



# Deep Speech Synthesis from MRI-Based Articulatory Representations

Peter Wu<sup>1</sup>, Tingle Li<sup>1</sup>, Yijing Lu<sup>3</sup>, Yubin Zhang<sup>3</sup>, Jiachen Lian<sup>1</sup>, Alan W Black<sup>2</sup>, Louis Goldstein<sup>3</sup>, Shinji Watanabe<sup>2</sup>, Gopala K. Anumanchipalli<sup>1</sup>

<sup>1</sup>University of California, Berkeley, United States

<sup>2</sup>Carnegie Mellon University, United States

<sup>3</sup>University of Southern California, United States

peterw1@berkeley.edu

## Abstract

In this paper, we study articulatory synthesis, a speech synthesis method using human vocal tract information that offers a way to develop efficient, generalizable and interpretable synthesizers. While recent advances have enabled intelligible articulatory synthesis using electromagnetic articulography (EMA), these methods lack critical articulatory information like excitation and nasality, limiting generalization capabilities. To bridge this gap, we propose an alternative MRI-based feature set that covers a much more extensive articulatory space than EMA. We also introduce normalization and denoising procedures to enhance the generalizability of deep learning methods trained on MRI data. Moreover, we propose an MRI-to-speech model that improves both computational efficiency and speech fidelity. Finally, through a series of ablations, we show that the proposed MRI representation is more comprehensive than EMA and identify the most suitable MRI feature subset for articulatory synthesis.

**Index Terms:** speech synthesis, articulatory synthesis

## 1. Introduction

Deep speech synthesis technology has made significant advancements in recent years, leading to high-performing models for tasks such as text-to-speech [1, 2, 3], voice conversion [4, 5], and speech translation [6, 7]. However, the development of brain-to-speech devices [8, 9] still poses significant challenges, requiring faster and more data-efficient models. Articulatory synthesis [10, 11, 12, 13, 14, 15] offers a potential solution by synthesizing speech from a compact, smooth, and interpretable articulatory space [16, 17, 18, 19, 20, 21].

Electromagnetic articulography (EMA) is a commonly used articulatory representation [13], but it only contains 6 x-y points, making it challenging to comprehensively capture articulatory movements. Real-time magnetic resonance imaging (MRI) is a state-of-the-art tool that captures dynamic information about vocal tract movements and shaping during human speech production, offering a feature-rich alternative to EMA. It contains hundreds of x-y points, including positional information for the hard palate, pharynx, epiglottis, velum, and larynx, all of which are important for speech production but not directly described in raw EMA data. Moreover, recent advances in image acquisition and reconstruction techniques have enabled sufficient temporal and spatial resolutions (*e.g.*, 12ms and  $2.4 \times 2.4$  mm<sup>2</sup>) that allow researchers to study the intricate and dynamic interactions during speech production [22, 23].

However, since MRI dataset participants speak from inside a tube-shaped MRI machine, there is a noticeable amount of reverberation in the collected utterances, resulting in an unsatisfactory performance on MRI-to-speech synthesis. To overcome

this problem, we enhance the utterances and propose a generative adversarial network (GAN) based method that directly synthesizes waveform from articulatory features, which produces noticeably more intelligible speech than the baselines.

We summarize our contributions as follows:

- We propose a novel MRI-based representation for articulatory synthesis, along with effective preprocessing strategies for such data.
- We demonstrate that our proposed model outperforms baselines across several evaluation metrics.
- We quantitatively and qualitatively identify the advantages of MRI over EMA and the most important MRI features for articulatory synthesis.

Code and additional related information will be available at <https://github.com/articulatory/articulatory>.

## 2. MRI Dataset

We utilize the real-time MRI and its corresponding audio recordings of one native American English speaker (female, 25-year-old), with a total speech duration of approximately 17 minutes and a sampling rate of 20 kHz, which are acquired from a publicly available multispeaker MRI dataset [19]. This dataset includes midsagittal real-time MRI videos with a spatial resolution of  $2.4 \times 2.4$  mm<sup>2</sup> ( $84 \times 84$  pixels) and a temporal resolution of 12-ms (83 frames per second), capturing the vocal tract movements during the production of a comprehensive set of scripted and spontaneous speech materials.

To prepare the MRI data for our model, we use a semi-automatic method [24] to track the contours of vocal tract air-tissue boundaries in each raw MRI frame (Figure 1) and segmented the contours into anatomical components, as shown in Figure 2 and Figure 3. To mitigate the problem of overfitting, we pruned the MRI feature set by discarding segments that did not contribute much to understanding how speech production varies across utterances. Figure 3 presents the full set of segments, while Figure 2 shows the reduced set. The original set comprises 170 x-y coordinates, whereas the reduced set contains only 115. We then concatenate and flatten the 115 x-y coordinates into a 230-dimensional vector, which we used as input for our MRI-to-speech synthesis task.

Since the raw MRI data is composed of long utterances with lots of silences, we first segment the utterances into sentence-long pieces. Then we employ a pre-trained BERT-based model<sup>1</sup> to estimate sentence boundaries, and align the audio recordings as well as the orthographic and phonological transcriptions using Montreal force aligner [25]. The resulting alignments are

<sup>1</sup><https://huggingface.co/felflare/bert-restore-punctuation>



Figure 1: One MRI frame during the utterance “apa”.

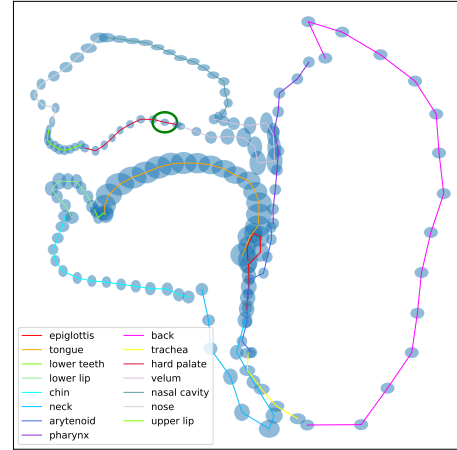


Figure 3: Standard deviation of each MRI feature (Sec. 2).




Figure 2: Extracted MRI features for the utterance “apa”. Lighter is earlier in time. The labeled points are the estimated EMA features (Sec. 2).

manually calibrated by professional phoneticians. By utilizing the estimated sentence boundaries and word alignments, we split the audio recordings and MRI data into 236 utterances, totaling 11 minutes. Finally, these utterances are randomly split into a 0.85-0.05-0.10 train-val-test split, resulting in 200, 11, and 25 utterances in the train, val, and test sets, respectively.

Furthermore, the head location is not fixed within the MRI data, which can negatively impact the ability of such models to generalize to unseen positions. Thus, we adopt a centering approach that centers each frame around a relatively fixed point to improve generalizability. Specifically, we calculate the standard deviation  $(\sigma_x, \sigma_y)$  of each of the 170 points across the training set and center every frame at the point with the lowest  $\sqrt{\sigma_x^2 + \sigma_y^2}$ . This center point is located on the hard palate, circled in green in Figure 3. We note that the standard deviation results reflect human speech production behavior, as the hard palate is relatively still across utterances whereas the tongue varies noticeably, highlighting the interpretability of our MRI-based articulatory features. Another preprocessing step that we found useful was denoising, which we detail in Section 3.4.

### 3. Models

#### 3.1. Intermediate-Representation Baselines

Currently, a popular speech synthesis approach is to first synthesize an intermediate representation from the input and then map the intermediate representation to the waveform domain [26, 27, 14]. Wu *et al.* [13] showed that directly synthesizing speech from EMA outperformed a spectrum-intermediate approach in terms of computational efficiency and yielded comparable synthesis quality. Intuitively, omitting the spectrum intermediate reflects how the human speech production process does not perform this intermediate mapping [13]. In this work, we also compare using intermediate representations with directly synthesizing from inputs. We observe two popular types of intermediate representations in the literature: (1) spectrums [26, 27, 14], and (2) deep representations [4]. To compare our proposed direct modelling approach in Section 3.3 with both intermediate modelling methods, we experiment with Mel spectrogram calculation and HuBERT [28] intermediates. For the Mel spectrogram calculation, we use size-240 hops, size-1024 FFTs, Hann windows, and 80 Mels. With HuBERT, we use the output of the model’s last hidden layer, linearly interpolated to match the MRI input sampling rate. We denote spectrum-intermediate models with “Spe.” and HuBERT ones with “Hub.” in our results below for readability. In our MRI-to-speech task, direct modeling is both more computationally efficient and more high-fidelity than the intermediate approaches, as discussed in Sections 4 and 5. We detail the model architectures of our intermediate-representation baselines in Section 3.2.

#### 3.2. CNN-BiLSTM Baseline

As per Yu *et al.* [14], we employ the CNN-BiLSTM architecture as the baseline method. This method involves processing each MRI frame through a sequence of four CNN layers, with two max-pooling layers incorporated in the middle. The extracted features are then aggregated along the time axis and fed to a BiLSTM layer to generate the mel-spectrogram. Since the inputs in our MRI-to-speech task are sequences of vectors rather than the MRI video inputs used in Yu *et al.* [14], we use 1D convolutions instead of 2D and 3D. Finally, a neural vocoder is used to reconstruct the waveform signal. For this vocoder, we use HiFi-CAR [13], which outperforms the WaveGlow architecture [29] used by Yu *et al.* [14]. HiFi-CAR is an autoregressive

version of the HiFi-GAN convolutional network [30], detailed in Section 3.3. It is worth noting that they used the original speech data, without any denoising, resulting in unsatisfactory performance. For a fair comparison, we also train this model using enhanced speech. In our experiments in Sections 4 and 5 below, we refer to this model as CBL for readability.

### 3.3. HiFi-CAR Model

Similar to the method used by Wu *et al.* [13], our model directly synthesizes waveforms from articulatory features without the need for an intermediate representation. Specifically, we build on their HiFi-CAR model, which is a HiFi-GAN convolutional neural network [30] modified to be autoregressive using the CAR-GAN audio encoder [31]. To our knowledge, training models to directly synthesize waveforms from MRI data has not yielded successful results previously. However, we observe that this model outperforms our baselines in terms of both computational efficiency and fidelity, as discussed in Sections 4 and 5. We also use the HiFi-CAR vocoder to map intermediate features to waveforms for our intermediate-representation baselines. For all HiFi-CAR models, we initialize their weights with those of a HiFi-GAN spectrum-to-waveform vocoder pre-trained on LibriTTS.<sup>2</sup> We note that this initialization approach noticeably improves performance compared to Wu *et al.* [13]. Further modeling details can be found in the accompanying codebase.

### 3.4. Speech Enhancement Model

The currently available dataset [19] suffers from poor quality due to significant reverberation and noise, which poses a significant challenge for accurate modeling of the relationship between MRI and speech. To circumvent this issue, we employed an off-the-shelf Adobe Podcast toolkit<sup>3</sup>, which processes speech recordings to enhance their quality and makes them sound as if they were recorded in a professional studio. Therefore the resulting speech is better suited for our purposes. Unfortunately, due to its proprietary nature, we do not have access to its technical details. Through our observation, however, we conjecture that it may contain a pipeline of bandwidth extension [32] and speech enhancement [33]. Specifically, we hypothesize that the toolkit up-samples the speech to 48kHz and leverages HiFi-GAN [34] to generate high-quality speech. We downsample the enhanced speech to the waveform sampling rate of our MRI dataset to keep model output lengths the same. In our MRI-to-speech task, we use  $0.9 * y_e + 0.1 * y_o$  as our target waveform, where  $y_e$  is the enhanced waveform and  $y_o$  is the original one. Using this weighted sum yields more intelligible MRI-to-speech models than using just  $y_e$ , which may be due to how deep speech enhancers add irregular noise that can be smoothed to more learnable targets by adding the original, more natural waveforms.

## 4. Computational Efficiency

Given the importance of computational efficiency for real-time, on-device speech synthesizers we compare the number of parameters and inference times between our model and the baselines, summarized in Table 1. GPU trials use one RTX A5000 GPU, and CPU trials use none. Like Wu *et al.* [13], we report inference time as the mean and standard deviation of five tri-

<sup>2</sup><https://github.com/kan-bayashi/ParallelWaveGAN>

<sup>3</sup><https://podcast.adobe.com/enhance>

Table 1: Average inference time and number of parameters for MRI-to-speech models. See Section 4 for details.

Model	CPU (s) ↓	GPU (s) ↓	Params. ↓
CBL (Spe.) [14]	.66 ± .05	.081 ± .009	$1.9 * 10^7$
CBL (Hub.)	.69 ± .04	.090 ± .016	$2.3 * 10^7$
HiFi-CAR	<b>.58 ± .03</b>	<b>.061 ± .015</b>	<b><math>1.5 * 10^7</math></b>

Table 2: AB test results. See Section 5.1 for details.

Baseline Type	AB Test Votes		
	Baseline	Ours	Same
CBL (Spe.) [14]	1	<b>53</b>	0
+ Denoising	18	<b>33</b>	3

als, each calculating the average time to synthesize an utterance in our test. Our model is faster and uses less parameters than both intermediate-representation baselines, reinforcing the idea that directly mapping articulatory features to speech is more efficient than relying on an intermediate representations.

## 5. Synthesis Quality

### 5.1. Subjective Fidelity Evaluation

We perform a subjective AB preference test on Amazon Mechanical Turk (MTurk). In this evaluation, each participant is asked to distinguish between the utterances generated by our method and the baselines in terms of naturalness. We compared our model with two baselines: (1) Yu *et al.* [14], detailed in Section 3.2, and (2) [14] trained with our denoised waveforms described in Section 3.4 as targets. For each of the two AB tests, we asked 6 native English listeners to rank a total of 9 random samples from the test set in this study. To prevent listeners from randomly submitting results, we added an audio pair consisting of one audio sample consisting entirely of noise and another audio sample containing high-fidelity speech. In Table 2, we summarize the total number of votes for each of the three options for both AB tests. Our model outperforms both baselines, receiving the most votes. The almost unanimous vote for our model in the AB test with the non-denoised baseline highlights the importance of denoising waveforms accompanying MRI data. Our model also noticeably outperforms the denoised baseline, suggesting that direct synthesis approach described in Section 3.3 is well-suited for articulatory synthesis.

### 5.2. Objective Fidelity Evaluation

We perform an objective evaluation of synthesis quality by analyzing the mel-cepstral distortions (MCD) [35] between ground truths and synthesized samples, as in Wu *et al.* [13]. Table 3 summarizes these results, reporting the mean and standard deviation of the MCDs across utterances. Our HiFi-CAR approach outperforms both intermediate-representation baselines, suggesting that our direct modelling method is suitable for the MRI-to-speech task.

### 5.3. Transcription

We also compare our method with the baselines in terms of speech intelligibility. Specifically, we compute the character error rate (CER) for speech transcription. We use Whisper [36],

Table 3: *MCD between MRI-to-speech model outputs and denoised ground truths. See Section 5.2 for details.*

Model	MCD ↓
CBL (Spe.) [14] + Denoising	7.31 ± 0.45
CBL (Hub.) + Denoising	8.84 ± 1.00
HiFi-CAR	<b>6.64 ± 0.64</b>

Table 4: *ASR CER for MRI-to-speech model outputs. See Section 5.3 for details.*

Model	CER ↓
CBL (Spe.) [14] + Denoising	84.7% ± 36.4%
CBL (Hub.) + Denoising	84.2% ± 15.7%
HiFi-CAR	<b>69.2% ± 28.1%</b>

a state-of-the-art automatic speech recognition (ASR) model, to generate text from the speech synthesized using each method and all test set utterances. Table 4 summarizes these results. Like in Table 3, our model outperforms both baselines, reinforcing the suitability of our model for MRI-to-speech synthesis.

## 6. Comparing MRI and EMA Features

As mentioned in Section 1, MRI provides much more information about the vocal tract than EMA. MRI is a superset of EMA. Specifically, EMA has one x-y coordinate for each of the following locations: upper lip, lower lip, lower incisor, tongue tip, tongue body, and tongue dorsum. Points at all of these locations are present in the MRI data, so we can actually approximate EMA features from MRI by choosing one MRI point at each EMA location, as visualized in Figure 2. In this figure, segments are the connected MRI points and the shaped dots are the estimated EMA locations. We compare these two articulatory feature sets by comparing the outputs of our proposed MRI-to-speech model with those of this model trained to synthesize speech from our estimated 12-dimensional EMA features. The test set predictions of this EMA-to-speech model yielded an MCD of  $6.986 \pm 0.587$  and ASR CER of  $73.2\% \pm 6.7\%$ . Both of these values are worse than those of our MRI-to-speech model, summarized in Tables 3 and 4. This suggests that MRI features are more complete representations of the human vocal tract than EMA features. Thus, articulatory synthesis models should incorporate features beyond EMA in order to achieve human-like fidelity across all utterances, with MRI features being a potential feature set to extend towards. We identify which of the MRI features would be the most valuable to add to the articulatory feature set in Section 7.

## 7. Identifying Important MRI Features

We also study which of the MRI features are the most useful for synthesis in order to provide insight into which features should be present in an ideal articulatory feature set for articulatory synthesis. Specifically, we created 50 subsets of our 230-dimensional MRI feature set, each composed of a random 90% subset of the 230 features. With each feature subset, we masked the 23 MRI features not in the subset to 0.0 and synthesized the test set utterances. Then, we computed the average MCD between the test set ground truths and the synthesized waveforms. For each MRI feature, we assign it a score equal

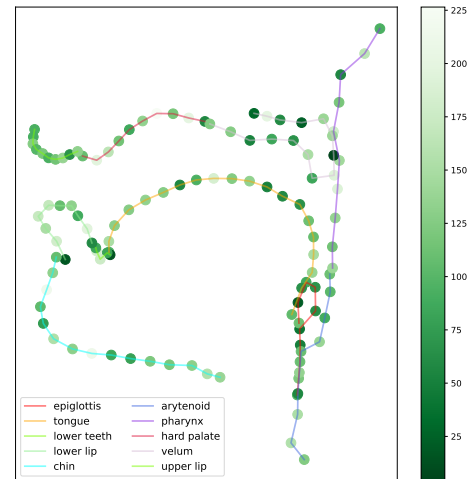


Figure 4: *Importance of each MRI feature for MRI-to-speech synthesis. See Section 6 for details.*

to the average of the MCD values corresponding to experiments where that feature was unmasked. Since our subsets are chosen randomly, we try each feature an equal number of times in expectation. We rank the MRI features by score, with lower rank values being better and corresponding to a lower score and average MCD. Figure 4, with darker green points corresponding to better MRI features. We note that each of the six EMA locations described in Section 6 have an MRI point that is ranked as important. Moreover, for these six locations, besides the upper lip, the number of points ranked as important is fairly sparse. This suggests that six points chosen in the EMA feature set are all very valuable for articulatory synthesis. The important MRI features also correspond well to the phonetic constriction task variables, *e.g.*, those used in [37] to model articulatory synergies from real-time MRI images. Beyond the corresponding EMA locations, points around and in the pharyngeal region (*e.g.*, between tongue root or epiglottis and rear pharyngeal wall) and velic region are also ranked as important. This suggests that these features are also essential for fully-specified, high-fidelity articulatory synthesis. The pharyngeal features are relevant to the production of various speech sounds [38], like /a/ [39] and some variants of /t/ in English [40]. The velic features are crucial to the production of nasal sounds. Both of these features are not available from EMA. Thus, moving forward, we plan to incorporate pharyngeal and velic features in all of our articulatory synthesis models. Points around constriction locations, whether at the lips, tongue, or throat, are generally ranked as important. Thus, when designing sparse articulatory feature sets, it may be useful to prioritize these constriction locations.

## 8. Conclusion and Future Directions

In this work, we devise a new articulatory synthesis method using MRI-based features, providing preprocessing and modelling strategies for working with such data. Based on MCD, ASR, human evaluation, timing, and memory measurements, our model achieves noticeably better fidelity and computational efficiency than the prior intermediate-representation approach. Through speech synthesis ablations, we also show the advantages of MRI over EMA and identify the most important MRI features for articulatory synthesis. Moving forward, we will extend our work to multi-speaker synthesis and inversion tasks [19, 17].

## 9. References

- [1] T. Hayashi, R. Yamamoto, T. Yoshimura, P. Wu, J. Shi, T. Saeki, Y. Ju, Y. Yasuda, S. Takamichi, and S. Watanabe, “Espnet2-tts: Extending the edge of tts research,” *arXiv preprint arXiv:2110.07840*, 2021.
- [2] P. Wu, J. Shi, Y. Zhong, S. Watanabe, and A. W. Black, “Cross-lingual transfer for speech processing using acoustic language similarity,” in *ASRU*, 2021.
- [3] D. Lim, S. Jung, and E. Kim, “Jets: Jointly training fastspeech2 and hifi-gan for end to end text to speech,” 09 2022, pp. 21–25.
- [4] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhota, W.-N. Hsu, A. Mohamed, and E. Dupoux, “Speech Resynthesis from Discrete Disentangled Self-Supervised Representations,” in *Interspeech*, 2021.
- [5] K. Lakhota, E. Kharitonov, W.-N. Hsu, Y. Adi, A. Polyak, B. Bolte, T.-A. Nguyen, J. Copet, A. Baevski, A. Mohamed *et al.*, “On generative spoken language modeling from raw audio,” *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 1336–1354, 2021.
- [6] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [7] K. Deng, S. Watanabe, J. Shi, and S. Arora, “Blockwise streaming transformer for spoken language understanding and simultaneous speech translation,” *Interspeech*, 2022.
- [8] G. K. Anumanchipalli, J. Chartier, and E. F. Chang, “Speech synthesis from neural decoding of spoken sentences,” *Nature*, vol. 568, no. 7753, pp. 493–498, 2019.
- [9] S. L. Metzger, J. R. Liu, D. A. Moses, M. E. Dougherty, M. P. Seaton, K. T. Littlejohn, J. Chartier, G. K. Anumanchipalli, A. Tu-Chan, K. Ganguly *et al.*, “Generalizable spelling using a speech neuroprosthesis in an individual with severe limb and vocal paralysis,” *Nature Communications*, vol. 13, no. 1, p. 6510, 2022.
- [10] G. Fant, “What can basic research contribute to speech synthesis?” *Journal of Phonetics*, vol. 19, no. 1, pp. 75–90, 1991.
- [11] P. Rubin, T. Baer, and P. Mermelstein, “An articulatory synthesizer for perceptual research,” *The Journal of the Acoustical Society of America*, vol. 70, no. 2, pp. 321–328, 1981.
- [12] C. Scully, “Articulatory synthesis,” in *Speech production and speech modelling*. Springer, 1990, pp. 151–186.
- [13] P. Wu, S. Watanabe, L. Goldstein, A. W. Black, and G. K. Anumanchipalli, “Deep speech synthesis from articulatory representations,” in *Interspeech*, 2022.
- [14] Y. Yu, A. H. Shandiz, and L. Tóth, “Reconstructing speech from real-time articulatory mri using neural vocoders,” in *EUSIPCO*, 2021, pp. 945–949.
- [15] G. Beguš, A. Zhou, P. Wu, and G. K. Anumanchipalli, “Articulation gan: Unsupervised modeling of articulatory learning,” *ICASSP*, 2023.
- [16] T. Toda, A. Black, and K. Tokuda, “Acoustic-to-articulatory inversion mapping with gaussian mixture model,” in *ICSLP*, 2004.
- [17] P. Wu, L.-W. Chen, C. J. Cho, S. Watanabe, L. Goldstein, A. W. Black, and G. K. Anumanchipalli, “Speaker-independent acoustic-to-articulatory speech inversion,” in *ICASSP*, 2023.
- [18] C. J. Cho, P. Wu, A. Mohamed, and G. K. Anumanchipalli, “Evidence of vocal tract articulation in self-supervised learning of speech,” *ICASSP*, 2023.
- [19] Y. Lim, A. Toutios, Y. Bliesener, Y. Tian, S. G. Lingala, C. Vaz, T. Sorensen, M. Oh, S. Harper, W. Chen *et al.*, “A multispeaker dataset of raw and reconstructed speech production real-time mri video and 3d volumetric images,” *Scientific data*, vol. 8, no. 1, pp. 1–14, 2021.
- [20] J. Lian, A. W. Black, L. Goldstein, and G. K. Anumanchipalli, “Deep neural convolutive matrix factorization for articulatory representation decomposition,” in *Interspeech*, 2022.
- [21] J. Lian, A. W. Black, Y. Lu, L. Goldstein, S. Watanabe, and G. K. Anumanchipalli, “Articulatory representation learning via joint factor analysis and neural matrix factorization,” in *ICASSP*, 2023.
- [22] S. G. Lingala, B. P. Sutton, M. E. Miquel, and K. S. Nayak, “Recommendations for real-time speech mri,” *Journal of Magnetic Resonance Imaging*, vol. 43, no. 1, pp. 28–44, 2016.
- [23] S. G. Lingala, Y. Zhu, Y.-C. Kim, A. Toutios, S. Narayanan, and K. S. Nayak, “A fast and flexible mri system for the study of dynamic vocal tract shaping,” *Magnetic resonance in medicine*, vol. 77, no. 1, pp. 112–125, 2017.
- [24] E. Bresch and S. Narayanan, “Region segmentation in the frequency domain applied to upper airway real-time magnetic resonance images,” *IEEE transactions on medical imaging*, vol. 28, no. 3, pp. 323–338, 2008.
- [25] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldii,” in *Interspeech*, 2017.
- [26] T. G. Csapó, C. Zainkó, L. Tóth, G. Gosztolya, and A. Markó, “Ultrasound-based articulatory-to-acoustic mapping with waveglow speech synthesis,” in *Interspeech*, 2020.
- [27] M.-A. Georges, P. Badin, J. Diard, L. Girin, J.-L. Schwartz, and T. Hueber, “Towards an articulatory-driven neural vocoder for speech synthesis,” in *International Seminar on Speech Production*, 2020.
- [28] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, 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*, vol. 29, pp. 3451–3460, 2021.
- [29] R. Prenger, R. Valle, and B. Catanzaro, “Waveglow: A flow-based generative network for speech synthesis,” in *ICASSP*, 2019, pp. 3617–3621.
- [30] J. Kong, J. Kim, and J. Bae, “Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis,” in *NeurIPS*, 2020.
- [31] M. Morrison, R. Kumar, K. Kumar, P. Seetharaman, A. Courville, and Y. Bengio, “Chunked autoregressive gan for conditional waveform synthesis,” in *ICLR*, 2022.
- [32] J. Su, Y. Wang, A. Finkelstein, and Z. Jin, “Bandwidth extension is all you need,” in *ICASSP*, 2021, pp. 696–700.
- [33] J. Su, Z. Jin, and A. Finkelstein, “Hifi-gan: High-fidelity denoising and dereverberation based on speech deep features in adversarial networks,” in *Interspeech*, vol. 2020, 2017, pp. 4506–4510.
- [34] ———, “Hifi-gan-2: Studio-quality speech enhancement via generative adversarial networks conditioned on acoustic features,” in *WASPAA*, 2021.
- [35] A. W. Black, “CMU wilderness multilingual speech dataset,” in *ICASSP*, 2019.
- [36] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.04356*, 2022.
- [37] T. Sorensen, A. Toutios, L. Goldstein, and S. S. Narayanan, “Characterizing vocal tract dynamics across speakers using real-time mri,” in *Interspeech*, 2016, pp. 465–469.
- [38] S. R. Moisiuk, J. H. Esling, L. Crevier-Buchman, and P. Halimi, “Putting the larynx in the vowel space: Studying larynx state across vowel quality using mri,” in *ICPhS*, 2019.
- [39] M. I. Proctor, C. Y. Lo, and S. S. Narayanan, “Articulation of english vowels in running speech: A real-time mri study,” in *ICPhS*, 2015.
- [40] A. Alwan, S. Narayanan, and K. Haker, “Toward articulatory-acoustic models for liquid approximants based on mri and epg data. part ii. the rhotics,” *The Journal of the Acoustical Society of America*, vol. 101, no. 2, pp. 1078–1089, 1997.