognition of the improvised utterances.

H1 is motivated by real human-agent interactions. Many widely-deployed speech agents (e.g. at home, cars etc.) are likely used in specific target domains (e.g. “weather”, “music” etc.), in which most commands have relatively fixed and specialized content. Thus, the potential of leveraging a larger corpus with free content (e.g. public emotional utterances collected from Youtube) for target domain adaptation is explored under source-target setting in **H1**. Currently, a large number of emotional utterances can be easily collected from online resources, such as online videos [25] and call recordings [26, 27], however, copyright issues, privacy, and lack of controlled recording settings (background noise, etc.) might cause serious issues in commercial products or usage. Thus, **H2** is motivated in that the many available curated emotion datasets are collected in carefully designed settings with controlled conditions, including with user consent, and leveraging these data for improving performances on natural sessions can be highly impactful.

To alleviate the domain mismatch, traditional transfer learning approaches, such as fine-tuning on target domain, often require hundreds of labeled samples, which can be challenging in the ER area considering the scarcity and cost of labeled data. Recently, domain adversarial training [28] has been employed in ER to remove domain information [21, 29, 30]. In our work, to verify the two proposed hypotheses above, we build an emotion recognition model with the combination of adversarial iterative training strategy and softlabel loss from work [21] and [31]. The adversarial iterative training strategy, on the one hand, is to maximize the capture of emotion-related information; on the other hand, is to remove domain information (the difference between various emotion elicitation approaches). The introduction of softlabel loss ensures that domain transfer learning can consider both domains and emotion categories simultaneously. This is motivated by constraints in real-world use settings, where it is not feasible to have enough target domain labeled data to conduct supervised tuning for each customer or each new usage scenario.

* Work performed while interning at Amazon.




Figure 1: Our model includes the representation encoder (ENC), one domain classifier (DC), and one emotion classifier (EC). We use a domain classifier loss to optimize the DC and use domain confusion loss, emotion classifier loss and softlabel loss to remove the domain information and optimize the emotion classification accuracy.

2. Methodology

Our work aims to (1) investigate domain adaptation across utterances obtained through different emotion elicitation approaches and (2) apply domain adversarial learning and softlabel loss with limited target samples. The goal of adversarial iterative training strategy is to predict emotion class labels while simultaneously finding representations that make the domains as similar as possible. Technically, a max-entropy adversarial network model from [21] is utilized by optimizing over a loss which includes both classification error on the labeled samples as well as a combined gradient reversal technique with an entropy loss function to make the domains indistinguishable.

However, while maximizing domain confusion pulls the marginal distributions of the domains together, it doesn't necessarily align the different emotion categories of the target with those in the source. The correlations across different emotions are related to the properties of human emotion expression and perception. For example, from the acoustic expression aspect, happiness and anger often result in a similar higher pitch. In addition, in emotion perception, sadness can be easily confused with neutral states regardless of utterance type as shown in [6]. Motivated by work [31], besides a single emotion label, we also employ softlabel loss to perform the emotion category adaptation. We explicitly transfer the similarity of emotion categories from the source to the target and further optimize our representation to produce the same structure in the target domain using a few target labeled examples as reference points.

2.1. Max-Entropy Adversarial Network

Our proposed model takes acoustic and visual features as inputs, and contains three modules: the representation encoder (ENC), the domain classifier (DC), and the emotion classifier (EC) as shown in Figure 1.

For the acoustic ENC, we use stacked 1D convolutional layers and followed by sequential GRU layers [32] to obtain the acoustic representation for each utterance. For the visual ENC, we extract the intermediate embedding from a pre-trained inception-ResNet [33] model on a face recognition task, then use this embedding as the input feature to another GRU model to generate the visual representation. Then, we concatenate the acoustic and visual representations and add fully connected layers to generate the emotion representation. This fixed dimension representation, as the output of the ENC, is further connected to two classification modules: the emotion classifier (EC) and the domain classifier (DC). We name the parameters of each component as θ_{enc} , θ_{EC} and θ_{DC} , respectively.

The model takes as input the labeled source data $\{x_S, y_S\}$, and a few target labeled data $\{x_T, y_T\}$. The goal is to train an emotion category classifier that operates on the emotion repre-

sentation $enc(x)$, and can correctly classify the target samples during the evaluation. The training strategy has the iterative update process: On the one hand, we attempt to optimize the DC to estimate the domain information; on the other hand, we attempt to correctly estimate the emotion label and remove the domain information for ENC output.

2.1.1. Training of DC

For a given representation v , the elicitation domain classifier (DC) can be regarded as a “discriminator”, which is trained to distinguish the domain identity. Mathematically, this can be done by optimizing a cross entropy objective function as follow:

$$\min_{\theta_{DC}} L_D(x_S, x_T, \theta_{enc}; \theta_{DC}) = -\sum_d \mathbb{1}[y_D = d] \log P(d|enc(x)) \quad (1)$$

where d is the domain label for sample x , and $P(d|enc(x))$ is the softmax output of domain classifier. In this training step, only the parameters of DC, i.e. θ_{DC} , are optimized.

2.1.2. Training of ENC and EC

For the training of ENC and EC, we need to optimize the model to improve the emotion classification accuracy. Simultaneously, we need to optimize ENC to increase the uncertainty or randomness of DC's outputs. Mathematically, we maximize the entropy value of DC's output [21] to promote the equal likelihood for all domains.

$$\max_{\theta_{enc}} L_{conf}(x_S, x_T, \theta_{DC}; \theta_{enc}) = -\sum_d P(d|enc(x)) \log P(d|enc(x)) \quad (2)$$

where $P(d|enc(x))$ is the softmax output of domain classifier, d refers to two emotion elicitation approaches considered in our work (acted/scripted vs. natural/spontaneous). Only the parameters of ENC are optimized based on this loss function.

In addition, the performance of the emotion classifier is optimized by minimizing cross entropy loss from EC's output:

$$\min_{\theta_{enc}, \theta_{EC}} L_{emo}(x, y, \theta_{enc}, \theta_{EC}) = -\sum_k \mathbb{1}[y = k] \log P_k \quad (3)$$

where the P is the softmax output of emotion classifier and k is the emotion label of the training sample $(x, y) \in (x_S, y_S)$ or (x_T, y_T) .

To combine these two objective functions together, we flip the sign of L_{conf} to do a gradient reversal and minimize the weighted overall loss sums:

$$L_{total}(x_S, y_S, x_T, y_T, \theta_{DC}; \theta_{enc}, \theta_{EC}) = L_{emo}(x, y; \theta_{enc}, \theta_{EC}) - \lambda_{conf} L_{conf}(x_S, x_T, \theta_{DC}; \theta_{enc}) \quad (4)$$

We notice the loss function (1) and (2) are in direct opposition to one another: learning a good domain classifier means

Utterance category	Description
Natural Interaction (NI)	Recording natural interaction during the breaks
Other-improvised (OI)	All the actors' turns during the improvisation sessions (not just the target sentences)
Target-improvised (TI)	Sentences with fixed lexical content during improvised dialog that expresses target emotions
Target-read (TR)	Actors read the sentences with fixed lexical content portraying the four target emotions
Whole MSP-IMPROV	Combination of all utterances above

Table 1: Emotional utterance type descriptions per emotional class in MSP-Improv dataset

the ENC output is not domain invariant, and obtaining a domain invariant ENC output can lead to a lower performance of DC. Thus, we perform the iterative update for losses (1) and (4) given the parameters from the previous iteration rather than globally optimizing all the parameters.

2.2. Softlabel loss

Although the adversarial domain confusion training tries to align the marginal distribution across different domains, the alignment of classes between each domain is not guaranteed. Thus, we further use the softlabel loss to align the source and target emotion classes, in which we try to ensure the relationship between emotion categories is preserved across source and target. This relationship originates from the human expression regardless of the elicitation approaches.

The softlabel is generated by averaging the per-category activation of source training examples using the source model. Thus, if in total we have K emotion classes, for each emotion category k , we have one softlabel $l^{(k)}$ with K dimensions to depict relative similarity across all emotion categories. The softlabel loss is defined as the cross-entropy loss between the soft activation of a target sample and the soft label corresponding to that emotion category of the sample.

$$L_{soft}(x_T, y_T, \theta_{enc}, \theta_{EC}) = - \sum_i l_i^{(y_T)} \log P_i \quad (5)$$

where P_i is the softmax output of emotion classifier with a high temperature of τ to make sure related emotion classes have enough probability mass to calculate the softlabel loss.

Lastly, we include this softlabel loss (5) to the total loss (4) to formulate the updated total loss during the training of ENC and EC as follows:

$$L_{total}(x_S, y_S, x_T, y_T, \theta_{DC}; \theta_{enc}, \theta_{EC}) = L_{emo}(x, y; \theta_{enc}, \theta_{EC}) - \lambda_{conf} L_{conf}(x_S, x_T, \theta_{DC}; \theta_{enc}) + \lambda_{soft} L_{soft}(x_T, y_T; \theta_{DC}, \theta_{enc}) \quad (6)$$

3. Data and Feature Extraction

The MSP-IMPROV [6] corpus is a multi-modality emotional database, which is designed to have a controlled data collection process over lexical content and emotion while also promoting naturalness in the recordings. It includes four types of utterances under different emotion elicitation scenarios as shown in Tables 1 and 2. Within MSP-IMPROV, utterances are classified as four types based on different elicitation scenarios: Natural Interaction (NI), Other-Improvised (OI), Target-Improvised (TI) and Target-Read (TR). To the best of our knowledge, currently, the MSP-IMPROV is the only dataset that was designed with such different emotion elicitation strategies for audio-visual emotion recognition. In addition, the ‘‘emotion elicitation approach’’ category label was provided for each utterance.

Utterance category	Angry	Sad	Happy	Neutral	Total
NI	38	66	1372	1164	2640
OI	470	633	1048	1789	3940
TI	115	106	136	283	640
TR	169	80	88	241	578

Table 2: Number of utterances per emotional class in MSP-Improv

For each utterance, we extract 40 dimensions of Mel-filter bank coefficients and energy as acoustic features. For visual features, we use an embedding of 512 dimensions from a pre-trained face recognition model, which has a similar structure of inception-ResNet [33]. The extracted embedding captures the facial information and is used for emotion recognition.

4. Experiment setup

We use utterances combined from different categories as source and target domain to verify our hypotheses, and further investigate the possibility of domain transfer learning with limited samples of target domain. In terms of the utterance content, TR and TI contain the designed scripts, while NI and OI are mostly from natural interactions and improvisation sessions. To verify our proposed hypotheses, we design the source and target split with the following combinations:

1. *Using NI and OI as source domain*

This split is used to test hypothesis (H1). Under this combination, source domain mainly contains utterances from natural or improvised conversations while for the target domain, lexical content of the utterances is fixed.

2. *Using TR and TI as source domain*

This split is used to test hypothesis (H2). Under this setting, the lexical content is fixed in the source domain while in the target domain, the lexical content can be variable.

To alleviate the emotion class imbalance issue, we utilize a weighted balanced data sampler to make sure samples from less representative categories have larger weights. In addition, within each training batch, we ensure the number of samples is balanced between the source and target domains. As we mentioned before, in a real usage scenario, the amount of labeled data of the target domain is often limited. Considering this situation, domain adaptation is performed with a very limited number of labeled target samples under our experimental setting. For samples of the target domain, we use only 10% as the training set, and 40% is used as a development set and the remaining is for evaluation.

Considering the imbalance across different emotions, we report the average recall value, which is obtained by computing the per-class classification accuracy independently and then averaging over 4 different emotion categories. For each experiment, we perform the experiments five times with exclusive training samples, and report the average 4-class emotion recognition recall value. As a comparison, we also conduct the following experiments:

Source only: the model is trained using source domain data only, and using the target domain data for evaluation.

Source + Target: The model is trained using samples both from source and target domains without any adversarial techniques.

The configuration of our model is shown in Table 3.

Training details:	Adam optimizer (lr=0.001) + L2 regularization (weight=0.01) batch size=32; epochs=100; $\tau = 2$.
ENC	Acoustic part: Conv1D(in_ch=41,out_ch=64, kernel size=10, stride=2), PReLU Conv1D(in_ch=64,out_ch=128, kernel size=5, stride=2), PReLU GRU(in_size=128, hidden_size=128, num_layers=1) Visual part: GRU(in_size=512, hidden_size=512, num_layers=1) Merge part: Linear(in=128+512, out=128), PReLU, Dropout Linear(in=128, out=128)
EC	Linear(in=128, out=32) PReLU Linear(in=32, out=10) PReLU Linear(in=10, out=4)
DC	Linear(in=128, out=32) PReLU Linear(in=32, out=10) PReLU Linear(in=10, out=2)

Table 3: Model structure and training configuration details

5. Results and Discussion

5.1. Using NI and OI as source domain

We use the NI and OI data as the source domain, and TR or TI as the target domain. Compared with TR, samples in TI have same fixed lexical content, but are recorded during an improvised dialog.

Experiment	Source Domain \rightarrow Target Domain	
	NI+OI \rightarrow TR	NI+OI \rightarrow TI
Source only	0.5005 (0.021)	0.7330 (0.020)
Source + Target	0.5837 (0.031)	0.7288 (0.034)
Domain adversarial loss only	0.6225 (0.018)	0.7002 (0.007)
Domain adversarial loss + softlabel loss	0.6339 (0.023)	0.7092 (0.022)

Table 4: Mean (std) of 4-class emotion recognition average recall value using NI+OI as source domain

The experiment results are shown in Table 4. We first notice the domain mismatch exists between NI+OI and TR: the performance of “source only” model has the lowest performance among all experiments. The “Source + Target” experiment shows that adding target domain training samples helps in model domain adaptation as expected. Even though only a few samples (10% of the samples from the target domain) are employed during training, the performance improves from 0.5005 to 0.5837 directly. More importantly, the results of our proposed model show that the domain adversarial loss can help the target domain adaptation. The softlabel loss, considering the relationships across different emotion categories, also contributes to the domain adaptation under limited target domain samples with the highest performance of 0.6339.

However, in the experiment of NI+OI \rightarrow TI, there is no improvement by adding target training samples. Surprisingly, in contrast to the previous findings, the domain adversarial loss and softlabel loss even worsen the performance. A closer inspection reveals that many samples within OI are exactly the same scripted utterances spoken in an improvised way as those in TI. As described in Table 1, OI and TI were both recorded during the improvisation sessions. We find the actors often practiced for a while and repeated scripted utterances multiple times before they were ready for the Target-improvised (TI) recording in many improvisation sessions. During dataset creation, those practicing utterances were largely kept and included in the group of OI.

Thus, OI utterances overlap with TI utterances in terms of the emotion elicitation manner and content information, which might consequently lead to the fact that adding new samples from target domain doesn’t provide extra information of target

domain. Moreover, this indistinct boundary between source and target can potentially lead to the domain confusion loss back propagating unintended bias during adaptation training, which can lead to a lower performance after adding the adversarial loss. In addition, we observe a large absolute increase of accuracy from using TR to TI as target domain (from 0.5005 to 0.7330), which is likely caused due to the existing overlap between OI and TI.

Apart from this domain overlapping issue caused by the data collection process, the results from our experiment verify the first proposed hypothesis: improvised utterances can be used as source domain to help the emotion recognition for the utterances with the scripted utterances.

5.2. Using TR and TI as source domain

We further reverse the source and target domain: TR and TI are used as the source domain and we select NI or OI as target domain. Under this setting, we are investigating whether we can apply the domain transfer learning from the acted dataset with fixed lexical content to those in natural conversation with variable lexical content. Compared with the combinations in Sec. 5.1, we have fewer samples in the source domain, but more samples in the target domain.

Experiment	Source Domain \rightarrow Target Domain	
	TI+TR \rightarrow NI	TI+TR \rightarrow OI
Source only	0.4296(0.028)	0.5499 (0.010)
Source + Target	0.4440 (0.032)	0.5886 (0.010)
Domain adversarial loss only	0.4601 (0.013)	0.5904 (0.024)
Domain adversarial loss + softlabel loss	0.4653 (0.023)	0.5920 (0.017)

Table 5: Mean (std) of 4-class emotion recognition average recall value using TI+TR as source domain

From Table 5, similarly, we notice the domain mismatch between source and target from the lower performance of “Source only” models. A better emotion classification accuracy of “Source+Target” model indicates that extra target training samples can help the model adaptation for the target domain. For both experiments, we observe the domain adversarial loss and softlabel loss can help the model to achieve better performance with limited training samples of the target domain. By comparing the results across two target domains, we notice NI has a lower classification accuracy. A reasonable explanation is that NI has a larger domain mismatch since both emotion elicitation manners and content information are different from the source.

All results show that we can improve the emotion recognition for daily natural conversation by leveraging domain adversarial loss and softlabel loss on the content fixed corpus. The results verify the hypotheses we proposed in the beginning.

6. Conclusion

The goal of this work is to focus on the domain adaptation across emotion utterances under different elicitation approaches. We proposed two hypotheses and investigated transfer learning between utterances with fixed scripted lexical content and utterances collected from spontaneous, improvised production, particularly with limited labeled target samples. Our adversarial training framework extracts domain invariant emotion representations. Moreover, we utilize the softlabel loss to take the correlations across different emotion categories into consideration. Our results confirm our proposed hypotheses and further provide new insights into elicitation approaches in emotion data collection, as well as the importance of domain adaptation in emotion recognition.

7. References

- [1] B. W. Schuller, "Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends," *Communications of the ACM*, vol. 61, no. 5, pp. 90–99, 2018.
- [2] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. W. Schuller, and S. Zafeiriou, "End-to-end multimodal emotion recognition using deep neural networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1301–1309, 2017.
- [3] P. Zhan and M. Westphal, "Speaker normalization based on frequency warping," in *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 2. IEEE, 1997, pp. 1039–1042.
- [4] B. Schuller, D. Arsic, F. Wallhoff, and G. Rigoll, "Emotion recognition in the noise applying large acoustic feature sets," in *Proc. Speech Prosody 2006, Dresden*, 2006.
- [5] C. Busso and S. S. Narayanan, "Scripted dialogs versus improvisation: Lessons learned about emotional elicitation techniques from the IEMOCAP database," in *Proceedings of InterSpeech*, Brisbane, Australia, sep 2008, pp. 1670–1673.
- [6] 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.
- [7] P. C. Petrantonakis and L. J. Hadjileontiadis, "A novel emotion elicitation index using frontal brain asymmetry for enhanced eeg-based emotion recognition," *IEEE Transactions on information technology in biomedicine*, vol. 15, no. 5, pp. 737–746, 2011.
- [8] E. Parada-Cabaleiro, G. Costantini, A. Batliner, M. Schmitt, and B. W. Schuller, "Demos: An italian emotional speech corpus," *Language Resources and Evaluation*, vol. 54, no. 2, pp. 341–383, 2020.
- [9] N. Lim, "Cultural differences in emotion: differences in emotional arousal level between the east and the west," *Integrative medicine research*, vol. 5, no. 2, pp. 105–109, 2016.
- [10] J. Zhao, R. Li, J. Liang, S. Chen, and Q. Jin, "Adversarial domain adaption for multi-cultural dimensional emotion recognition in dyadic interactions," in *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 37–45.
- [11] C.-H. Wu, J.-C. Lin, and W.-L. Wei, "Survey on audiovisual emotion recognition: databases, features, and data fusion strategies," *APSIPA transactions on signal and information processing*, vol. 3, 2014.
- [12] K. Feng and T. Chaspari, "A review of generalizable transfer learning in automatic emotion recognition," *Frontiers in Computer Science*, vol. 2, p. 9, 2020.
- [13] B. J. Abbaschian, D. Sierra-Sosa, and A. Elmaghraby, "Deep learning techniques for speech emotion recognition, from databases to models," *Sensors*, vol. 21, no. 4, p. 1249, 2021.
- [14] R. Milner, M. A. Jalal, R. W. Ng, and T. Hain, "A cross-corpus study on speech emotion recognition," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 304–311.
- [15] P. Song, "Transfer linear subspace learning for cross-corpus speech emotion recognition," *IEEE Transactions on Affective Computing*, vol. 10, no. 2, pp. 265–275, 2019.
- [16] S. Latif, R. Rana, S. Younis, J. Qadir, and J. Epps, "Transfer learning for improving speech emotion classification accuracy," *arXiv preprint arXiv:1801.06353*, 2018.
- [17] J. Gideon, M. McInnis, and E. M. Provost, "Improving cross-corpus speech emotion recognition with adversarial discriminative domain generalization (addog)," *IEEE Transactions on Affective Computing*, 2019.
- [18] S.-w. Lee, "Domain generalization with triplet network for cross-corpus speech emotion recognition," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 389–396.
- [19] B.-C. Chiou and C.-P. Chen, "Speech emotion recognition with cross-lingual databases," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [20] S. Latif, J. Qadir, and M. Bilal, "Unsupervised adversarial domain adaptation for cross-lingual speech emotion recognition," in *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2019, pp. 732–737.
- [21] 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)*. IEEE, 2020, pp. 7144–7148.
- [22] Y. Yin, B. Huang, Y. Wu, and M. Soleymani, "Speaker-invariant adversarial domain adaptation for emotion recognition," in *Proceedings of the 2020 International Conference on Multimodal Interaction*, 2020, pp. 481–490.
- [23] S. Mariooryad and C. Busso, "Exploring cross-modality affective reactions for audiovisual emotion recognition," *IEEE Transactions on affective computing*, vol. 4, no. 2, pp. 183–196, 2013.
- [24] G.-Y. Chao, C.-M. Chang, J.-L. Li, Y.-T. Wu, and C.-C. Lee, "Generating fmri-enriched acoustic vectors using a cross-modality adversarial network for emotion recognition," in *Proceedings of the 20th ACM International Conference on Multimodal Interaction*, 2018, pp. 55–62.
- [25] L.-P. Morency, R. Mihalcea, and P. Doshi, "Towards multimodal sentiment analysis: Harvesting opinions from the web," in *Proceedings of the 13th international conference on multimodal interfaces*, 2011, pp. 169–176.
- [26] C. M. Lee and S. S. Narayanan, "Toward detecting emotions in spoken dialogs," *IEEE transactions on speech and audio processing*, vol. 13, no. 2, pp. 293–303, 2005.
- [27] L. Devillers and L. Vidrascu, "Real-life emotions detection with lexical and paralinguistic cues on human-human call center dialogs," in *Ninth International Conference on Spoken Language Processing*, 2006.
- [28] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [29] 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.
- [30] R. Peri, H. Li, K. Somandepalli, A. Jati, and S. Narayanan, "An empirical analysis of information encoded in disentangled neural speaker representations," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 194–201.
- [31] 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*, 2015, pp. 4068–4076.
- [32] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [33] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.