

Mixed-EVC: Mixed Emotion Synthesis and Control in Voice Conversion

Kun Zhou¹, Berrak Sisman², Carlos Busso², Bin Ma¹, Haizhou Li^{3,4}

¹ Institute of Intelligent Computing, Alibaba Group, Singapore

² The University of Texas at Dallas, United States of America

³ The Chinese University of Hong Kong, Shenzhen, China

⁴ National University of Singapore, Singapore

Abstract

Emotional voice conversion (EVC) traditionally targets the transformation of spoken utterances from one emotional state to another, with previous research mainly focusing on discrete emotion categories. We propose a novel EVC framework, Mixed-EVC, which only leverages discrete emotion training labels. We construct an attribute vector that encodes the relationships among these discrete emotions, which is predicted using a ranking-based support vector machine and then integrated into a sequence-to-sequence (seq2seq) EVC framework. Mixed-EVC not only learns to characterize the input emotional style but also quantifies its relevance to other emotions during training. As a result, users have the ability to assign these attributes to achieve their desired rendering of mixed emotions. Objective and subjective evaluations confirm the effectiveness of our approach in terms of mixed emotion synthesis and control while surpassing traditional baselines in the conversion between discrete emotions.

1. Introduction

Human speech often encompasses a blend of emotions, resulting in the emergence of complex emotional expressions, as evidenced in prior studies [1, 2, 3]. Emotional voice conversion (EVC) aims to manipulate the emotional state of a spoken utterance while keeping speaker identity and linguistic content unchanged [4]. This paper represents a progressive step in the field of EVC, with a unique focus on infusing a quantifiable mixed emotion rendering into a human voice. The primary objective is to enhance the naturalness of human-computer interactions [5], for example, enriching the emotional responses within a dialogue system [6, 7].

EVC poses unique challenges due to the complex structure of emotions [8]. People use different words to describe the emotions that they feel, showing that there are nearly 34,000 distinct emotions that a human may experience [9]. To understand how these emotions correlate with each other, scientists analyze them in a valence-arousal space [8]. The evidence for the valence-arousal view comes from statistical analysis of how people report feelings [10]. However, the analysis from a valence-arousal view has not always been able to tell us the real difference between emotions [11]. Plutchik’s emotion wheel [12] provides a more straightforward way to describe emotions. 8 primary emotions: anger, fear, sadness, disgust, surprise, anticipation, trust, and joy, are arranged in an emotion wheel. The diverse amount of emotions could be produced either by changing the intensity or by adding up the primary emotions. Although preliminary studies [13, 14] have explored synthesizing mixed emotions for text-to-speech systems, we observe a



Figure 1: An illustration of mixed emotion synthesis: nuanced emotion types (Excited, Outrage and Disappointment) could be regarded as the mixed states of primary emotions (Happy, Sad, Angry and Surprise) [12].

lack of study on mixed emotion synthesis in the literature of EVC, with existing studies mostly focusing on the conversion between discrete emotions. In this research, we draw inspiration from the emotion wheel theory and introduce an approach that employs voice conversion techniques to manipulate human emotions into a mixed emotional state.

Speech emotions are inherently supra-segmental and intricate, complex with multiple acoustic cues such as speech quality, pitch, energy, and speaking rate [15]. Addressing these complexities in EVC calls for the modeling of both spectral and prosodic variations at the same time, leading to the study of sequence-to-sequence (seq2seq) architecture for EVC [16, 17, 18, 19]. To learn emotion information, existing EVC frameworks mostly leverage pre-defined discrete emotion labels as the supervision to emotion training, either learning a translation model between emotion pairs [20, 21, 22, 23, 24], or disentangling emotional elements with auto-encoders [25, 26, 27]. Emotion conversion can be achieved by assigning an emotion label or transferring from an emotional speech. These methods restrict learning richer style descriptions of emotions but rather produce a stereotypical emotional pattern [28]. Consequently, they confine emotions within specific categories, posing challenges when it comes to examining the connections between different emotional states, encompassing the entirety of human emotions, and creating a mixed emotional profile. This paper aims to fill these gaps.

This paper presents the first investigation of mixed emo-

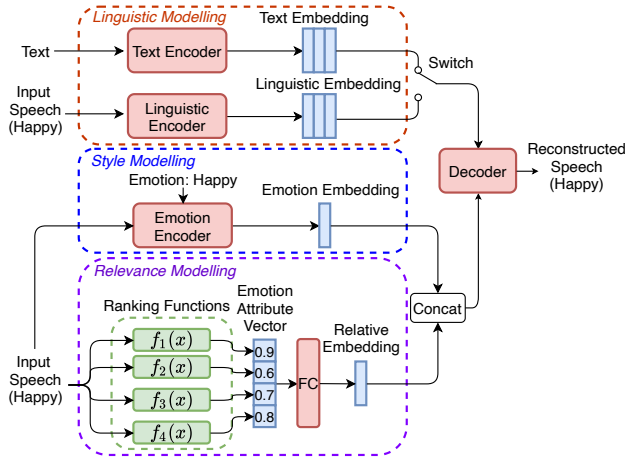


Figure 2: The training diagram of Mixed-EVC, where the red boxes represent the modules that are involved in the training. The pre-trained ranking functions automatically predict emotion attributes that indicate the degree of relevance between the input emotional style (Happy) and other emotional styles (Angry, Sad, Surprise and Neutral).

tion synthesis and control for EVC, denoted as “Mixed-EVC”, which aims to address two challenges: (1) how to describe and quantify the combination of emotions; and (2) how to assess the produced mixed outcomes. Mixed-EVC introduces a novel approach to explicitly quantify and encode the relationships among discrete emotions into an attribute vector and distinguishes itself by infusing diverse emotional behaviors into the human voice leveraging limited discrete emotion labels. Mixed-EVC also allows users to intuitively and quantifiably control the emotion rendering through categorical classes, offering a more user-friendly alternative than manipulating continuous emotional attributes. Our key contributions can be outlined as follows:

- We propose a novel formulation to quantify the mixture of emotions. We construct a ranking function for each pair of emotions, where the ranking values represent the degree of relevance with respect to an emotion;
- During training, the attribute can be precisely predicted by the ranking function, guiding the decoder to quantify the relevance between the input emotional style and all other emotions. At run-time, the users can define those attributes to generate various emotional mixtures;
- We design evaluation metrics to assess the effectiveness of our approach in terms of synthesizing mixed emotions and enabling precise control. We also present an interactive demo to the readers, showing the capability of our model in terms of mixed emotion synthesis and control.

The rest of the paper is organized as follows: We motivate our work by introducing the background in Section 2. We then introduce our proposed method in Section 3. In Section 4, we report our experimental results, and introduce our interactive demo in Section 5. Section 6 concludes the paper.

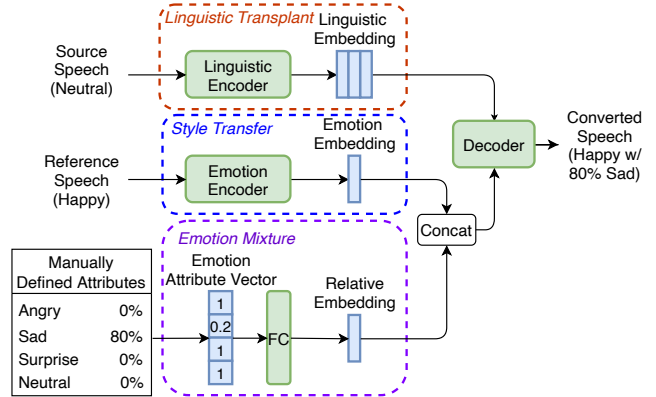


Figure 3: The run-time diagram of the proposed Mixed-EVC, where the green boxes represent pre-trained modules. By assigning manually defined attributes, Mixed-EVC transfers the source emotion to a target emotional mixture while preserving the source linguistic content.

2. From Emotion Wheel to Mixed Emotion Synthesis

The Emotion Wheel [12], proposed by Robert Plutchik in 1980, illustrates the vast range of human emotions and their inter-relationships. It suggests that combinations of eight primary emotions—joy, trust, fear, surprise, sadness, anticipation, anger, and disgust—can give rise to nuanced emotions. For example, emotions such as excitement, outrage, and disappointment can emerge from combinations of surprise and happiness, anger, and sadness, respectively, as illustrated in Figure 1. The Emotion Wheel acknowledges the complexity of emotional blends, bridging the gap between psychological emotion studies and computational speech emotion modeling. It highlights the challenges in speech emotion recognition and synthesis, emphasizing the need to consider a wider spectrum of emotions than traditionally accounted for. In our research, we integrate insights from the Emotion Wheel into voice conversion models, proposing a method to quantifiably render mixed emotions in human voice, which we will discuss next.

3. Mixed-EVC

In this section, we present the details of our proposed emotional voice conversion framework, Mixed-EVC. We begin by introducing a novel approach that facilitates the synthesis of mixed emotions (Sec. 3.1). Following that, we introduce the core components and the training diagram of Mixed-EVC (Sec. 3.2). Lastly, we describe how to render and control the desired mixed emotions (Sec. 3.3).

3.1. Modeling Emotion Attribute with Pairwise Ranking

Most emotional speech databases group utterances into several discrete categories thus restricting to learning of richer descriptions of the emotional style [4]. Instead of predicting the presence of a specific emotion, we propose to model the relative difference between a pair of emotions, which we define as “*Emotion Attribute*”. During the training, we characterize the input emotional style by measuring its relevance to other emotions by explicitly modeling emotion attributes. For example, while it is hard to give a consensus emotion label to an utterance, espe-

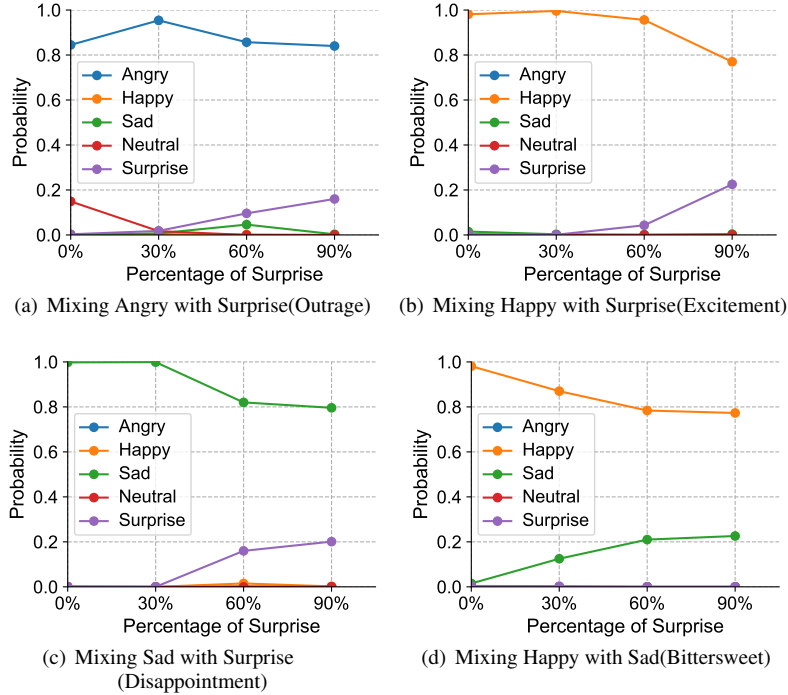


Figure 4: The classification probabilities utilizing a pre-trained SER model. Each data point corresponds to an averaged probability value obtained by analyzing 20 converted utterances displaying mixed emotions, all of which are drawn from the evaluation set.

cially when the emotions are ambiguous, we can agree that it sounds less happy than A but angrier than B [29].

We employ the idea of relative attributes [30] to encode the relationship between discrete emotions into a quantifiable vector. Relative attributes were first proposed in computer vision [30] and later applied in speech processing [31, 32, 33]. Prior studies show that relative attributes provide a more detailed description of an image, thus surpassing traditional classifiers on image recognition tasks [30]. Inspired by that, we propose to learn a ranking function for each emotion attribute given relative similarity constraints on paired emotional samples. More specifically, we are given a training set $T = \{\mathbf{x}_n\}$, where \mathbf{x}_n is the acoustic feature vector of the n^{th} training sample with M emotion labels. Our goal is to learn a ranking function f_m given below:

$$f_m(\mathbf{x}_n) = \mathbf{W}_m \mathbf{x}_n, \quad (1)$$

where $m = 1, \dots, M$ and \mathbf{W} is a weighting matrix. We construct an ordered set O_m and an unordered set U_m as the supervision sets, which satisfy the following constraints:

$$\forall (x_i, x_j) \in O_m : \mathbf{W}_m \mathbf{x}_i > \mathbf{W}_m \mathbf{x}_j \quad (2)$$

$$\forall (x_i, x_j) \in U_m : \mathbf{W}_m \mathbf{x}_i = \mathbf{W}_m \mathbf{x}_j, \quad (3)$$

The sample x_i should have a stronger presence of emotion attribute a_m than x_j in the ordered set O_m , while their presence should be similar in the unordered set U_m . The weighting matrix \mathbf{W} is estimated by solving a support vector machine (SVM) problem [34]. The emotion attribute thus corresponds to the distance from the decision boundary. We generate pairs for all emotions and iteratively conduct the aforementioned ranking training in a pairwise manner. During inference, the trained ranking function can automatically predict the emotion attribute, which signifies the relevance or similarity

between two emotional styles. Each emotion attribute is standardized within the $[0, 1]$ range to be easily quantifiable, where a lower value indicates a greater similarity in emotional style. All these emotion attributes together compose an emotion attribute vector, quantifying the degree of relevance between an input emotional style and all other emotions. Next, we explain how those attributes contribute to the EVC training.

3.2. Seq2Seq Emotional Voice Conversion Training

We integrate our emotion attribute vector into a seq2seq-based emotional voice conversion framework, as illustrated in Figure 2. Our approach, distinct from traditional frame-wise modeling [19], jointly models feature mapping and alignment and automatically predicts speech duration during run-time.

Given an input speech, the linguistic encoder is trained to predict a sequence of linguistic embeddings (“Linguistic Modeling”), while the emotion encoder focuses on encoding the emotional style into an utterance-level emotion embedding (“Style Modeling”). The pre-trained relative functions further predict an emotion attribute vector that signifies the relevance of the input emotional style to other emotions (“Relevance Modeling”). This emotion attribute vector is then transformed using a fully connected (FC) layer, resulting in a relative embedding. Ultimately, the decoder learns to reconstruct the input emotional style by combining information from both the emotion and relative embeddings. Our framework not only explicitly characterizes the input emotional style but also establishes a relationship with other emotions.

To overcome the instability issues in seq2seq models, we conduct emotional training with the following strategies: (1) introducing text supervision [35], and (2) pre-training with a large neutral-speaking corpus [16]. We use text transcriptions

Table 1: MOS with 95% confidence interval to evaluate the speech quality of synthesized mixed emotions.

Configuration		MOS
Mixing Angry with Surprise	Ground truth (Angry)	4.78 ± 0.17
	+ 0% Surprise	3.50 ± 0.24
	+ 30% Surprise	3.34 ± 0.26
	+ 60% Surprise	3.15 ± 0.34
	+ 90% Surprise	3.20 ± 0.31
Mixing Happy with Surprise	Ground truth (Happy)	4.85 ± 0.12
	+ 0% Surprise	3.63 ± 0.27
	+ 30% Surprise	3.53 ± 0.19
	+ 60% Surprise	3.17 ± 0.30
	+ 90% Surprise	3.05 ± 0.34
Mixing Sad with Surprise	Ground truth (Sad)	4.74 ± 0.12
	+ 0% Surprise	3.63 ± 0.27
	+ 30% Surprise	3.53 ± 0.19
	+ 60% Surprise	3.17 ± 0.30
	+ 90% Surprise	3.05 ± 0.34
Mixing Happy with Sad	Ground truth (Happy)	4.84 ± 0.12
	+ 0% Sad	3.17 ± 0.35
	+ 30% Sad	3.48 ± 0.30
	+ 60% Sad	3.42 ± 0.32
	+ 90% Sad	3.08 ± 0.36

to assist the framework in disentangling the linguistic information from the speech. A text encoder predicts a sequence of text embeddings from the input text. The text and linguistic embeddings are then fed into the decoder in an alternative manner. Following the previous literature [35], we employ a contrastive loss to ensure the similarity between text and linguistic embeddings.

3.3. Mixed Emotion Rendering and Control

During run-time, Mixed-EVC carries out the conversion of the source emotion into an emotional mixture, illustrated in Figure 3. This process involves several steps: first, the linguistic encoder transforms the source linguistic content into internal representations. Next, the emotion encoder encapsulates a set of reference speech into an utterance-level emotion embedding. Additionally, attributes of other emotion types can be introduced by manually specifying an attribute vector. This flexibility enables us to adjust the proportions of each emotion type, thereby enabling the creation of diverse emotional mixtures.

4. Experiments

4.1. Experimental Setup

We conduct experiments with the ESD dataset [4], where we randomly choose one female speaker (“0019”) with five emotions (“*Neutral*”, “*Happy*”, “*Sad*”, “*Angry*” and “*Surprise*”). We follow the data partition in the ESD dataset, and for each emotion, we use 300, 30, and 20 utterances for training, testing, and evaluation, respectively. We train a universal EVC model for all emotions. At run-time, we convert *Neutral* to 4 different mixtures of emotions that are:

- **Outrage, Excitement, Disappointment:** where we convert *Neutral* to *Angry*, *Happy* and *Sad* (Neu-Ang, Neu-Hapy, Neu-Sad) respectively while introducing different percent-

Table 2: BWS test results to evaluate the perception of the mixed feelings (*Outrage*, *Excitement*, *Disappointment* and *Bittersweet*) in the converted mixed emotions.

Configuration	Best (%)	Worst (%)	
(a) Perception of Outrage			
Mixing Angry with Surprise	+ 0% Surprise	15.4	40.0
	+ 30% Surprise	15.4	24.6
	+ 60% Surprise	13.8	21.6
	+ 90% Surprise	55.4	13.8
(b) Perception of Excitement			
Mixing Happy with Surprise	+ 0% Surprise	3.1	69.2
	+ 30% Surprise	12.3	10.8
	+ 60% Surprise	23.1	9.2
	+ 90% Surprise	61.5	10.8
(c) Perception of Disappointment			
Mixing Sad with Surprise	+ 0% Surprise	13.8	53.8
	+ 30% Surprise	16.9	12.3
	+ 60% Surprise	21.5	27.7
	+ 90% Surprise	47.8	6.2
(d) Perception of Bittersweet			
Mixing Happy with Sad	+ 0% Sad	4.6	41.5
	+ 30% Sad	20.0	12.3
	+ 60% Sad	30.8	10.8
	+ 90% Sad	44.6	35.4

ages of *Surprise* into the mixture. These 3 mixtures have been studied in emotion theory [12];

- **Bittersweet:** where we convert *Neutral* to *Happy* while introducing different percentages of *Sad*. Even though happiness and sadness are two oppositely valenced emotions in Russell’s model [36], there are some debates that agree with the co-existence of happiness and sadness [2];

The training pipeline is described below: We first pre-train the relative ranking function between each emotion pair. Each utterance is represented by a 384-dimensional feature vector defined by the INTERSPEECH Emotion Challenge [37], which is used to train the ranking functions. The ranking functions achieved 98% accuracy in classifying the emotion categories on the test set of the ESD dataset. We then follow a 2-stage training strategy [16] to train our proposed framework: (1) Style Pre-training with the VCTK Corpus [38], and (2) Seq2Seq EVC Training with the ESD dataset. The inputs to the EVC model are acoustic features represented by an 80-dimensional Mel-spectrogram extracted every 12.5 ms with a frame size of 50 ms for the short-time Fourier transform (STFT) and phoneme sequences converted by the Festival [39] G2P tool. It should be noted that we only use acoustic features as inputs during the conversion.

Our proposed framework has a similar structure to the model presented in [35]. The linguistic encoder comprises an encoder, a 2-layer 256-cell BLSTM, and a decoder employing a 1-layer 512-cell BLSTM with an attention layer followed by a fully connected (FC) layer with an output channel of 512. The decoder has the same model architecture as that of Tacotron [40]. The text encoder is a 3-layer 1D CNN with a kernel size of 5 and a channel of 512. The style encoder is a 2-layer, 128-cell




Figure 5: Preference test results of emotion similarity between the proposed Mixed-EVC and the baseline Seq2Seq-EVC over one-to-one emotion conversion.

BLSTM followed by an FC layer with an output channel of 64. The classifier is a 4-layer FC with 512, 512, 512, and 5 channels. For training, we specify a batch size of 64 for style pre-training and 4 for emotion training. During style pre-training, the learning rate is set at 0.001, and a weight decay of 0.0001 is applied. For emotion pre-training, we adopt a learning rate halving strategy every 7 epochs.

4.2. Objective Evaluation

We assess the quality of our synthesized mixed emotions by employing a speech emotion recognition (SER) model pre-trained on the ESD dataset. This SER model shares a similar structure to the one presented in [41], encompassing components such as a 3-D CNN layer, a BLSTM layer, an attention layer, and an FC layer. Our analysis leverages the classification probabilities obtained from the final softmax layer of the SER. We contend that these classification probabilities serve as a concise summary of emotional cues aggregated from preceding layers, aiding in decision-making and providing us with a valuable tool for investigating emotional blends. We report the classification probabilities in Figure 4. We first evaluate three different mixed effects that are *Outrage*, *Excitement* and *Disappointment*, where we convert *Neutral* to *Angry*, *Neutral* to *Happy* and *Neutral* to *Sad* respectively. These transformations are done while gradually increasing the percentage (0%, 30%, 60% and 90%) of *Surprise*. As shown in Figure 4(a), (b), and (c), we observe that the probability of *Surprise* consistently increases when we increase the percentage of *Surprise* from 0% to 90%. The likelihoods of *Angry*, *Happy*, and *Sad* consistently maintain the highest in three different emotional mixtures. This pattern arises from their distinct characterization by the emotion encoder, establishing them as the predominant emotional components within the mixture. We then evaluate the mixed effect of *Bittersweet* as shown in Figure 4(d), where *Sad* is further introduced when we convert *Neutral* to *Happy*. As shown in Figure 4(d), we find a similar observation as in Figure 4(a), (b) and (c). These observations indicate that mixed emotions can be objectively recognized by a pre-trained SER.

4.3. Subjective Evaluation

We then conduct listening experiments to evaluate our synthesized results in terms of speech quality and emotion perception. 15 subjects participated in all the experiments and each of them listened to 112 synthesized utterances in total.

We first report mean opinion scores (MOS) results for speech quality, where all participants are asked to listen to the reference speech (“Ground truth”) and the synthesized speech with mixed emotions and score the “quality” of each speech




Figure 6: The illustration of the interactive demo - an emotion triangle. Each pot combines basic emotions - happy, angry, sad and surprise in various proportions. Each corner of the triangle represents the emotional blend of excitement, outrage, and disappointment.

sample on a 5-point scale (‘5’ for excellent, ‘4’ for good, ‘3’ for fair, ‘2’ for poor, and ‘1’ for bad). As shown in Table 1, the audio quality of synthesized speech slightly decreases as we increase the percentage of emotional mixture. Nonetheless, the quality remains between fair and good across the board.

We then conduct best-worst scaling (BWS) tests to evaluate the emotion perception of the synthesized mixed emotions. All participants are asked to choose the best and the worst emotion according to their perception of the mixed emotion (*Outrage*, *Excitement*, *Disappointment* and *Bittersweet*). As shown in Table 2(a), (b) and (c), we observe that most participants can perceive the mixed feelings and choose those with 90% of *Surprise* as the “Best”, and those of 0% of *Surprise* as the “Worst”. We take one step further to evaluate the perception of *Bittersweet* as shown in Table 2(d). Most participants chose those with 90% *Sad* as the “Best” confirming the strength of our approach. The results of *Bittersweet* are not as distinguishable as those of the other three mixed emotions, which suggests that it could be more challenging for listeners to perceive a *Bittersweet* feeling. These results show that we are able to create new emotional feelings that are subtle but do not exist in the database. We also show the effectiveness of controllability by varying the percentages of primary emotions to synthesize different emotional mixtures.

We further conduct a preference test to evaluate the performance of Mixed-EVC on one-to-one emotion conversion. We choose Seq2Seq-EVC [16] as the baseline. Participants are asked to choose the best one in terms of the emotion similarity with the reference emotion. As shown in Figure 5, our proposed Mixed-EVC outperforms the baseline Seq2Seq-EVC in all conversion pairs. These results show that Mixed-EVC could not only synthesize emotional mixtures but also have superior performance on emotion conversion.

5. An Interactive Demo - Emotion Triangle

Mixed emotion synthesis is a key method to model transitions between discrete emotional states. In our application study, we developed an “Emotion Triangle” as an interactive demo for users to explore these transitions. This model challenges us to accurately represent the internal states that occur between distinct emotions. Our approach, the Mixed-EVC, enables the synthesis of these intermediary states by blending various percentages of basic emotions—happiness, anger, sadness, and surprise. This tool not only aids in understanding how different feelings can combine to form complex, mixed emotional states

but also demonstrates the effectiveness of Mixed-EVC in controlling these blends. As shown in Figure 6, each corner of the triangle illustrating a specific emotional mixture, representing excitement, outrage, and disappointment, thereby further validating the controllability of Mixed-EVC. For a practical exploration of these concepts, we encourage readers to visit the online demo.¹

6. Conclusion

We introduce Mixed-EVC, a seq2seq emotional voice conversion (EVC) framework, to address the existing research gap in mixed emotion synthesis within the context of EVC. We formulate emotional styles as an attribute and explicitly model the degree of relevance between different emotions through a ranking-based SVM. By manually adjusting relevance at runtime, Mixed-EVC could produce different emotional mixtures. Both objective and subjective evaluations show the effectiveness of synthesizing different mixed emotions. The speech samples are publicly available².

7. References

- [1] Larry A Herzberg, “To blend or to compose: A debate about emotion structure,” 2012.
- [2] Patti Williams and Jennifer L Aaker, “Can mixed emotions peacefully coexist?,” *Journal of consumer research*, vol. 28, no. 4, pp. 636–649, 2002.
- [3] Huang-Cheng Chou, Chi-Chun Lee, and Carlos Busso, “Exploiting co-occurrence frequency of emotions in perceptual evaluations to train a speech emotion classifier,” in *Proc. Interspeech*, 2022, vol. 2022.
- [4] Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li, “Emotional voice conversion: Theory, databases and esd,” *Speech Communication*, vol. 137, pp. 1–18, 2022.
- [5] Dagmar Schuller and Björn W Schuller, “The age of artificial emotional intelligence,” *Computer*, vol. 51, no. 9, pp. 38–46, 2018.
- [6] Mostafa M Amin, Rui Mao, Erik Cambria, and Björn W Schuller, “A wide evaluation of chatgpt on affective computing tasks,” *arXiv preprint arXiv:2308.13911*, 2023.
- [7] Weixiang Zhao, Yanyan Zhao, Xin Lu, Shilong Wang, Yanpeng Tong, and Bing Qin, “Is chatgpt equipped with emotional dialogue capabilities?,” *arXiv preprint arXiv:2304.09582*, 2023.
- [8] James A Russell, “Core affect and the psychological construction of emotion.,” *Psychological review*, vol. 110, no. 1, pp. 145, 2003.
- [9] Robert Plutchik, “The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice,” *American scientist*, vol. 89, no. 4, pp. 344–350, 2001.
- [10] Joe LeDoux, Liz Phelps, and Cristina Alberini, “What we talk about when we talk about emotions,” *Cell*, vol. 167, no. 6, pp. 1443–1445, 2016.
- [11] Alan Cowen, “How many different kinds of emotion are there,” *Age*, vol. 12, pp. 13, 2018.
- [12] Robert Plutchik and Henry Kellerman, *Theories of emotion*, vol. 1, Academic Press, 2013.
- [13] Haobin Tang, Xulong Zhang, Jianzong Wang, Ning Cheng, and Jing Xiao, “EmoMix: Emotion Mixing via Diffusion Models for Emotional Speech Synthesis,” in *Proc. INTERSPEECH 2023*, 2023, pp. 12–16.
- [14] Kun Zhou, Berrak Sisman, Rajib Rana, Björn W. Schuller, and Haizhou Li, “Speech synthesis with mixed emotions,” *IEEE Transactions on Affective Computing*, vol. 14, no. 4, pp. 3120–3134, 2023.
- [15] ZHOU KUN, “Emotion modelling for speech generation,” 2022.
- [16] Kun Zhou, Berrak Sisman, and Haizhou Li, “Limited Data Emotional Voice Conversion Leveraging Text-to-Speech: Two-Stage Sequence-to-Sequence Training,” in *Proc. Interspeech 2021*, 2021, pp. 811–815.
- [17] Carl Robinson, Nicolas Obin, and Axel Roebel, “Sequence-to-sequence modelling of f0 for speech emotion conversion,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6830–6834.
- [18] Heejin Choi and Minsoo Hahn, “Sequence-to-sequence emotional voice conversion with strength control,” *IEEE Access*, vol. 9, pp. 42674–42687, 2021.
- [19] Zijiang Yang, Xin Jing, Andreas Triantafyllopoulos, Meishu Song, Ilhan Aslan, and Björn W Schuller, “An overview & analysis of sequence-to-sequence emotional voice conversion,” *Interspeech 2022*.
- [20] Junichi Yamagishi, Koji Onishi, Takashi Masuko, and Takao Kobayashi, “Modeling of various speaking styles and emotions for hmm-based speech synthesis,” in *Eighth European Conference on Speech Communication and Technology*, 2003.
- [21] Ryo Aihara, Ryoichi Takashima, Tetsuya Takiguchi, and Yasuo Ariki, “Gmm-based emotional voice conversion using spectrum and prosody features,” *American Journal of Signal Processing*, 2012.
- [22] Shumin An, Zhenhua Ling, and Lirong Dai, “Emotional statistical parametric speech synthesis using lstm-rnns,” in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2017, pp. 1613–1616.
- [23] Kun Zhou, Berrak Sisman, and Haizhou 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.
- [24] Zhaojie Luo, Jinhui Chen, Tetsuya Takiguchi, and Yasuo Ariki, “Emotional voice conversion using dual supervised adversarial networks with continuous wavelet transform f0 features,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 10, pp. 1535–1548, 2019.
- [25] Kun Zhou, Berrak Sisman, and Haizhou Li, “Vaw-gan for disentanglement and recombination of emotional elements in speech,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 415–422.
- [26] Kun Zhou, Berrak Sisman, Mingyang Zhang, and Haizhou Li, “Converting Anyone’s Emotion: Towards

¹**Interactive Demo:** https://demo9646.github.io/Emotion_Triangle/

²**Speech Demo:** https://demo9646.github.io/Mixed_EVC/

- Speaker-Independent Emotional Voice Conversion,” in *Proc. Interspeech 2020*, 2020, pp. 3416–3420.
- [27] Tae-Ho Kim, Sungjae Cho, Shinkook Choi, Sejik Park, and Soo-Young Lee, “Emotional voice conversion using multitask learning with text-to-speech,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7774–7778.
- [28] Andreas Triantafyllopoulos, Björn W Schuller, Gökçe İymen, Metin Sezgin, Xiangheng He, Zijiang Yang, Panagiotis Tzirakis, Shuo Liu, Silvan Mertes, Elisabeth André, et al., “An overview of affective speech synthesis and conversion in the deep learning era,” *Proceedings of the IEEE*, 2023.
- [29] Georgios N Yannakakis, Roddy Cowie, and Carlos Busso, “The ordinal nature of emotions: An emerging approach,” *IEEE Transactions on Affective Computing*, vol. 12, no. 1, pp. 16–35, 2018.
- [30] Devi Parikh and Kristen Grauman, “Relative attributes,” in *2011 International Conference on Computer Vision*. IEEE, 2011.
- [31] Xiaolian Zhu, Shan Yang, Geng Yang, and Lei 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.
- [32] Yi Lei, Shan Yang, and Lei 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.
- [33] Kun Zhou, Berrak Sisman, Rajib Rana, Björn W. Schuller, and Haizhou Li, “Emotion intensity and its control for emotional voice conversion,” *IEEE Transactions on Affective Computing*, vol. 14, no. 1, pp. 31–48, 2023.
- [34] Olivier Chapelle, “Training a support vector machine in the primal,” *Neural computation*, vol. 19, no. 5, pp. 1155–1178, 2007.
- [35] Jing-Xuan Zhang, Zhen-Hua Ling, and Li-Rong Dai, “Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 540–552, 2019.
- [36] James A Russell, “A circumplex model of affect,” *Journal of personality and social psychology*, vol. 39, no. 6, pp. 1161, 1980.
- [37] Björn Schuller, Stefan Steidl, and Anton Batliner, “The interspeech 2009 emotion challenge,” in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [38] Christophe Veaux, Junichi Yamagishi, Kirsten MacDonal, et al., “Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit,” 2016.
- [39] Alan Black, Paul Taylor, Richard Caley, Rob Clark, Korin Richmond, Simon King, Volker Strom, and Heiga Zen, “The festival speech synthesis system, version 1.4.2,” *Unpublished document available via <http://www.cstr.ed.ac.uk/projects/festival.html>*, 2001.
- [40] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al., “Tacotron: Towards end-to-end speech synthesis,” *arXiv preprint arXiv:1703.10135*, 2017.
- [41] Mingyi Chen, Xuanji He, Jing Yang, and Han Zhang, “3-d convolutional recurrent neural networks with attention model for speech emotion recognition,” *IEEE Signal Processing Letters*, vol. 25, no. 10, pp. 1440–1444, 2018.