

Textless Speech Emotion Conversion using Discrete & Decomposed Representations

Felix Kreuk^{1,2,*}, Adam Polyak², Jade Copet², Eugene Kharitonov², Tu-Anh Nguyen², Morgane Rivière², Wei-Ning Hsu², Abdelrahman Mohamed², Emmanuel Dupoux^{2,3}, Yossi Adi²

Bar-Ilan University, Ramat-Gan, Israel¹

Meta AI Research² EHESS, Paris³

{felixkreuk, adiyoss}@fb.com

Abstract

Speech emotion conversion is the task of modifying the perceived emotion of a speech utterance while preserving the lexical content and speaker identity. In this study, we cast the problem of emotion conversion as a spoken language translation task. We use a decomposition of the speech signal into discrete learned representations, consisting of phonetic-content units, prosodic features, speaker, and emotion. First, we modify the speech content by translating the phonetic-content units to a target emotion, and then predict the prosodic features based on these units. Finally, the speech waveform is generated by feeding the predicted representations into a neural vocoder. Such a paradigm allows us to go beyond spectral and parametric changes of the signal, and model non-verbal vocalizations, such as laughter insertion, yawning removal, etc. We demonstrate objectively and subjectively that the proposed method is vastly superior to current approaches and even beats text-based systems in terms of perceived emotion and audio quality. We rigorously evaluate all components of such a complex system and conclude with an extensive model analysis and ablation study to better emphasize the architectural choices, strengths and weaknesses of the proposed method. Samples are available under the following link: [\[samples\]](#).

1 Introduction

Generating spoken utterances and dialogue that sound natural is a fundamental requirement to improve human-computer interaction (Tits et al., 2019). One of the main roadblock in improving naturalness in speech generation is the modeling of expressive and emotional states. The difficulty is that emotion is a phenomenon affecting all linguistic levels simultaneously: when one goes from a happy

* Work done while Felix Kreuk was an Intern at Meta AI Research.




Figure 1: An illustration of the proposed system. The input signal is first encoded as a discrete sequence of content units (E_c). Next, a sequence to sequence ($S2S$) model is applied to translate between the sequences corresponding to different emotions. Then we predict the duration (E_{dur}) and F0 (E_{F0}) before feeding these signals to a vocoder (G). The last 4 components are conditioned by emotion (z_{emo}). We use pink to denote models and green to denote representations.

to an angry state, one may use different vocabulary, insert non-verbal vocalizations (cries, grunts, etc), modify prosody (intonation and rhythm), and change voice quality due to stress. Vice versa each of the levels contributes to the perception of the emotional state of the speaker, where the non verbal aspects can often override the lexical content (Mehrabian and Wiener, 1967).

Existing emotion generation or emotion conversion techniques have a hard time producing convincing results because they only manage to tackle a subset of these levels. In a nutshell, signal-based approaches are mainly focused on manipulating parameters of the speech signal (Inanoglu and Young, 2009; Aihara et al., 2012; Gao et al., 2018; Robinson et al., 2019; Polyak et al., 2020a; Rizos et al., 2020; Zhou et al., 2020a,b, 2021a,d,b), and can only address changes at the level of voice and prosody. In contrast, text based approaches (Skerry-Ryan et al., 2018a; Ren et al., 2020; Kim et al., 2021; Sorin et al., 2020; Hono et al., 2020;

Tan et al., 2021; Polyak et al., 2020b) can generate expressive speech, but struggle with nonverbal vocalizations because they are typically not annotated in speech corpora.

In this work we focus on the task of speech emotion conversion under the parallel dataset setting, modifying the perceived emotion of a speech utterance while preserving the speaker identity and the lexical content (Zhou et al., 2021c).

The proposed pipeline is comprised of four main blocks: speech tokenizer, content translation model, prosody prediction model, and a neural vocoder. We start by extracting discrete representation of the speech signal. We translate these representation to a target emotion (e.g., removing laughter, inserting yawning) while preserving the lexical content. Then, we predict prosodic features based on the translated representations. A neural vocoder synthesizes the speech waveform from the translated phonetic-content, predicted prosody, speaker label and target emotion label. The overall system is depicted in Figure 1.

Our contribution is twofold: We propose a novel textless approach by casting the task of speech emotion conversion as a *spoken language translation* problem. We demonstrate how such a paradigm can be used to model expressive non-verbal communication cues as well as generating high-quality speech samples. Finally, we demonstrate for the first time the coverage of all levels of expressive speech modeling simultaneously. Results show that the proposed method is vastly superior to emotion conversion techniques based on signal only, and also beats text-based approaches in terms of generation quality and perceived emotion.

2 Related work

Speech emotion conversion. There is an increasing number of studies for speech emotion conversion recently (Zhou et al., 2021c). Early studies on speech emotional conversion include Gaussian Mixture Model (GMM) technique (Aihara et al., 2012), sparse representation technique (Aihara et al., 2014), or an incorporated framework of Hidden Markov Model (HMM), GMM and fundamental frequency segment selection method (Inanoglu and Young, 2009). Recently, speech emotion conversion has seen a great leap in performance following the advances in deep nets. This include using multi-layer perceptron (Luo et al., 2016; Lorenzo-Trueba et al.,

2018), highway networks (Shankar et al., 2019a), deep bi-directional LSTM networks (Ming et al., 2016), and sequence-to-sequence model (Robinson et al., 2019). When considering non-parallel data, new techniques have been proposed to learn the translation between emotional domains with CycleGAN (Zhou et al., 2020a) and StarGAN (Rizos et al., 2020), to disentangle the emotional elements from speech with auto-encoders (Cao et al., 2020; Zhou et al., 2021b), and to leverage text-to-speech (Kim et al., 2020) or Automatic Speech Recognition (ASR) (Liu et al., 2020). Studies have also revealed that the emotions can be expressed through universal principles that are shared across different individuals and cultures (Ekman, 1992; Manokara et al., 2021). This motivates research on multispeaker (Shankar et al., 2019b, 2020), and speaker-independent emotion conversion (Zhou et al., 2020b; Choi and Hahn, 2021).

Unsupervised speech representation learning aims to distill features useful for downstream tasks, such as phone discrimination (Kharitonov et al., 2021b; Schneider et al., 2019) and semantic prediction (Lai et al., 2021; Wu et al., 2020), by constructing pretext tasks that can exploit large quantities of unlabeled speech. Pretext tasks in the literature can be roughly divided into two categories: reconstruction and prediction. Reconstruction is often implemented in the form of auto-encoding (van den Oord et al., 2017), where speech is first encoded into a low-dimensional space, and then decoded back to speech. Various constraints can be imposed on the encoded space, such as temporal smoothness (Ebbers et al., 2017; Glarner et al., 2018; Khurana et al., 2020), discreteness (Ondel et al., 2016; van den Oord et al., 2017), presence of hierarchy (Hsu et al., 2017), and information bottlenecks for speech representation decomposition (Qian et al., 2020, 2021). Prediction-based approaches task a model with predicting information of unseen speech based on its context. Examples of information include spectrograms (Wang et al., 2020; Chi et al., 2021; Chung and Glass, 2020), cluster indices (Hsu et al., 2021), derived signal processing features (Ravanelli et al., 2020), and binary labels of whether a candidate is the target unseen spectrogram (Schneider et al., 2019; Kharitonov et al., 2021b; Baevski et al., 2020).

Spoken Language Modeling. The closest related work to ours concerns with modeling spo-

ken language without textual supervision. [Lakhotia et al. \(2021\)](#) suggested utilizing recent success of self-supervised speech representation learning for discovering discrete units and modeling them. Then, conditional and unconditional speech generation can be achieved by sampling unit sequences from a unit-Language-Model, and synthesizing them to speech using a units-to-Speech model ([Lakhotia et al., 2021](#)). Moreover, the authors proposed a set of evaluation functions to assess the quality of the overall system. Next, [Polyak et al. \(2021\)](#) evaluated the robustness and disentanglement properties of several Speech-to-units models and demonstrated the ability to perform voice conversion, and lightweight speech codec. Lastly, [Kharitonov et al. \(2021a\)](#) proposed a multi-stream transformer model to jointly process "pseudo-text" units and prosodic features (i.e., duration and F0). This was shown beneficial in improving content modeling capabilities. In this work, we leverage a similar approach and speech representation scheme but for speech emotion conversation via translation.

3 Model

As emotion manifests itself in multiple aspects of spoken language, to optimally convert emotion one needs to consider all aspects in the conversion process. For example, emotion can be expressed via a prosodic features (high pitch, slow speaking rate, etc.), speaking style (yelling, whispering, etc.), and non-verbal vocalizations (laughing, yawning, etc.).

We suggest to use a decomposed representation of the speech signal to synthesize speech in the target emotion. We consider four components in the decomposition: phonetic-content, prosodic features (i.e., F0 and duration), speaker identity, and emotion-label, denoted by $\mathbf{z}_c, (\mathbf{z}_{dur}, \mathbf{z}_{F0}), \mathbf{z}_{spk}, \mathbf{z}_{emo}$ respectively.

Specifically, we propose the following cascaded pipeline: (i) extract \mathbf{z}_c from the raw waveform using a Self-Supervised Learning (SSL) model; (ii) translate non-verbal vocalizations in \mathbf{z}_c while preserving the lexical content (e.g., when converting from amused to sleepy, we remove laughter and insert yawning); (iii) predict the prosodic features of the target emotion based on the translated content; (iv) synthesize the speech from the translated content, predicted prosody, target speaker identity and target emotion-label. See [Figure 1](#) for a visual description of the method.

3.1 Speech Input Representation

Phonetic-content representation. To represent speech phonetic-content we extract a discrete representation of the audio signal using a pre-trained SSL model, namely HuBERT. We use a SSL representation for phonetic-content in order to capture non-verbal vocalizations (unlike text where they are often not annotated). We discretize this representation for better modeling and sampling (as opposed to regressing on continuous variables). This paradigm allows us to benefit from all recent advances in Natural Language Processing (NLP). We chose HuBERT for the phonetic-content units as it was shown to better disentangle between speech content and both speaker and prosody compared to other SSL-based models ([Polyak et al., 2021](#)).

Denote the domain of audio samples by $\mathcal{X} \subset \mathbb{R}$. The representation for an audio waveform is therefore a sequence of samples $\mathbf{x} = (x^1, \dots, x^T)$, where each $x^i \in \mathcal{X}$ for all $1 \leq t \leq T$. The content encoder E_c is a HuBERT model ([Hsu et al., 2021](#)) pre-trained on the LibriSpeech corpus ([Panayotov et al., 2015](#)). HuBERT is a self-supervised model trained on the task of masked prediction of continuous audio signals, similarly to BERT ([Devlin et al., 2019](#)). During training, the targets are obtained via clustering of MFCCs features or learned representations from earlier iterations. The input to the content encoder E_c is an audio waveform \mathbf{x} , and the output is a spectral representation sampled at a lower frequency $\mathbf{z}' = (z_c'^1, \dots, z_c'^L)$ where $L < T$. Since HuBERT outputs continuous representations, an additional k-means step is needed in order to quantize these representations into a discrete unit sequence denoted by $\mathbf{z}_c = (z_c^1, \dots, z_c^L)$ where $z_c^i \in \{1, \dots, K\}$ and K is the size of the vocabulary. For the rest of the paper, we refer to these discrete representations as "units". We extracted representations from the 9th layer of HuBERT model and set $K = 200$. Following [Lakhotia et al. \(2021\)](#), repeated units were omitted (e.g., 0, 0, 0, 1, 1, 2 \rightarrow 0, 1, 2). We denote such sequences by "deduped".

Speaker representation. Our goal is to convert speech emotion while keeping the speaker identity fixed. To that end, we construct a speaker-representation \mathbf{z}_{spk} , and include it as an additional conditioner during the waveform synthesis phase. To learn \mathbf{z}_{spk} we optimize the parameters of a fixed size look-up-table. Although such modeling limits our ability to generalize to new and unseen speak-

ers, it produces higher quality generations (Polyak et al., 2021). We additionally experimented with representing \mathbf{z}_{spk} as a d -vector using a method similar to Heigold et al. (2016). However, we observed that such approach keeps source emotion prosodic features, resulting in inferior disentanglement during the waveform synthesis phase.

Emotion-label representation. We represent the emotion-label using a categorical variable represented by a 1-hot vector. We observed that this component controls for timbre characteristics of the generated speech signal (e.g., roughness, smoothness, etc.) during the synthesis phase.

3.2 Speech Emotion Conversion

Using the above representations we propose to synthesize the speech signal in the target emotion. We use a translation model to convert between phonetic-content units of a source emotion to phonetic-content units of the target emotion. This serves as a learnable insertion/deletion/substitution mechanism for non-verbal vocalizations, while preserving the lexical content (e.g., removing yawning while preserving the verbal content). Next, we predict the prosodic features (duration and F0) based on the translated phonetic-content units and target emotion-label, and inflate the sequence according to the predicted durations. This will later be used as a conditioning for the waveform synthesis phase.

Unit translation. To translate the speech content units, we use a sequence-to-sequence Transformer model (Vaswani et al., 2017) denoted by E_{s2s} . The input to E_{s2s} is a sequence of phonetic-content units representing a speech utterance in the source emotion \mathbf{z}_c^{src} . The model is trained to output a sequence of phonetic-content units \mathbf{z}_c^{tgt} containing the same lexical content with the addition/deletion/substitution of speech cues related to emotion expression (e.g., inserting laughter units). The optimization minimizes the cross-entropy (CE) loss between the predicted units $\hat{\mathbf{z}}_c = E_{s2s}(\mathbf{z}_c^{src}, \mathbf{z}_{emo})$ and ground-truth units for each location in the sequence,

$$L_{s2s} = \sum_{i=1}^L CE(E_{s2s}(\mathbf{z}_c^{src}, \mathbf{z}_{emo})_i, \mathbf{z}_c^{tgt}_i). \quad (1)$$

In our experiments, we observed that directly optimizing the above model for translation captures emotion transfer, but fails to maintain the same lexical content, producing expressive yet unintelligible




Figure 2: An illustration of the sequence-to-sequence emotion translation component E_{s2s} . Here different encoders and decoders are used for each emotions, but we also tested shared architectures in Appendix A.2.

speech utterances (e.g., the model adds laughter but corrupts the sentence by removing needed syllables). To mitigate this, we pre-train the translation model on the task of language denoising auto-encoder similarly to BART (Lewis et al., 2020). To better support translation between all emotions we use a dedicated encoder and decoder for each emotion (see Figure 2). We additionally evaluated a shared-encoder and shared-decoder model, in which we condition on the target emotion, as well as a model where only the encoder is shared. We evaluated all approaches, results are summarized in Table 3 Appendix A.2.

Prosody prediction. Next, using the translated phonetic-content unit sequence we predict the prosodic features corresponding to the target emotion. We consider the prosodic representation as a tuple of content unit durations and F0.

We start by describing the duration prediction process. Due to working on deduped sequences, we first need to predict the duration of each phonetic-content unit. We follow a similar approach to the one proposed in (Ren et al., 2020) and use a Convolutional Neural Network (CNN) to learn the mapping between content units to durations. We denote this model by E_{dur} . During training of E_{dur} , we input the deduped phonetic-content units \mathbf{z}_c and use the ground-truth phonetic-content unit durations as supervision. We minimize the Mean Squared Error (MSE) between the network’s output and the target durations. We also evaluated n-gram-based duration prediction models. The n-gram models were trained by counting the mean frequency μ and the standard deviation σ of each n-gram in the training set. During inference, we predict the duration of each n-gram by sampling from $\mathcal{N}(\mu, \sigma)$. For unseen n-grams we back-off to a smaller n-gram model. Results can be found on

Table 5 in the Appendix A.2.

We now turn to describe the F0 prediction process. We use a F0 estimation model to predict the pitch from a sequence of phonetic-content units \mathbf{z}_c . Our model, denoted by E_{F0} , is a CNN, followed by a linear layer projecting the output to \mathbb{R}^d . The final activation layer is set to be a sigmoid such that the network outputs a vector in $[0, 1]^d$. We extract the F0 using the YAAPT (Kasi and Zahorian, 2002) algorithm to serve as targets during training. Next, we normalize the F0 values using the mean and standard deviation per speaker. We follow a similar approach to Kim et al. (2018) and discretize the range of F0 values into d bins represented by one-hot encodings. Next, we apply Gaussian-blur on these encodings to get the final supervision targets denoted by $\mathbf{z}_{F0} = (z_{F0}^1, \dots, z_{F0}^d)$ where each $z_{F0}^i \in [0, 1]^d$, and $d = 50$. Formally, we minimize the binary-cross-entropy (BCE) for each coordinate of the target and the network output,

$$L_{F0} = \sum_{i=1}^d BCE(E_{F0}(\mathbf{z}_c, \mathbf{z}_{emo}^{tgt})_i, z_{F0}^i). \quad (2)$$

During inference multiple frequency bins can be activated to a different extent, similarly to Kim et al. (2018) we output the F0 value corresponding to the weighted-average of the activated bins. This modeling allows for a better output range when converting bins back to F0 values, as opposed to a single representative F0 value per bin. For F0-to-bin conversion, we use an adaptive binning strategy such that the probability mass of each bin is the same, similarly to Kharitonov et al. (2021a). Results are summarized on Table 4 Appendix A.2 together with additional comparisons to log-F0 estimation models, uniform binning, and argmax decoding.

Notice, the mapping between discrete unit sequences and prosodic feature (F0 and durations) is one-to-many, as it depends on the target emotion. Hence, we additionally condition both E_{dur} and E_{F0} on the target emotion denoted by \mathbf{z}_{emo}^{tgt} . Both model were trained independently.

Speech synthesis. We follow Polyak et al. (2021), and use a variation of the HiFi-GAN neural vocoder (Kong et al., 2020). The architecture of HiFi-GAN consists of a generator G and a set of discriminators D . We adapt the generator component to take as input a sequence of predicted phonetic-content units inflated using the predicted

durations, predicted F0, target speaker-embedding, and a target emotion-label. The above features are concatenated along the temporal axis and fed into a sequence of convolutional layers that output a 1-dimensional signal. The sample rates of unit sequence and F0 are matched by means of linear interpolation, while the speaker-embedding and emotion-label are replicated.

The discriminators are comprised of two sets: Multi-Scale Discriminators (MSD) and Multi-Period Discriminators (MPD). The first type operates on different sizes of a sliding window over the signal (2, 4), while the latter samples the signal at different rates (2,3,5,7,11). Overall, each discriminator D_i is trained by minimizing the following loss function, $L_{adv}(D_i, G) = \sum_{\mathbf{x}} \|1 - D_i(\hat{\mathbf{x}})\|_2^2$ and $L_D(D_i, G) = \sum_{\mathbf{x}} \|1 - D_i(\mathbf{x})\|_2^2 + \|D_i(\hat{\mathbf{x}})\|_2^2$, where $\hat{\mathbf{x}} = G(\hat{\mathbf{z}}_c, E_{F0}(\hat{\mathbf{z}}_c, \mathbf{z}_{emo}), \mathbf{z}_{spk}, \mathbf{z}_{emo})$ is the time-domain signal reconstructed from the decomposed representation. There are two additional loss terms used for optimizing G . The first is a Mean-Absolute-Error (MAE) reconstruction loss in the log-mel frequency domain $L_{recon}(G) = \sum_{\mathbf{x}} \|\phi(\mathbf{x}) - \phi(\hat{\mathbf{x}})\|_1$. where ϕ is the spectral operator computing the Mel-spectrogram. The second loss term is a feature matching loss, which penalizes for large discrepancies in the intermediate discriminator representations, $L_{fm}(D_i, G) = \sum_{\mathbf{x}} \sum_{j=1}^R \|\xi_j(\mathbf{x}) - \xi_j(\hat{\mathbf{x}})\|_1$ where ξ_j is the operator that extracts the intermediate representation of the j th layer of discriminator D_i with R layers. The overall objective for optimizing the system is:

$$L_G(D, G) = \left[\sum_{i=1}^J L_{adv}(D_i, G) + \lambda_{fm} L_{fm}(D_i, G) \right] + \lambda_r L_{recon}(G),$$

and $L_D(D, G) = \sum_{i=1}^J L_D(D_i, G)$, where $\lambda_{fm} = 2$ and $\lambda_r = 45$.

4 Experimental Setup

We use the Emotional Voices Database (EmoV) for training and evaluating our model. EmoV consists of 7000 speech utterances based on transcripts from the CMU Arctic Database (Kominek and Black, 2004). Each transcript was recorded in multiple acted emotions (neutral, amused, angry, sleepy, disgusted) by multiple native speakers (two males speakers and two females speakers).

This allows us to create a dataset of utterance pairs for the task of translation. Specifically, we

create pairs of utterances that are based on the same transcript but are recorded with different acted emotions. Due to the small size of this dataset (~ 9 hours), we further augment it by creating additional parallel pairs from different speakers. Overall, the size of the entire dataset is 78,324 pairs. We split the data to train/validation/test sets with a ratio of 90/5/5 such that there is no overlap of utterances between the sets. In our experiments, splitting randomly (e.g., overlapping transcripts) led to a memorization of the utterances and failed to generalize to unseen data. We explore two data regimes for pre-training stage: (1) large scale pre-training on a mix of Librispeech, Blizzard2013 (Chalamandaris et al., 2013) and EmoV (Adigwe et al., 2018). Denoted by OURS⁺; (2) small scale setup in which we pre-train on VCTK (Yamagishi et al., 2019) and EmoV. We explore the latter for a fair comparison with the evaluated baselines. We denote this setup by OURS⁻. A full description of architecture hyperparameters and pre-training details can be found in the Appendix A.1.

4.1 Baselines

We compare the proposed method to a textless speech emotion conversion method, VAWGAN (Zhou et al., 2020b), as well as a State-of-The-Art (SOTA) text-based emotional voice conversion model, Seq2seq-EVC (Zhou et al., 2021a). We also evaluate an expressive Text-to-Speech (TTS) system based on Tacotron2 (Shen et al., 2018).¹

For the text-based approach we use the Tacotron2 and Seq2seq-EVC models. The input to Tacotron2 is the ground-truth text representing the speech content. We modify the Tacotron2 architecture by adding a Global-Style-Token (Skerry-Ryan et al., 2018b) to control for the target emotion. The inputs to Seq2seq-EVC are the ground-truth phonemes coupled with the source speech utterance. The output of both systems is the Mel-spectrogram of the speech utterance in the target emotion. To reconstruct the time-domain signal we use the HiFi-GAN vocoder.

All baselines were trained and evaluated on the EmoV dataset. Seq2seq-EVC and Tacotron2 were first pre-trained on VCTK.

4.2 Evaluation

Subjective Evaluation. Recall, our goal is to perform speech emotion conversion. To that end,

¹We consider text-based systems as ones using textual annotations during training.

# units	layer no.	pre-trained	WER	PER
100	6	✗	88.85	72.17
100	6	✓	36.01	29.67
100	9	✗	83.72	66.63
100	9	✓	31.69	27.90
200	6	✗	91.06	77.14
200	6	✓	31.97	27.49
200	9	✗	84.92	71.45
200	9	✓	23.92	25.95

Table 1: Evaluation of different token extraction configurations and the effect of pre-training the translation model. # units denotes the vocabulary size, layer no. denotes the layer index used in HuBERT.

we propose a new subjective metric called Emotion-Mean-Opinion-Classification (eMOC). In an eMOC study, a human rater is presented with a speech utterance and a set of emotion categories. The rater is instructed to select the emotion that best fits the speech utterance. The eMOC score is the percentage of raters that selected the target emotion given a speech recording. The final score is averaged over all raters and utterances in the study. Additionally, we measure the perceived audio-quality using the Mean-Opinion-Score (MOS). When comparing MOS and eMOC against baselines, we report results for emotion conversion from Neutral as all baselines are constrained to Neutral as the source emotion. We provide additional results of our system in the any-to-any conversion setting. The CrowdMOS package (Ribeiro et al., 2011) was used in all subjective experiments with the recommended recipes for outliers removal. Participants were recruited using a crowd-sourcing platform.

ASR Based. Our goal is to perform speech emotion conversion while preserving the lexical content of the speech signal. However, to compare different translation models with different vocabularies, one cannot simply use metrics such as BLEU, hence, we report Word Error Rate (WER) and Phoneme Error Rate (PER) metrics extracted using a pre-trained SOTA ASR system². As ASR models suffer from performance degradation when evaluated on expressive speech, WER and PER metrics are reported on emotions converted to neutral: {amused, angry, sleepy, disgusted}-neutral.

5 Results

We first tune independently the different components of our system using objective metrics: the

²We use a BASE wav2vec 2.0 phoneme detection model trained on LibriSpeech-960h with CTC loss from scratch.

content units extraction configuration, and the prosodic modeling modules (F0 and duration estimators). Next we conduct a subjective evaluation for our best system in terms of audio-quality (MOS) and perceived emotions (eMOC) and compare the proposed method against the baselines (Section 4.1). Finally, we run an ablation study where we measure the impact of each component on the perceived emotion.

5.1 Model Tuning

Yang et al. (2021) found that intermediate HuBERT representations obtained from different layers have an impact on the downstream task at hand. Specifically, when considering spoken language modeling, Lakhotia et al. (2021) used the 6th layer, while Kharitonov et al. (2021a) used layer 9. Additionally, Lakhotia et al. (2021) showed that the number of units (k) also has an impact on the overall performance of the system. To better understand the effect of these architectural configurations in our setting, we experimented with extracting HuBERT features from layers 6 and 9, using 100 and 200 clusters for the k-means post-processing step. Additionally, we measure the impact that pre-training the translation model has on performance. As we compare models with different vocabulary sizes, we cannot use metrics such as BLEU, hence, we report WER and PER.

For emotion conversion, we find that using the 9th HuBERT layer and 200 tokens performs best. An evaluation of different model design architectures can be found in Table 1.

5.2 Subjective Evaluation

Figure 3 depicts the MOS, and eMOC scores for the proposed method and the evaluated baselines. For a fair comparison against baselines, all emotion were converted from Neutral. Results suggest that both OURS⁻ and OURS⁺ surpass the baselines in terms of both MOS and eMOC, with the exception MOS score for Angry and eMOC for Sleepy where OURS⁻ is comparable to Seq2seq-EVC.

While Tacotron2 and Seq2seq-EVC succeed in conveying the target emotion, they produce less natural expressive speech utterances, which is reflected in lower MOS. Unlike our method, both Tacotron2 and Seq2seq-EVC are text-based systems hence they attempt to learn an alignment (an attention map) between text inputs and audio targets. This task is particularly challenging as non-verbal cues (e.g., laughter, breathing) are

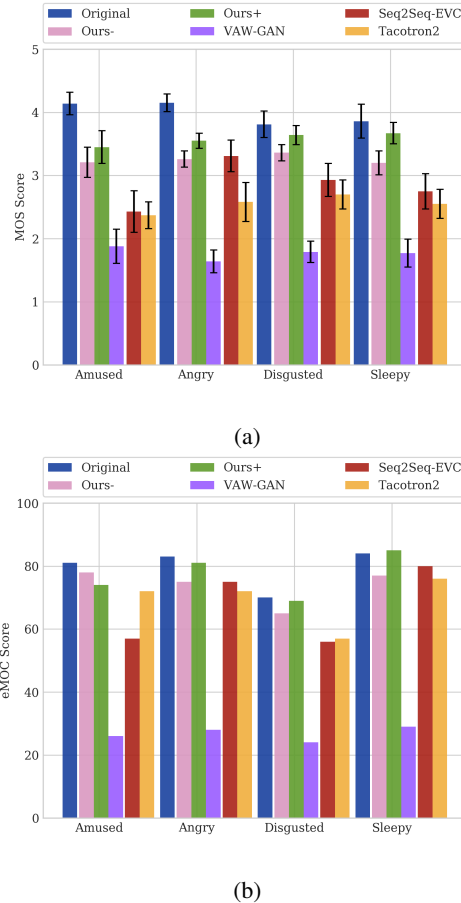



Figure 3: Sub-figure (a) reports the MOS, measuring the perceived audio-quality. We report mean scores with a Confidence Interval of 95%. Sub-figure (b) reports the eMOC score, measuring the perceived emotion. We report mean scores for each emotion (chance level: 25%). OURS⁺ was pre-trained on Librispeech, Blizzard2013 and EmoV, and OURS⁻ was pre-trained on VCTK and EmoV.


not annotated in the text inputs. We hypothesize that this misalignment leads to less natural expressive speech production. We provide additional eMOC results for the proposed method, evaluated baselines, and ground truth recordings in the Appendix A.2.

The evaluated baselines considered only Neutral as the source emotion. As the proposed method supports many-to-many³ emotion voice conversion, we additionally report MOS and eMOC results when converting from any emotion to any emotion. Figure 4 summarizes the MOS and eMOC scores for the many-to-many setting. We additionally provide the full eMOC results for the many-to-many setting in the Table 6 of the Appendix A.2. In terms of MOS, results suggest that converting to Neutral

³from a set of pre-defined emotions



(a) MOS



(b) eMOC

Figure 4: MOS and eMOC results for many-to-many emotion conversion. We report MOS score for each pair of source and target emotions. For eMOC we report percentage of raters that chose the target emotion.

and from Neutral produced the most natural sounding utterances. In terms of eMOC, having Amused as the source emotion surprisingly yielded the highest overall accuracy, with $\sim 80\%$ of conversions successful (average of first row in Figure 4b). On the other hand, the Sleepy emotion was the hardest to convert, with 12% of raters still choosing Sleepy after conversion. The pair of emotions that were hardest to distinguish are Disgusted and Neutral.

5.3 Ablation

Recall, we use a decomposed speech representation consisting of four feature sets. Here, we gauge the effect of each feature by gradually adding different components and evaluating their impact on the eMOC metric. Specifically, we start by evaluating the source features and replacing the emotion-token. Then, we predict the unit durations and F0 for the target emotion. Lastly, the full effect of our method is achieved by incorporating the unit translation model. For reference, we report results for

Tokens	Emotion	F0	Duration	eMOC			
				Amused	Angry	Disgusted	Sleepy
Original – Neutral				20.11	21.66	7.59	17.34
src	src	src	src	20.24	18.62	8.37	17.72
src	tgt	src	src	27.66	25.76	10.83	25.19
src	tgt	src	pred	30.89	35.71	32.11	70.31
src	tgt	pred	pred	79.12	86.11	69.21	79.12
pred	tgt	pred	pred	85.16	90.61	75.89	84.23
Original – Emotion				86.31	89.92	76.18	88.01

Table 2: Effect of components in our system on perceived emotion. First row denotes the ground-truth neutral recordings while last row denotes the ground-truth emotional recording. “src”, “tgt”, and “pred” denote features extracted from source speech, target speech, and predicted by our system respectively.

the original recording, and resynthesized one (i.e., using source features only) and the target recording. Results are summarized in Table 2.

Results suggest that the proposed method is comparable to ground-truth recordings in terms of the perceived emotion. Interestingly, for the Sleepy emotion, modifying the timbre using the target emotion-token and adjusting the unit durations is enough to reach 70.31%, while the rest of the emotions require further processing. For Amused, Angry and Disgusted modifying the F0 reaches performance of 5% below the ground-truth recordings. When applying the entire pipeline results are on par with the ground-truth.

5.4 Out-of-distribution samples

Due to the size of the EmoV dataset (7000 samples overall) the number of unique utterance is small. As a result, the model can memorize the utterances and might fail to generalize to unseen sentences. Hence, we experimented with converting out-of-domain recordings. To that end, we input our system with recordings from the LibriSpeech dataset. As LibriSpeech consists of non-expressive samples, we treat them as “neutral”. We convert these samples to different emotions and evaluate the performance. For evaluation we randomly sampled 20 utterances converted to each emotion. The proposed method reaches an average eMOC score of $82.25\% \pm 7.32$ and an average MOS score of 3.69 ± 0.26 across all four emotions (amused, angry, disgusted, sleepy). When considering the lexical content, the WER between the ground-truth text and ASR based transcriptions of the generated audio is 27.92. These results are similar to the ones reported for the EmoV dataset.

6 Conclusion

We presented a novel textless method for speech emotion conversion using decomposed and discrete representations. Using such representations we suggested casting speech emotion conversion as a spoken language translation problem. We learn a mapping to translate between discrete speech units from one emotion to another. Results suggest the proposed method outperforms the evaluated baselines by a great margin. We demonstrated how the proposed system is able to model expressive non-verbal vocalizations as well as generate high-quality expressive speech. We conclude with an ablation study and analysis of the different components composing our system.

This study serves as the foundation for improving speech emotion conversion and building general textless expressive speech generation models. For future work, we would like to design and evaluate an end-to-end system where we jointly model content and prosody, as well as extending the current work to the non-parallel dataset setting.

References

- Adaeze Adigwe, Noé Tits, Kevin El Haddad, Sarah Ostadabbas, and Thierry Dutoit. 2018. The emotional voices database: Towards controlling the emotion dimension in voice generation systems. *arXiv preprint arXiv:1806.09514*.
- Ryo Aihara, Ryoichi Takashima, Tetsuya Takiguchi, and Yasuo Ariki. 2012. Gmm-based emotional voice conversion using spectrum and prosody features. *American Journal of Signal Processing*, 2(5):134–138.
- Ryo Aihara, Reina Ueda, Tetsuya Takiguchi, and Yasuo Ariki. 2014. Exemplar-based emotional voice conversion using non-negative matrix factorization. In *Signal and Information Processing Association Annual Summit and Conference (APSIPA), 2014 Asia-Pacific*, pages 1–7. IEEE.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework for self-supervised learning of speech representations](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, volume 33, pages 12449–12460.
- Yuexin Cao, Zhengchen Liu, Minchuan Chen, Jun Ma, Shaojun Wang, and Jing Xiao. 2020. Nonparallel emotional speech conversion using vae-gan. In *INTERSPEECH*, pages 3406–3410.
- Aimilios Chalamandaris, Pirros Tsiakoulis, Sotiris Karabetsos, Spyros Raptis, and I LTD. 2013. The ilsp/innoetics text-to-speech system for the blizzard challenge 2013. In *Blizzard Challenge Workshop*. Citeseer.
- Po-Han Chi, Pei-Hung Chung, Tsung-Han Wu, Chun-Cheng Hsieh, Shang-Wen Li, and Hung-yi Lee. 2021. [Audio ALBERT: A lite BERT for self-supervised learning of audio representation](#). In *IEEE Spoken Language Technology Workshop (SLT)*, pages 344–350.
- Heejin Choi and Minsoo Hahn. 2021. Sequence-to-sequence emotional voice conversion with strength control. *IEEE Access*, 9:42674–42687.
- Yu-An Chung and James Glass. 2020. [Improved speech representations with multi-target autoregressive predictive coding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2353–2358. Association for Computational Linguistics.
- Jacob Devlin et al. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL*.
- Janek Ebbers, Jahn Heymann, Lukas Drude, Thomas Glarner, Reinhold Haeb-Umbach, and Bhiksha Raj. 2017. [Hidden Markov Model variational autoencoder for acoustic unit discovery](#). In *Proc. INTERSPEECH*, pages 488–492.
- Paul Ekman. 1992. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200.
- Jian Gao, Deep Chakraborty, Hamidou Tembine, and Olaitan Olaleye. 2018. Nonparallel emotional speech conversion. *arXiv preprint arXiv:1811.01174*.
- Thomas Glarner, Patrick Hanebrink, Janek Ebbers, and Reinhold Haeb-Umbach. 2018. [Full Bayesian Hidden Markov Model variational autoencoder for acoustic unit discovery](#). In *Proc. INTERSPEECH*, pages 2688–2692.
- Georg Heigold, Ignacio Moreno, Samy Bengio, and Noam Shazeer. 2016. End-to-end text-dependent speaker verification. In *ICASSP*.
- Yukiya Hono, Kazuna Tsuboi, Kei Sawada, Kei Hashimoto, Keiichi Oura, Yoshihiko Nankaku, and Keiichi Tokuda. 2020. Hierarchical multi-grained generative model for expressive speech synthesis. *arXiv preprint arXiv:2009.08474*.
- Wei-Ning Hsu, Yao-Hung Hubert Tsai, Benjamin Bolte, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. [HuBERT: How much can a bad teacher benefit ASR pre-training?](#) In *Neural Information Processing Systems Workshop on Self-Supervised Learning for Speech and Audio Processing Workshop*, pages 6533–6537.

- Wei-Ning Hsu, Yu Zhang, and James Glass. 2017. [Unsupervised learning of disentangled and interpretable representations from sequential data](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 1878–1889. Curran Associates, Inc.
- Zeynep Inanoglu and Steve Young. 2009. Data-driven emotion conversion in spoken english. *Speech Communication*, 51(3):268–283.
- K. Kasi and S. A. Zahorian. 2002. Yet another algorithm for pitch tracking. *ICASSP*.
- Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhota, Tu-Anh Nguyen, Morgane Rivière, Abdelrahman Mohamed, Emmanuel Dupoux, et al. 2021a. Text-free prosody-aware generative spoken language modeling. *arXiv preprint arXiv:2109.03264*.
- Eugene Kharitonov, Morgane Rivière, Gabriel Synnaeve, Lior Wolf, Pierre-Emmanuel Mazaré, Matthijs Douze, and Emmanuel Dupoux. 2021b. [Data augmenting contrastive learning of speech representations in the time domain](#). *arXiv preprint arXiv:2007.00991*, pages 215–222.
- Sameer Khurana, Antoine Laurent, Wei-Ning Hsu, Jan Chorowski, Adrian Lancucki, Ricard Marxer, and James Glass. 2020. [A convolutional deep Markov Model for unsupervised speech representation learning](#). In *Proc. INTERSPEECH*, pages 3790–3794.
- Jong Wook Kim, Justin Salamon, Peter Li, and Juan Pablo Bello. 2018. Crepe: A convolutional representation for pitch estimation. In *ICASSP*.
- Minchan Kim, Sung Jun Cheon, Byoung Jin Choi, Jong Jin Kim, and Nam Soo Kim. 2021. Expressive text-to-speech using style tag. *arXiv preprint arXiv:2104.00436*.
- Tae-Ho Kim, Sungjae Cho, Shinkook Choi, Sejik Park, and Soo-Young Lee. 2020. Emotional voice conversion using multitask learning with text-to-speech. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7774–7778. IEEE.
- John Kominek and Alan W Black. 2004. The cmu arctic speech databases. In *Fifth ISCA workshop on speech synthesis*.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. [HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, volume 33, pages 17022–17033.
- Cheng-I Lai, Yung-Sung Chuang, Hung-Yi Lee, Shang-Wen Li, and James Glass. 2021. [Semi-supervised spoken language understanding via self-supervised speech and language model pretraining](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7468–7472.
- Kushal Lakhota, Evgeny Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, et al. 2021. Generative spoken language modeling from raw audio. *arXiv preprint arXiv:2102.01192*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880. Online. Association for Computational Linguistics.
- Songxiang Liu, Yuewen Cao, and Helen Meng. 2020. Multi-target emotional voice conversion with neural vocoders. *arXiv preprint arXiv:2004.03782*.
- Jaime Lorenzo-Trueba, Gustav Eje Henter, Shinji Takaki, Junichi Yamagishi, Yosuke Morino, and Yuta Ochiai. 2018. Investigating different representations for modeling and controlling multiple emotions in dnn-based speech synthesis. *Speech Communication*, 99:135–143.
- Zhaojie Luo, Tetsuya Takiguchi, and Yasuo Arik. 2016. Emotional voice conversion using deep neural networks with mcc and f0 features. In *2016 IEEE/ACIS 15th International Conference on Computer and Information Science (ICIS)*, pages 1–5. IEEE.
- Kunalan Manokara, Mirna Đurić, Agneta Fischer, and Disa Sauter. 2021. Do people agree on how positive emotions are expressed? a survey of four emotions and five modalities across 11 cultures. *Journal of Nonverbal Behavior*, pages 1–32.
- Albert Mehrabian and Morton Wiener. 1967. Decoding of inconsistent communications. *Journal of personality and social psychology*, 6(1):109.
- Huaiping Ming, Dong-Yan Huang, Lei Xie, Jie Wu, Minghui Dong, and Haizhou Li. 2016. Deep bidirectional lstm modeling of timbre and prosody for emotional voice conversion. In *Interspeech*, pages 2453–2457.
- Lucas Ondel, Lukáš Burget, and Jan Černocký. 2016. [Variational inference for acoustic unit discovery](#). *Procedia Computer Science*, 81:80–86.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [Fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT*, pages 48–53.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [LibriSpeech: an ASR corpus based on public domain audio books](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210. IEEE.

- Kishore Papineni et al. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proc. of ACL*.
- Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharitonov, Kushal Lakhota, Wei-Ning Hsu, Abdelrahman Mohamed, and Emmanuel Dupoux. 2021. Speech resynthesis from discrete disentangled self-supervised representations. *arXiv preprint arXiv:2104.00355*.
- Adam Polyak, Lior Wolf, Yossi Adi, and Yaniv Taigman. 2020a. [Unsupervised cross-domain singing voice conversion](#). In *Proc. INTERSPEECH*, pages 801–805.
- Adam Polyak, Lior Wolf, and Yaniv Taigman. 2020b. [TTS skins: Speaker conversion via ASR](#). In *Proc. INTERSPEECH*, pages 786–790.
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Mark Hasegawa-Johnson, and David Cox. 2020. Unsupervised speech decomposition via triple information bottleneck. In *International Conference on Machine Learning*, pages 7836–7846. PMLR.
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Jinjun Xiong, Chuang Gan, David Cox, and Mark Hasegawa-Johnson. 2021. Global rhythm style transfer without text transcriptions. *arXiv preprint arXiv:2106.08519*.
- M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio. 2020. [Multi-task self-supervised learning for robust speech recognition](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6989–6993.
- Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2020. FastSpeech 2: Fast and high-quality end-to-end text to speech. *arXiv preprint arXiv:2006.04558*.
- F. Ribeiro, D. Florêncio, C. Zhang, and M. Seltzer. 2011. [CROWDMOS: An approach for crowdsourcing mean opinion score studies](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2416–2419.
- Georgios Rizos, Alice Baird, Max Elliott, and Björn Schuller. 2020. Stargan for emotional speech conversion: Validated by data augmentation of end-to-end emotion recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3502–3506. IEEE.
- Carl Robinson, Nicolas Obin, and Axel Roebel. 2019. Sequence-to-sequence modelling of f0 for speech emotion conversion. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6830–6834. IEEE.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019. [wav2vec: Unsupervised pre-training for speech recognition](#). In *Proc. INTERSPEECH*, pages 3465–3469.
- Ravi Shankar, Hsi-Wei Hsieh, Nicolas Charon, and Archana Venkataraman. 2019a. Automated emotion morphing in speech based on diffeomorphic curve registration and highway networks. *Proc. Interspeech 2019*.
- Ravi Shankar, Hsi-Wei Hsieh, Nicolas Charon, and Archana Venkataraman. 2020. Multi-speaker emotion conversion via latent variable regularization and a chained encoder-decoder-predictor network. *arXiv preprint arXiv:2007.12937*.
- Ravi Shankar, Jacob Sager, and Archana Venkataraman. 2019b. A multi-speaker emotion morphing model using highway networks and maximum likelihood objective. In *INTERSPEECH*, pages 2848–2852.
- J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu. 2018. [Natural TTS synthesis by conditioning WaveNet on MEL spectrogram predictions](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4779–4783.
- RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron Weiss, Rob Clark, and Rif A Saurous. 2018a. Towards end-to-end prosody transfer for expressive speech synthesis with tacotron. In *international conference on machine learning*, pages 4693–4702. PMLR.
- RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron Weiss, Rob Clark, and Rif A Saurous. 2018b. Towards end-to-end prosody transfer for expressive speech synthesis with tacotron. In *international conference on machine learning*, pages 4693–4702. PMLR.
- Alexander Sorin, Slava Shechtman, and Ron Hoory. 2020. Principal style components: Expressive style control and cross-speaker transfer in neural tts. In *INTERSPEECH*, pages 3411–3415.
- Xu Tan, Tao Qin, Frank Soong, and Tie-Yan Liu. 2021. A survey on neural speech synthesis. *arXiv preprint arXiv:2106.15561*.
- Noé Tits, Kevin El Haddad, and Thierry Dutoit. 2019. The theory behind controllable expressive speech synthesis: A cross-disciplinary approach. In *Human 4.0-From Biology to Cybernetic*. IntechOpen.
- Aaron van den Oord, Oriol Vinyals, et al. 2017. [Neural discrete representation learning](#). In *Advances in Neural Information Processing Systems*, pages 6306–6315.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- W. Wang, Q. Tang, and K. Livescu. 2020. [Unsupervised pre-training of bidirectional speech encoders via masked reconstruction](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6889–6893.
- Anne Wu, Changhan Wang, Juan Pino, and Jiatao Gu. 2020. [Self-supervised representations improve end-to-end speech translation](#). In *Proc. INTERSPEECH*, pages 1491–1495.
- Junichi Yamagishi, Christophe Veaux, Kirsten MacDonald, et al. 2019. Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92).
- Shu-wen Yang et al. 2021. Superb: Speech processing universal performance benchmark. *arXiv preprint arXiv:2105.01051*.
- Kun Zhou, Berrak Sisman, and Haizhou Li. 2020a. [Transforming spectrum and prosody for emotional voice conversion with non-parallel training data](#). *arXiv preprint arXiv:2002.00198*.
- Kun Zhou, Berrak Sisman, and Haizhou Li. 2021a. [Limited data emotional voice conversion leveraging text-to-speech: Two-stage sequence-to-sequence training](#). *arXiv preprint arXiv:2103.16809*.
- Kun Zhou, Berrak Sisman, and Haizhou Li. 2021b. [Vaw-gan for disentanglement and recombination of emotional elements in speech](#). In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 415–422. IEEE.
- Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li. 2021c. [Emotional voice conversion: Theory, databases and esd](#). *arXiv preprint arXiv:2105.14762*.
- Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li. 2021d. [Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset](#). In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 920–924. IEEE.
- Kun Zhou, Berrak Sisman, Mingyang Zhang, and Haizhou Li. 2020b. [Converting anyone’s emotion: Towards speaker-independent emotional voice conversion](#). *arXiv preprint arXiv:2005.07025*.

A Appendix

A.1 Implementation Details

We use the sequence-to-sequence Transformer model as implemented in `fairseq` (Ott et al.,

2019). The model contains 3 layers for both encoder and decoder modules, 4 attention heads, embedding size of 512, FFN size of 512, and dropout probability of 0.1.

For pre-training, we explore two settings: (1) large scale – using a mix of LibriSpeech, Blizzard2013 and EmoV, denoted by OURS⁺; and (2) small scale – using a mix of VCTK and EmoV, denoted by OURS⁻. We stop training after 3M update steps. We use the following input augmentations: *Infilling* using $\lambda = 3.5$ for the Poisson distribution, *Token masking* with probability of 0.3, *Random masking* with probability of 0.1, and *Sentence permutation*. Finally, we fine-tune this model on the task of translation using paired utterances of different emotions from EmoV and early-stop using the loss from Equation 1.

Our F0 prediction model E_{F0} consists of six 1-D convolutional layers, where the number of kernels per layer is 256 and the respective kernel size is 5. For non-linearity we used the ReLU activation function followed by layer-norm and dropout ($p = 0.1$).

Our duration prediction model E_{dur} consists of two convolutional layers, where the number of kernels per layer is 256 and the respective kernel size is 3. For non-linearity we used the ReLU activation function followed by layer-norm and dropout ($p = 0.5$).

Computational Resources. All experiments done in this study are conducted using 8 NVIDIA V100 GPUs with 32GB memory each.

A.2 Additional Results

Model Architecture We evaluated three weight-sharing schemes for this model: (i) all emotions use the same encoder and decoder components of the Transformer architecture. In this case, we condition the model on the target emotion. This can be done by prepending a special target emotion token at the beginning of the decoding procedure. We denote this approach by “share-all”. (ii) All emotions share the same encoder but have separate dedicated decoders. In this case, no target emotion conditioning is needed. We denote this approach by “share-enc”. Finally, (iii) each emotion has a dedicated encoder and decoder. We denote this approach by “share-none”. We evaluated all three approaches and summarized the results in Table 3.

We report the BLEU score (Papineni et al., 2002) and Unit Error Rate (UER). UER measures relative

edit-distance between the target unit sequence and the generated unit sequence. We additionally report the WER and PER metrics extracted using a pre-trained SOTA ASR system⁴. Results are reported on emotions converted to neutral: {amused, angry, sleepy, disgusted}-neutral. The BLEU and UER metrics are computed using all emotion pairs from Neutral and to Neutral.

Architecture	BLEU	UER	WER	PER
Share-all	31.97	42.56	25.47	26.44
Share-enc	32.22	41.16	23.91	25.41
Share-none	32.37	41.38	23.91	25.95

Table 3: Evaluation of three weight-sharing schemes for the translation model.

It can be seen that the share-enc and share-none architectures are comparable, while the share-all configuration is inferior. Although both share-enc and share-none configurations are similar in terms of lexical reconstruction, in our listening tests share-none generated more expressive speech.

F0 estimation module We evaluate the F0 estimation model using the Mean Absolute Error (MAE) between ground-truth F0 and predicted F0. In this experiment, we explore a number of configurations for training such an estimator. Specifically, we evaluate different binning strategies, normalization methods and prediction rules. For binning strategies, we explore adaptive binning vs. uniform binning. Under normalization, we explore no normalization, mean normalization and mean & standard deviation normalization⁵. Finally, in addition to the weighted-average prediction rule described in Section 3.2, we also evaluate an argmax prediction rule where the highest scoring bin is predicted. Results are summarized in Table 4, “Log” denotes applying the logarithm function before normalization.

Results suggest that the weighted-average prediction rule is preferable to argmax, especially when used in conjunction with adaptive binning. This can be explained by large-range bins in the adaptive case, leading to larger MAE when selecting a single bin using the argmax operator. Although adaptive quantization reaches the best performance, under specific settings uniform quantization can reach comparable results. For normalization, it is

⁴We use a BASE wav2vec 2.0 phoneme detection model trained on LibriSpeech-960h with CTC loss from scratch.

⁵The mean and standard deviation are computed using the F0 values per speaker.

Log	Quantization	Norm.	Prediction	MAE
✗	uniform	✗	argmax	83.51
✗	uniform	✗	w-avg	44.92
✗	uniform	mean	argmax	53.74
✗	uniform	mean	w-avg	35.63
✗	uniform	mean-std	argmax	63.01
✗	uniform	mean-std	w-avg	35.69
✗	adaptive	✗	argmax	129.7
✗	adaptive	✗	w-avg	45.21
✗	adaptive	mean	argmax	127.8
✗	adaptive	mean	w-avg	36.06
✗	adaptive	mean-std	argmax	155.7
✗	adaptive	mean-std	w-avg	35.38
✓	uniform	✗	argmax	62.69
✓	uniform	✗	w-avg	67.40
✓	uniform	mean	argmax	52.24
✓	uniform	mean	w-avg	63.67
✓	uniform	mean-std	argmax	50.95
✓	uniform	mean-std	w-avg	51.25
✓	adaptive	✗	argmax	127.9
✓	adaptive	✗	w-avg	67.40
✓	adaptive	mean	argmax	110.4
✓	adaptive	mean	w-avg	54.21
✓	adaptive	mean-std	argmax	144.9
✓	adaptive	mean-std	w-avg	51.25

Table 4: Evaluation of different F0 estimation configurations. The MAE is reported for voiced frames only.

Model	MAE	Acc@0ms	Acc@20ms	Acc@40ms
CNN	0.77	51.12	86.24	94.08
1-gram	1.47	29.41	67.34	83.26
3-gram	1.16	36.78	76.04	88.32
5-gram	1.02	37.32	81.36	90.46

Table 5: Evaluation of four duration prediction models. CNN architecture is as described in (Ren et al., 2020).

preferable to normalize, the specific normalization method has little impact on performance.

Duration prediction module. We evaluate the duration prediction models using the MAE between target and predicted durations. For a more complete analysis, we also report the accuracy using thresholds of 0ms, 20ms and 40ms. We explore a CNN duration predictor similarly to (Ren et al., 2020) and three n-gram based models. The results are summarized in Table 5. As expected, the CNN outperforms n-gram models, with ~94% accuracy when considering a tolerance level of 40ms.

eMOC Confusion Matrices We provide the eMOC confusion matrices for samples generated by each of the evaluated models and the ground truth recordings. See Figures 5, 6, 7, 8, 9, 10. We additionally provide the full eMOC results for the many-to-many setting in the Table 6.

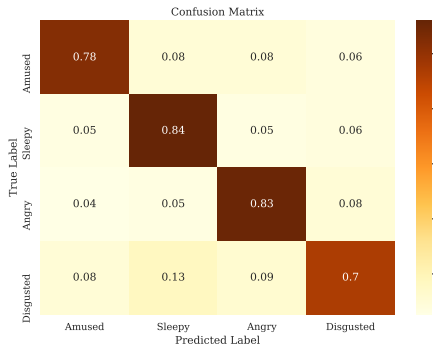


Figure 5: Ground Truth Recordings




Figure 9: Tacotron2

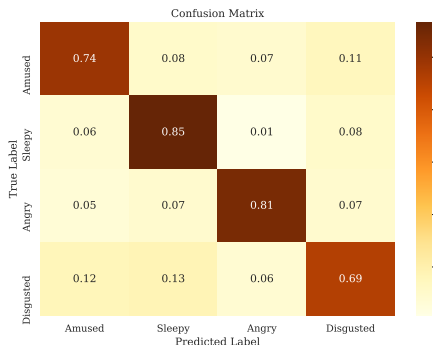


Figure 6: Ours⁺




Figure 10: VAW-GAN




Figure 7: Ours⁻




Figure 8: Seq2Seq-EVC

A.3 Subjective Evaluation

For the crowd-sourced subjective evaluations we computed the MOS and eMOC metrics. For the MOS metric raters were asked: “Rate the quality

and naturalness of the given speech utterance on a scale of 1 to 5 (1 being of low quality and naturalness and 5 being of high quality and naturalness)”. For the eMOC raters were asked: “Select the emotion for the given emotion categories that best suits the given speech utterance”.

All raters are native English speakers located in the United-States.

B Risks & Limitations

The main risks in the proposed approach (as in any generative model) is the development of a high quality and natural speech synthesis model. Such technology might be used to alter emotion in speech recordings. To deal with that we limit the number of speakers that can be synthesized with the proposed approach using a look-up-table. Another potential risk of the proposed method which is also one of its limitations is that the generated speech content is not always perfect, hence might lead to wrong pronunciations. As described in the results section, the WER of the proposed system is ~ 27.92 which means we still have a big room for improvement. Another limitation of the current approach is the need for parallel corpus to convert between different emotions.

Source	Target	Amused	Angry	Disgusted	Sleepy	Neutral
Amused	Amused	-	-	-	-	-
	Angry	0.06	0.81	0.04	0.04	0.05
	Disgusted	0.04	0.04	0.77	0.01	0.14
	Sleepy	0.06	0.02	0.03	0.85	0.04
	Neutral	0.06	0.07	0.09	0.03	0.75
Angry	Amused	0.83	0.07	0.04	0.02	0.04
	Angry	-	-	-	-	-
	Disgusted	0.05	0.12	0.59	0.07	0.17
	Sleepy	0.04	0.06	0.03	0.82	0.05
	Neutral	0.06	0.08	0.04	0.03	0.79
Disgusted	Amused	0.71	0.08	0.1	0.05	0.06
	Angry	0.03	0.78	0.05	0.09	0.05
	Disgusted	-	-	-	-	-
	Sleepy	0.09	0.06	0.05	0.72	0.08
	Neutral	0.08	0.02	0.1	0.03	0.77
Sleepy	Amused	0.63	0.02	0.12	0.14	0.09
	Angry	0.06	0.77	0.03	0.09	0.05
	Disgusted	0.03	0.07	0.66	0.13	0.11
	Sleepy	-	0	-	-	-
	Neutral	0.05	0.08	0.12	0.12	0.63
Neutral	Amused	0.7	0.08	0.06	0.08	0.08
	Angry	0.06	0.77	0.06	0.07	0.04
	Disgusted	0.08	0.07	0.63	0.1	0.12
	Sleepy	0.06	0.02	0.03	0.79	0.1
	Neutral	-	-	-	-	-

Table 6: Subjective evaluation for any-to-any emotion conversion. The “source” and “target” columns represent the source and target emotion respectively. Each emotion column (Amused, Angry, Disgusted, Sleepy and Neutral columns) represent the percentage of human raters that chose the said emotion when converting from source to target emotion.