



# Attention-based Interactive Disentangling Network for Instance-level Emotional Voice Conversion

Yun Chen<sup>1</sup>, Lingxiao Yang<sup>1</sup>, Qi Chen<sup>1</sup>, Jian-Huang Lai<sup>1,2,3</sup> and Xiaohua Xie<sup>1,2,3\*</sup>

<sup>1</sup>School of Computer Science and Engineering, Sun Yat-Sen University, China

<sup>2</sup>Guangdong Province Key Laboratory of Information Security Technology, China

<sup>3</sup>Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

{cheny2259, chenq377}@mail2.sysu.edu.cn, {yanglx9, stsljh, xiexiaoh6}@mail.sysu.edu.cn

## Abstract

Emotional Voice Conversion aims to manipulate a speech according to a given emotion while preserving non-emotion components. Existing approaches cannot well express fine-grained emotional attributes. In this paper, we propose an *Attention-based Interactive Disentangling Network (AINN)* that leverages instance-wise emotional knowledge for voice conversion. We introduce a two-stage pipeline to effectively train our network: Stage I utilizes inter-speech contrastive learning to model fine-grained emotion and intra-speech disentanglement learning to better separate emotion and content. In Stage II, we propose to regularize the conversion with a multi-view consistency mechanism. This technique helps us transfer fine-grained emotion and maintain speech content. Extensive experiments show that our AINN outperforms state-of-the-arts in both objective and subjective metrics.

**Index Terms:** emotional voice conversion, instance-level, feature disentanglement, strength learning

## 1. Introduction

Emotional Voice Conversion (EVC) aims to transfer the emotion of a speech while preserving emotion-independent components. EVC has garnered much attention due to its potential applications in areas such as virtual assistants [1], voice cloning [2], and human-computer interaction [3]. The recent advancement in EVC is driven by the use of Generative Adversarial Networks (GANs) [4, 5, 6], Autoencoder [7, 8, 9, 10, 11], and Sequence-to-sequence model [12, 13, 14]. In general, most of existing EVC methods take an emotional domain descriptor  $\phi$  (emotion label or a pre-defined emotion embedding) as a condition to guide the input speech  $\mathbf{a}_s$  for emotion conversion (Fig. 1(a)). However, since emotions have varying degrees in speech, it is difficult to capture fine-grained emotions using only the domain descriptor. We define this task as **Domain-level emotional voice conversion (D-EVC)**:  $\tilde{\mathbf{a}}_s = \mathcal{M}(\mathbf{a}_s, \phi)$ .

To achieve the fine-grained expression, we propose to extract instance-level emotional representation (e.g. emotional category, strength) from a reference speech  $\mathbf{a}_r$ , and then incorporate it with the emotion-independent components of source speech  $\mathbf{a}_s$ :  $\tilde{\mathbf{a}}_s = \mathcal{M}(\mathbf{a}_s, \mathbf{a}_r)$ . We refer to this task as **Instance-level emotional voice conversion (I-EVC)** (Fig. 1(b)). Due to the representation gap between the emotional domain and speech instance, I-EVC faces two challenges: 1) Unlike D-EVC which provides pure emotional category, I-EVC needs to capture the emotion-related information from a reference speech and avoid the interference of emotion-irrelevant components. 2) Unlike the well-defined emotion category, emotional strength is continuous and highly subjective, making it

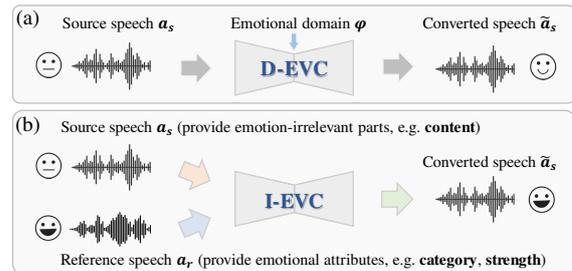


Figure 1: Illustration of (a) Domain-level emotional voice conversion, and (b) Instance-level emotional voice conversion.

hard to be modeled without supervised signals. StarGANv2-VC [15] achieves instance-level conversion by involving a reference speech for style representation. Nevertheless, this approach lacks strength-related modeling, resulting in uncontrolled strength variations in the converted speech. Recently, Emovox [16] extracts strength value by the relative function [17] and uses it to help strength transfer. However, it does not explicitly establish the relationship between pre-calculated strength and high-level emotional speech features, limiting the strength conversion. Besides, Emovox requires additional text transcriptions to disentangle the emotional feature, which is less flexible for low-resourced languages [18].

To address instance-level emotional voice conversion, we propose an *Attention-based Interactive Disentangling Network (AINN)* with a two-stage training pipeline. The first stage aims to disentangle emotion and emotion-irrelevant features under the auto-encoder paradigm. Specifically, we contrastively learn the relation of inter-speech to model emotional strength and construct complete emotion representation. Meanwhile, we establish an attention-based method to effectively decompose the emotion and content of a speech. In the second stage, multi-view consistency is applied to enable the model to learn more task-adaptive knowledge. Our method ensures that the speech content remains nearly unchanged, while the fine-grained emotion is improved. The main contributions of this paper include:

- We propose an Attention-based Interactive Disentangling Network (AINN) with an efficient two-stage training strategy for instance-level emotional voice conversion.
- We propose a text-free feature disentanglement method that utilizes attention-based emotional cues to better decompose emotion and content.
- We incorporate contrastive learning to effectively model emotional strength, and leverage it to improve speech consistency on fine-grained emotion.
- Extensive experiments demonstrate the effectiveness of our method in both objective and subjective metrics.

\*Corresponding Author

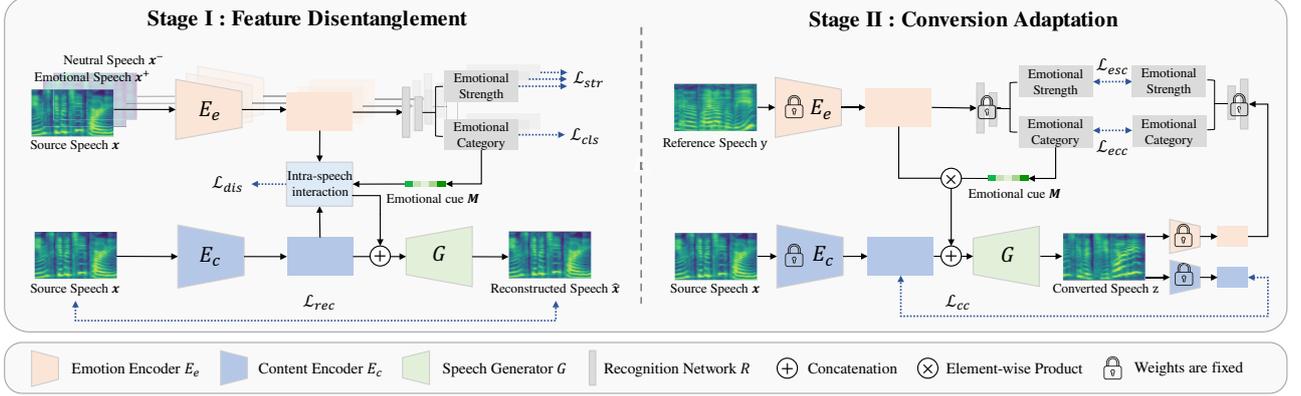


Figure 2: The framework of the proposed method. It consists of two stages: Feature disentanglement and Conversion adaptation.

## 2. Method

Given a source speech, our goal is to generate a new speech, where we can adjust the emotional category and strength of it towards a reference speech, and ensure its content remains unchanged. We expect the emotional attributes to be only guided by reference speech and the content aligns with the source speech. To this end, we design a two-stage training pipeline: *Feature disentanglement* and *Conversion adaptation*. The overview of our framework is shown in Fig. 2. In Stage I, We train the whole network to disentangle the emotional and content features from the input speech mel-spectrogram in an auto-encoder manner. Specifically, we introduce an *inter-speech interaction* method to model instance-level emotion and an *intra-speech interaction* method to enforce independent content and calibrated emotion. Stage II only fine-tunes the generator to fuse two parts of features and synthesize reasonable speech in the conversion progress. It is achieved by *multi-view consistency* including emotional category, strength, and content.

### 2.1. Stage I: Feature disentanglement

#### 2.1.1. Inter-speech interactive emotion representation

Given a source speech mel-spectrogram  $\mathbf{x}$ , we can obtain an emotional feature  $\mathbf{e} = E_e(\mathbf{x})$ . To make the emotion encoder  $E_e$  capture rich fine-grained emotional features, we feed  $\mathbf{e}$  into the two-branch recognition network  $R = \{R_c, R_s\}$  and force it to achieve *emotion classification* and *strength assessment*.

**Emotional classification.** For  $K$  emotional category, we adopt widely used cross-entropy loss, which is defined as follows:

$$\mathcal{L}_{cls} = \mathbb{E}_{\mathbf{x}} \left[ - \sum_{k=1}^K \mathbf{p}_k \log(\mathbf{q}_k) \right], \quad (1)$$

where  $\mathbf{p}$  represents the one-hot ground truth and  $\mathbf{q} = R_c(\mathbf{e})$  indicates the classification probability distribution.

**Strength assessment.** Due to lack of annotations, it is difficult to directly apply supervision to emotional strength. Therefore, we adopt a self-supervised approach to assess strength by mining relative relations [17] between speeches. In previous works [16, 19, 20], neutral speech is often treated as non-emotion speech, and emotional strength is considered as relative difference between neutral speech and emotional speech. Based on that prior, we build relative relations with two hypotheses: 1) The strength of neutral speech approaches zero, and the strength of emotional speech is significantly higher than that of neutral

speech. 2) The strength difference between emotional speeches is small. Therefore, we model the relative relation as a triplet  $\{\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-\}$ , where the positive pair  $\{\mathbf{x}, \mathbf{x}^+\}$  shares the same emotional category, and the negative pair  $\{\mathbf{x}, \mathbf{x}^-\}$  contains an emotional speech mel-spectrogram  $\mathbf{x}$  and a neutral speech mel-spectrogram  $\mathbf{x}^-$ . We set the strength of  $\mathbf{x}^-$  to zero, then the triplet contrastive loss is applied to narrow the strength difference of the positive pair, and enlarge that of the negative pair:

$$\begin{aligned} \mathcal{L}_{str} = \mathbb{E}_{\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-} [ & \max\{|s^x - s^{x^+}|, \delta_1\} \\ & + \max\{1 - (s^x - s^{x^-}), \delta_2\}], \end{aligned} \quad (2)$$

where  $s = R_s(\mathbf{e})$  is emotional strength, and  $s^x, s^{x^+}, s^{x^-}$  are the strengths of  $\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-$ , respectively. Two margins  $\delta_1$  and  $\delta_2$  are employed to prevent the network from overfitting to extreme strength. The incorporation of strength learning allows the network to acquire knowledge related to strength, enhancing personalized information in emotional feature. This learned network is also used to predict reasonable strength values for consistency learning in Stage II (Sec. 2.2).

#### 2.1.2. Intra-speech interactive disentanglement

To achieve effective disentanglement over emotional and content features, it is crucial to identify meaningful frames in a speech. These frames should have two attributes: 1) They should not be silent frames as they do not contain useful information. 2) The emotion expressed in these frames should be clearly perceivable so that the network can focus on the significant difference between emotional and content features. In this section, we first introduce an attention mechanism that helps locate these meaningful frames in the speech. And then, the feature disentanglement can be realized by minimizing the similarity between emotional feature  $\mathbf{e} = E_e(\mathbf{x})$  and content feature  $\mathbf{c} = E_c(\mathbf{x})$ . This process is shown in Fig. 3.

Firstly, inspired by the idea of class activation mapping (CAM) [21, 22], we estimate an emotional cue  $\mathbf{M}$  to indicate the meaningful frames:

$$\mathbf{M} = \text{ReLU}(\boldsymbol{\theta}^T \cdot \mathbf{e}), \quad (3)$$

where  $\boldsymbol{\theta}$  is the classifier weight of  $R_c$ , relating to the corresponding emotional category.  $\mathbf{M}$  is further normalized to  $[0, 1]$  by the maximum value along the temporal axes. Since  $\mathbf{M}$  is class-aware attention, it naturally filters silent frames that are not related to emotion, and concentrates on meaningful frames.

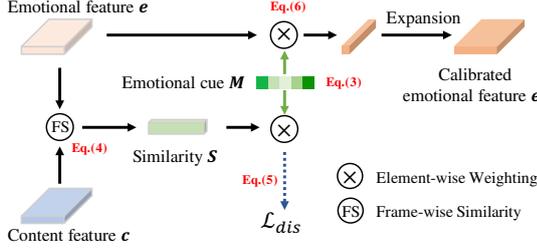


Figure 3: Illustration of Intra-speech interactive method.

Secondly, we calculate the frame-wise similarity  $\mathcal{S}$  for emotional feature  $e$  and content feature  $c$ :

$$\mathcal{S}_t = \frac{|e_t \cdot c_t|}{\|e_t\|_2 \cdot \|c_t\|_2}, \quad (4)$$

where  $t \in [0, T]$  indexes the  $t$ -th frame of the feature. Thus the feature disentanglement loss is defined as:

$$\mathcal{L}_{dis} = \mathbb{E}_{\mathbf{x}} \left[ \frac{1}{T} \sum_{t=0}^T M_t \cdot \mathcal{S}_t \right]. \quad (5)$$

Since emotional cues give high responses on frames with the discriminative emotional tone, we calibrate the emotional feature by emotional cues:

$$\mathbf{e}^* = \frac{1}{T} \sum_{t=0}^T M_t \cdot e_t. \quad (6)$$

To train the network, we employ the auto-encoder manner to reconstruct the speech mel-spectrogram  $\mathbf{x}$  by a generator  $G$  and define a reconstruction loss as:

$$\mathcal{L}_{rec} = \mathbb{E}_{\mathbf{x}} [\|G(\mathbf{c}, \mathbf{e}^*) - \mathbf{x}\|_1]. \quad (7)$$

**Full objectives.** The total loss of Stage I is as follows:

$$\mathcal{L}_{stage1} = \mathcal{L}_{rec} + \mathcal{L}_{cls} + \lambda_{dis} \mathcal{L}_{dis} + \lambda_{str} \mathcal{L}_{str}, \quad (8)$$

where  $\lambda_{dis}$  and  $\lambda_{str}$  are hyperparameters for balancing losses.

## 2.2. Stage II: Conversion adaptation

Given the fixed emotion encoder  $E_e$ , content encoder  $E_c$ , and recognition network  $R$  pre-trained in Stage I, our second stage is to fine-tune only the generator  $G$  to adapt to conversion using a style transfer paradigm [23, 24]. Owing to the pre-trained networks that provide a good feature disentanglement for speech, our framework incorporates a multi-view consistency using well-disentangled features, improving the expression of fine-grained emotions and preserving the integrity of speech content. As shown in the right part of Fig. 2, our network takes a source speech mel-spectrogram  $\mathbf{x}$  and a reference speech mel-spectrogram  $\mathbf{y}$  as input. The network extracts their content feature  $E_c(\mathbf{x})$  and the calibrated emotional feature  $E_e(\mathbf{y})^*$ . Then the generator  $G$  produces the converted mel-spectrogram  $\mathbf{z} = G(E_c(\mathbf{x}), E_e(\mathbf{y})^*)$ . We constrain  $\mathbf{z}$  to be similar to  $\mathbf{x}$  and  $\mathbf{y}$  in terms of content and emotion, respectively.

**Content consistency.** The converted speech mel-spectrogram  $\mathbf{z}$  is required to preserve the content feature of  $\mathbf{x}$ . Following the commonly used constraint in voice conversion [25], we employ a consistency loss  $\mathcal{L}_{cc}$  to prevent the content from being lost:

$$\mathcal{L}_{cc} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\|E_c(\mathbf{z}) - E_c(\mathbf{x})\|_1]. \quad (9)$$

**Emotion consistency.** Moreover, compared to content information, the emotional feature tends to be more global in a speech. Therefore, we apply two consistency losses to regulate the emotional attributes from a global perspective.

First, we choose KL-divergence as emotional category consistency loss  $\mathcal{L}_{ecc}$  to align two distributions:

$$\mathcal{L}_{ecc} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left[ - \sum R_c(E_e(\mathbf{y})) \log \frac{R_c(E_e(\mathbf{z}))}{R_c(E_e(\mathbf{y}))} \right]. \quad (10)$$

In addition, an MSE-based emotional strength consistency loss  $\mathcal{L}_{esc}$  is adopted to regress the assessed value, shown as:

$$\mathcal{L}_{esc} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\|R_s(E_e(\mathbf{z})) - R_s(E_e(\mathbf{y}))\|_2^2], \quad (11)$$

**Full objectives.** The total loss of Stage II is defined as:

$$\mathcal{L}_{stage2} = \mathcal{L}_{rec} + \lambda_c \mathcal{L}_{cc} + \lambda_c \mathcal{L}_{ecc} + \lambda_c \mathcal{L}_{esc}. \quad (12)$$

where  $\mathcal{L}_{rec}$  is the reconstruction loss to ensure the stability of converted speech.  $\lambda_c$  is a hyperparameter for balancing loss.

## 3. Experiments

### 3.1. Experimental Setup

**Dataset.** We train our model on 10 English speakers from the Emotional Speech Dataset (ESD) [11] with five emotions. Each speaker includes 300 training, 20 validation, and 30 test speeches. In our experiments, all speech data is utilized in the form of an 80-dimensional mel-spectrogram extracted from the audio with a sampling rate of 16kHz. We also use a pre-trained HiFi-GAN [27] to synthesize audio from mel-spectrogram.

**Data preparation.** To adapt to the I-EVC task, we construct source-reference speech pairs as input. The source and reference speeches are sampled from different emotional categories of the same speaker. We create 20 conversion types between five emotions, such as neutral to happy and sad to angry. For each conversion type, we randomly sample 750 speech pairs for training, 50 pairs for validation, and 15 pairs for testing.

**Network architecture.** The architecture of the content encoder and generator are based on AutoVC [25], and the emotion encoder consists of six convolution layers and an LSTM layer. The recognition network is a  $1 \times 1$  convolution layer. More details are given in the supplementary material.

**Implementation details.** We adopt Adam optimizer with a learning rate of  $1e-4$  for Stage I and  $1e-5$  for Stage II. The network is trained with a batch size of 64 for 50000 iterations in both stages. The total training time is about nine hours on an NVIDIA A100 GPU. For the loss functions in Eq. (8) and Eq. (12), we set  $\lambda_{dis} = 0.2$ ,  $\lambda_{str} = 1$ ,  $\lambda_c = 0.0002$ . To avoid strength collapsing, those two margins in Eq. (2) are set to 0.5. Our code and converted samples are available at <https://ainn-evc.github.io/>.

### 3.2. Comparison with previous methods

Our method is compared with three state-of-the-arts methods, including domain-level EmotionalStarGAN [5] and Seq2Seq-EVC [14], and instance-level Emovox [16].

**Objective Evaluation.** Following prior works [5, 11], we adopt Mel-cepstral distortion (MCD) [28] as a basic evaluation metric for evaluating content preservation. Besides, we use a pre-trained emotion classifier ACRNN [29] to evaluate the accuracy ( $\text{ACC}_{cls}$ ) of converted speech. Furthermore, we evaluate the RMSE of strength between reference speech and converted speech. The strength value is obtained by our strength assess-

Table 1: *Quantitative comparisons of converted speech with previous methods. The best results are shown in bold. The \* denotes methods with text supervision and pretrained on VCTK dataset [26].*

Methods	Type	Objectiveness			Subjectiveness	
		MCD ↓	ACC <sub>cls</sub> ↑	RMSE <sub>str</sub> ↓	Naturalness ↑	Similarity ↑
EmotionalStarGAN [5]	Domain-level	5.219	0.743	0.312	2.78 ± 0.56	48.37%
Seq2Seq-EVC* [14]	Domain-level	5.323	0.550	0.215	2.62 ± 0.54	55.82%
Emovox* [16]	Instance-level	5.088	0.460	0.199	2.96 ± 0.53	47.67%
AINN (Ours)	Instance-level	<b>4.596</b>	<b>0.830</b>	<b>0.117</b>	<b>3.43 ± 0.59</b>	<b>76.12%</b>

Table 2: *Ablation study of the methods.*

Methods	MCD ↓	ACC <sub>cls</sub> ↑	RMSE <sub>str</sub> ↓
baseline ( $\mathcal{L}_{rec} + \mathcal{L}_{cls}$ )	4.922	0.027	0.446
+ Stage I (w/o EC)	4.766	0.323	0.283
+ Stage I (w/ EC)	<b>4.493</b>	0.647	0.146
+ Stage I & II (Full)	4.596	<b>0.830</b>	<b>0.117</b>

ment network  $R_s$ . Note that the strength assessment network is frozen in Stage II, it would not be impacted by our converted speech, thereby a fair comparison can be drawn. As shown in Tab. 1, the proposed AINN achieves the best score in MCD, which indicates that our model effectively transforms emotion while maintaining the content. Similarly, our AINN improves the classification accuracy by up to 83%, outperforming all compared methods. For RMSE<sub>str</sub>, Emovox obtains the best value 0.199 among all compared methods. Our method, however, largely outperforms Emovox due to the proposed instance-level emotion modeling and multi-view consistency.

**Subjective Evaluation.** We further conduct a user study to evaluate the naturalness and emotional similarity of converted speech. Following previous works [11, 18, 30], we recruit 15 subjects from the social network and they are proficient in English. Each of them is asked to listen to 70 converted speeches. From Tab. 1, we note that the proposed method receives notable preference for both emotion similarity and overall performance.

### 3.3. Ablation Study

We conduct ablation experiments to verify the contribution of each component in our method. As shown in Tab. 2, we define the model trained using only reconstruction and classification losses as the baseline method. Compared to the baseline, the introduction of the feature disentanglement in Stage I lead to significant improvements across all evaluation metrics. Furthermore, the use of emotional cues (EC) largely boosts the model’s performance, with the ACC<sub>cls</sub> increasing from 0.323 to 0.647 and RMSE<sub>str</sub> decreasing to 0.146. (The emotional cues are visualized in the supplementary material.) This indicates that the attention mechanism successfully captures emotion-related information, resulting in better emotion representation. Our full model with two stages increases the classification accuracy and narrows the RMSE of strength a lot with the contribution of multi-view consistency. These results demonstrate that the two-stage training strategy with the well-designed disentanglement method and multi-view consistency is highly effective.

### 3.4. Feature Disentanglement

Following prior work [18], we conduct emotion classification on the content feature of source speech to demonstrate the effect of feature disentanglement. As shown in Tab. 3, we observe a decline in the classification accuracy of content feature from 56.3% to 34.3% with the incorporation of the proposed methods. This proves the content feature does not contain emotion-related information. By contrast, the classification accuracy of emotional feature remains consistently high, confirming that

Table 3: *Emotion classification accuracy on content feature and emotional feature of source speech.*

Methods	Content feature ↓	Emotional feature ↑
baseline	0.563	0.903
Stage I (w/o EC)	0.395	0.904
Stage I (Full)	<b>0.343</b>	<b>0.915</b>

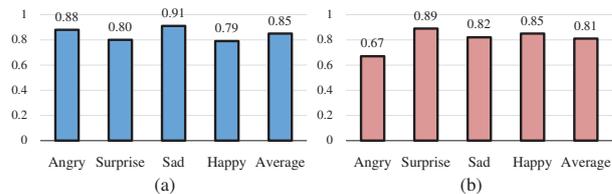


Figure 4: *The strength accuracy of (a) source speech and (b) converted speech.*

emotional information is preserved in feature disentanglement.

### 3.5. Strength Learning

We perform a subjective experiment to investigate the effectiveness of the proposed strength learning method. We create two groups of speech data, one including 20 pairs of real speeches and the other including 20 pairs of converted speeches. Each pair contains two emotional speeches sharing the same emotion with varying emotional strengths. Subjects are asked to identify the more expressive speech in each pair. The results show that the strength accuracy achieves 85% in the real speech (Fig. 4 (a)) and 81% in the converted speech (Fig. 4 (b)). This consistently high performance provides strong evidence for the correctness of our strength assessment. The pitch contour of converted speech is shown in the supplementary material. Furthermore, we conduct an objective comparison between our full model and the model trained without strength learning. In this case, the model without strength learning degrades to domain-level EVC, resulting in a slight increase (+0.027) in RMSE<sub>str</sub>. In summary, the effectiveness of strength learning is demonstrated from both subjective and objective aspects.

## 4. Conclusions

We propose an *Attention-based Interactive disEntangling Network (AINN)* and a two-stage training pipeline for instance-level emotional voice conversion. Our method effectively models emotional strength, and disentangles emotional and content features with the help of emotional cues. Moreover, we introduce multi-view consistency to enhance fine-grained emotions and preserve the integrity of speech content. Experimental results achieve state-of-the-art in all metrics. Future work includes improving the quality of converted speech and achieving cross-speaker instance-level emotional voice conversion.

**Acknowledgements.** This project is supported by Guangdong-Hong Kong-Macao Greater Bay Area International Science and Technology Innovation Cooperation Project (No.2021A0505030080).

## 5. References

- [1] M. Elgaar, J. Park, and S. W. Lee, "Multi-speaker and multi-domain emotional voice conversion using factorized hierarchical variational autoencoder," in *ICASSP*. IEEE, 2020, pp. 7769–7773.
- [2] Q. Chen, M. Tan, Y. Qi, J. Zhou, Y. Li, and Q. Wu, "V2c: Visual voice cloning," in *CVPR*, 2022, pp. 21 242–21 251.
- [3] K. Zhou, B. Sisman, R. Liu, and H. Li, "Emotional voice conversion: Theory, databases and esd," *Speech Commun.*, vol. 137, pp. 1–18, 2022.
- [4] K. Zhou, B. Sisman, and H. Li, "Transforming Spectrum and Prosody for Emotional Voice Conversion with Non-Parallel Training Data," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 230–237.
- [5] G. Rizos, A. Baird, M. Elliott, and B. Schuller, "Stargan for emotional speech conversion: Validated by data augmentation of end-to-end emotion recognition," in *ICASSP*. IEEE, 2020, pp. 3502–3506.
- [6] R. Shankar, J. Sager, and A. Venkataraman, "Non-parallel emotion conversion using a deep-generative hybrid network and an adversarial pair discriminator," in *Interspeech*, 2020, pp. 3396–3400.
- [7] J. Gao, D. Chakraborty, H. Tembine, and O. Olaleye, "Nonparallel emotional speech conversion," in *Interspeech*, 2019, pp. 2858–2862.
- [8] K. Zhou, B. Sisman, M. Zhang, and H. Li, "Converting anyone's emotion: Towards speaker-independent emotional voice conversion," in *Interspeech*, 2020, pp. 3416–3420.
- [9] Y. Cao, Z. Liu, M. Chen, J. Ma, S. Wang, and J. Xiao, "Nonparallel emotional speech conversion using vae-gan," in *Interspeech*, 2020, pp. 3406–3410.
- [10] K. Zhou, B. Sisman, and H. Li, "Vaw-gan for disentanglement and recombination of emotional elements in speech," in *SLT*. IEEE, 2021, pp. 415–422.
- [11] 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*. IEEE, 2021, pp. 920–924.
- [12] C. Robinson, N. Obin, and A. Roebel, "Sequence-to-sequence modelling of f0 for speech emotion conversion," in *ICASSP*. IEEE, 2019, pp. 6830–6834.
- [13] T.-H. Kim, S. Cho, S. Choi, S. Park, and S.-Y. Lee, "Emotional voice conversion using multitask learning with text-to-speech," in *ICASSP*. IEEE, 2020, pp. 7774–7778.
- [14] K. Zhou, B. Sisman, and H. Li, "Limited Data Emotional Voice Conversion Leveraging Text-to-Speech: Two-Stage Sequence-to-Sequence Training," in *Interspeech*, 2021, pp. 811–815.
- [15] Y. A. Li, A. Zare, and N. Mesgarani, "Starganv2-vc: A diverse, unsupervised, non-parallel framework for natural-sounding voice conversion," in *Interspeech*, 2021, pp. 1349–1353.
- [16] K. Zhou, B. Sisman, R. Rana, B. W. Schuller, and H. Li, "Emotion intensity and its control for emotional voice conversion," *IEEE Trans. Affect. Comput.*, pp. 1–18, 2022.
- [17] D. Parikh and K. Grauman, "Relative attributes," in *ICCV*. IEEE, 2011, pp. 503–510.
- [18] S.-H. Lee, J.-H. Kim, H. Chung, and S.-W. Lee, "Voicemixer: Adversarial voice style mixup," in *NeurIPS*, vol. 34, 2021, pp. 294–308.
- [19] 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.
- [20] T. Li, X. Wang, Q. Xie, Z. Wang, and L. Xie, "Cross-speaker emotion disentangling and transfer for end-to-end speech synthesis," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 30, pp. 1448–1460, 2022.
- [21] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *CVPR*, 2016, pp. 2921–2929.
- [22] Q. Chen, L. Yang, J.-H. Lai, and X. Xie, "Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation," in *CVPR*, 2022, pp. 4288–4298.
- [23] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *ECCV*. Springer, 2016, pp. 694–711.
- [24] P. Zhang, L. Yang, J.-H. Lai, and X. Xie, "Exploring dual-task correlation for pose guided person image generation," in *CVPR*, 2022, pp. 7713–7722.
- [25] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "Autovc: Zero-shot voice style transfer with only autoencoder loss," in *ICML*. PMLR, 2019, pp. 5210–5219.
- [26] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," *University of Edinburgh. The Centre for Speech Technology Research (CSTR)*, 2017.
- [27] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS*, vol. 33, 2020, pp. 17 022–17 033.
- [28] R. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," in *PACRIM*, vol. 1. IEEE, 1993, pp. 125–128.
- [29] 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, 2018.
- [30] C.-B. Im, S.-H. Lee, S.-B. Kim, and S.-W. Lee, "Emoq-tts: Emotion intensity quantization for fine-grained controllable emotional text-to-speech," in *ICASSP*. IEEE, 2022, pp. 6317–6321.