
OpenVoice: Versatile Instant Voice Cloning

Zengyi Qin *
MIT & MyShell.ai

Wenliang Zhao
Tsinghua University

Xumin Yu
Tsinghua University

Xin Sun
MyShell.ai

Abstract

We introduce OpenVoice, a versatile instant voice cloning approach that requires only a short audio clip from the reference speaker to replicate their voice and generate speech in multiple languages. OpenVoice represents a significant advancement in addressing the following open challenges in the field: **1) Flexible Voice Style Control.** OpenVoice enables granular control over voice styles, including emotion, accent, rhythm, pauses, and intonation, in addition to replicating the tone color of the reference speaker. The voice styles are not directly copied from and constrained by the style of the reference speaker. Previous approaches lacked the ability to flexibly manipulate voice styles after cloning. **2) Zero-Shot Cross-Lingual Voice Cloning.** OpenVoice achieves zero-shot cross-lingual voice cloning for languages not included in the massive-speaker training set. Unlike previous approaches, which typically require extensive massive-speaker multi-lingual (MSML) dataset² for all languages, OpenVoice can clone voices into a new language without any massive-speaker training data for that language. OpenVoice is also **computationally efficient**, costing tens of times less than commercially available APIs that offer even inferior performance. To foster further research in the field, we have made the source code³ and trained model publicly accessible. We also provide qualitative results in our demo website⁴. OpenVoice has served more than 2M users worldwide as the voice engine of MyShell.ai.

1 Introduction

Instant voice cloning (IVC) in text-to-speech (TTS) synthesis means the TTS model can clone the voice of any reference speaker given a short audio sample without additional training on the reference speaker. It is also referred to as Zero-shot TTS. IVC enables the users to flexibly customize the generated voice and exhibits tremendous value in a wide variety of real-world applications, such as media content creation, customized chatbots, and multi-modal interaction between humans and computers or large language models.

An abundant of previous work has been done in IVC. Examples of auto-regressive approaches include VALLE [16] and XTTS [3], which extract the acoustic tokens or speaker embedding from the reference audio as a condition for the auto-regressive model. Then the auto-regressive model sequentially generate acoustic tokens, which are then decoded to raw audio waveform. While these methods can clone the tone color, they do not allow users to flexibly manipulate other important style parameters such as emotion, accent, rhythm, pauses and intonation. Also, auto-regressive models are relatively computationally expensive and has relatively slow inference speed. Examples of non-autoregressive approach include YourTTS [2] and the recently developed Voicebox [8], which demonstrate significantly faster inference speed but are still unable to provide flexible control over

*Corresponding email: qinzy@mit.edu

²We use the term **MSML dataset** to denote a speech dataset with multiple languages, and each language has a vast amount of speakers. It is a typical requirement for training cross-lingual instant voice cloning models.

³<https://github.com/myshell-ai/OpenVoice>

⁴<https://research.myshell.ai/open-voice>

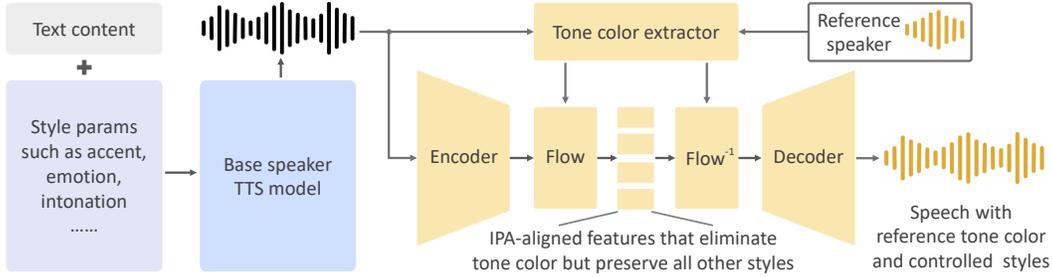


Figure 1: Illustration of the OpenVoice framework. We use a base speaker model to control the styles and languages, and a converter to embody the tone color of the reference speaker into the speech.

style parameters besides tone color. Another common disadvantage of the existing methods is that they typically require a huge MSML dataset in order to achieve cross-lingual voice clone. Such combinatorial data requirement can limit their flexibility to include new languages. In addition, since the voice cloning research [8, 16] by tech giants are mostly closed-source, there is not a convenient way for the research community to step on their shoulders and push the field forward.

We present OpenVoice, a flexible instant voice cloning approach targeted at the following key problems in the field:

- In addition to cloning the tone color, how to have flexible control of other important style parameters such as emotion, accent, rhythm, pauses and intonation? These features are crucial for generating in-context natural speech and conversations, rather than monotonously narrating the input text. Previous approaches [2, 3, 16] can only clone the monotonous tone color and style from the reference speaker but do not allow flexible manipulation of styles.
- How to enable zero-shot cross-lingual voice cloning in a simple way. We put forward two aspects of zero-shot capabilities that are important but not solved by previous studies:
 - If the language of the reference speaker is not presented in the MSML dataset, can the model clone their voice?
 - If the language of the generated speech is not presented in the MSML dataset, can the model clone the reference voice and generate speech in that language?

In previous studies [18, 8], the language of the reference speaker and the generated language by the model should both exist in great quantity in the MSML dataset. But what if neither of them exist?

- How to realize super-fast speed real-time inference without downgrading the quality, which is crucial for massive commercial production environment.

To address the first two problems, OpenVoice is designed to decouple the components in a voice as much as possible. The generation of language, tone color, and other important voice features are made independent of each other, enabling flexible manipulation over individual voice styles and language types. This is achieved without labeling any voice style in the MSML training set. We would like to clarify that **the zero-shot cross-lingual task in this study is different from that in VALLE-X [18]**. In VALLE-X, data for all languages need to be included in the MSML training set, and the model cannot generalize to an unseen language outside the MSML training set. By comparison, OpenVoice is designed to generalize to completely unseen languages outside the MSML training set. The third problem is addressed by default, since the decoupled structure reduces requirement on model size and computational complexity. We do not require a large model to learn everything. Also, we avoid auto-regressive or diffusion components to speed up the inference.

Our internal version of OpenVoice before this public release has been used tens of millions of times by users worldwide between May and October 2023. It powers the instant voice cloning backend of MyShell.ai and has witnessed several hundredfold user growth on this platform. To facilitate the research progress in the field, we explain the technology in great details and make the source code with model weights publicly available.

2 Approach

The technical approach is **simple** to implement but surprisingly effective. We first present the intuition behind OpenVoice, then elaborate on the model structure and training.

2.1 Intuition

The Hard. It is obvious that **simultaneously** cloning the tone color for any speaker, enabling flexible control of all other styles, and adding new language with little effort could be very challenging. It requires a huge amount of combinatorial datasets where the controlled parameters intersect, and pairs of data that only differ in one attribute, and are well-labeled, as well as a relatively large-capacity model to fit the dataset.

The Easy. We also notice that in regular single-speaker TTS, as long as voice cloning is not required, it is relatively easy to add control over other style parameters and add a new language. For example, recording a single-speaker dataset with 10K short audio samples with labeled emotions and intonation is sufficient to train a single-speaker TTS model that provides control over emotion and intonation. Adding a new language or accent is also straightforward by including another speaker in the dataset.

The intuition behind OpenVoice is to decouple the IVC task into separate subtasks where every subtask is much easier to achieve compared to the coupled task. The cloning of tone color is fully decoupled from the control over all remaining style parameters and languages. **We propose to use a base speaker TTS model to control the style parameters and languages, and use a tone color converter to embody the reference tone color into the generated voice.**

2.2 Model Structure

We illustrate the model structure in Figure. 1. The two main components of OpenVoice are the base speaker TTS model and the tone color converter. The base speaker TTS model is a single-speaker or multi-speaker model, which allows control over the style parameters (e.g., emotion, accent, rhythm, pauses and intonation), accent and language. The voice generated by this model is passed to the tone color converter, which changes the tone color of the base speaker into that of the reference speaker.

Base Speaker TTS Model. The choice of the base speaker TTS model is very flexible. For example, the VITS [6] model can be modified to accept style and language embedding in its text encoder and duration predictor. Other choices such as InstructTTS [17] can also accept style prompts. It is also possible to use commercially available (and cheap) models such as Microsoft TTS, which accepts speech synthesis markup language (SSML) that specifies the emotion, pauses and articulation. One can even skip the base speaker TTS model, and read the text by themselves in whatever styles and languages they desire. In our OpenVoice implementation, we used the VITS [6] model by default, but other choices are completely feasible. We denote the outputs of the base model as $\mathbf{X}(L_I, S_I, C_I)$, where the three parameters represent the language, styles and tone color respectively. Similarly, the speech audio from the reference speaker is denoted as $\mathbf{X}(L_O, S_O, C_O)$.

Tone Color Converter. The tone color converter is an encoder-decoder structure with a invertible normalizing flow [12] in the middle. The encoder is an 1D convolutional neural network that takes the short-time Fourier transformed spectrum of $\mathbf{X}(L_I, S_I, C_I)$ as input. All convolutions are single-strided. The feature maps outputted by the encoder are denoted as $\mathbf{Y}(L_I, S_I, C_I)$. The tone color extractor is a simple 2D convolutional neural network that operates on the mel-spectrogram of the input voice and outputs a single feature vector that encodes the tone color information. We apply it on $\mathbf{X}(L_I, S_I, C_I)$ to obtain vector $\mathbf{v}(C_I)$, then apply it on $\mathbf{X}(L_O, S_O, C_O)$ to obtain vector $\mathbf{v}(C_O)$.

The normalizing flow layers take $\mathbf{Y}(L_I, S_I, C_I)$ and $\mathbf{v}(C_I)$ as input and outputs a feature representation $\mathbf{Z}(L_I, S_I)$ that eliminates the tone color information but preserves all remaining style properties. **The feature $\mathbf{Z}(L_I, S_I)$ is aligned with International Phonetic Alphabet (IPA) [1] along the time dimension.** Details about how such feature representation is learned will be explained in the next section. Then we apply the normalizing flow layers in the inverse direction, which takes $\mathbf{Z}(L_I, S_I)$ and $\mathbf{v}(C_O)$ as input and outputs $\mathbf{Y}(L_I, S_I, C_O)$. This is a critical step where the tone color C_O from the reference speaker is embodied into the feature maps. Then the $\mathbf{Y}(L_I, S_I, C_O)$ is decoded into raw waveforms $\mathbf{X}(L_I, S_I, C_O)$ by HiFi-Gan [7] that contains a stack of transposed 1D convolutions. The entire model in our OpenVoice implementation is feed-forward without any auto-regressive compo-

ment. The tone color converter is conceptually similar to voice conversion [14, 11], but with different emphasis on its functionality, inductive bias on its model structure and training objectives. The flow layers in the tone color converter are structurally similar to the flow-based TTS methods [6, 5] but with different functionalities and training objectives.

Alternative Ways and Drawbacks. Although there are alternative ways [4, 9, 14] to extract $\mathbf{Z}(L_I, S_I)$, we empirically found that the proposed approach achieves the best audio quality. One can use HuBERT [4] to extract discrete or continuous acoustic units [14] to eliminate tone color information, but we found that such method also eliminates emotion and accent from the input speech. When the input is an unseen language, this type of method also has issues preserving the natural pronunciation of the phonemes. We also studied another approach [9] that carefully constructs information bottleneck to only preserve speech content, but we observed that this method is unable to completely eliminate the tone color.

Remark on Novelty. OpenVoice does not intend to *invent* the submodules in the model structure. Both the base speaker TTS model and the tone color converter borrow the model structure from existing work [5, 6]. The contribution of OpenVoice is the **decoupled framework that separates the voice style and language control from the tone color cloning**. This is very simple, but very effective, especially when one wants to control styles, accents or generalize to new languages. If one wanted to have the same control on a coupled framework such as XTTS [3], it could require a tremendous amount of data and computing, and it is relatively hard to fluently speak every language. In OpenVoice, as long as the single-speaker TTS speaks fluently, the cloned voice will be fluent. Decoupling the generation of voice styles and language from the generation of tone color is the core philosophy of OpenVoice. We also provided our insights of using flow layers in tone color converter, and the importance of choosing a universal phoneme system in language generalization in our experiment section.

2.3 Training

In order to train the base speaker TTS model, we collected audio samples from two English speakers (American and British accents), one Chinese speaker and one Japanese speaker. There are 30K sentences in total, and the average sentence length is 7s. The English and Chinese data has emotion classification labels. We modified the VITS [6] model and input the emotion categorical embedding, language categorical embedding and speaker id into the text encoder, duration predictor and flow layers. The training follows the standard procedure provided by the authors of VITS [6]. The trained model is able to change the accent and language by switching between different base speakers, and read the input text in different emotions. We also experimented with additional training data and confirmed that rhythm, pauses and intonation can be learned in exactly the same way as emotions.

In order to train the tone color converter, we collected 300K audio samples from 20K individuals. Around 180K samples are English, 60K samples are Chinese and 60K samples are Japanese. This is what we called the MSML dataset. The training objectives of the tone color converter is two-fold. First, we require the encoder-decoder to produce natural sound. During training, we feed the encoder output directly to the decoder, and supervised the generated waveform using the original waveform with mel-spectrogram loss and HiFi-GAN [7] loss. We will not detail here as it has been well explained by previous literature [7, 6].

Second, we require flow layers to eliminate as much tone color information as possible from the audio features. During training, for each audio sample, its text is converted to a sequence of **phonemes in IPA** [1], and each phoneme is represented by a learnable vector embedding. The sequence of vector embedding is passed to a transformer [15] encoder to produce the feature representation of the text content. Denote this feature as $\mathbf{L} \in \mathbb{R}^{c \times l}$, where c is the number of feature channels and l is the number of phonemes in the input text. The audio waveform is processed by the encoder and flow layers to produce the feature representation $\mathbf{Z} \in \mathbb{R}^{c \times t}$, where t is the length of the features along the time dimension. Then we align \mathbf{L} with \mathbf{Z} along the time dimension using dynamic time warping [13, 10] (an alternative is monotonic alignment [5, 6]) to produce $\bar{\mathbf{L}} \in \mathbb{R}^{c \times t}$, and minimize the KL-divergence between $\bar{\mathbf{L}}$ and \mathbf{Z} . Since $\bar{\mathbf{L}}$ does not contain any tone color information, the minimization objective would encourage the flow layers to remove tone color information from their output \mathbf{Z} . The flow layers are conditioned on the tone color information from the tone color encoder, which further helps the flow layers to identify what information needs to be eliminated. In addition, we do not provide any style or language information for the flow layers to condition on,

which prevents the flow layers to eliminate information other than tone color. Since the flow layers are invertible, conditioning them on a new piece of tone color information and running its inverse process can add the new tone color back to the feature representations, which are then decoded to the raw waveform with the new tone color embodied.

3 Experiment

The evaluation of voice cloning is hard to be objective for several reasons. First, different research studies (e.g., [8], [2]) usually have different training and test sets. The numerical comparison could be intrinsically unfair. Even their metrics such as Mean Opinion Score can be evaluated by crowdsourcing, the diversity and difficulty of the test set would significantly influence the results. For example, if many samples in the test set are neural voices that concentrate on the mean of human voice distributions, then it is relatively easy for most methods to achieve good voice cloning results. Second, different studies usually have different training sets, where the scale and diversity would have considerable influence of the results. Third, different studies can have different focus on their core functionalities. OpenVoice mainly aims at tone color cloning, flexible control over style parameters, and making cross-lingual voice clone easy even without massive-speaker data for a new language. These are different from the objectives of previous work on voice cloning or zero-shot TTS. Therefore, instead of comparing numerical scores with existing methods, we mainly focus on analyzing the qualitative performance of OpenVoice itself, and make the audio samples publicly available for relevant researchers to freely evaluate.

Accurate Tone Color Cloning. We build a test set of reference speakers selected from celebrities, game characters and anonymous individuals. The test set covers a wide voice distributions including both expressive unique voices and neutral samples in human voice distribution. With any of the 4 base speakers and any of the reference speaker, OpenVoice is able to accurately clone the reference tone color and generate speech in multiple languages and accents. We invite the readers to this website⁵ for qualitative results.

Flexible Control on Voice Styles. A premise for the proposed framework to flexibly control the speech styles is that the tone color converter is able to only modify the tone color and preserves all other styles and voice properties. In order to confirm this, we use both our base speaker model and the Microsoft TTS with SSML to generate a speech corpus of 1K samples with diverse styles (emotion, accent, rhythm, pauses and intonation) as the base voices. After converting to the reference tone color, we observed that all styles are well-preserved. In rare cases, the emotion will be slightly neutralized, and one way that we found to solve this problem is to replace the tone color embedding vector of this particular sentence with the average vector of multiple sentences with different emotions from the same base speaker. This gives less emotion information to the flow layers so that they do not eliminate the emotion. Since the tone color converter is able to preserve all the styles from the base voice, controlling the voice styles becomes very straightforward by simply manipulating the base speaker TTS model. The qualitative results are publicly available on this website⁶.

Cross-Lingual Voice Clone with Ease. OpenVoice achieves near zero-shot cross-lingual voice cloning without using any massive-speaker data for an unseen language. It does require a base speaker of the language, which can be achieved with minimum difficulty with the off-the-shelf models and datasets. On our website⁷, we provide an abundant of samples that demonstrates the cross-lingual voice clone capabilities of the proposed approach. The cross-lingual capabilities are two-fold:

- When the language of the reference speaker is unseen in the MSML dataset, the model is able to accurately clone the tone color of the reference speaker.
- When the language of the generated speech is unseen in the MSML dataset, the model is able to clone the reference voice and speak in that language, as long as the base speaker TTS supports that language.

Fast Inference with Low Cost. Since OpenVoice is a feed-forward structure without any auto-regressive component, it achieves very high inference speed. Our experiment shows that a slightly

⁵<https://research.myshell.ai/open-voice/accurate-tone-color-cloning>

⁶<https://research.myshell.ai/open-voice/flexible-voice-style-control>

⁷<https://research.myshell.ai/open-voice/zero-shot-cross-lingual-voice-cloning>

optimized version of OpenVoice (including the base speaker model and the tone converter) is able to achieve $12\times$ real-time performance on a single A10G GPU, which means it only takes 85ms to generate a one second speech. Through detailed GPU usage analysis, we estimate that the upper bound is around $40\times$ real-time, but we will leave this improvement as future work.

Importance of IPA. We found that using IPA as the phoneme dictionary is crucial for the tone color converter to perform cross-lingual voice cloning. As we detailed in Section 2.3, in training the tone color converter, the text is first converted into a sequence of phonemes in IPA, then each phoneme is represented by a learnable vector embedding. The sequence of embedding is encoded with transformer layers and compute loss against the output of the flow layers, aiming to eliminate the tone color information. IPA itself is a cross-lingual unified phoneme dictionary, which enables the flow layers to produce a language-neutral representation. Even if we input a speech audio with unseen language to the tone color converter, it is still able to smoothly process the audio. We also experimented with other types of phoneme dictionaries but the resulting tone color converter tend to mispronounce some phonemes in unseen languages. Although the input audio is correct, there is a high likelihood that the output audio is problematic and sounds non-native.

4 Discussion

OpenVoice demonstrates remarkable instance voice cloning capabilities and is more flexible than previous approaches in terms of voice styles and languages. The intuition behind the approach is that it is relatively easy to train a base speaker TTS model to control the voice styles and languages, as long as we do not require the model to have the ability to clone the tone color of the reference speaker. Therefore, we proposed to decouple the tone color cloning from the remaining voice styles and the language, which we believe is the foundational design principle of OpenVoice. In order to facilitate future research, we make the source code and model weights publicly available.

References

- [1] I. P. Association. *Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet*. Cambridge University Press, 1999.
- [2] E. Casanova, J. Weber, C. D. Shulby, A. C. Junior, E. Gölge, and M. A. Ponti. Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone. In *International Conference on Machine Learning*, pages 2709–2720. PMLR, 2022.
- [3] CoquiAI. Xtts taking text-to-speech to the next level. *Technical Blog*, 2023.
- [4] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.
- [5] J. Kim, S. Kim, J. Kong, and S. Yoon. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. *Advances in Neural Information Processing Systems*, 33:8067–8077, 2020.
- [6] J. Kim, J. Kong, and J. Son. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, pages 5530–5540. PMLR, 2021.
- [7] J. Kong, J. Kim, and J. Bae. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033, 2020.
- [8] M. Le, A. Vyas, B. Shi, B. Karrer, L. Sari, R. Moritz, M. Williamson, V. Manohar, Y. Adi, J. Mahadeokar, et al. Voicebox: Text-guided multilingual universal speech generation at scale. *arXiv preprint arXiv:2306.15687*, 2023.
- [9] J. Li, W. Tu, and L. Xiao. Freevc: Towards high-quality text-free one-shot voice conversion. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [10] M. Müller. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84, 2007.
- [11] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W.-N. Hsu, A. Mohamed, and E. Dupoux. Speech resynthesis from discrete disentangled self-supervised representations. *arXiv preprint arXiv:2104.00355*, 2021.
- [12] D. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.

- [13] P. Senin. Dynamic time warping algorithm review. *Information and Computer Science Department University of Hawaii at Manoa Honolulu, USA*, 855(1-23):40, 2008.
- [14] B. van Niekerk, M.-A. Carbonneau, J. Zaïdi, M. Baas, H. Seuté, and H. Kamper. A comparison of discrete and soft speech units for improved voice conversion. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6562–6566. IEEE, 2022.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [16] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, et al. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*, 2023.
- [17] D. Yang, S. Liu, R. Huang, G. Lei, C. Weng, H. Meng, and D. Yu. Instructtts: Modelling expressive tts in discrete latent space with natural language style prompt. *arXiv preprint arXiv:2301.13662*, 2023.
- [18] Z. Zhang, L. Zhou, C. Wang, S. Chen, Y. Wu, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, et al. Speak foreign languages with your own voice: Cross-lingual neural codec language modeling. *arXiv preprint arXiv:2303.03926*, 2023.