

# Preset-Voice Matching for Privacy Regulated Speech-to-Speech Translation Systems

Daniel Platnick<sup>\*1,2</sup>, Bishoy Abdelnour<sup>\*1</sup>, Eamon Earl<sup>\*1</sup>,  
Rahul Kumar<sup>1</sup>, Zahra Rezaei<sup>1</sup>, Thomas Tsangaris<sup>1</sup>, Faraj Lagum<sup>1</sup>

<sup>1</sup>Vosyn, Etobicoke, Canada

<sup>2</sup>Vector Institute for Artificial Intelligence, Toronto, Canada

daniel.platnick@torontomu.ca,

{bishoynour,eamon.sc.earl}@gmail.com

## Abstract

In recent years, there has been increased demand for speech-to-speech translation (S2ST) systems in industry settings. Although successfully commercialized, cloning-based S2ST systems expose their distributors to liabilities when misused by individuals and can infringe on personality rights when exploited by media organizations. This work proposes a regulated S2ST framework called *Preset-Voice Matching (PVM)*. *PVM* removes cross-lingual voice cloning in S2ST by first matching the input voice to a similar *prior consenting* speaker voice in the target-language. With this separation, *PVM* avoids cloning the input speaker, ensuring *PVM* systems comply with regulations and reduce risk of misuse. Our results demonstrate *PVM* can significantly improve S2ST system run-time in multi-speaker settings and the naturalness of S2ST synthesized speech. To our knowledge, *PVM* is the first explicitly regulated S2ST framework leveraging similarly-matched preset-voices for dynamic S2ST tasks.

## 1 Introduction

Progress in deep learning and voice cloning technology has enhanced public access to robust AI-driven voice cloning systems. These systems can help solve complicated speech-to-speech translation (S2ST) tasks like automated video dubbing (auto-dubbing) by generating audio deepfakes (Brannon et al., 2022; Shoaib et al., 2023; Amezaga and Hajek, 2022). Cloning systems are desirable for dynamic speech tasks because they can generate a clone from an input voice given an audio sample as short as a few seconds (Arik et al., 2018). Currently, voice cloning technology is highly unregulated and can be harmful if misused or commercialized irresponsibly (Liu et al., 2023a).

As voice cloning systems can clone an arbitrary voice and do not require permission, they raise

several privacy concerns (Baris, 2024). Risks related to voice cloning technology include lack of informed consent, biometric privacy, and the spread of misinformation through deepfakes (Frankovits and Mirsky, 2023; Okolie, 2023). Robust regulations are necessary to mitigate these risks, protect individual rights, and prevent misuse (Baris, 2024; Moreno, 2024; Sudhakar and Shanthi, 2023).

The risks of unregulated voice cloning technologies are compounded by a high demand for voice cloning-based products. Pressure to capitalize on a newly budding market of cloning-based products can lead businesses to emphasize speed over careful and tested development. Since voice cloning technology is so new, regulatory measures are required and in the process of being implemented, but not yet fully in place. Given these challenges, it is crucial to integrate privacy regulations into AI-powered voice cloning systems (Liu et al., 2023b; Tee and Murugesan, 2021; Tariq et al., 2023).

To address the need for regulated voice cloning technology, we propose *Preset-Voice Matching (PVM)*, a regulated S2ST framework. *PVM* bakes regulatory precautions into the S2ST process by removing the explicit training objective of cloning an unknown input speaker’s voice, and instead cloning a *similar* preset-voice of a consenting speaker. *PVM* can be easily installed on top of existing cascaded S2ST pipelines, improving regulatory compliance. We find this process also decreases system run-time in multi-speaker auto-dubbing scenarios and improves speaker naturalness relative to state-of-the-art voice cloning systems when translating across our tested languages.

The intention of this paper is to put forward a regulated *PVM* S2ST framework that is robust against legislative changes and future liability concerns. We demonstrate *PVM* is desirable for S2ST over current benchmark voice cloning frameworks due to its inherent safety, lower run-time in multi-speaker scenarios, and enhanced speaker natural-

<sup>\*</sup>These authors contributed equally and share co-first authorship

ness. We show this by providing and testing a *PVM* algorithm which we call *GEMO-Match*. We hope this work inspires others to develop and tune the framework for different high-performance environments. Our main contributions are as follows:

1. We propose *PVM*, a novel privacy-regulated S2ST framework which leverages consented preset-voices to clone a preset-voice similar to the input voice.
2. We provide a gender-emotion based *PVM* algorithm, *GEMO-Match*, and use it to demonstrate *PVM* in multilingual settings.
3. We empirically analyze *GEMO-Match* in terms of robustness, multilingual capability, and run-time, on two speech emotion datasets and discuss the implications of our system.
4. We create and provide a *Combined Gender-Dependent Dataset (CGDD)*, which combines various benchmark speech-emotion datasets for training future gender-dependent *PVM* algorithms.

The rest of this paper is organized as follows. Background information is provided in Section 2. Our *PVM* framework and *GEMO-Match* algorithm are detailed in sections 3 and 4. Relevant datasets are described in Section 5. Section 6 explains our experimental setup as well as the techniques, algorithms, and parameters used in the study. Section 7 includes experimental results and analysis. We discuss potential future work towards *PVM* and conclude the paper in sections 8 and 9. We address *PVM* limitations in Section 10.

## 2 Background Information

Speech-to-speech translation (S2ST) is typically achieved by direct translation or cascaded approaches (Etchegoyhen et al., 2022). Direct translation approaches use speech and linguistic encoder/decoders (Jia et al., 2019) to directly translate speech signals from one language to another. Cascading architectures split S2ST into three sub-tasks, using separate but connected speech-to-text (STT), text-to-text (TTT), and text-to-speech (TTS) modules (Huang et al., 2023). Cascading architectures have been the traditional method for S2ST.

Two common approaches for synthesizing speech from text are concatenative and parametric TTS. Concatenative TTS combines pre-recorded

clips from a database to form a final speech output (Gujarathi and Patil, 2021). Parametric TTS attempts to model and predict speech variations given text and a reference voice (King, 2011). Parametric deep learning methods have shown ubiquitous success spanning various industries from computer vision to text synthesis (Lecun and Bengio, 1995; Fayyaz et al., 2022; Platnick et al., 2024; Ning et al., 2019). As deep neural network (DNN) based TTS methods can lead to natural and expressive synthesized voices, they are desirable for many speech tasks (Barakat et al., 2024).

Wavenet is a benchmark DNN-based TTS model (van den Oord et al., 2016). Since its creation, there have been many advancements in sequence-to-sequence TTS models trained to produce human-like speech (Wang et al., 2017). Wavenet performs speech synthesis by training on a set of human voices, conditioning on their unique speaker ID to generate natural-sounding utterances in the voice of a selected speaker (van den Oord et al., 2016). Recently, there have been models which aim to extend this behavior by cloning voices unseen in training, resulting in zero-shot voice cloning (Zhang et al., 2023).

Cross-lingual voice cloning is difficult due to complexities in discriminating between language-specific and speaker-specific features within a singular waveform, and mapping these features across different languages (Eren and Team, 2023). Additionally, training robust multilingual speech generation models requires vast amounts of processed language and speech data in multiple languages with a variety of utterances and speakers. The performance of these models depends on the data they are trained on (Rebai et al., 2017).

Preset-voice TTS methods generate speech from stored options of preset or pre-recorded voices. Preset-voice methods are typically used in static or repetitive systems which do not require dynamic adaptive functionality. Examples include pre-programmed transit operator dispatch messages, medical alert systems in healthcare, and emergency flight announcements (Strathman et al., 2001; Eyesan and Okuboyejo, 2013; Samaras and Ferreira, 2019).

Due to the static nature of current preset-voice methods, they have not previously been used for dynamic S2ST tasks like auto-dubbing. Such dynamic tasks require modelling different speakers across languages based on incoming media data (Brannon et al., 2022). In addition to providing a

regulated *PVM* framework, this work aims to extend the application of preset-voice TTS methods to more dynamic settings.

### 3 Preset-Voice Matching Framework

This section explains our privacy regulated *Preset-Voice Matching (PVM)* framework.

*PVM* bakes privacy regulations into the S2ST process by cloning a similar and prior consenting preset-voice, instead of the voice originally input to the S2ST system. The *PVM* framework connects to cascading S2ST architectures, performing additional computations alongside the STT, TTT, and TTS modules. The *PVM* framework consists of 3 sub-modules.

Module 1, the *Similarity Feature Extraction* module, extracts features from the inputted voice. It then uses the extracted features to match the input voice to the most similar preset-voice from the *Preset-Voice Library*. Module 2, the *Preset-Voice Library*, contains a collection of consented target-language preset-voices, partitioned by discrete feature codes depending on the *PVM* implementation. Module 3, the *TTS Module*, generates TTS in the target-language using the matched preset-voice from the *Preset-Voice Library*.

We describe these 3 modules below in detail.

#### 3.1 Feature Extraction and Voice Matching

The *Similarity Feature Extraction* module extracts meaningful features from the input voice. These features are used to determine the most similar consented preset-voice in the target-language from our preset-voice library. This module takes in speech signals as input and outputs similarity feature encodings (gender-emotion pair combinations in the case of *GEMO-Match*) to match a consented similar preset-voice.

#### 3.2 Target-language Preset-Voice Libraries

Module 2, the *Preset-Voice Library*, contains a collection of preset-voices in desired target-languages. The *Preset-Voice Library* acts as a feature codebook, informing the mapping between feature encodings and target-language preset-voice samples. This module takes in a feature code as input, and outputs a matched consenting speaker preset-voice sample.

#### 3.3 Text-to-Speech with Matched Preset-Voice

As input, the *TTS Module* takes in the matched consented preset-voice and target-language text (from

an auxiliary TTT module). The *TTS Module* outputs a clone of the most similar preset-voice in a desired language relative to the features extracted in the *Similarity Feature Extraction* module. Any voice cloning TTS model supporting the desired target-languages can be used in the *TTS Module*. Therefore, *PVM* is a general framework and is easily modifiable for many industry settings.

### 4 GEMO-Match Algorithm

In this section we describe *GEMO-Match*, an example *PVM* framework implementation.

Following a similar notion to (Singh and Prasad, 2023), *GEMO-Match* employs a hierarchical gender-dependent emotion classifier architecture trained with a gender-dependent training method. The process of splitting gender and emotion in emotion classification simplifies the emotion classification problem. As *GEMO-Match* is a *PVM* framework, it contains the 3 *PVM* modules: the *Similarity Feature Extraction* module, the *Preset-Voice Library*, and the *TTS Module*.

These modules and their process are described below.

#### 4.1 GEMO-Match Modules

The *GEMO-Match Similarity Feature Extraction* module contains 3 classifiers in two stages. The first stage contains the *gender classifier*, and the second stage includes both the *male-emotion classifier*, and the *female-emotion classifier*. The *Similarity Feature Extraction* classifiers are trained in the source language (English).

In *GEMO-Match*, the *Preset-Voice Library* contains previously consenting speakers in desired target-languages for a given S2ST task. The *Preset-Voice Library* partitions target-language preset-voices by language, gender, and emotion. The number of target-languages supported by *GEMO-Match* depends on the ability to gather preset-voices in each desired target-language. The *Preset-Voice Library* in our provided implementation includes two target-languages, French and German. Therefore, the *GEMO-Match* implementation can translate from English to either French or German.

The *GEMO-Match TTS Module* performs TTS. The *TTS Module* is straightforward and performs TTS given a matched preset-voice and a text prompt in the desired target-language. We implement *GEMO-Match* with two distinct TTS models, discussed in 6.2 and 6.3.

## 4.2 GEMO-Match Algorithm Flow

First, source language speech is input to the *Similarity Feature Extraction* module. The *gender classifier* then classifies the input voice as male or female. Next, given the gender classification result, the source speech is input to the corresponding gender-dependent emotion classifier. The appropriate gender-dependent emotion classifier will then classify the source language speech as happy, angry, sad, disgust, or neutral. The two-stage classifier output pair is then concatenated (i.e., Female - Sad).

The resulting concatenation is used alongside the intended target-language to query the most similar preset-voice in the *Preset-Voice Library*. Finally, the feature-matched preset-voice is passed alongside a text prompt to the voice cloning TTS model. This algorithm assumes that the intended target-language will be an input to the system. The performance of *GEMO-Match* depends primarily on the robustness of the *Similarity Feature Extraction* classifiers.

## 5 Dataset Descriptions

In this section, we describe the datasets used to test our framework.

We experimented with two speech-emotion datasets: the Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) (Livingstone and Russo, 2018), and the *Combined Gender-Dependent Dataset (CGDD)*, which we curated by combining four benchmark speech datasets. To ensure compatibility with our gender-emotion based *GEMO-Match* algorithm, we split the RAVDESS dataset by gender and relabeled it with gender-emotion pairs. Further details on RAVDESS and *CGDD* are outlined in 5.1 and 5.2.

### 5.1 RAVDESS Dataset

RAVDESS is a benchmark emotional speech dataset containing 1440 audio files of 24 professional actors (12 female and 12 male) with the emotions calm, happy, sad, angry, fearful, surprise, and disgust (Livingstone and Russo, 2018). As *GEMO-Match* requires consistent labeling across source and target-language data, we focus on a subset of 5 common emotions: happy, angry, sad, disgust, and calm (neutral). Each speech sample was originally provided with two intensities, normal and strong. We filtered the speech files to include only strong intensities as the emotion is more apparent in those

samples. After filtering, the RAVDESS subset contains a total of 5 speech recordings per actor per emotion.

### 5.2 Combined Gender-Dependent Dataset

Training a robust gender-emotion classifier requires numerous samples of speakers from various demographics, speaking a variety of utterances with different emotional intensities. We found that many available speech-emotion datasets have limited variance in regards to at least one of these features. To help facilitate gender-dependent training research, we provide a *Combined Gender-Dependent Dataset (CGDD)*, made from combining four benchmark emotional speech datasets: RAVDESS, CREMA-D, SAVEE, and TESS (Livingstone and Russo, 2018; Cao et al., 2014; Phukan et al., 2023; Pichora-Fuller and Dupuis, 2020).

The RAVDESS dataset is explained in section 5.1. CREMA-D is comprised of 7,442 audio recordings of 91 actors. These clips include 48 male actors and 43 female actors, with ages ranging from 20 to 74. SAVEE database includes four English male speakers aged between 27 and 31, totaling 480 files. The TESS database contains two female speakers, one aged 26 and the other aged 64, with a total of 2800 files.

The *CGDD* dataset is processed for gender-dependent training, useful for hierarchical emotion detection algorithms like *GEMO-Match*. We further processed the audio based on pitch frequency and loudness to obtain a higher-quality dataset. As pitch and loudness are crucial attributes of speech, we filter the data to ensure the files are within a suitable range for speech recognition (Zaïdi et al., 2021). Additionally, we use RMS loudness to eliminate excessively quiet or loud files. The best quality was found with a pitch frequency range of 75 Hz to 3000 Hz. We removed audio samples with RMS loudness less than -23 dBFS and greater than -20 dBFS.

### 5.3 Data Pre-processing

We processed the RAVDESS and *CGDD* datasets to be compatible with the hierarchical gender-dependent emotion classification architecture of the *GEMO-Match Similarity Feature Extraction* module. For both datasets, we partitioned the speech signal files on gender and further organized them into five gender-emotion directories. We then converted the speech signals to mel-spectrograms using the Fast Fourier Transform. Next, the mel-



spectrograms were converted to image representation (PNG format) to be processed by a pre-trained ResNet50 model initialized with ImageNet weights (Deng et al., 2009). Our data pre-processing methodology is similar to the procedures outlined in (Sinha et al., 2020). The Python library Librosa was used to convert speech signal files to mel-spectrogram signals.

## 6 Experimental Setup

This section details the setup of each experiment, which show additional strengths of the *PVM* framework, beyond its inherent regulatory benefits.

We demonstrate the effectiveness of *PVM* for S2ST with *GEMO-Match* in terms of robustness, multilingual capability, and run-time. Our experiments were run on a single Tesla T4 GPU with 40 cores. We discuss each experiment in detail below.

### 6.1 GEMO-Match Robustness

For this test, we assess the robustness of *GEMO-Match*. The performance of *GEMO-Match* depends on the three *Similarity Feature Extraction* classifiers. We fine-tuned and evaluated these classifiers on the RAVDESS and our *CGDD* dataset in terms of accuracy and precision. Each classifier was implemented as a ResNet50 previously pre-trained on ImageNet. The results of the six classifiers are shown in tables 1 and 2.

The same approach was used to train each ResNet50. The *gender classifiers* were trained for 20 epochs, while the *male-emotion* and *female-emotion classifiers* required 30 epochs to converge. Each emotion classifier was trained using a dynamic learning rate schedule: 0.01 for the first 20 epochs, reduced to 0.001 for the remaining 10.

We used the Adam optimizer, and the Pytorch ImageDataGenerator function for data augmentation (Kingma and Ba, 2017). The classifiers were trained using a batch size of 32 and a train-test-validation split of 60-20-20. The models were optimized using categorical cross entropy as the loss function, incorporating batch normalization and dropout layers for regularization. The activation functions used were ReLU for internal layers and softmax for the output layer.

### 6.2 GEMO-Match Multilingualism

We test *GEMO-Match* in terms of speaker naturalness on the task of translating English speech into French and German speech. *GEMO-Match* is

implemented within a cascaded S2ST system using SeamlessM4T for TTT, and XTTS as the TTS module (Communication et al., 2023; Eren and Team, 2023). XTTS is a state-of-the-art TTS model which supports zero-shot voice cloning across 17 languages. Instead of performing STT, we provide ground truth source-language (English) text directly to the TTT model (SeamlessM4T) to measure the isolated performance of *GEMO-Match* across multiple languages. We measured speaker naturalness using the standard metric Non-intrusive Objective Speech Quality Assessment (NISQA) (Mittag et al., 2021; Yi et al., 2022).

We show *PVM* algorithms lead to higher naturalness in S2ST outputted speech by alleviating the need to perform *cross-lingual* voice cloning. We compare two cases of S2ST. The first case is when XTTS performs cross-lingual cloning from an English voice input to the target-languages German and French. In the second case, *GEMO-Match* performs the cross-lingual matching, allowing XTTS to run monolingual TTS given the matched target-language voice as input.

The French and German preset-voices used in this experiment are sourced from the CAFE, and EmoDB datasets respectively (Gournay et al., 2018; Burkhardt et al., 2005). For each target-language in both experimental pipelines, we used 150 English text transcriptions from the CREMA-D dataset alongside emotive English audios from RAVDESS as input (Cao et al., 2014). We ensured that our RAVDESS audios had an average NISQA (3.54) similar to the preset-voices in our target-languages. For additional context, we included the average preset-voice NISQA scores for both target-languages in Table 3.

### 6.3 GEMO-Match Run-time

We compared the run-time of *GEMO-Match* to state-of-the-art TTS models VALL-E X, XTTS, SeamlessM4T, and OpenVoice, as shown in Figure 1 (Qin et al., 2024). The *gender*, *male-emotion*, and *female-emotion classifiers* were implemented using the same lightweight ResNet50 models as in 6.1. Each model was given 10 identical utterances with their respective transcriptions, and average inference run-times were calculated. The inputs were each 15 seconds and varied in tone, emotion, pacing, and vocabulary.

We compared *PVM* (using *GEMO-Match*) with OpenVoice as they are both cascaded TTS frameworks that decouple voice-cloning from voice syn-

Emotions	RAVDESS Precision		CGDD Precision	
	Male-Emo	Female-Emo	Male-Emo	Female-Emo
Happy	0.78	0.56	0.51	0.78
Angry	0.78	1.00	0.82	0.87
Sad	0.50	0.40	0.59	0.66
Disgust	0.30	0.40	0.78	0.72
Neutral	0.80	0.90	0.72	0.85

Table 1: Precision of *GEMO-Match* gender-dependent emotion classifiers (ResNet50 pre-trained) on 5 emotions from RAVDESS and *CGDD*. Training the ResNet50 on the *CGDD* dataset results in better generalization across emotions in terms of precision.

Classifier	RAVDESS Accuracy	CGDD Accuracy
Gender	94.00	<b>97.00</b>
Male-Emotion	62.00	<b>63.21</b>
Female-Emotion	65.00	<b>71.29</b>

Table 2: Test set accuracies of *GEMO-Match* classifiers.

thesis. OpenVoice uses a variation of VITS for TTS in its open-source implementation (Kim et al., 2021). For consistent comparisons with OpenVoice, we use StyleTTS2 for TTS with *GEMO-Match* (Li et al., 2023). StyleTTS2 and VITS are both styling-based models and display similar run-times. StyleTTS2 is a monolingual TTS model, and we use it to show the run-time benefits of *PVM* removing cross-lingual voice cloning in cascaded S2ST systems.

Figure 2 compares *GEMO-Match* with the OpenVoice framework in terms of run-time scaling in multi-speaker scenarios. We plotted the number of times each system must re-run auxiliary modules while performing TTS over time in multi-speaker instances. The plots were generated using Python.

## 7 Experimental Results and Analysis

In this section, we discuss and analyze our experimental results.

Section 7.1 describes the results of the *GEMO-Match* robustness experiment, contained in tables 1 and 2. Next, section 7.2 provides an analysis on the results in Table 3. Section 7.3 then highlights our run-time experiment results.

### 7.1 GEMO-Match Robustness Results

Tables 1 and 2 show the precision and accuracy of the *Similarity Feature Extraction* module classifiers. Testing *GEMO-Match* on RAVDESS across emotions, the *Male-Emotion Classifier* performs

best on happy, angry, and neutral, which have precision scores of 78%, 78%, and 80%, respectively. The *Female-Emotion Classifier* performs well on angry and neutral, achieving 100% and 90% precision, respectively. We find *GEMO-Match* overfits to certain gender-emotion classes when trained on RAVDESS. This is prevalent in the *Female-Emotion Classifier* performance, as it classifies angry emotions with perfect precision, but classifies sad and disgust with 40% precision.

As illustrated in Table 1, *GEMO-Match* generalizes more consistently across emotions when trained on *CGDD* compared to *RAVDESS*. In the cases of both datasets shown in Table 1, *GEMO-Match* tends to classify angry and neutral effectively. The improvements in generalization described in Table 1 when using *CGDD* instead of *RAVDESS* showcases that some benchmarks are currently lacking variation. *CGDD* can remedy this, as it has higher variance compared to *RAVDESS*, comprising of multiple benchmark datasets as described in section 5.2.

Table 2 shows the accuracy of *GEMO-Match* on RAVDESS and *CGDD*. The *GEMO-Match* gender classifier scored 94% accuracy on the RAVDESS dataset, and 97% on *CGDD*. The best *GEMO-Match* emotion classifier results are found when training and testing on *CGDD*, which results in 63% accuracy for the *Male-Emotion Classifier* and 71% for the *Female-Emotion Classifier*. Therefore, our proposed *CGDD* dataset can improve model generalization compared to benchmark datasets like RAVDESS.

### 7.2 GEMO-Match Multilingual Results

The results in Table 3 show *PVM* implementations can significantly improve the output naturalness of S2ST systems by enabling monolingual TTS within S2ST. We find this trend holds across the two tested languages, French and German. When

Target Language	XTTS Input NISQA	XTTS Output NISQA
<i>Cross-lingual Cloning (English prompt)</i>		
French	3.54	3.54
German	3.54	3.41
<i>Monolingual Cloning (PVM-matched preset)</i>		
French	<b>3.39</b>	<b>3.43</b>
German	<b>3.47</b>	<b>3.69</b>

Table 3: Speech quality behavior when decoupling multilingual transformation and voice cloning in S2ST. XTTS performs significantly better when cloning in a monolingual context. Inputs are shown in parentheses.

XTTS performs cross-lingual TTS from English to German, NISQA values decrease from 3.54 (English) to 3.41 (German). Similarly, when XTTS cross-lingually clones from English to French, the input-output NISQA values are 3.54 and 3.54, respectively. Overall, XTTS either maintained or degraded the input naturalness when performing cross-lingual cloning in our experiments.

We find XTTS performs much better in a monolingual setting, which can significantly enhance S2ST quality. The average NISQA score when XTTS cloned from German preset-voices to German outputs increased from 3.47 to 3.69. The same increase is seen with French, though to a lesser degree. For our tested language pairs, *GEMO-Match* consistently improves output naturalness by allowing S2ST pipelines to clone in a monolingual context while maintaining cross-lingual behavior.

### 7.3 GEMO-Match Run-time Results

The run-time results of different TTS approaches are shown in Figure 1. VALL-E X and XTTS, deep multilingual voice cloning models, are slowest on average. SeamlessM4T offers multilingualism in multiple modalities, but does not clone voices, and has significantly lower runtime than the aforementioned models. This underscores additional complexities inherent to achieving speech translation and voice cloning in a single embedding space.

In our experiments, the lowest run-times were achieved by our *PVM* implementation (*GEMO-Match* with StyleTTS2) and OpenVoice. Both of these frameworks are not strictly limited to a specific TTS module for processing. As such, the runtime of their auxiliary, decoupled systems are noted separately in Figure 1. OpenVoice uses the post-processing tone extractor described in (Qin et al., 2024), and *PVM* uses *GEMO-Match*. For




Figure 1: Comparative processing times of different models. OpenVoice’s *tone extractor* and *GEMO-Match* are distinguished from their TTS processing times.

these isolated auxiliary modules, we achieved an average runtime of 0.52 for OpenVoice and 0.61 seconds for *GEMO-Match*.

Figure 2 compares these auxiliary modules under sequential inference on long multi-speaker inputs. For this comparison, we focus on the run-time of the entire S2ST system. Figure 2 shows that *GEMO-Match* need only run when a new speaker is presented in the input, while OpenVoice must always post-process the TTS output to achieve the desired result. Therefore, *PVM* offers favourable scaling properties, making it desirable for many commercial use-cases.

## 8 Future Work

*PVM* is a general framework for regulated S2ST that can be integrated into pre-existing cascaded S2ST pipelines. The performance of *PVM* is directly dependent on the quality of the individual swappable components of the pipeline. Consequently, the efficacy of any *PVM* implementation is expected to increase with general advancements in TTS technology. There are many ways to improve the *PVM* framework, and we propose some ideas for future work.

For future work, we propose a cascaded voice cloning TTS system which uses an initial *vocal*

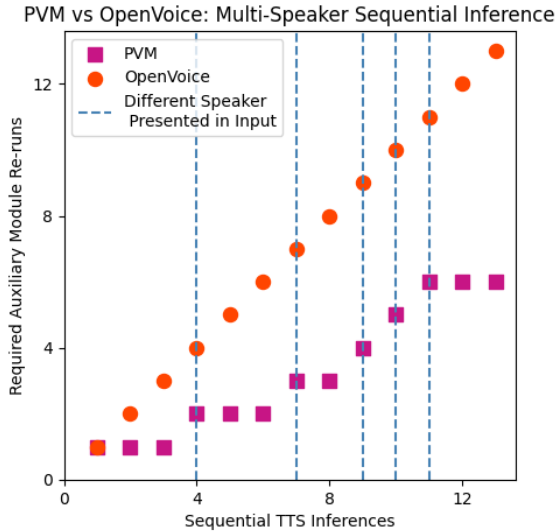


Figure 2: The OpenVoice tone extractor post-processes every TTS output. *GEMO-Match* only needs to re-run on the arrival of a different speaker from the one present in the previous input.

*encoder* with learned weights to extract and compress relevant features from the input voice. The system would perform the classical cloning tasks on this encoded voice in a downstream, decoupled TTS model. This would allow voices to be stored in the *Preset-Voice Library* in their encoded formats rather than speech signals, likely decreasing run-time complexity. Using a cascaded learning process, the TTS module would learn to effectively clone and only synthesize voices encoded by the *vocal encoder*. During distribution of the system, the *vocal encoder* would not be published. In this way, the system could not be used to clone a voice outside of the pre-encoded preset-voices in the *Preset-Voice Library*.

*GEMO-Match* uses classifiers which depend on labeled data. This dependency motivates the development of alternative *PVM* instances capable of voice-matching without relying on labeled data. We posit that learned encodings can be used, akin to self-supervised learning mechanisms employed by transformer architectures, to extract robust internal representations of speech inputs (Devlin et al., 2019; Babu et al., 2021). This would require a new training pipeline with an objective function for maximizing speaker similarity between the input voice and the matched voice. The resulting *PVM* system could use latent feature representations to perform voice matching, and training would not require labeled speech datasets.

## 9 Conclusion

We proposed *Preset-Voice Matching*, a novel framework that bakes regulatory precautions into the S2ST process. *PVM* achieves this by removing the explicit objective of cloning an unknown input speaker’s voice, and instead cloning a similar preset-voice of a consenting speaker. This paradigm is extensible to a variety of industry settings to regulate the behavior of S2ST systems. Quantitative experiments show *PVM* is a desirable framework compared to the tested benchmarks in terms of run-time and naturalness of multilingual translation output. Additionally, we provided *CGDD*, a gender-dependent speech-emotion dataset. We then showed *CGDD* leads to better model generalization and robustness in terms of accuracy and precision compared to the benchmark RAVDESS dataset. We hope this work inspires others to create more privacy regulated S2ST systems using the *PVM* framework.

## 10 PVM Limitations

In this section, we discuss the limitations of *GEMO-Match* and the *PVM* framework.

*GEMO-Match* requires training 3 unique classifiers for every source-language supported by the system. Specifically, the three *Feature Extraction Module* classifiers need to be trained on language specific emotional speech datasets processed into 3 versions: the entire dataset labeled by gender, and two subsets containing the gender-dependent labeled data. Gathering and processing data as described for each desired source-language may be complicated depending on data availability.

We acknowledge that the three features language, gender, and emotion alone are inadequate to fully capture the breadth of speaker variance across human speech. There are scenarios which demand more closely matched consented speakers in terms of vocal characteristics of the input speaker. *GEMO-Match* has strong limitations in this respect, which necessitates systems with more granularity in terms of speech feature extraction than what is offered by *GEMO-Match*.

Additionally, *PVM* makes no attempt to mimic background ambience or environmental noise in the inputted audios, as it loses this information when matching to a preset-voice. This is a drawback of *PVM*, as maintaining background audio noise information is highly important in some settings. However, many modern S2ST systems denoise in-



put audio to improve model performance, and add the noise back during post-processing. *PVM* would not be limited in such an environment, and can ensure high-quality voice inputs to the TTS module by always mapping to high-quality consenting speaker audios.

Lastly, we consider the drawback of error propagation in the *PVM* framework, inherent to cascaded architectures with separate modules. Ultimately, using a set of separate modules introduces additional points of failure, causing inaccuracies which are passed to downstream tasks.

## 11 Appendix

### 11.1 Industry Applications

In this section, we include some examples of cases where *PVM* can be applied to industry settings.

APIs are a common avenue for controlled public access to ML models and pipelines. These access points are commonly subjected to adversarial attacks, where imperceptible artefacts are injected into inputs to produce undesirable results. In the *PVM* framework, the audio input by our user is not directly passed to the TTS model, and is only matched to a consented speaker using feature similarity. This limits the scope of poor results that could be triggered by an adversarial user by negating direct access to the TTS model. Additionally, propagating audio input data from a genuine user through fewer modules in the pipeline limits opportunities for sensitive bio-metric data to be extracted by malicious third parties. Ultimately, removing direct control over synthesis of the input voice prevents bad actors from cloning a non-consenting speaker for nefarious goals.

We also consider how *PVM* can be extended to help regulate open-source models. As mentioned in Section 8, an autoencoder could be applied to derive robust latent space representations of the preset-voices. Matching based on similarity would still occur on the raw preset-voice audios, but their corresponding preset encodings would be passed as input to the voice cloning TTS model. The encoder/decoder models would not be published alongside the rest of the system. As the TTS model would have only been trained on the latent embeddings, the published system could not be hijacked to clone non-consenting voices.

In content localization systems, media content is leased by distributing platforms, while rights to the reproduction of the likenesses of individuals

present in the content is not readily available. Not only can *PVM* secure these systems in the manners mentioned above, but its regulated application can help bring this budding market to life by efficiently producing translated content in only the voices of consenting speakers. We believe *PVM* provides feasibility to the commercialization of such systems while being robust against future industry regulations.

We hope these examples give insight into the vast extensibility of the *PVM* framework.

## 12 Acknowledgements

We thank the anonymous reviewers for their feedback on this work. We also thank Joy Christian, Chao-Lin Chen, Sina Pordanesh, Akash Lakhani, Kaivil Brahmhatt, and Darshan Sarkale of Vosyn’s *PVM* team for their contributions on the early stages of this work.

## References

- Naroa Amezaga and Jeremy Hajek. 2022. [Availability of voice deepfake technology and its impact for good and evil](#). In *Proceedings of the 23rd Annual Conference on Information Technology Education*, SIGITE ’22. Association for Computing Machinery.
- Sercan O. Arik, Jitong Chen, Kainan Peng, Wei Ping, and Yanqi Zhou. 2018. [Neural voice cloning with a few samples](#). *Preprint*, arXiv:1802.06006.
- Arun Babu, Changan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2021. [Xls-r: Self-supervised cross-lingual speech representation learning at scale](#). *Preprint*, arXiv:2111.09296.
- H. Barakat, O. Turk, and C. Demiroglu. 2024. [Deep learning-based expressive speech synthesis: a systematic review of approaches, challenges, and resources](#). *EURASIP Journal on Audio, Speech, and Music Processing*, 11:11.
- Antonios Baris. 2024. [Ai covers: legal notes on audio mining and voice cloning](#). In *AI covers: legal notes on audio mining and voice cloning*.
- William Brannon, Yogesh Virkar, and Brian Thompson. 2022. [Dubbing in practice: A large scale study of human localization with insights for automatic dubbing](#). *Preprint*, arXiv:2212.12137.
- Felix Burkhardt, Astrid Paeschke, Miriam Rolfes, Walter F Sendlmeier, and Benjamin Weiss. 2005. [A database of German emotional speech](#). In *Proceedings of the Annual Conference of the International Speech Communication Association (INTER-SPEECH)*, volume 5, pages 1517–1520, Lisbon, Portugal. ISCA.

- Houwei Cao, David G. Cooper, Michael K. Keutmann, Ruben C. Gur, Ani Nenkova, and Ragini Verma. 2014. [Crema-d: Crowd-sourced emotional multimodal actors dataset](#). *IEEE Transactions on Affective Computing*, 5(4):377–390.
- Seamless Communication, Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, Christopher Klaiber, Pengwei Li, Daniel Licht, Jean Maillard, Alice Rakotoarison, Kaushik Ram Sadagopan, Guillaume Wenzek, Ethan Ye, Bapi Akula, Peng-Jen Chen, Naji El Hachem, Brian Ellis, Gabriel Mejia Gonzalez, Justin Haaheim, Prangthip Hansanti, Russ Howes, Bernie Huang, Min-Jae Hwang, Hirofumi Inaguma, Somya Jain, Elahe Kalbassi, Amanda Kallet, Iliia Kulikov, Janice Lam, Daniel Li, Xutai Ma, Ruslan Mavlyutov, Benjamin Peloquin, Mohamed Ramadan, Abinesh Ramakrishnan, Anna Sun, Kevin Tran, Tuan Tran, Igor Tufanov, Vish Vogeti, Carleigh Wood, Yilin Yang, Bokai Yu, Pierre Andrews, Can Balioglu, Marta R. Costa-jussà, Onur Celebi, Maha Elbayad, Cynthia Gao, Francisco Guzmán, Justine Kao, Ann Lee, Alexandre Mourachko, Juan Pino, Sravya Popuri, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Paden Tomasello, Changhan Wang, Jeff Wang, and Skyler Wang. 2023. [SeamlessM4T: Massively multilingual & multimodal machine translation](#). *Preprint*, arXiv:2308.11596.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. [Imagenet: A large-scale hierarchical image database](#). In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *Preprint*, arXiv:1810.04805.
- Gölge Eren and The Coqui TTS Team. 2023. [Coqui tts](#).
- Thierry Etchegoyhen, Haritz Arzelus, Harritxu Gete, Aitor Alvarez, Iván G. Torre, Juan Manuel Martín-Doñas, Ander González-Docasal, and Edson Benites Fernandez. 2022. [Cascade or direct speech translation? a case study](#). *Applied Sciences*, 12(3).
- Omatseyin Eyesan and Senanu Okuboyejo. 2013. [Design and implementation of a voice-based medical alert system for medication adherence](#). In *Design and Implementation of a Voice-Based Medical Alert System for Medication Adherence*, volume 9.
- Zeshan Fayyaz, Daniel Platnick, Hannan Fayyaz, and Nariman Farsad. 2022. [Deep unfolding for iterative stripe noise removal](#). In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7.
- Guy Frankovits and Yisroel Mirsky. 2023. [Discussion paper: The threat of real time deepfakes](#). *Preprint*, arXiv:2306.02487.
- Philippe Gournay, Olivier Lahaie, and Roch Lefebvre. 2018. [A canadian french emotional speech dataset](#).
- Priyanka Gujarathi and Sandip Raosaheb Patil. 2021. [Review on unit selection-based concatenation approach in text to speech synthesis system](#). In *Cybernetics, Cognition and Machine Learning Applications*, Algorithms for Intelligent Systems (AIS), pages 191–202.
- Wen-Chin Huang, Benjamin Peloquin, Justine Kao, Changhan Wang, Hongyu Gong, Elizabeth Salesky, Yossi Adi, Ann Lee, and Peng-Jen Chen. 2023. [A holistic cascade system, benchmark, and human evaluation protocol for expressive speech-to-speech translation](#). *Preprint*, arXiv:2301.10606.
- Ye Jia, Ron J. Weiss, Fadi Biadisy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu. 2019. [Direct speech-to-speech translation with a sequence-to-sequence model](#). *Preprint*, arXiv:1904.06037.
- Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021. [Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech](#). *Preprint*, arXiv:2106.06103.
- Simon King. 2011. [An introduction to statistical parametric speech synthesis](#). *c Indian Academy of Sciences*, 36:837–852.
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#). *Preprint*, arXiv:1412.6980.
- Yann Lecun and Y. Bengio. 1995. Convolutional networks for images, speech, and time-series. In *Convolutional Networks for Images, Speech, and Time-Series*.
- Yinghao Aaron Li, Cong Han, Vinay S. Raghavan, Gavin Mischler, and Nima Mesgarani. 2023. [Styletts 2: Towards human-level text-to-speech through style diffusion and adversarial training with large speech language models](#). *Preprint*, arXiv:2306.07691.
- Chang Liu, Jie Zhang, Tianwei Zhang, Xi Yang, Weiming Zhang, and Nenghai Yu. 2023a. [Detecting voice cloning attacks via timbre watermarking](#). *Preprint*, arXiv:2312.03410.
- Zihao Liu, Yan Zhang, and Chenglin Miao. 2023b. [Protecting your voice from speech synthesis attacks](#). In *Proceedings of the 39th Annual Computer Security Applications Conference, ACSAC '23*, page 394–408. Association for Computing Machinery.
- Steven R. Livingstone and Frank A. Russo. 2018. [The Ryerson Audio-Visual Database of Emotional Speech and Song \(RAVDESS\)](#).
- Gabriel Mittag, Babak Naderi, Assmaa Chehadi, and Sebastian Möller. 2021. [Nisqa: A deep cnn-self-attention model for multidimensional speech quality prediction with crowdsourced datasets](#). In *InterSpeech 2021*. ISCA.

- Felipe Romero Moreno. 2024. [Generative ai and deep-fakes: a human rights approach to tackling harmful content](#). *International Review of Law, Computers & Technology*, 0(0):1–30.
- Yishuang Ning, Sheng He, Zhiyong Wu, Chunxiao Xing, and Liang-Jie Zhang. 2019. [A review of deep learning based speech synthesis](#). *Applied Sciences*, 9(19).
- Chidera Okolie. 2023. Artificial intelligence-altered videos (deepfakes) and data privacy concerns. *Journal of International Women’s Studies*, 25:13.
- Orchid Chetia Phukan, Arun Balaji Buduru, and Rajesh Sharma. 2023. [A comparative study of pre-trained speech and audio embeddings for speech emotion recognition](#). *Preprint*, arXiv:2304.11472.
- M. Kathleen Pichora-Fuller and Kate Dupuis. 2020. [Toronto emotional speech set \(TESS\)](#).
- Daniel Platnick, Sourena Khanzadeh, Alireza Sadeghian, and Richard Valenzano. 2024. [Ganssemble for Small and Imbalanced Data Sets: A Baseline for Synthetic Microplastics Data](#). *Proceedings of the Canadian Conference on Artificial Intelligence*. <https://caiac.pubpub.org/pub/0hhra7j6>.
- Zengyi Qin, Wenliang Zhao, Xumin Yu, and Xin Sun. 2024. [Openvoice: Versatile instant voice cloning](#). *Preprint*, arXiv:2312.01479.
- Ilyes Rebai, Yessine BenAyed, Walid Mahdi, and Jean-Pierre Lorré. 2017. [Improving speech recognition using data augmentation and acoustic model fusion](#). *Procedia Computer Science*, 112:316–322. Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 21st International Conference, KES-20176-8 September 2017, Marseille, France.
- Panos Samaras and Michael Ferreira. 2019. Emergency communication systems effectiveness in an airport environment. *Journal of business continuity & emergency planning*, 12:242–252.
- Mohamed R. Shoaib, Zefan Wang, Milad Taleby Ahvanooy, and Jun Zhao. 2023. [Deepfakes, misinformation, and disinformation in the era of frontier ai, generative ai, and large ai models](#). *Preprint*, arXiv:2311.17394.
- Vandana Singh and Swati Prasad. 2023. [Speech emotion recognition system using gender dependent convolution neural network](#). *Procedia Computer Science*, 218:2533–2540. International Conference on Machine Learning and Data Engineering.
- Harsh Sinha, Vinayak Awasthi, and Pawan K. Ajmera. 2020. [Audio classification using braided convolutional neural networks](#). *IET Signal Processing*, 14(7):448–454.
- James Strathman, Thomas Kimpel, Kenneth Dueker, Richard Gerhart, Kenneth Turner, David Griffin, Steve Callas, Wan Ahmad, Teri Curtis, Wendy Martin, John Alldrin, Robert Mcalister, Brian Shields, Kathryn Saviers, Brook Bradford, Michael Seid, Twyla Ferguson, and Lillian Wilkinson. 2001. [Bus transit operations control: Review and an experiment involving tri-met’s automated bus dispatching system](#). *Journal of Public Trans.*, 4.
- K. N Sudhakar and M.B Shanthi. 2023. [Deepfake: An endanger to cyber security](#). In *2023 International Conference on Sustainable Computing and Smart Systems (ICSCSS)*, pages 1542–1548.
- Shahroz Tariq, Alsharif Abuadbbba, and Kristen Moore. 2023. [Deepfake in the metaverse: Security implications for virtual gaming, meetings, and offices](#). In *The 2nd Workshop on the security implications of Deepfakes and Cheapfakes*. ACM.
- Wee Jing Tee and R.K. Murugesan. 2021. [Protecting Data Privacy and Prevent Fake News and Deepfakes in Social Media via Blockchain Technology](#), pages 674–684. Springer.
- Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. [Wavenet: A generative model for raw audio](#). *Preprint*, arXiv:1609.03499.
- Yuxuan Wang, R. J. Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc V. Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous. 2017. [Tacotron: A fully end-to-end text-to-speech synthesis model](#). *CoRR*, abs/1703.10135.
- Gaoxiong Yi, Wei Xiao, Yiming Xiao, Babak Naderi, Sebastian Möller, Wafaa Wardah, Gabriel Mittag, Ross Cutler, Zhuohuang Zhang, Donald S. Williamson, Fei Chen, Fuzheng Yang, and Shidong Shang. 2022. [Conferencingspeech 2022 challenge: Non-intrusive objective speech quality assessment \(nisqa\) challenge for online conferencing applications](#). *Preprint*, arXiv:2203.16032.
- Julian Zaidi, Hugo Seuté, Benjamin van Niekerk, and Marc-André Carbonneau. 2021. [Daft-exprt: Robust prosody transfer across speakers for expressive speech synthesis](#). *CoRR*, abs/2108.02271.
- Ziqiang Zhang, Long Zhou, Chengyi Wang, Sanyuan Chen, Yu Wu, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, Lei He, Sheng Zhao, and Furu Wei. 2023. [Speak foreign languages with your own voice: Cross-lingual neural codec language modeling](#). *Preprint*, arXiv:2303.03926.