

# EmoReg: Directional Latent Vector Modeling for Emotional Intensity Regularization in Diffusion-based Voice Conversion

Ashishkumar Prabhakar Gudmalwar<sup>1</sup>, Ishan Darshan Biyani<sup>1</sup>, Nirmesh J. Shah<sup>1</sup>,  
Pankaj Wasnik<sup>1</sup>, Rajiv Ratn Shah<sup>2</sup>

<sup>1</sup>Media Analysis Group, Sony Research India, Bangalore

<sup>2</sup>Indraprastha Institute of Information Technology (IIIT), Delhi, India

{ashish.gudmalwar1,ishan.biyani,nirmesh.shah,pankaj.wasnik}@sony.com, rajivrtn@iiitd.ac.in

## Abstract

The Emotional Voice Conversion (EVC) aims to convert the discrete emotional state from the source emotion to the target for a given speech utterance while preserving linguistic content. In this paper, we propose regularizing emotion intensity in the diffusion-based EVC framework to generate precise speech of the target emotion. Traditional approaches control the intensity of an emotional state in the utterance via emotion class probabilities or intensity labels that often lead to inept style manipulations and degradations in quality. On the contrary, we aim to regulate emotion intensity using self-supervised learning-based feature representations and unsupervised directional latent vector modeling (DVM) in the emotional embedding space within a diffusion-based framework. These emotion embeddings can be modified based on the given target emotion intensity and the corresponding direction vector. Furthermore, the updated embeddings can be fused in the reverse diffusion process to generate the speech with the desired emotion and intensity. In summary, this paper aims to achieve high-quality emotional intensity regularization in the diffusion-based EVC framework, which is the first of its kind work. The effectiveness of the proposed method has been shown across state-of-the-art (SOTA) baselines in terms of subjective and objective evaluations for the English and Hindi languages.

**Demo Page** — <https://nirmesh-sony.github.io/EmoReg/>

## Introduction

Despite significant progress in the field of Generative AI, speech synthesis models still encounter several challenges when it comes to the AI-based dubbing of entertainment content such as movies and serials (Brannon et al. 2023; Wu et al. 2023; Hu et al. 2021; Mhaskar et al. 2024; Sahipjohn et al. 2020). AI-based dubbing involves replicating input speech emotion and controlling its intensity depending on the context and emotion of the scene (Brannon et al. 2023; Zhou et al. 2022c,a; Amiriparian et al. 2023). Most of today’s text-to-speech (TTS) systems can produce high-quality, high-fidelity, natural speech output, but they still lack expressiveness and fine control over emotional states (Barakat et al. 2024; Gudmalwar et al. 2024). Hence, professional voice-over/dubbing artists are still preferred in

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

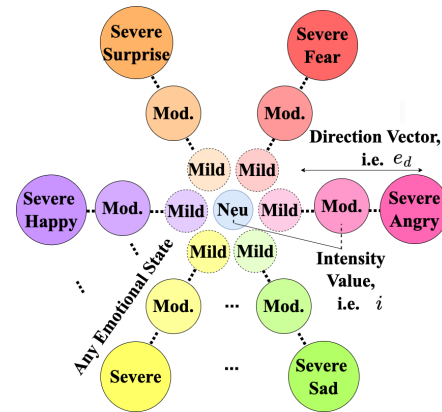


Figure 1: Conceptual representation of emotional intensity regularization based on direction vector and intensity value. Here, Mod. refers to Moderate and Neu refers to Neutral.

the dubbing industry to tackle the complex demands of generating emotionally engaging speech (Gutentag et al. 2017). This poses challenges in terms of dubbing at a large scale, turnaround time, and operational costs.

One of the major roadblocks in building emotional speech synthesis is the unavailability of a large emotional speech database with diverse emotional expressions and a wide range of emotion intensities. Even annotating accurate emotional states in the speech requires expert knowledge and is costly and labor-intensive as labeling each speech file with the correct emotion and intensity involves detailed subjective assessments. Hence, emotional intensity control or regularization is still an under-explored open research problem. To this end, we primarily focus on the emotional voice conversion (EVC) task, a sub-topic of emotional speech synthesis for emotional intensity regularization (Zhou et al. 2022c, 2021, 2022a; Shah et al. 2023). In particular, we employ a self-supervised learning (SSL)-based framework to tackle the issues related to the unavailability of a large annotated emotional speech database.

The EVC aims to alter the discrete emotional state of a speech utterance while preserving linguistic content and the speaker’s identity (Zhou et al. 2021). Furthermore, emotional intensity regularization involves fine control over in-

tensity associated with the target emotional state. For example, a neutral utterance can be converted into a mild, moderate, or severe level of anger emotion, as shown in Figure 1. Current EVC methods usually tackle this scenario with discrete (Zhou et al. 2022b; Shah et al. 2023; Zhou et al. 2022c; Ekman and Friesen 1971) or continuous emotion representations (Russell 1980; Posner, Russell, and Peterson 2005; Cho et al. 2024). Discrete emotion representations contain categorical labels such as happy, angry, sad, etc. On the other hand, continuous emotion representations are obtained from the circumplex models and contain continuous-independent dimensions, such as arousal and valence (Russell 1980; Posner, Russell, and Peterson 2005). Achieving emotion intensity control in continuous emotion representation is relatively easier than discrete emotion representations (Gunes et al. 2011; Prabhu, Lehmann-Willenbrock, and Gerkmann 2023; Zhou et al. 2022c,a; Cho et al. 2024). However, obtaining such continuous emotion representations-based annotated data is challenging, as discussed earlier. Hence, this paper aims to achieve emotion intensity regularization for the discrete emotion representations-based EVC methods.

In summary, we propose a method to achieve emotion intensity regularization by means of self-supervised learning and direction latent vector modeling using a diffusion-based EVC. Our key contributions can be summarized as follows:

- To the best of the authors' knowledge this is the first attempt that achieves a high-quality emotional intensity regularization in the diffusion-based EVC framework.
- We propose a novel direction latent vector modeling-based approach for obtaining fine control over intensity while transitioning across different emotional states.
- The proposed EmoReg utilizes the SSL-based audio feature representations, which are obtained after finetuning the SSL-based framework for a downstream task related to emotions classification.
- Effectiveness of the proposed EmoReg approach has been shown against SOTA baselines with relevant subjective and objective evaluation techniques across two languages, namely English and Hindi.

## Related Work

EVC methods are usually categorized into parallel and non-parallel approaches depending on the nature of the training data. Parallel means each speaker has spoken the same utterance in different emotional states and non-parallel means recorded sentences are different depending on the emotional states. Since emotion is also dependent on the content that is being spoken, hence, researchers have primarily focused on non-parallel EVC approaches (Shah et al. 2023). Traditionally, researchers have explored various generative models for the discrete emotion representation-based EVC tasks, namely, GMM (Aihara et al. 2012), HMM (Inanoglu and Young 2007), DNN (Lorenzo-Trueba et al. 2018), Sequence-to-sequence models (Ming et al. 2016), Variational Auto Encoder (VAE) (Zhou, Sisman, and Li 2021b), Generative Adversarial Networks (GAN) and its variants (Shah et al. 2023; Li, Zare, and Mesgarani 2021), Diffusion Model-based ap-

proaches (Prabhu et al. 2024), etc. However, none of these approaches tackles emotion intensity regularization tasks.

Broadly, emotional intensity regularization is achieved from the emotional labels and the reference speech utterance. The emotion label-based approach is utilized to manipulate auxiliary features such as attention weights, saliency maps, ranking function, or signal processing attributes to achieve emotion intensity regularization (Zhou et al. 2022c,a; Matsumoto, Hara, and Abe 2020; Schnell and Garner 2021; Choi and Hahn 2021; Um et al. 2020). Among these, one of the most prominent techniques is relative attribute (Zhu et al. 2019; Zhou et al. 2022c,a), which describes the distinctions between binary classes using a learned ranking function. However, these methods fail to capture high-level complex abstract representations, which results in artifacts in the converted output and leads to significant degradation in the quality. Also, emotional manipulations often involve inter-dependencies across different types of features, hence manipulating any single feature in isolation fails in achieving emotional intensity regularization. On the other hand, some simple approaches manipulate learned emotion representations via scaling (Choi and Hahn 2021) or interpolations (Um et al. 2020) to achieve emotion regularization. However, these approaches do not work well since emotional embedding space does not align well with the assumption of linear interpolation. Furthermore, high-dimensional emotional embedding space is often sparsely populated between two styles. Thus, such approaches result in inept style manipulations.

Another emotional intensity control-based strategy leverages the extended emotional dimensions, namely, Arousal, Valence, and Dominance (AVD) (Russell 1980). AVD-based emotional dimensions provide continuous and fine-grained emotion representations, which can be utilized to alter emotional states in a continuous manner more appropriately than discrete emotion representations. However, obtaining such annotations is difficult due to inherent subjectivity-related issues associated with different annotators and the high costs associated with preparing such data. Hence, emotional intensity regularization is an under-explored research area.

In summary, all previous discrete and continuous emotional intensity control methods produce low-quality noisy converted output. To tackle quality-related issues recently, diffusion-based models have seen great success as a potent and versatile generative AI technology in computer vision, audio, and reinforcement learning (Ho, Jain, and Abbeel 2020). Diffusion models serve as samplers in these applications, producing new samples while actively guiding them toward task-desired features (Ho, Jain, and Abbeel 2020). They also offer versatile high-dimensional data modeling (Ho, Jain, and Abbeel 2020). The diffusion model-based approach has also obtained remarkable success in the voice conversion (Popov et al. 2021a) and EVC task (Prabhu et al. 2024). However, it has not been explored for achieving emotion intensity control-related tasks. Hence, we exploit the diffusion-based EVC method with the proposed DVM-based strategy for the emotion intensity regularization task. The next section presents details about the proposed diffusion-based EmoReg architecture.

## Proposed Methodology

### Problem Formulation

Given input speech  $x_0$  along with reference emotion speech  $x_r$  and intensity value  $i$ , the goal of EVC is to generate emotional intensity regularized converted speech  $y$  such that  $Y = G(X_0, e_s, e_r, i, \theta)$ , where  $X_0$  and  $Y$  denotes Mel-spectrogram features corresponding to the input speech  $x_0$  and the converted emotional speech  $y$ , respectively. In addition,  $e_s, e_r$  denotes emotional feature representation for source and reference speech, respectively.  $\theta$  is model parameters corresponding to the considered conditional diffusion probabilistic model-based EVC architecture. Details of the proposed architecture is as follows.

### Proposed EmoReg Architecture

The proposed EmoReg architecture uses diffusion-based model to achieve EVC while controlling the emotion intensity  $i$ . It comprises a diffusion-based decoder and a set of encoders illustrated in Figure 3. The diffusion decoder is responsible for emotion-controllable speech synthesis, while the encoders individually encode the emotion and speech representation that require disentanglement. The decoder is conditioned on the proposed Direction Vector Modeling (DVM) based features  $e_{ir}$ , which facilitate the transition between different emotions. At the output of the diffusion decoder, we obtain the Mel-spectrogram  $Y$ , which can be converted to the output speech signal  $y = H(Y)$ , with  $H(\cdot)$  representing the HiFiGAN vocoder. Subsequent sections delve deeper into the architectural components of the proposed method.

### Encoders

Our EmoReg model consists of two encoders. The first encoder is used to encode the lexical content, while the second one is used to obtain emotion representations.

**Phoneme Encoder:** In this part, we encode the lexical content at the phoneme level by using speaker- and emotion-independent average phoneme-level Mel characteristics. To achieve this, we utilize the transformer-based encoder as described in (Popov et al. 2021a), which has previously been used in speech conversion applications.

$$\bar{\mathbf{X}} = \phi(\mathbf{X}_0) \quad (1)$$

Here,  $\bar{\mathbf{X}}$  is average Mel-spectrogram of source audio,  $\mathbf{X}_0$  is source speech Mel-spectrogram and  $\phi(\cdot)$  represents the pre-trained average phoneme encoder.

**SSL Emotion Embedding Network:** Due to the unavailability of a large annotated emotional speech database, we plan to utilize SSL-based representations. We fine-tune the pre-trained SSL emotion2vec (Ma et al. 2023) embedding network  $E(\cdot)$  using English and Hindi emotional speech databases to learn emotional embedding representations.

We used a pretrained emotion2vec model (Ma et al. 2023) to generate 768-dimensional embeddings, which we then fine-tuned for classifying target emotions. Similar to the diffusion-based text generation models (Gao et al. 2022), we also find that in the high dimensional embedding space, the

insufficient noise results in a simple denoising task, which leads to the deterioration of the model. In the process, we first passed these embeddings through a fully connected layer to reduce them to 256 dimensions. We then passed these reduced embeddings through an output layer to classify them into the target emotions (neutral, happy, sad and angry). The training was conducted in a supervised manner. After training, we used the network with the newly added fully connected layer to obtain 256-dimensional emotion embeddings. These emotional embeddings are conditioned in the decoder of the diffusion model in order to achieve target emotional state in the converted output. Additionally, these emotional embeddings are manipulated via proposed directional vector modelling to obtain a fine control over intensities associated with the target emotional states.

### Direction Vector Modeling

During training phase, the 256-dimension emotional embeddings from SSL Emotion Embedding Network are used as an input to the DVM module as shown in Figure 2. We modeled the emotion embedding space using a 64-component Gaussian Mixture Model (GMM) (Reynolds et al. 2009) to derive the local mean vector for each emotion, i.e., Angry, Happy, Neutral, and Sad. After extracting the GMM features for each emotion, pairwise subtraction is performed between the embeddings of Angry, Happy, and Sad with those of Neutral to derive the emotional direction vector matrix for all possible transitions from local mean of one emotional state to another. Subsequently, Principal Component Analysis (PCA) (Abdi and Williams 2010) is applied to the emotional direction matrix corresponding to each emotion, reducing it to 128 components since not all 256 components are reflecting changes related to emotional states, they may also be affected due to content variability. Ablation analysis for selecting number of principle components is presented in Table 5.

During the inference phase, given a source sample in the Neutral emotion and a reference sample in another emotion (e.g., Angry), the emotional direction vector is calculated by subtracting the source’s neutral embedding ( $e_s$ ) from the reference’s emotional embedding ( $e_r$ ). This direction vector is then transformed using the PCA to obtain a 128-dimensional vector, which is then converted back to the original 256-dimensional space using PCA inverse transformation, result-




Figure 2: Three key steps of the proposed DVM approach. 1) Fitting local GMM to each emotional state. 2) computing directional vectors for all possible transitions from the local mean of one emotional state to another. 3) Applying PCA to find relevant direction for emotional transition.

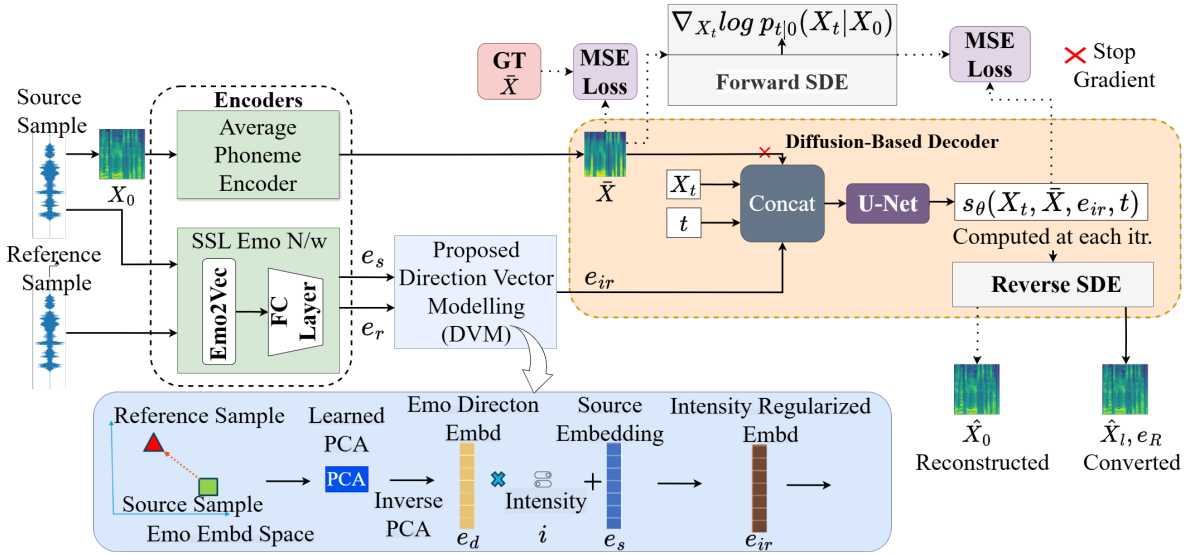


Figure 3: Block diagram of the proposed DVM-based Emotion Intensity Regularized EVC architecture. Dotted arrows represents operations performed only during training. Also, GT  $\bar{X}$  are derived by replacing each phoneme Mel-spectrogram feature in the input with its corresponding pre-calculated average feature.

ing in the final emotional reference embedding  $e_d$ . The emotional reference embedding is then scaled by a specified factor corresponding to the intensity value  $i$ , which ranges between 0 to 1, and added to the source embedding (i.e.,  $e_s$ ) to shift it towards the reference emotion (i.e.,  $e_{ir}$ ) in the embedding space. The diffusion model’s reverse Stochastic Differential Equation (SDE) process is conditioned on the resulting scaled embeddings, which effectively transform the source sample into the output speech with a reference target emotion regulated by the emotion intensity value.

### Diffusion-based Decoder

The diffusion-based decoder follows the SDE formalism as described in (Popov et al. 2021b). Consider the continuous diffusion time-step variable  $t$ , which characterizes the progression of the diffusion process. The forward SDE for  $0 < t \leq 1$  is given by:

$$d\mathbf{X}_t = \frac{1}{2}\beta_t(\bar{\mathbf{X}} - \mathbf{X}_t)dt + \sqrt{\beta_t}d\mathbf{w} \quad (2)$$

where  $\mathbf{X}_t$  represents the current process state with initial condition  $\mathbf{X}_0$  and  $\mathbf{w}$  is the typical Wiener process (Karatzas et al. 2014). The noise schedule is represented by a non-negative function  $\beta_t$ . The mean evolution here represents an interpolation that ends roughly at the distribution of average voice phoneme characteristics  $\bar{\mathbf{X}}$  at  $t = 1$  and begins at the distribution of source  $\mathbf{X}_0$  at  $t = 0$ . There is an associated reverse SDE with the forward SDE as given in Eq. 2.

$$d\mathbf{X}_t = \left[ -\frac{1}{2}\beta_t(\bar{\mathbf{X}} - \mathbf{X}_t) + \beta_t\nabla_{\mathbf{X}_t}\log p_t(\mathbf{X}_t|\bar{\mathbf{X}}) \right] dt + \beta_t d\tilde{\mathbf{w}} \quad (3)$$

where  $d\tilde{\mathbf{w}}$  is the Wiener process which at each diffusion time-step moving backward. Here noise  $\beta_t = \beta_0 + t(\beta_1 - \beta_0)$  follows linear schedule for both  $\beta_0$  and  $\beta_1$  such that  $e^{-\int_0^1(\beta_s ds)}$  tends to zero (Popov et al. 2021a). Additionally, the reverse SDE has the same trajectory as the forward SDE; that is, it begins roughly with the distribution of average-voice and ends at  $t = 0$  with the distribution of source-targets. Here, we employ the score model  $s_\theta$  as the U-Net architecture from (Ronneberger, Fischer, and Brox 2015). The score model takes  $\mathbf{X}_t, t$  and regulated emotion intensity features  $e_{ir} = f(e_s, e_r, i)$ , where  $f(\cdot)$  represents function for the DVM approach. Hence, given the emotion intensity regularized embedding  $e_{ir}$ , we can then utilize the reverse SDE with the learned  $s_\theta(\mathbf{X}_t, \bar{\mathbf{X}}, t, e_{ir})$  to estimate the reconstructed  $\mathbf{X}_0$  from the average voice  $\bar{\mathbf{X}}$ .

While doing inference, a source sample in a neutral emotion is converted to a specified target emotion by utilizing a reference sample in the corresponding target emotion. Both the source emotional embedding and the reference emotional embedding are given as input to the DVM, which models the direction vector and produces an emotion intensity regularized embedding. This embedding and the learned score function are then used to transform the source sample into the reference sample emotion.

## Experimental Analysis

### Dataset

Our performance evaluation primarily uses the Emotional Speech Database (ESD) (Zhou et al. 2021). The ESD dataset contains 29 hours of 350 parallel recorded utterances in five distinct emotions, i.e., Neutral, Happy, Angry, Sad, and Sur-

prise, from ten English speakers, respectively. In addition, we utilized 36 hours of internally developed non-parallel Hindi emotional speech database recorded by 9 native Hindi speakers for evaluating performance across languages because emotional speech database is not publicly available for other Indian languages. To record this database, we have collected Hindi texts from stories and categorized these texts into four emotions, namely, Neutral, Happy, Angry, and Sad. Each speaker has recorded one hour of data on each emotion. We used a 90:10 train-test split for both ESD-English (excluding Chinese) and Hindi data. We presented our analysis for three emotional state conversion scenarios along with intensity control: neutral-to-happy, neutral-to-sad, and neutral-to-angry. Note, we use nonparallel reference emotional samples from different speaker for evaluation.

## Implementation Details

Initially, emotion-independent speech representation is achieved by using average phoneme-level Mel-spectrogram characteristics. The process begins with aligning speech frames with phonemes in the ESD dataset using the Montreal Forced Aligner (MFA) (McAuliffe et al. 2017). Once aligned, the Mel features of each phoneme are aggregated across the entire dataset to obtain their average Mel-spectrogram features. The Phoneme Encoder is then trained to minimize the mean square error between the output Mel-spectrograms and the ground truth average voice Mel-spectrograms. These ground truth spectrograms are derived by replacing each phoneme Mel-spectrogram feature in the input with its corresponding average feature calculated earlier. The encoder is trained for 300 epochs with a batch size of 32 and a base learning rate of  $5e-5$  using the Adam optimizer. Our decoder is built on top of DiffVC architecture (Popov et al. 2021a). The number of parameters in the Encoder is 8.5mn and the decoder is 118mn. The decoder is trained for 126 epochs with a batch size of 32 using the Adam optimizer with a base learning rate of  $1e-4$ . The noise schedule parameters  $\beta_0$  and  $\beta_1$  are set to 0.05 and 20.0, respectively. The model mainly operates on Mel-spectrograms with 80 Mel features and a sampling rate of 22.05 kHz. The entire training process takes approximately 4 hours for the encoder and 12 hours for the decoder, both trained on an NVIDIA L40S GPU.

## Baseline Methods

- **EmoVox (Zhou et al. 2022c):** This approach explicitly controls the emotion intensity in the EVC task by using characteristic features to express fine-grained emotion intensity and further learn the actual emotion encoder from an emotion-labeled database.
- **Mixed Emotion (Zhou et al. 2022a):** It is a two stages seq-to-seq training method for emotional voice conversion. It uses a multi-speaker TTS corpus to perform style initialization in stage 1 to separate content from speaking style. Stage 2 uses a limited amount of emotional speech data for emotion training to learn how to disentangle the speech’s linguistic and emotional style information.

- **Other baselines:** We have also considered CycleGAN-EVC (Fu et al. 2021), StarGAN-EVC (Rizos et al. 2020), StyleVC (Du et al. 2021), DISSC (Maimon and Adi 2022) and Seq2Seq-EVC (Zhou, Sisman, and Li 2021a) to compare performance of proposed EmoReg approach w.r.t emotion voice conversion. However, we have excluded them from intensity control related evaluations as these approaches do not support emotion regularization.
- **Ablation EmoReg w/o DVM:** Here, we utilized the proposed EmoReg architecture without directional vector modeling. In particular, we applied to scale in the direction vector learned between the global mean of target emotion embedding reference and source emotion embedding along with intensity-related scaling function, i.e., via traditional interpolation-based strategy.

## Experimental Results

**Objective Evaluation** Since our key goal is obtaining control over the emotion intensity, we primarily considered the emotion similarity score as an objective evaluation metric. We use it to measure the effectiveness of the proposed approach for the emotion conversion and intensity regularization tasks. It is obtained by computing the cosine distance between the generated and ground truth speech emotion embeddings from the pre-trained emotion classifiers. Furthermore, Word Error Rate (WER) and Character Error Rate (CER) are considered to measure the intelligibility of the output while performing emotion intensity regularization on a continuous scale. Here, we used Whisper-small model (Radford et al. 2023). The emotion similarity score for emotion voice conversion is calculated for Neutral-to-Angry, Neutral-to-Sad, and Neutral-to-Happy emotion conversion scenarios for the proposed approach and baseline methods, as shown in Table 1. These are denoted by the abbreviations: Neu-Ang, Neu-Sad, and Neu-Hap, respectively. Similarly, we achieved improvements of 2 to 11% with the proposed DVM-based approach in AutoPCP-based objective evaluations for emotional similarity (Barrault et al. 2023) and scale-wise subjective evaluations for emotional controllability. Detail results are omitted due to space constraints.

Table 1 illustrates the comprehensive evaluation of the proposed EmoReg with DVM, which outperforms all the SOTA approaches in the EVC task. The performance of the EmoReg approach is further investigated on different intensity scores using emotion similarity score, as shown in Figure 4. We used an intensity scale range of 0 to 1 with a 0.2 step size and we could only consider EmoVox (Zhou

Methods	Neu-Ang $\uparrow$	Neu-Sad $\uparrow$	Neu-Hap $\uparrow$	Average $\uparrow$
Emovox	0.94 $\pm$ 0.004	0.94 $\pm$ 0.004	0.95 $\pm$ 0.004	0.94 $\pm$ 0.004
Mixed Emotion	0.94 $\pm$ 0.004	0.92 $\pm$ 0.004	0.90 $\pm$ 0.004	0.92 $\pm$ 0.004
CycleGAN-EVC	0.96 $\pm$ 0.004	0.92 $\pm$ 0.004	0.91 $\pm$ 0.004	0.93 $\pm$ 0.004
StarGAN-EVC	0.95 $\pm$ 0.004	0.91 $\pm$ 0.004	0.91 $\pm$ 0.004	0.93 $\pm$ 0.004
Seq2Seq-EVC	0.96 $\pm$ 0.004	0.93 $\pm$ 0.004	0.87 $\pm$ 0.004	0.92 $\pm$ 0.004
StyleVC	0.96 $\pm$ 0.004	0.92 $\pm$ 0.004	0.91 $\pm$ 0.004	0.93 $\pm$ 0.004
DISSC	0.88 $\pm$ 0.004	0.91 $\pm$ 0.004	0.87 $\pm$ 0.004	0.89 $\pm$ 0.004
Ablation	0.96 $\pm$ 0.004	0.93 $\pm$ 0.004	0.95 $\pm$ 0.004	0.94 $\pm$ 0.004
Proposed	<b>0.97 <math>\pm</math> 0.003</b>	<b>0.96 <math>\pm</math> 0.003</b>	<b>0.95 <math>\pm</math> 0.003</b>	<b>0.96 <math>\pm</math> 0.003</b>

Table 1: Analysis of emotion similarity scores along with margin of error corresponding to the 95% CI.




Figure 4: Analysis of emotion similarity score with respect to incremental emotion intensity scale.

et al. 2022c) and MixedEmotion (Zhou et al. 2022a) models, as others i.e., CycleGAN-StarGAN-EVC and Seq2Seq-EVC lack support for emotion intensity regularization due to their architectural limitations.

Additionally, from Figure 4, it is apparent that the emotion similarity score increases with an increase in emotion intensity scale which shows that the proposed EmoReg with DVM can achieve fine control over emotion intensity. Whereas, the emotion similarity score of the baselines and ablation does not vary with an increase in intensity scale and hence, fails to achieve fine control over emotion intensity.

To evaluate the intelligibility of EVC-generated speech, we calculated the WER and CER for baselines and proposed EmoReg. From Table 2, we can see that our approach achieves lower WER for the considered intensity scales and, on average, compared to EmoVox and Mixed Emotion. It further signifies that the proposed EmoReg effectively maintains speech intelligibility during EVC while allowing fine control over emotion intensity. Interestingly, the ablation model of EmoReg without DVM achieves lower WER at intensity scales of 0.0, 0.2, and 0.4. This could be due to the fact that, within a small range of emotional intensity, traversing within a closely related emotional space often leads to similar emotional state representations. However, our approach with DVM estimates the optimum direction for transitioning from one emotional state to another. Hence, after 0.6 intensity-related distance, the model achieves higher emotional similarity scores and lower WER/CER scores.

**Subjective Evaluation** For the subjective evaluation, we considered the Mean Opinion Score (MOS) metric on overall speech quality. A total of 25 individuals with no known hearing impairment between 25 and 35 years of age participated in the subjective evaluation. We asked subjects to rate each sample based on quality on a scale of 0 to 100, where 100 indicates the highest and 0 lowest quality. The evaluation of the proposed EmoReg, compared to baseline methods on the ESD dataset, is depicted in Figure 5. The figure shows that the proposed approach yields an absolute increment of 13.56% and 43.13% in the MOS value for quality compared to the baselines EmoVox and Mixed Emotion, respectively. This noticeable quality gap is evident in the samples we provided and samples from the baselines. Diffusion model-based architectures are one key reason for achiev-

Methods	0	0.2	0.4	0.6	0.8	1	Avg
WER ↓							
Emovox	34.82	34.52	34.43	35.86	39.15	38.57	36.23
Mixed Emotion	72.22	55.56	75.89	70	71.12	77.78	70.43
Ablation	<b>13.86</b>	<b>13.34</b>	<b>13.75</b>	14.28	13.76	14.29	13.88
<b>Proposed</b>	14.6	13.65	13.89	<b>13.09</b>	<b>13.25</b>	<b>13.65</b>	<b>13.69</b>
CER ↓							
Emovox	21.82	20.57	22.72	21.63	24.5	23.49	22.46
Mixed Emotion	36.57	30.1	38.65	36.11	38.89	43.06	37.23
Ablation	<b>8.27</b>	8.03	8.18	7.94	8.18	8.12	8.12
<b>Proposed</b>	8.7	<b>7.98</b>	<b>8.03</b>	<b>7.78</b>	<b>7.93</b>	<b>8.02</b>	<b>8.07</b>

Table 2: Comparison of WER and CER for proposed approach with baseline methods for English language.

Methods	Neu-Ang ↑	Neu-Sad ↑	Neu-Hap ↑	Avg ↑
English				
Ablation	0.96 ± 0.004	0.93 ± 0.004	0.95 ± 0.004	0.94 ± 0.004
<b>Proposed</b>	<b>0.97 ± 0.003</b>	<b>0.96 ± 0.003</b>	<b>0.95 ± 0.003</b>	<b>0.96 ± 0.003</b>
Hindi				
Ablation	0.89 ± 0.003	0.86 ± 0.003	<b>0.89 ± 0.003</b>	<b>0.88 ± 0.003</b>
<b>Proposed</b>	<b>0.91 ± 0.003</b>	<b>0.87 ± 0.003</b>	0.87 ± 0.003	<b>0.88 ± 0.003</b>

Table 3: Emotion Similarity scores across languages along with 95 % confidence interval.

ing high-quality output. Moreover, the proposed EmoReg architecture with DVM achieves a 5.52% absolute increment in MOS compared to the ablation model without DVM. This indicates that the proposed DVM is helping in traversing emotion embedding space with improved-quality samples. In addition, We also obtained a similar improvement in the performance (1%) with the proposed DVM in the non-diffusion-based architecture, namely VITS. Detail results are omitted due to space constraint.

**Performance Across Languages:** The effectiveness of the proposed approach is evaluated across different databases using similar objective and subjective assessments for both English and Hindi languages. Table 3 shows the emotion similarity scores for Neutral-to-Angry, Neutral-to-Sad, and Neutral-to-Happy emotion voice conversion for ablation and the proposed EmoReg approach for both languages. It is evident from Table 3 that the proposed approach also performs well for the Hindi language. Additionally, the intelligibility of converted speech in both languages is evaluated using WER and CER along with MOS score and is presented in Table 4.




Figure 5: MUSHRA-based MOS scores for speech quality for proposed EmoReg approach and baseline methods.




Figure 6: Visualization of Mel-spectrogram and pitch variation for emotion conversion from Neutral to Angry, Happy and Sad emotion with different intensity scale.

We conducted an ablation analysis of emotional similarity scores using different numbers of PCA components for both English and Hindi languages. PCA is primarily applied to the direction vectors obtained among all possible combinations of transitions from local mean vectors presented in two different emotional states. While these direction vectors primarily reflect changes related to emotional states, they may also be influenced by content and gender differences, as emotional representations are not entirely independent of language or gender. To identify direction vectors corresponding to emotional state differences, we selected three different numbers of principal components (64, 128, and 256) to capture variability, focusing primarily on emotional state transitions. As shown in Table 5, 128 components yield the best emotion similarity scores for both languages. Therefore, 128 principal components were chosen for the proposed DVM-based approach.

Methods	WER ↓		CER ↓		MOS ↑	
	English	Hindi	English	Hindi	English	Hindi
Ablation	13.88	33.33	8.12	<b>17.74</b>	71.15	80.07
Proposed	<b>13.69</b>	<b>32.93</b>	<b>8.07</b>	18.20	<b>76.67</b>	<b>80.58</b>

Table 4: Comparison of WER, CER and MUSHRA based MOS scores for the ablation and the proposed approaches.

Methods	English			Hindi		
	64	128	256	64	128	256
Neu-Ang ↑	<b>0.938</b>	<b>0.938</b>	0.937	0.889	<b>0.898</b>	0.896
Neu-Sad ↑	0.949	<b>0.950</b>	0.944	0.889	0.889	<b>0.896</b>
Neu-Hap ↑	0.911	<b>0.934</b>	0.928	0.849	<b>0.858</b>	0.851

Table 5: Ablation analysis of emotional similarity score for different number of PCA components.

## Visual Analysis

Additionally, we present the spectrographic visual analysis for various emotions, Neutral-to-Angry, Neutral-to-Sad, and Neutral-to-Happy in Figure 6. For Neutral-to-Angry emotion, we can see that with an increase in emotion intensity scale, pitch frequency range and dynamics also increase which is indicated by a white box. As harmonics are an integer multiple of the fundamental frequency, we can also see that for scales 0.6 and 1.0, the harmonics are wider than those on scale 0. It indicates that the proposed approach can control the intensity of emotion. Similar observations are seen in the case of Neutral-to-Happy emotion, where the pitch range increases with an increase in intensity scale. For the Neutral-to-Sad condition, the harmonics in the spectrogram are getting narrower with an increase in intensity scale because sadness is a low pitch-flattish emotion where the pitch range is usually lower as compared to high pitch with high variance emotions like anger and happiness.

## Conclusions

In this paper, we introduced the EmoReg model for emotion voice conversion with emotion intensity regularization. By leveraging SSL-based emotion embeddings, we achieved effective emotion representation from speech. We proposed a direction vector modeling (DVM) to transition between emotional states while controlling emotion intensity. We evaluated our approach against the SOTA architectures for both English and Hindi languages. Our approach demonstrated significant improvements of 13.56% over EmoVox and 43.13% over Mixed Emotion in speech quality. Additionally, the proposed EmoReg model outperformed existing methods in various objective evaluations. Moving forward, we aim to extend this work to enhance emotional intensity regularization in challenging environments, such as those with background music or noise.

## References

- Abdi, H.; and Williams, L. J. 2010. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4): 433–459.
- Aihara, R.; Takashima, R.; Takiguchi, T.; and Ariki, Y. 2012. GMM-based emotional voice conversion using spectrum and prosody features. *American Journal of Signal Processing*, 2(5): 134–138.
- Amiriparian, S.; Schuller, B. W.; Asghar, N.; Zen, H.; and Burkhardt, F. 2023. Guest editorial: special issue on affective speech and language synthesis, generation, and conversion. *IEEE Trans. on Affect. Comput.*, 14(1): 3–5.
- Barakat, H.; et al. 2024. Deep learning-based expressive speech synthesis: a systematic review of approaches, challenges, and resources. *EURASIP Journal on Audio, Speech, and Music Processing*, 2024(1): 11.
- Barrault, L.; Chung, Y.-A.; Meglioli, M. C.; Dale, D.; Dong, N.; Duppenhaler, M.; Duquenne, P.-A.; Ellis, B.; Elsahar, H.; Haaheim, J.; et al. 2023. Seamless: Multilingual Expressive and Streaming Speech Translation. *arXiv preprint arXiv:2312.05187*.
- Brannon, W.; et al. 2023. Dubbing in practice: A large scale study of human localization with insights for automatic dubbing. *Trans. of the Assoc. for Comput. Ling.*, 11: 419–435.
- Cho, D.-H.; Oh, H.-S.; Kim, S.-B.; Lee, S.-H.; and Lee, S.-W. 2024. EmoSphere-TTS: Emotional Style and Intensity Modeling via Spherical Emotion Vector for Controllable Emotional Text-to-Speech. In *INTERSPEECH*.
- Choi, H.; and Hahn, M. 2021. Sequence-to-sequence emotional voice conversion with strength control. *IEEE Access*, 9: 42674–42687.
- Du, Z.; Sisman, B.; Zhou, K.; and Li, H. 2021. Disentanglement of emotional style and speaker identity for expressive voice conversion. *arXiv preprint arXiv:2110.10326*.
- Ekman, P.; and Friesen, W. V. 1971. Constants across cultures in the face and emotion. *Journal of personality and social psychology*, 17(2): 124.
- Fu, C.; Liu, C.; Ishi, C. T.; and Ishiguro, H. 2021. Cycletransgan-vec: A cyclegan-based emotional voice conversion model with transformer. *arXiv preprint arXiv:2111.15159*.
- Gao, Z.; Guo, J.; Tan, X.; Zhu, Y.; Zhang, F.; Bian, J.; and Xu, L. 2022. Difformer: Empowering diffusion models on the embedding space for text generation. *arXiv preprint arXiv:2212.09412*.
- Gudmalwar, A.; Shah, N.; Akarsh, S.; Wasnik, P.; and Shah, R. R. 2024. VECL-TTS: Voice identity and Emotional style controllable Cross-Lingual Text-to-Speech. In *INTERSPEECH*.
- Gunes, H.; Schuller, B.; Pantic, M.; and Cowie, R. 2011. Emotion representation, analysis and synthesis in continuous space: A survey. In *IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, 827–834.
- Gutentag, T.; Halperin, E.; Porat, R.; Bigman, Y. E.; and Tamir, M. 2017. Successful emotion regulation requires both conviction and skill: Beliefs about the controllability of emotions, reappraisal, and regulation success. *Cognition and Emotion*, 31(6): 1225–1233.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, C.; et al. 2021. Neural dubber: Dubbing for videos according to scripts. *Advances in neural information processing systems*, 34: 16582–16595.
- Inanoglu, Z.; and Young, S. J. 2007. A system for transforming the emotion in speech: combining data-driven conversion techniques for prosody and voice quality. In *INTERSPEECH*, 490–493.
- Karatzas, I.; et al. 2014. *Brownian motion and stochastic calculus*, volume 113. springer.
- Li, Y. A.; Zare, A.; and Mesgarani, N. 2021. Starganv2-vc: A diverse, unsupervised, non-parallel framework for natural-sounding voice conversion. *arXiv preprint arXiv:2107.10394*.
- Lorenzo-Trueba, J.; Henter, G. E.; Takaki, S.; Yamagishi, J.; Morino, Y.; and Ochiai, Y. 2018. Investigating different representations for modeling and controlling multiple emotions in DNN-based speech synthesis. *Speech Communication*, 99: 135–143.
- Ma, Z.; et al. 2023. emotion2vec: Self-supervised pre-training for speech emotion representation. *arXiv preprint arXiv:2312.15185*.
- Maimon, G.; and Adi, Y. 2022. Speaking style conversion in the waveform domain using discrete self-supervised units. *arXiv preprint arXiv:2212.09730*.
- Matsumoto, K.; Hara, S.; and Abe, M. 2020. Controlling the Strength of Emotions in Speech-Like Emotional Sound Generated by WaveNet. In *INTERSPEECH*, 3421–3425.
- McAuliffe, M.; Socolof, M.; Mihuc, S.; Wagner, M.; and Sonderegger, M. 2017. Montreal forced aligner: Trainable text-speech alignment using kald. In *Interspeech*, volume 2017, 498–502.
- Mhaskar, S.; Shah, N.; Zaki, M.; Gudmalwar, A.; Wasnik, P.; and Shah, R. 2024. Isometric Neural Machine Translation using Phoneme Count Ratio Reward-based Reinforcement Learning. In *Findings of the Association for Computational Linguistics: NAACL 2024*, 3966–3976. Mexico City, Mexico: Association for Computational Linguistics.
- Ming, H.; Huang, D.-Y.; Xie, L.; Wu, J.; Dong, M.; and Li, H. 2016. Deep Bidirectional LSTM Modeling of Timbre and Prosody for Emotional Voice Conversion. In *Interspeech*, volume 2016, 2453–2457.
- Popov, V.; et al. 2021a. Diffusion-based voice conversion with fast maximum likelihood sampling scheme. *arXiv preprint arXiv:2109.13821*.
- Popov, V.; et al. 2021b. Grad-tts: A diffusion probabilistic model for text-to-speech. In *International Conference on Machine Learning*, 8599–8608. PMLR.
- Posner, J.; Russell, J. A.; and Peterson, B. S. 2005. The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and psychopathology*, 17(3): 715–734.

- Prabhu, N. R.; Lay, B.; Welker, S.; Lehmann-Willenbrock, N.; and Gerkmann, T. 2024. EMOCONV-Diff: Diffusion-Based Speech Emotion Conversion for Non-Parallel and in-the-Wild Data. In *ICASSP*, 11651–11655.
- Prabhu, N. R.; Lehmann-Willenbrock, N.; and Gerkmann, T. 2023. In-the-wild speech emotion conversion using disentangled self-supervised representations and neural vocoder-based resynthesis. In *Speech Communication; 15th ITG Conference*, 176–180. VDE.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, 28492–28518. PMLR.
- Reynolds, D. A.; et al. 2009. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663).
- Rizos, G.; Baird, A.; Elliott, M.; and Schuller, B. 2020. Star-gan for emotional speech conversion: Validated by data augmentation of end-to-end emotion recognition. In *ICASSP*, 3502–3506.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Russell, J. A. 1980. A circumplex model of affect. *Journal of personality and social psychology*, 39(6): 1161.
- Sahipjohn, N.; Gudmalwar, A.; Shah, N.; Wasnik, P.; and Shah, R. R. 2020. DubWise: Video-Guided Speech Duration Control in Multimodal LLM-based Text-to-Speech for Dubbing. In *INTERSPEECH*. Kos Island, Greece.
- Schnell, B.; and Garner, P. N. 2021. Improving emotional TTS with an emotion intensity input from unsupervised extraction. In *Proc. 11th ISCA Speech Synth. Workshop*, 60–65.
- Shah, N.; Singh, M.; Takahashi, N.; and Onoe, N. 2023. Nonparallel emotional voice conversion for unseen speaker-emotion pairs using dual domain adversarial network & virtual domain pairing. In *ICASSP 2023*, 1–5.
- Um, S.-Y.; Oh, S.; Byun, K.; Jang, I.; Ahn, C.; and Kang, H.-G. 2020. Emotional speech synthesis with rich and granularized control. In *ICASSP*, 7254–7258.
- Wu, Y.; et al. 2023. Videodubber: Machine translation with speech-aware length control for video dubbing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 13772–13779.
- Zhou, K.; Sisman, B.; Busso, C.; Ma, B.; and Li, H. 2022a. Mixed-EVC: Mixed Emotion Synthesis and Control in Voice Conversion. *arXiv preprint arXiv:2210.13756*.
- Zhou, K.; Sisman, B.; and Li, H. 2021a. Limited data emotional voice conversion leveraging text-to-speech: Two-stage sequence-to-sequence training. *arXiv preprint arXiv:2103.16809*.
- Zhou, K.; Sisman, B.; and Li, H. 2021b. Vaw-gan for disentanglement and recombination of emotional elements in speech. In *IEEE spoken language technology workshop (SLT)*, 415–422.
- Zhou, K.; Sisman, B.; Liu, R.; and Li, H. 2021. Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 920–924.
- Zhou, K.; Sisman, B.; Liu, R.; and Li, H. 2022b. Emotional voice conversion: Theory, databases and ESD. *Speech Communication*, 137: 1–18.
- Zhou, K.; Sisman, B.; Rana, R.; Schuller, B. W.; and Li, H. 2022c. Emotion intensity and its control for emotional voice conversion. *IEEE Trans. on Affect. Comput.*, 14(1): 31–48.
- Zhu, X.; Yang, S.; Yang, G.; and Xie, L. 2019. Controlling emotion strength with relative attribute for end-to-end speech synthesis. In *IEEE ASRU*, 192–199.