

# VoiceMe: Personalized voice generation in TTS

Pol van Rijn<sup>1</sup>, Silvan Mertes<sup>2</sup>, Dominik Schiller<sup>2</sup>, Piotr Dura<sup>3</sup>, Hubert Siuzdak<sup>3</sup>, Peter M. C. Harrison<sup>4</sup>, Elisabeth André<sup>2</sup>, Nori Jacoby<sup>1</sup>

<sup>1</sup>Max-Planck-Institute for Empirical Aesthetics, Frankfurt, Germany

<sup>2</sup>Human-Centered Artificial Intelligence, Augsburg, Germany

<sup>3</sup>Charactr Inc., Los Angeles, US

<sup>4</sup>University of Cambridge, Cambridge, UK

pol.van-rijn@ae.mpg.de

## Abstract

Novel text-to-speech systems can generate entirely new voices that were not seen during training. However, it remains a difficult task to efficiently create personalized voices from a high dimensional speaker space. In this work, we use speaker embeddings from a state-of-the-art speaker verification model (SpeakerNet) trained on thousands of speakers to condition a TTS model. We employ a human sampling paradigm to explore this speaker latent space. We show that users can create voices that fit well to photos of faces, art portraits, and cartoons. We recruit online participants to collectively manipulate the voice of a speaking face. We show that (1) a separate group of human raters confirms that the created voices match the faces, (2) speaker gender apparent from the face is well-recovered in the voice, and (3) people are consistently moving towards the real voice prototype for the given face. Our results demonstrate that this technology can be applied in a wide number of applications including character voice development in audiobooks and games, personalized speech assistants, and individual voices for people with speech impairment.

**Index Terms:** voice, personalization, speech, avatar, characters, human-computer interaction

## 1. Introduction

In the last few years, multi-speaker text-to-speech (TTS) models have been proposed that can create entirely new high-quality voices [1, 2]. While these approaches can generate unique voices that are distinct from voices seen during training, it is unclear how to create a personalized voice that fits one’s own mental representation. The existence of such a tool would open the door to numerous creative and practical applications, such as developing customized voices for digital agents and robots, personalized speech assistants, bringing fictional characters or paintings to life, or developing individualized voices for speech impaired people.

Here we use humans to search the speaker latent space in a trained TTS model that uses speaker embeddings from a state-of-the-art speaker verification model [3]. This allows participants to efficiently find voices that match their perception of faces in realistic photographs, art portraits and cartoons. Subjective evaluations from an independent group of human raters confirm that the created voices match the faces.

---

Audio samples and the supplementary figure can be found here: <https://polvanrijn.github.io/VoiceMe/> All code is here: <https://github.com/polvanrijn/VoiceMe>

## 2. Background

Previous human-in-the-loop approaches have shown that humans can iteratively build feature representations for personalizing sound and speech. For example, Ritschel et al. [4] used an interactive evolutionary algorithm to allow users to create sounds that express intentions and emotions for a social robot. Van Rijn et al. [5] used an adaptive method to sample from latent semantic human representations [6] to find prototypes of emotional speech by adapting the latent representation of a GST Tacotron model [7].

In the past years, various approaches have been proposed to create novel voices. Jia et al. [1] demonstrated that an independently trained speaker encoder network trained on a speaker verification task is able to produce useful conditioning for a multi-speaker text-to-speech model. By sampling random points from the obtained speaker embedding space, the authors generated fictitious voices that were not seen during the training. Another approach was proposed by Stanton et al. [2] that does not rely on transfer learning from the speaker verification task, but jointly learns a distribution over speaker embeddings, also allowing for sampling a novel voice. However, all of these papers focus on speaker generation, but not on speaker personalization. To our knowledge, this is the first paper using speaker generation models for voice personalization.

In the present paper, we adopt the method proposed by Jia et al. [1] and used embeddings from a speaker verification network as speaker representation: because they can be trained on noisy speech of thousands of speakers, do not require transcripts, can extract speaker embeddings for unseen voices, and the obtained voice prototypes can be reused in future models if trained on the same pretrained speaker verification network. While this work uses this particular state-of-the-art architecture for voice synthesis (Figure 1A), our human-sampling approach (Figure 1B) can be extended to work with a large class of speaker generation models so that they can be used for voice personalization.

## 3. Methods

### 3.1. TTS Architecture

Here, we use a modified version of the state-of-the-art TTS model VITS [8]. VITS is an end-to-end architecture with the following components (see Figure 1A): (i) text frontend composed of text normalization followed by a grapheme-to-phoneme model utilizing IPA characters (International Phonetic Alphabet) [9], (ii) transformer-based Text Encoder with a projection layer used to construct prior distribution, (iii) Normalizing Flow greatly improving the flexibility of prior latent space,

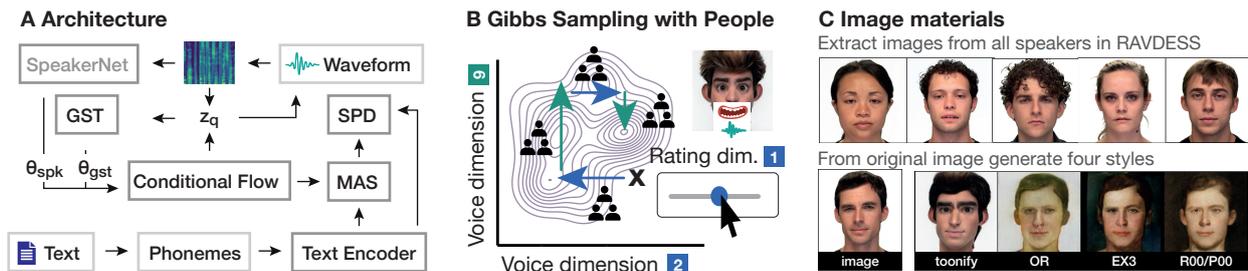


Figure 1: *Overview.* (A) The architecture used in this paper is a modified version of the VITS model using SpeakerNet and GST embeddings. (B) Participants move a slider adjusting a single dimension of the voice. The same slider is presented to three participants. The mean answer is passed to the next iteration. This allows us to gradually optimize the voice which fits to a face over the course of iterations. (C) Image materials. Images are extracted for all speakers in the RAVDESS corpus. We use deep-learning style-transfer to convert the images to cartoons and paintings.

(iv) convolutional Posterior Encoder, (v) Monotonic Alignment Search (MAS) and Stochastic Duration Predictor (SDP) modules learning to align input characters to encoded spectrogram frames, (vi) high-fidelity GAN vocoder based on the HiFi-GAN [10] architecture.

We use the pretrained speaker verification network SpeakerNet-M [3] as speaker embeddings. The model was trained on 7,205 speakers and has a satisfactory trade-off between the quality on a speaker recognition task and small computational footprint. We furthermore found that SpeakerNet mainly encoded voice pitch and timbre information and not prosody. In order to learn prosodic variation, we extend our model with a bank of Global Style Tokens (GST) [7], which extracts style embeddings from encoded spectrogram frames. We initialized 16 Global Style Tokens, 8 attention heads and the resulting embedding size was set to 256. As GST encoder an 8-layer convolutional network was used with the same architecture as the Posterior Encoder. During training the speaker and style embedding were separately L2-normalized and concatenated. During the experiments, we use the same zero embedding to keep prosody approximately constant across samples.

To prevent the GSTs from learning speaker-dependent features in presence of SpeakerNet embeddings, an additional adversarial loss is proposed as follows. Alongside discriminators a separate shallow feed-forward neural network is trained to reconstruct the speaker embeddings from extracted style embeddings. During the discriminator step, this network minimizes a cosine distance between real and reconstructed speaker embeddings using cosine embedding loss:  $(1 - \cos(x, \hat{x}))$ . Conversely, during the generator step, the style extractor is penalized if this network succeeded in reconstructing speaker embeddings – the loss function is then:  $(\max(0, \cos(x, \hat{x})))$ . The full architecture is depicted in Figure 1A.

We applied transfer learning from the publicly available VCTK checkpoint<sup>1</sup> and the training was continued using two NVIDIA V100 GPUs. For the first 400k iterations only the discriminators were allowed to train due to the lack of a published discriminator checkpoint, then normal training continued for additional 2M iterations with learning rate lowered to 1e-4.

### 3.2. Gibbs Sampling with People

‘Gibbs Sampling with People’ (GSP) is an adaptive technique [6], where many participants collaborate to navigate a stimu-

lus space and identify regions associated with a given semantic concept, in our case optimize voice parameters so that they match a target face animation. The participants’ responses are organized into sequences of iterations called “chains”. In each trial, the participant is presented with a stimulus (a synthesized voice) and a slider, where the slider is coupled to a particular dimension of the stimulus space that changes from trial to trial. The participant is instructed to move the slider to find the stimulus most associated with the target face. In our implementation, three different participants contribute trials for a given iteration in a given chain, and their responses are averaged. The resulting stimulus (based on the average response) is then passed along the chain of participants, with each successive generation of participants optimizing a different dimension of the stimulus (Figure 1B). This procedure is repeated multiple times, cycling through each of the dimensions of the sample space. Harrison et al. ([6]) demonstrated that the emergent process corresponds to a Gibbs sampler that maps the relationship between the target face and the participants’ internalized representations of the corresponding voice parameters

In the current experiment, participants change the first ten principal components of the SpeakerNet embeddings. Initial piloting suggested that these principal components had the desired property of intuitive interpretability (e.g., PC2 has a strong gender effect), and prior research with related models suggested that 10 principal components should be enough to achieve meaningful control over the stimuli [6]. The principal components were computed on SpeakerNet embeddings extracted on a single utterance of the 45,825 speakers present in the train, dev and test partitions of the English CommonVoice dataset [11] and account for 25.4 % of the variance. The participants are prompted to adjust a slider which corresponds to one principal component to make the voice maximally similar to a face (see Figure 1C for some of the faces). For practical reasons, every slider contains a finite resolution of 31 equally-spaced slider positions. As opposed to using static images, we use Wav2Lip [12] to synchronize the lips to the voice so that the resulting stimulus sounds more natural.

## 4. Materials

### 4.1. Stimuli

In order to compare the personalized voice with a ground truth we extracted stills and speaker embeddings from recordings of

<sup>1</sup><https://github.com/jaywalnut310/vits>

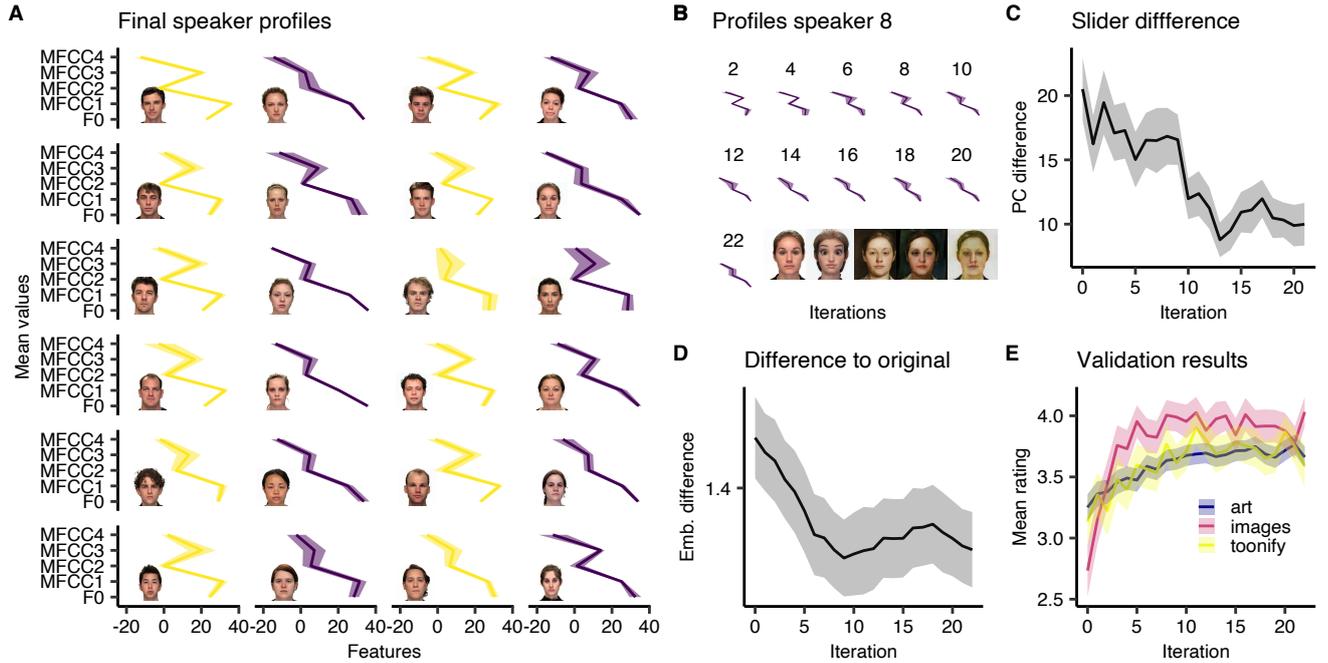


Figure 2: *Results.* (A) Final speaker profiles on mean values for F0 and MFCC1-4. The shaded areas in all plots refer to 95% confidence intervals. (B) Example profiles for speaker 8 at different iterations. (C) The Euclidean distance between consecutive iterations is larger for earlier iterations compared to later iterations. (D) The Euclidean distance to the original reference drops over the first ten dimensions and slightly increases and decreases. (E) The mean opinion score increases over the course of iterations.

real speakers. It was important to use speakers who are unknown to the participants, because this might constrain possible voices attributed to the face. We used the same neutral utterance of all 24 speakers in the RAVDESS corpus [13].

To demonstrate our approach does not only work for real faces, but also for fictional characters, we created for each original image four fictional characters based on style transfer. We used toonify [14] and three additional art portrait styles from Ai Gahaku<sup>2</sup>: OR, EX3 and ROO or P00 (see example in figure 1C). We selected the images in the following way. We start by creating 12 art portraits and one toonified version and then select four styles with the highest perceptual similarity to the real photo [15] (see Figure 1C). Thus, we select toonify, OR, and EX3 styles, but in 22 of 24 cases we select R00 and in all other cases we select P00.

For all 24 speakers, we use the extracted images and four styles with the highest perceptual similarity totaling 120 chains. To each chain we randomly assign one of the 720 phonetically balanced and semantically neutral Harvard sentences [16].

#### 4.2. Participants

All participants were recruited from Amazon Mechanical Turk (AMT) and provided informed consent in accordance with the Max Planck Society Ethics Council approved protocol (application 2021\_42). Participants were paid \$9/hour. Requirements for participation include a minimal age of 18 years, 99% or higher approval rate on at least 5,000 previous tasks on AMT, residency in the US and wearing headphones [17]. Participant recruitment was managed by PsyNet [6], an under-development framework for implementing complex experimental paradigms.

<sup>2</sup><https://ai-art.tokyo>

This framework builds on the Dallinger platform<sup>3</sup> for experiment hosting and deployment.

## 5. Results and discussion

### 5.1. Main experiment

180 US participants (69 female, 3 prefer not to say, 68 male) engaged in the main experiment. The age ranged from 19 to 78 years old ( $M = 41$ ,  $SD = 12$ ). We terminated the experiment after 48 hours, after which 99 out of the 120 chains were full (22 iterations). In the last iteration, we obtain unique profiles for all 24 speakers (Figure 2A) and while each voice is randomly initialized, the profiles rapidly develop towards the final prototypes (Figure 2B). Quantitatively, we show that the Euclidean distance between consecutive iterations within a chain decreases over the course of iterations (Figure 2C), stabilizing toward the final 15 iterations. This means that participants move the sliders to a lesser extent at later iterations, suggesting convergence. We also compared the created speaker embedding at each iteration with the original speaker embedding of the real speaker. As we can see in Figure 2D, the difference to the reference is dropping for the first 10 iteration and then mildly increases and decreases again.

### 5.2. Validation

In a separate validation experiment, participants ( $N = 110$ , 50 female, 1 prefer not to say, 59 male; age  $M = 35$ ,  $SD = 10$ ) rated how well the voice matches the moving face on a 5-point Mean Opinion Score (MOS): ‘Excellent’, ‘Good’, ‘Fair’, ‘Poor’, and ‘Bad match’. The validation included all stimuli generated in

<sup>3</sup><https://github.com/Dallinger/Dallinger>

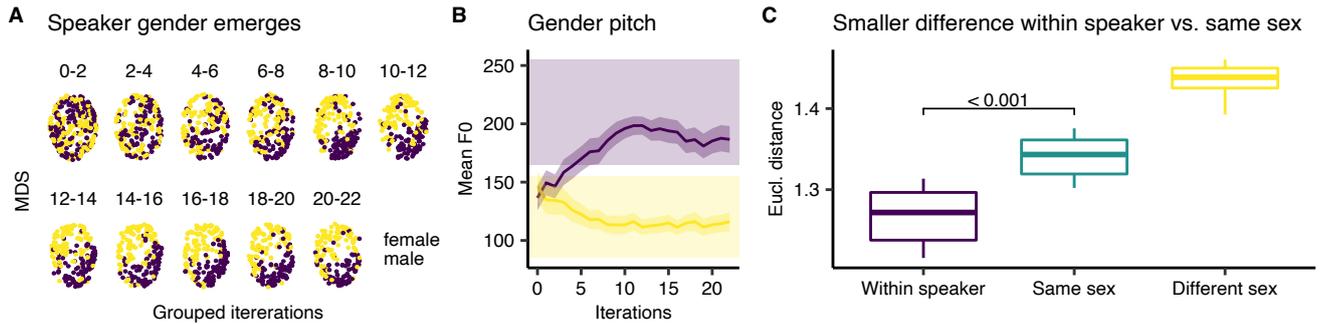


Figure 3: *Gender difference.* (A) Generated speaker embeddings put into MDS space in which voices for male and female pictures occupy increasingly distinct areas in the voice latent space. Each row represents two iterations, read from left to right, top to bottom. (B) The pitch strongly increases for female faces and is lowered for male faces in the first ten iterations. The solid background reflects common pitch ranges for the sexes. F0 is expressed in Hz. (C) In the final iterations the difference in the generated voice is smaller within the different styles of the same face (purple), compared to a random same sex (green) or different sex face (yellow).

the first experiment, (overall 2,409 stimuli). Participants performed 200 ratings per experiment and consequently on average every stimulus was rated 9.1 times. As depicted in Figure 2E, the average MOS increases over the course of iterations for all styles. However, the increase is largest for the original faces moving from a 2.7 MOS at iteration 0 to a MOS of 4.0 (‘Good match’) in the later iterations (Wilcoxon rank sum test,  $Z = .42$ ,  $p < 0.001$ , Bonferroni-adjusted). The trend is followed by the cartoons (Wilcoxon rank sum test,  $Z = .18$ ,  $p < 0.001$ , Bonferroni-adjusted). For art portraits the improvement over iterations is smallest (Wilcoxon rank sum test,  $Z = .16$ ,  $p < 0.001$ , Bonferroni-adjusted).

### 5.3. Toward personalized voice characteristics

To further understand what kind of voice features were selected by the personalization process, we visualize the speaker latent space using Multi Dimensional Scaling on all voices created in the experiment. As shown in Figure 3A, over the course of iterations, male and female faces occupy increasingly distinct areas in the voice latent space. Furthermore, the average pitch starts at roughly the same point due to the random initialization of the voices and over the course of iterations is lowered for male and increased for female voices (Figure 3B). Voices for males and females converge in a pitch range common for the sex (85-155 and 165-255 Hz respectively) as indicated by the shaded areas [18].

Based on these results, we can state that speaker gender apparent from the face is well-recovered in the voice. However, do people only focus on the gender or also on other characteristics of the face? In order to address this question, we run another analysis. Here, we compute the Euclidean difference between the voices created for different styles of the same speaker versus a random speaker of the same sex. Using bootstrapping ( $n = 1,000$ ) we show that the voice differences within the same speaker are significantly smaller compared to a voice of a random speaker of the same sex (Figure 3C). The results show that the voice prototypes can capture face-specific image features in addition to gender.

### 5.4. Limitations and outlook

Since the ratings plateau at a MOS of 4 (‘good match’) there is still room for improving the match. One improvement would

be to also apply GSP on the prosodic style embeddings instead of fixing the prosody by constraining a zero embedding on all chains. Enforcing neutral prosody on all stimuli can negatively affect intelligibility, which might be amplified by the fact that participants should focus on the match between voice and face and not the intelligibility of the speech.

Another approach to disentangle prosody could be to initially obtain the speaker embeddings from a speaker verification network and at a later stage to learn them with the other components of the model. This initialization with speaker verification embeddings can be a reasonable starting point. In the present study, we used the first ten principal components computed on the SpeakerNet Embeddings, but future research can explore alternative parametrizations of the speaker latent space that potentially better align with human perception and might explain more variance or alternatively increase the number of dimensions manipulated by the GSP procedure.

## 6. Conclusion

In this paper, we presented a human-in-the-loop approach for generating speech personalized to a specific face. To this end, we used an embedding space of a speaker verification network as input to a state-of-the-art TTS model. The embeddings were modified collaboratively by the use of a GSP experiment in order to find voice that is consistent to shown images of realistic faces, portraits and cartoon-style faces. Our evaluation showed that over the course of our experiment, the MOS of the generated speech, representing the consistency between the speech and the shown faces, increased from 2.7 to 4.0. This indicates that our approach is promising for synthesizing highly personalized speech for a specific speaker’s visual appearance. Taken together our results open up a vast range of creative and practical applications including personalized voices in audiobooks and games, personalized speech assistants, and individualized voices for people with speech impairments.

## 7. Acknowledgments

This work has partially been funded by the European Union Horizon 2020 research and innovation programme, grant agreement 856879 (PRESENT).

## 8. References

- [1] Y. Jia, Y. Zhang, R. Weiss, Q. Wang, J. Shen, F. Ren, P. Nguyen, R. Pang, I. Lopez Moreno, Y. Wu *et al.*, “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” *Advances in neural information processing systems*, vol. 31, 2018.
- [2] D. Stanton, M. Shannon, S. Mariooryad, R. Skerry-Ryan, E. Battenberg, T. Bagby, and D. Kao, “Speaker generation,” *arXiv preprint arXiv:2111.05095*, 2021.
- [3] N. R. Koluguri, J. Li, V. Lavrukhin, and B. Ginsburg, “Speaker-er-net: 1d depth-wise separable convolutional network for text-independent speaker recognition and verification,” *arXiv preprint arXiv:2010.12653*, 2020.
- [4] H. Ritschel, I. Aslan, S. Mertes, A. Seiderer, and E. André, “Personalized synthesis of intentional and emotional non-verbal sounds for social robots,” in *2019 8th International conference on affective computing and intelligent interaction (ACII)*. IEEE, 2019, pp. 1–7.
- [5] P. van Rijn, S. Mertes, D. Schiller, P. M. C. Harrison, P. Larrouy-Maestri, E. André, and N. Jacoby, “Exploring emotional prototypes in a high dimensional TTS latent space,” *CoRR*, vol. abs/2105.01891, 2021. [Online]. Available: <https://arxiv.org/abs/2105.01891>
- [6] P. Harrison, R. Marjeh, F. Adolphi, P. van Rijn, M. Anglada-Tort, O. Tchernichovski, P. Larrouy-Maestri, and N. Jacoby, “Gibbs sampling with people,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 10 659–10 671, 2020.
- [7] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 5180–5189.
- [8] 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*. PMLR, 2021, pp. 5530–5540.
- [9] M. Bernard, “Phonemizer,” <https://github.com/bootphon/phonemizer>, 2021.
- [10] J. Kong, J. Kim, and J. Bae, “Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis,” *CoRR*, vol. abs/2010.05646, 2020. [Online]. Available: <https://arxiv.org/abs/2010.05646>
- [11] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” in *LREC*, 2020, pp. 4218–4222.
- [12] K. R. Prajwal, R. Mukhopadhyay, V. P. Namboodiri, and C. Jawahar, “A lip sync expert is all you need for speech to lip generation in the wild,” in *Proceedings of the 28th ACM International Conference on Multimedia*. ACM, Oct. 2020. [Online]. Available: <https://doi.org/10.1145/3394171.3413532>
- [13] S. R. Livingstone and F. A. Russo, “The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in North American English,” *PLOS ONE*, vol. 13, no. 5, p. e0196391, 2018.
- [14] J. N. M. Pinkney and D. Adler, “Resolution dependent GAN interpolation for controllable image synthesis between domains,” *CoRR*, vol. abs/2010.05334, 2020. [Online]. Available: <https://arxiv.org/abs/2010.05334>
- [15] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *CVPR*, 2018.
- [16] “IEEE recommended practice for speech quality measurements,” IEEE, Tech. Rep., 1969, ISBN: 9781504402743. [Online]. Available: <https://ieeexplore.ieee.org/document/7405210/>
- [17] K. J. Woods, M. H. Siegel, J. Traer, and J. H. McDermott, “Headphone screening to facilitate web-based auditory experiments,” *Attention, Perception, & Psychophysics*, vol. 79, no. 7, pp. 2064–2072, 2017.
- [18] J. L. Fitch and A. Holbrook, “Modal vocal fundamental frequency of young adults,” *Archives of Otolaryngology - Head and Neck Surgery*, vol. 92, no. 4, pp. 379–382, Oct. 1970. [Online]. Available: <https://doi.org/10.1001/archotol.1970.04310040067012>