

Coding Speech Through Vocal Tract Kinematics

Cheol Jun Cho , Peter Wu, Tejas S. Prabhune, Dhruv Agarwal , and Gopala K. Anumanchipalli 

Abstract—Vocal tract articulation is a natural, grounded control space of speech production. The spatiotemporal coordination of articulators combined with the vocal source shapes intelligible speech sounds to enable effective spoken communication. Based on this physiological grounding of speech, we propose a new framework of neural encoding-decoding of speech – Speech Articulatory Coding (SPARC). SPARC comprises an articulatory analysis model that infers articulatory features from speech audio, and an articulatory synthesis model that synthesizes speech audio from articulatory features. The articulatory features are kinematic traces of vocal tract articulators and source features, which are intuitively interpretable and controllable, being the actual physical interface of speech production. An additional speaker identity encoder is jointly trained with the articulatory synthesizer to inform the voice texture of individual speakers. By training on large-scale speech data, we achieve a fully intelligible, high-quality articulatory synthesizer that generalizes to unseen speakers. Furthermore, the speaker embedding is effectively disentangled from articulations, which enables accent-preserving zero-shot voice conversion. To the best of our knowledge, this is the first demonstration of universal, high-performance articulatory inference and synthesis, suggesting the proposed framework as a powerful coding system of speech.

Index Terms—Speech coding, speech synthesis, articulatory synthesis, speech inversion, acoustic-to-articulatory inversion, electromagnetic articulography.

I. INTRODUCTION

HUMANS naturally produce intelligible speech by controlling articulators on the vocal tract. Such vocal tract articulation has long been claimed to be the physiological ground of speech production in various aspects [1]. The source-filter theory of speech describes articulation as shaping the vocal cavity to implement filters on source, or glottal flow, to create speech sounds [2], [3]. Articulatory phonetics and phonology have explained the basis of speech in terms of the coordination of articulators, identifying some canonical articulators that can determine the phonetic properties [1], [4], [5]. In cognitive neuroscience, the speech sensorimotor cortex has been proven to represent continuous, real-time vocal tract articulation while

Received 30 April 2024; revised 20 August 2024 and 17 October 2024; accepted 24 October 2024. Date of publication 20 November 2024; date of current version 3 February 2025. This work was supported in part by the National Science Foundation under Grant 2106928, in part by the Society of Hellman Fellows, Rose Hills Innovator and UC Noyce Initiative Programs at UC Berkeley, and in part by the BAIR Commons-Meta AI Research, Google Research Scholar Award and Robert E. and Beverly A. Brooks Endowed Chair in EECS. The guest editor coordinating the review of this article and approving it for publication was Prof. Minje Kim. (Corresponding authors: Cheol Jun Cho; Gopala K. Anumanchipalli.)

The authors are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: cheoljun@berkeley.edu; peterw1@berkeley.edu; prabhune@berkeley.edu; dhruvagarwal@berkeley.edu; gopala@berkeley.edu).

Digital Object Identifier 10.1109/JSTSP.2024.3497655

naturally speaking, suggesting the vocal tract articulation as a cognitive basis of speech production [6], [7], [8].

Furthermore, the recent findings by Cho et al. [9], [10] suggest that articulatory inversion naturally emerges from self-supervised learning (SSL) of speech. When probed on articulatory kinematics measured by electromagnetic articulography (EMA), the feature representation of the recent speech SSL models (e.g., HuBERT [11]) is highly correlated with EMA, where high-fidelity articulation can be reconstructed by a simple linear mapping from speech SSL features [9]. This suggests that the articulatory inference is a natural solution of SSL of speech for abstracting speech information. This emergent property is further shown to be universal across speakers, dialects, and even languages [10]. Together, these suggest that the biophysical, articulatory representation of speech is a shared coding principle in both biological and artificial intelligence of speech. This convergence raises an interesting question – *Can we represent any arbitrary speech using articulatory features?*

Previous studies have demonstrated that intelligible speech can be synthesized from articulatory features [12], [13], [14], [15], [16], and combined with acoustic-to-articulatory inversion (AAI), resynthesis frameworks have shown the potential of articulatory features as viable intermediate for speech coding systems [16], [17]. However, these previous methods are limited to a fixed set of speakers and the quality is still far behind the commercial speech synthesis models. This absence of a universal, generalizable framework has significantly limited the practical utility of articulatory-based speech coding.

Here, we first demonstrate a high-performance, universal articulatory encoder and decoder that can scale and generalize across an indefinite number of speakers. We leverage the universal articulatory inference by speech SSL [10] to build a generalizable articulatory encoder that transforms speech into vocal tract kinematic traces in a template articulatory space. The template articulatory space is agnostic to individual anatomical differences which are compensated by a separate speaker identity encoder. By training a synthesis model, or a vocoder, on a large-scale dataset, we achieve a universal articulatory vocoder that can generate fully intelligible, high-quality speech from any speaker’s articulation. Furthermore, the speaker identity encoder successfully disentangles speaker identity from articulation, which is demonstrated by a zero-shot, accent-preserving voice conversion. By closing the loop of articulatory encoding and decoding, we propose a novel, speech science guided encoding-decoding framework of speech – Speech Articulatory Coding (SPARC) (Fig. 1).¹ The SPARC framework shows a minimal

¹The code is available at <https://github.com/Berkeley-Speech-Group/Speech-Articulatory-Coding>.

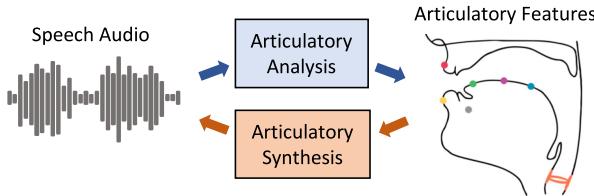


Fig. 1. The Speech Articulatory Coding (SPARC) framework. It encodes speech using articulatory features (analysis) and decodes these features back into speech (synthesis).

loss of intelligibility and quality compared to the original speech audio.²

Compared to existing neural coding of speech [18], [19], [20], [21], [22], representing speech as articulatory features has following benefits:

- *Low-dimensionality*: The articulatory features have only 14 channels with 50 Hz sampling rate.
- *Interpretability*: Each channel corresponds to the actual physical articulator on the vocal tract, which can be intuitively interpretable by visualization.
- *Controllability*: The features can be naturally controlled by the same principle as speech production.
- *Universality*: The articulatory encoding is universal across speakers and disentangled from individual anatomical variance.

With these unique benefits, we demonstrate empirical evidence of the promising potential of SPARC as a valid, novel coding framework for speech.

II. RELATED WORK

A. Electromagnetic Articulography

Electromagnetic articulography (EMA) measures time-varying displacements of vocal tract articulators synchronously while speaking. Typically, sensors are placed on the upper lip (UL), lower lip (LL), lower incisor (LI), tongue tip (TT), tongue blade (TB), and tongue dorsum (TD) (Fig. 2) [23]. A combination of displacements of these articulators on the mid-sagittal plane configures a place of articulation, and combined with source information, or manner of articulation, it shapes a phonetic content of speech. As the traces are continuously collected in real-time, the EMA data naturally reflect phoneme contextualization (coarticulation) and individual tendencies in pronunciations (accents). Given these properties, EMA has been widely accepted for studying articulatory bases of speech, providing biophysical evidence for many linguistic or cognitive theories of speech production [1], [6], [23]. However, EMA has been significantly limited to scale due to the complicated nature and high cost of the collection procedure.

B. Acoustic-to-Articulatory Inversion

To replace the complicated data collection procedure, acoustic-to-articulatory inversion (AAI) models have been

actively developed to predict EMA directly from speech audio [6], [7], [16], [17], [24], [25], [26], [27], [28]. However, the individual variance in vocal tract anatomy across speakers induces inconsistent placements of sensors, which has posed a significant barrier to developing a model that can generalize to unseen speakers [17], [27], [28], [29]. Despite such variability, a canonical basis of articulation is suggested to exist, which is agnostic to individual vocal tract anatomy [1], [10], [17], [30]. In fact, Cho et al. [10] demonstrated that a linear affine transformation can geometrically align one speaker's articulatory system to another's. This suggests that individual articulatory spaces are lying on the same linear space so that an articulatory space of any speaker can be a hypothetical universal template space of articulation. We empirically validate this statement by using a single-speaker AAI model as a universal articulatory encoder for our coding framework.

C. Articulatory Synthesis

Articulatory synthesis aims to generate speech audio from articulatory features. A century of efforts have been made to build articulatory synthesizers for basic research of speech [31], [32], [33], [34], [35], [36], [37], [38], [39]. Several methods have been proposed for improving intelligibility and quality, demonstrating broader use cases including text-to-speech (TTS) [13], prosody manipulation [40], [41], speech denoising [42], and speech brain-computer interfaces (BCIs) [7]. Some of these works utilize deep learning models to map articulatory features to acoustic features, which are then converted to audio using pretrained acoustic synthesizers [7], [15], [43], [44]. A recent study shows that a GAN-based generative model can directly synthesize speech waveform from articulatory features with high intelligibility [14]. However, to our knowledge, none of the existing approaches has achieved industrial-level performance, which requires high intelligibility, quality, and generalizability across unseen speakers.

D. Neural Coding of Speech

Many deep learning methods have been proposed to learn data-driven representations of speech. Various autoencoder-based frameworks have been suggested to jointly train encoders that compress audio into low-bitrate discrete units [20], [21], [45] or decompose speech into different factors [19], [22], [46], and decoders that reconstruct speech from encoded features with minimal loss of information. Also, pretrained speech SSL models have been utilized to extract rich linguistic content of speech, and synthesizers are trained to restore speech audio from those features [18], [19], [47], [48], [49], [50]. These SSL-based methods often utilize separate source modeling (e.g., pitch) and speaker encoding, since SSL model encoders tend to marginalize out acoustic and speaker information [51], [52]. We categorize all these kinds of closed-loop frameworks utilizing neural networks for both encoding and decoding as *neural coding of speech*. Though the existing methods achieve high fidelity in representing speech audio, the intermediate speech codes significantly lack interpretability, and embeddings of those codes are often high-dimensional.

²Audio samples: <https://berkeley-speech-group.github.io/sparc-demo>

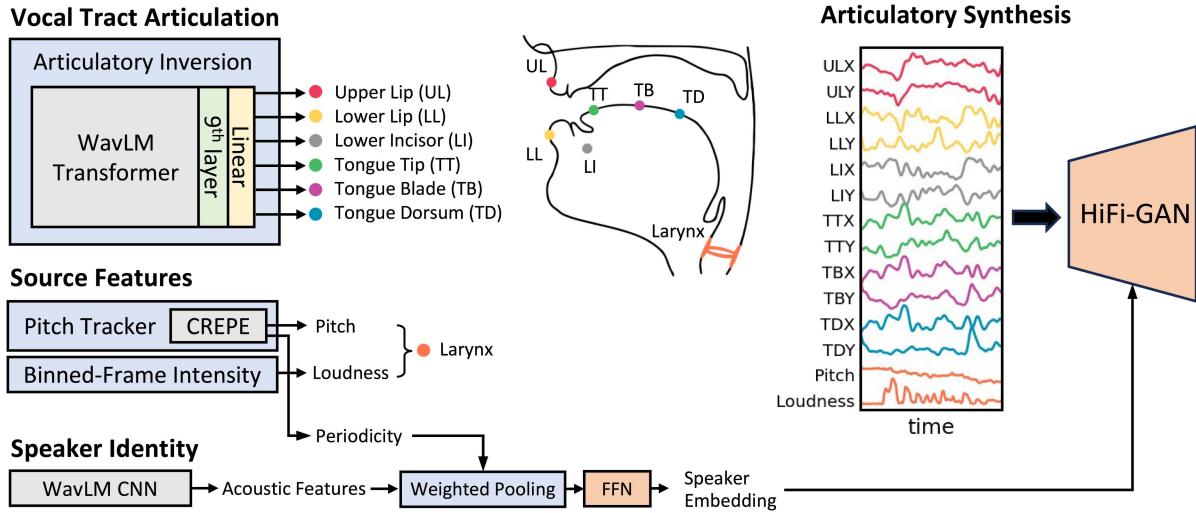


Fig. 2. Pipeline of articulatory analysis and synthesis. The articulatory analysis is composed of vocal tract articulation, source features, and speaker embedding, which are then fed to the synthesizer (HiFi-GAN) in the synthesis pipeline. The modules colored with orange (FFN and HiFi-GAN) are updated while training the synthesis model and other modules are fixed.

Here, we propose articulatory features as interpretable, controllable, and grounded coding of speech that can fully represent any arbitrary speech. A similar concurrent work is proposed towards interpretable speech representations, which utilizes sparse phonetic posteriograms instead of articulatory features [53].

III. METHODS

To bridge the interpretability and controllability gap in current neural speech coding systems, we propose neural articulatory inversion and synthesis as a new type of coding system that can provide an interpretable and controllable coding of speech.

A. Articulatory Analysis

In SPARC, speech is encoded as factors obtained by an analytic framework that infers three different components of speech: vocal tract articulation, source features, and speaker identity (Fig. 2). The first two are based on pre-trained analysis models that provide kinematic traces of the physical articulators, and vocal source features. The last component, speaker identity, is inferred by a model jointly trained with the synthesizer.

1) *Vocal Tract Articulation*: Based on the findings by [10] elaborated in Section II-B, we propose to use a single speaker's EMA as a template articulatory space to represent speaker-generic articulatory kinematics. We selected one of the largest single-speaker EMA datasets, MNGU0 [54], that includes 75 minutes of EMA collected while reading newspapers aloud. This dataset is widely accepted and verified in many studies, given a fine signal quality carefully controlled by the authors. We claim all speakers' articulations can be represented on this single-speaker EMA space without losing information that contributes to the intelligibility of speech. That is, EMA represents phonetic content in a way that can be detached from the variance of vocal tract anatomical structure across individuals.

We use the SSL-linear AAI approach proposed by [9], [10]. The SSL-linear model is built by training a linear mapping from SSL features to EMA, while keeping the SSL encoder weights frozen. This simple mapping can effectively find a linear subspace in the SSL feature space which is highly correlated with EMA, as shown by previous probing studies [9], [10]. We use the WavLM Large model [55], which shows the highest correlation amongst speech SSL models as reported in [9]. Note that the linear head is the only fitted part here, thus, maintaining the generalization capacity of the WavLM encoder that is attained by pretraining on large-scale speech data and adversarial data augmentation [55]. Furthermore, the speaker information tends to diminish after a few early layers [51], which indicates that the mapping can be speaker-agnostic, further contributing to multi-speaker generalizability.

The original 200 Hz EMA data is downsampled to 50 Hz to match the sampling rate of the SSL features, and each channel is z-scored within utterances. The 9th layer of the WavLM Transformer encoder is used to extract speech features for the inversion, where the input audio has 16000 Hz sampling frequency, zero mean, and unit variance. A low-pass filter is applied to the features to remove high-frequency noise using the Butterworth filter with order of 5, where the frequency threshold is set as 10 Hz. The linear inversion model is trained by ordinary least squares. All data in MNGU0 dataset are used for training the main model after selecting the best layer using cross-validation. See Appendix A.1 for the details of the layer selection procedure. The resulting AAI model outputs 12 channels of EMA (X and Y axis of each of 6 articulators) with 50 Hz frequency.

We claim the proposed AAI model can be universal based on two hypotheses. First of all, all individual articulatory spaces lie on the same linear space. Articulatory space is defined as a vector space where each basis corresponds to physical location of a specific articulator. The space is speaker dependent as individual has different vocal tract anatomy. However, an affine

transformation between two spaces exists that spatially aligns different vocal tract structures. As elaborated in Section II-B, empirical evidence is demonstrated by [10] and we also report in Section IV-A. This suggests that despite variable vocal tract anatomies there exists a common basis to which all individual articulations can be registered. All we need is a high quality measurement of any articulatory space, which we claim the MNGU0 is sufficient for this purpose.

The second hypothesis is that the SSL-linear AAI is speaker agnostic. This again requires the following propositions to hold: there exists a subspace in SSL features that is agnostic to voice identity, and a linear mapping used to project the features to the subspace is also speaker independent. This means that even if the SSL-linear AAI is trained only on the MNGU0 speaker's voice, the model should project speech from other speakers to the same articulatory space. We provide empirical evidence to support this hypothesis in a later section (Section IV-D), where we show that the inference remains consistent after converting between different voice identities.

2) *Source Features*: Though EMA has a full descriptive capacity of the place of articulation, it lacks source information generated by the glottal excitation, which is crucial to implementing the manner of articulation and expressing the prosody of speech. Therefore, we include pitch (or fundamental frequency, f0) and loudness features to represent the source features [18], [19], [56]. The loudness feature also informs non-larynx constriction, which is important for voiced fricatives such as "z" and "v". We use CREPE [57] to infer pitch from speech, and loudness is measured by the average of absolute magnitudes of waves for every 20 ms. Together with the EMA from AAI, we referred to these features as "articulatory features" that have 14 channels (12 EMA + 2 source) and a 50 Hz sampling rate.

3) *Speaker Identity*: Since we use a template space for the vocal tract, the articulatory features lack information about the individual structures of the vocal tract anatomy. However, this structural information is an important determinant of the voice texture of an individual speaker, which is crucial in defining the speaker's identity [58]. For example, the vocal tract length is known to be correlated with gender and age in voice. Note that our definition of the speaker identity does not include information about dialect or accent which is actually aimed to be disentangled from the speaker identity. Here, we compensate for this missing information with a separate speaker identity encoder which is jointly trained with the vocoder to extract the speaker-specific texture information. To this end, we propose a simple yet effective speaker encoder, which minimizes the trainable portion of the model. Based on the observation by [51], [59] that the speaker information is largely concentrated in the CNN outputs of speech SSL models, the encoder consists of the frozen CNN extractor from WavLM Large followed by a weighted pooling layer and a learnable feedforward network (FFN). The pooling layer weighted-averages the acoustic features from WavLM CNN across frames, where the weight is given by the periodicity inferred from CREPE. This allows more attention to the periodic signals which may encode more information about voice texture than non-periodic portion of the input. Then, the FFN transforms the averaged features to a speaker embedding with 64 channels.

TABLE I
COMPARISON OF MULTI-SPEAKER AAI SYSTEMS

System	Model	Dataset	
		MOCHA	HPRC
Tract Variable	Wu et al. [17]	0.678	0.784
	Siriwardena & Espy-Wilson [28]	—	0.757
Single-speaker	Ours	0.739 ± 0.017	0.761 ± 0.012

Performance is measured as pearson correlation coefficient (PCC).

This speaker identity encoding is indispensable to fully represent multi-speaker speech.

B. Articulatory Synthesis

We adopt HiFi-GAN as the vocoder for articulatory synthesis [14], [47], [60]. The vocoder is trained to synthesize speech audio with a 16 K Hz sampling rate from the articulatory features. To condition on the speaker embedding, we apply FiLM [61] to each convolution module in the HiFi-GAN architecture, which modulates the output channels of each module. We adopt the same loss functions as [60]: mel spectrogram loss for reconstruction and multi-period and multi-scale discriminator loss for GAN training.

C. Dataset

For training the vocoder and speaker encoder, we use LibriTTS-R [62], an enhanced version of LibriTTS [63]. The dataset is comprised of 585 hours of reading audiobooks (555 hours for training). The original 24 K Hz audio is downsampled to 16 K Hz. We use VCTK [64] to further evaluate the generalizability of the model to a broader range of speakers and accents. The entire VCTK dataset is unseen during training and only used for evaluation.

Note that the FFN in the speaker identity encoding and the HiFi-GAN vocoder are the only trainable modules (orange modules in Fig. 2), and the rest of the pipeline remains fixed while training. More details of implementation and training are in Appendix B.1–7.

IV. EXPERIMENTS & RESULTS

We evaluate the proposed SPARC in different aspects to validate each part of the framework – universality of articulatory encoding (IV-A), information conservation in encoding-decoding (IV-B), multilingual generalization (IV-C), and disentanglement of speaker identity encoding (IV-D).

A. Universality of Single-Speaker Articulatory Encoding

As we claim a single-speaker AAI can be used for universal articulatory encoding, we compare our model with the existing state-of-the-art (SOTA) multi-speaker articulatory encoding systems that utilize tract variables [17], [28]. The tract variables are a set of articulatory parameters that can be derived from EMA and are known to be more regularized across speakers [65]. Table I denotes the performance of the tract-variable-based AAI systems and ours on two multi-speaker EMA datasets, MOCHA [66] and HPRC [67]. MOCHA includes 7 speakers

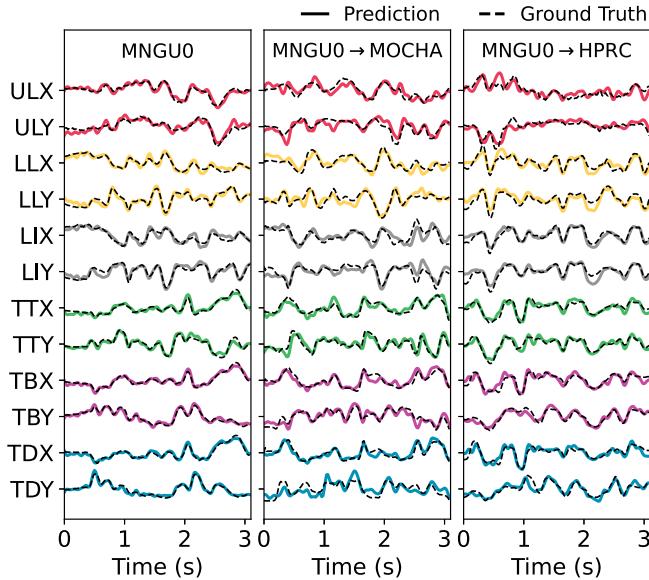


Fig. 3. Examples of SSL-linear prediction on MNGU0 (left), transformed prediction from MNGU0 to a female MOCHA speaker (middle) and to a male HPRC speaker (right). Predictions are denoted with the colored lines and ground truths are denoted with the black dotted lines.

with 27 minutes of data per speaker on average, and HPRC includes 8 speakers with 59 minutes of data per speaker on average. The performance is evaluated by the Pearson correlation coefficient (PCC), where results besides ours are retrieved from reference papers. (95% confidence interval is denoted for our case.) Since our AAI is trained on a single-speaker (MNGU0) articulatory space, the model outputs are not directly comparable to the EMA from MOCHA-TIMIT or HPRC. Therefore, we fit a linear model to spatially align MNGU0's articulatory space to another speaker's articulatory space to measure the correlation [10]. (See Appendix A.2. for details.) As a result, our single-speaker AAI system shows correlations higher in MOCHA and slightly lower in HPRC compared to Wu et al. [17], and slightly higher in HPRC compared to Siriwardena and Espy-Wilson [28]. Note that this is not a fair comparison since the scores are measured on different variables. However, this system-wise comparison suggests that the single-speaker approach can yield a similar level of consistency across multiple speakers. Moreover, our approach only fits a linear model, thus it is likely to be more generalizable than the previous systems which use non-linear, deep recurrent neural networks. Fig. 3 demonstrates prediction examples of our AAI model. The left panel shows a near-perfect prediction even with a simple linear mapping from WavLM, and even after the affine transformation from MNGU0 space to other speakers, the predictions show high correlations with the ground truths (middle and right panel). The prediction performance on MNGU0 shows 0.878 ± 0.012 average correlation, and additional analyses on AAI are reported in Appendix A.3-5.

B. Performance of Resynthesis by SPARC

We measure the intelligibility and quality of resynthesized speech audio to evaluate how well information is preserved in the

TABLE II
PERFORMANCE COMPARISON OF GROUND TRUTH SPEECH (GT) AND RESYNTHESIZED SPEECH BY SPARC (RESYNTH)

Dataset	Input	Intelligibility		Quality	
		WER↓	CER↓	MOS↑	UTMOS↑
LibriTTS-R	GT	4.22	2.24	4.01 ± 0.16	4.10 ± 0.58
	Resynth	5.43	2.90	3.82 ± 0.20	4.12 ± 0.58
VCTK	GT	1.66	0.74	3.65 ± 0.16	3.84 ± 0.54
	Resynth	3.73	1.96	3.67 ± 0.17	3.95 ± 0.52

encoding-decoding process. We use an out-of-the-box automatic speech recognition (ASR) model, Whisper ("openai/whisper-large-v3") [68], to evaluate the word error rate (WER) and character error rate (CER) of resynthesized speech.³ The quality of the audio is measured by a human-evaluated subjective metric, mean opinion score (MOS), and a machine-evaluated MOS, UTMOS [69]. We evaluate both ground truth speech and resynthesized speech on the test-clean subset of LibriTTS-R which includes 8.56 hours from 39 unseen speakers. To further evaluate the generalizability of the model, we evaluate the model on VCTK dataset that includes 107 speakers with 500 utterances in total. Table II summarizes the performance in comparison with the ground truth speech. For the quality metrics, we reported 95% confidence intervals. When tested on LibriTTS-R, the ASR on the resynthesized speech shows high intelligibility with WER of 5.43% and CER of 2.90%, which is marginally different from those of the ground truth. Moreover, our articulatory synthesis generates natural speech sounds showing decent quality with MOS of 3.82 which is a 0.19 decrease from the ground truth, and with UTMOS of 4.12 which is on par with the ground truth.

Furthermore, the results on VCTK also demonstrate high intelligibility with WER of 3.73% and CER of 1.96%. Though the gap from the ground truth is larger than LibriTTS-R, this is a remarkable level of performance given that the model is not exposed to the entire dataset during training. As LibriTTS-R has cleaner audio than VCTK, the model demonstrates some speech enhancement capacity which results in a higher MOS and UTMOS than the ground truth. While the speakers in LibriTTS-R are concentrated on American English, VCTK includes a variety of accents, primarily focusing on English speakers with different regional accents from the United Kingdom. Therefore, the high performance demonstrated on VCTK suggests that SPARC can robustly generalize to unseen speakers and accents.

C. Multilingual Generalization

We evaluate SPARC on speech corpora from other languages: 7 European languages (German, Dutch, Portuguese, Italian, Polish, Spanish, French) from multilingual LibriSpeech [70] and 3 East Asian languages (Korean, Japanese, and Chinese (Mandarin)) from KSS [71], JVS [72], and AISHELL [73], respectively. Except for KSS, all corpora include multiple speakers. We randomly sampled 200 utterances from the test split of each corpus to evaluate the resynthesized speech.⁴ Table III compares the intelligibility (WER and CER using Whisper) and

³The transcription is normalized before measuring the error rates.

⁴We excluded samples with numerals to avoid converting Arabic numerals.

TABLE III
PERFORMANCE COMPARISON ON MULTILINGUAL SPEECH DATA

Language	WER↓			CER↓			UTMOS↑		
	GT	EN-R	ML-R	GT	EN-R	ML-R	GT	EN-R	ML-R
German	5.23	15.32	9.70	1.89	6.66	3.95	2.96 ± 0.83	2.87 ± 0.58	2.83 ± 0.72
Dutch	9.90	20.20	15.39	2.82	8.83	5.70	2.91 ± 0.82	2.89 ± 0.65	2.77 ± 0.79
Portuguese	7.10	17.81	12.80	2.19	8.50	5.21	2.63 ± 1.34	2.81 ± 0.67	2.56 ± 0.92
Italian	9.63	22.09	15.53	2.43	7.40	4.98	2.78 ± 1.26	2.82 ± 0.61	2.67 ± 0.82
Polish	4.19	14.79	9.89	0.94	4.63	2.87	3.09 ± 0.76	3.00 ± 0.63	2.95 ± 0.76
Spanish	3.96	10.74	6.79	1.26	4.73	2.38	2.93 ± 0.86	2.88 ± 0.66	2.72 ± 0.81
French	4.62	16.67	8.04	2.65	9.90	4.54	2.79 ± 0.93	2.80 ± 0.65	2.69 ± 0.79
Korean	6.31	15.88	13.66	1.01	6.72	4.96	3.69 ± 0.63	3.47 ± 0.80	2.91 ± 0.95
Japanese	—	—	—	5.65	8.45	8.16	3.53 ± 0.80	3.49 ± 0.77	3.08 ± 0.83
Chinese	—	—	—	10.69	31.27	23.28	2.66 ± 1.21	3.06 ± 0.92	2.71 ± 1.12
Average	6.37	16.69	11.48	3.15	9.71	6.60	3.00 ± 0.66	3.01 ± 0.49	2.79 ± 0.29
w/o Chinese	6.37	16.69	11.48	2.32	7.31	4.75	—	—	—

Each metric is evaluated on ground truth speech (GT) and resynthesized speech by english-only trained model (EN-R) or multilingual model (ML-R).

quality (UTMOS). As no interword spacing is used for Japanese and Chinese orthographic systems, words in these languages are not separated as straightforwardly as other languages. Therefore, we omit the WER for these languages. We also evaluate the performance of a multilingual version of the model that is obtained by fine-tuning to the multilingual data (denoted as ML-R in Table III). The original English-only trained model is denoted as EN-R. See Appendix B.8. for details of the fine-tuning procedure.

As shown in Table III, the resynthesis by English-only model (EN-R) achieves average WER of 16.69 excluding Japanese and Chinese, average CER of 9.71. Roughly, if the conflated Chinese CER is excluded, the EN-R preserves a fair amount of information given that 93% of characters are correctly recognized. Despite the huge gap from the ground truth, this is remarkable generalizability given both encoding and decoding procedures have only seen English in training. When fine-tuned (ML-R), the average WER and CER are cut down to 11.48% and 6.60%, respectively. This indicates that some portion of the errors is induced by the out-of-domain application of the vocoder, thus that fine-tuning is able to yield a huge gain in performance. Yet, there are unresolved gaps from the ground truth, which would be attributed to the fact that we are only tuning the vocoder and keeping the encoder part intact. This is also aligned with the observation that there is a slight language bias in the articulatory representation in SSL [10]. In terms of synthesis quality, we find a minimal difference between ground truth (3.00 ± 0.66) and resynthesized audio from English-only trained model (3.01 ± 0.49), and the quality slightly drops to 2.79 ± 0.29 when the model is fine-tuned to multilingual dataset. (The scores are averaged across languages and the ranges denote 95% confidence intervals.) The drop in quality is likely due to the general inferiority in the quality of multilingual datasets compared to LibriTTS-R.

D. Speaker Recognition and Voice Control by Speaker Identity

The speaker identity encoding informs voice texture of speech, which is especially important as our articulatory features are speaker agnostic. On one hand, this suggests that the model can learn a speaker embedding that is disentangled from individual tendencies in articulation, or accents. Here, we evaluate

TABLE IV
FEW-SHOT SPEAKER IDENTIFICATION ACCURACY (SID ACC)

Model	SID ACC (%)↑	
	LibriTTS-R	VCTK
x-vector	91.7	83.4 (88.2)
r-vector	98.3	99.8 (99.8)
Ours	95.5	92.2 (94.6)

For VCTK, the accuracy of full-train data template embeddings is denoted in parentheses.

this claim by experiments on few-shot speaker identification and zero-shot voice conversion.

First, we build a few-shot, learning-free speaker identification (SID) by comparing the similarity between the speaker embeddings. For each test speaker, 10 first clips in the dataset are concatenated and then the speaker embedding is extracted to serve as a template embedding for the speaker. Then among potential test speakers, the identity is predicted by choosing the speaker with maximal similarity between the template and tested speaker embedding. For LibriTTS-R, we use speakers in the test-clean set with at least 10 clips, leaving 36 speakers. For VCTK, we use the train set clips to create the template embeddings for 107 test speakers. Additionally, we also evaluate SID accuracy using all data in the train set for creating the templates, which is denoted in parentheses in Table IV.

Our speaker encoding shows high discriminability in the few-shot SID, achieving 94.7% accuracy in LibriTTS-R and 92.2% accuracy in VCTK test sets. When all utterances in the train set are used, the accuracy for VCTK reaches 94.6%. Our embedding outperforms the widely adopted speaker embedding, x-vector [74] by a large margin, especially in VCTK which is a harder case with a larger number of speakers (Table IV).⁵ As a sanity check for the proposed SID task, we evaluate the accuracy with the current state-of-the-art speaker embedding, ResNet based r-vector by WeSpeaker [75], which shows a near-perfect SID accuracy.⁶ However, the r-vector embedding is extracted from very deep ResNet with 221 layers, which is way larger than ours and x-vector, therefore the model may incorporate more speaker information than voice texture (e.g., dialect and accent).

⁵We retrieved checkpoints from <https://huggingface.co/speechbrain/spkrec-xvect-voxceleb>.

⁶We use “resnet221_LM” checkpoint from <https://github.com/wenet-e2e/wespeaker>.

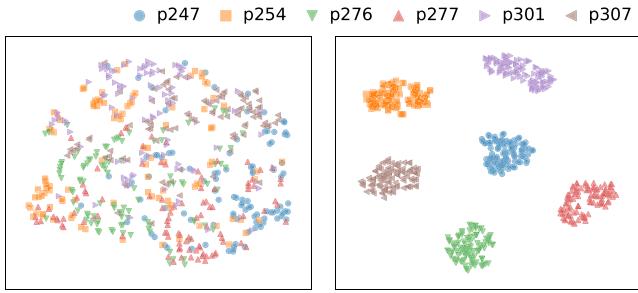


Fig. 4. Visualization by T-SNE of utterance-wise averaged articulations (left) and speaker embeddings (right) from 6 different VCTK speakers. The perplexity is set as 10.

Both x-vector and r-vector models are trained on VoxCeleb corpus [76], [77], [78]. Overall, this suggests that our speaker encoding can extract highly discriminable features for speaker identification.

We further evaluate the validity of our identity encoding by zero-shot voice cloning application. We selected 6 speakers from VCTK, where their mean pitch values are equidistant in the range from 80 Hz to 230 Hz. Fig. 4 visualizes the speaker embedding of the selected speakers, demonstrating that each speaker’s utterances are clustered distinctively. On the contrary, no speaker-relevant cluster is identifiable in the manifold of utterance-wise averaged articulations. For each target speaker, the first 10 clips in the training set are concatenated and then the speaker embedding is extracted. Then, the voice from the source speaker is converted by switching speaker embedding in the synthesis. Additionally, the pitch range of the source speaker is adjusted to match that of the target speaker, by z-scoring the source pitch trace; and shifting and rescaling to match the center and scale of the target pitch range (“P-rescale”). We convert all audio clips in test-clean set of LibriTTS-R with more than 2 seconds of duration, to each of 6 VCTK target speakers.

We evaluate “coding-recoding similarity” of the voice-converted speech. The coding-recoding similarity measures the correlation between the speech coding of original speech and that of synthesized speech. Note that our coding pipeline is also an analytic measurement that provides articulatory and source features of speech. Therefore, the metric indicates how much the articulation and source information is preserved in the voice conversion and speech synthesis process. For articulation, we measure the Pearson correlation of each of 12 channels and then average them, and for pitch and loudness, the correlation is separately reported. For speaker identity, the cosine similarity is measured between the conditioning target embedding and the embedding extracted from synthesized speech. Unlike articulatory and source features, speaker embedding is arbitrary and not interpretable. Therefore, we evaluate the accuracy in SID task, the same discriminability task designed above for VCTK. Since the task requires to pinpoint the target speaker from 107 speakers, the SID accuracy can inform how closely the voice-converted speech matches the target speaker while remaining distinct from others. Lastly, we measure WER and UTMOS to evaluate the intelligibility and quality of the converted audio.

As baselines, we evaluate existing voice conversion models: FreeVC [79] and QuickVC [50]. Both models encode speech contents using speech SSL models and have separate speaker encoders jointly trained with synthesizers. The former is more comparable to ours as the content encoder is also based on WavLM. However, none of the baseline models impose any grounded nor interpretable structures on the speech codes. We apply our framework to analyze articulatory and source features of the speech synthesized by baselines, while the correlations of speaker embeddings and SID scores are evaluated using their own speaker encoders, where we also denote the SID accuracies using our speaker encoder in parentheses (Table V). Furthermore, we evaluate the coding-recoding similarity of the self-targeted resynthesized speech (“Ours-Resynth” in Table V), to demonstrate consistency of the speech content in synthesis without voice conversion, which represents the upper bound of the scores. To avoid confusion, we do not report SID for resynthesis since the speaker identity is not imposed to be VCTK speakers. The results are summarized in Table V with 95% confidence intervals if applicable.

As a result, the voice-converted speech by our model can be accurately identified as the target speaker among 107 VCTK speakers, achieving 94.6% accuracy (Table V). Though the similarity in speaker embedding is lower in voice conversion than in the self-targeted resynthesis (0.879 and 0.962, respectively), the SID result suggests that this level of similarity is high enough to be discriminable from 107 unseen speakers. Furthermore, our approach shows higher SID accuracy than both of the baseline models. FreeVC shows inferior scores in both speaker similarity and SID scores. QuickVC shows higher speaker similarity than ours but shows significantly lower SID accuracy as 65% (66.7% using our speaker encoding). The latter case indicates that their speaker embedding is relatively simple to be more consistent through voice conversion, but lacks specificity to discriminate different speakers.

The SID accuracy and the coding-recoding similarity of pitch and speaker identity significantly drop when we ablate the pitch rescaling (“w/o P-rescale” in Table V), indicating an interplay between pitch range and speaker identity encoding. This may be induced by the natural correlation between speaker identity and pitch range, which is likely to be confounded by biological sex. However, the coding-recoding similarity of articulation and loudness remains intact, showing a marginal difference from the pitch-rescaled voice conversion.

The high coding-recoding similarity in articulation suggests that the articulation is well disentangled from speaker identity and pitch, which indicates that accent is preserved after the voice conversion. The correlation is as high as 0.944 and is marginally different from the self-targeted resynthesis, 0.954, which also indicates the stability of the coding by SPARC. This result further supports the claim that the proposed AAI is agnostic to speaker identity, yielding highly consistent inferences even with different voice identities. Compared to the baseline models, SPARC demonstrates superior similarities, especially in pitch, suggesting our model better disentangles pitch from the other features. Regarding the intelligibility, WER slightly increases after the voice conversion, but the difference is marginal and the

TABLE V
PERFORMANCE COMPARISON OF ZERO-SHOT VOICE CONVERSION

Model	Target Spk	Coding-Recoding Similarity (PCC)				SID↑	WER↓	UTMOS↑
		Articulation↑	Pitch↑	Loudness↑	Speaker↑			
FreeVC	VCTK	0.923 ± 0.021	0.881 ± 0.165	0.881 ± 0.062	0.839 ± 0.109	92.5 (81.4)	4.00	4.05 ± 0.51
QuickVC	VCTK	0.930 ± 0.021	0.868 ± 0.172	0.903 ± 0.063	0.971 ± 0.105	65.0 (66.7)	4.17	4.22 ± 0.36
Ours w/o P-rescale	VCTK	0.944 ± 0.015	0.944 ± 0.093	0.938 ± 0.037	0.879 ± 0.069	94.6	4.83	3.83 ± 0.58
Ours-Resynth	Self	0.954 ± 0.014	0.968 ± 0.109	0.936 ± 0.077	0.962 ± 0.038	—	4.40	4.00 ± 0.59

The SID accuracy using our speaker encoding is denoted in parenthesis.

speech remains highly intelligible with WER of 4.83%. Both baselines show lower WERs than our model (FreeVC: 4.00%, and QuickVC: 4.17%) indicating a potential internal process that might improve intelligibility by modifying the articulatory features of the original speech. Finally, the voice conversion shows some degraded UTMOS as 3.83, 0.17 decrease compared to the resynthesis, and QuickVC shows the highest UTMOS as 4.22. This suggests that our model has a potential trade-off between disentanglement and audio quality of the synthesized audio. However, this degrades in naturalness may be confounded by a natural correlation between accents and voice textures, so that some voice conversion cases may result in rare combinations of accent and speaker identity, harming the UTMOS score. To sum up, these results indicate that SPARC is a disentangled and stable coding system of speech.

V. DEMONSTRATION OF INTERPRETABILITY AND CONTROLLABILITY OF SPARC

The proposed articulatory coding is highly interpretable and controllable, which is very unique compared to previous neural coding approaches. Interpretability is naturally granted as each feature represents a physical articulator on the vocal tract. Furthermore, the articulatory features can be controlled with the same principles that govern speech production as vocal tract articulation is the control space of speech production. The control or modification of articulatory features then can be played back by the vocoder, which functions as a physical simulation of vocal tract articulation [4], [12], [16], [41], [80]. Here, we demonstrate these features with two example cases.

A. Case 1: Place of Articulation – “lock” vs “rock”

First of all, we demonstrate how SPARC can be used to interpret and control the place of articulation in speech. Here, we select two speech clips of speaking “lock (l-p-k)” and “rock (r-p-k)” as an example pair to illustrate the difference between alveolar and post-alveolar approximants. The articulatory traces of these two words extracted by SPARC are depicted in Fig. 5, along with snapshots of the vocal tract animation. These offer an intuitive interpretation of how the vocal tract is dynamically shaped while speaking. Furthermore, this successfully highlights the known phonological distinction between these two approximants: “l” approximates tongue to a more anterior part of the palate (alveolar) than “r” (post-alveolar).

Furthermore, we demonstrate a control simulation by interpolating tongue articulations between the words [14]. The mixing factor, α , is applied to weigh the sum of each of the tongue articulatory traces (TT, TB, and TD), i.e., $Art("lock") \cdot \alpha + Art("rock") \cdot (1 - \alpha)$, where $Art(\cdot)$ means articulatory encoding. The Fig. 5 bottom panel shows the interpolated traces with $\alpha \in \{1.0, 0.8, 0.6, 0.4, 0.2, 0.0, -0.2\}$, gradually changing pronunciation from “lock” to “rock”. Then, ASR is applied to synthesized speech to determine which word is perceived, drawing a perceptual boundary at $\alpha = 0.2$, which is denoted using different line styles in Fig. 5. This demonstration shows that SPARC provides interpretable control knobs, which can be manipulated to simulate speech sounds and observe the causal interaction between articulatory control and the generated speech.

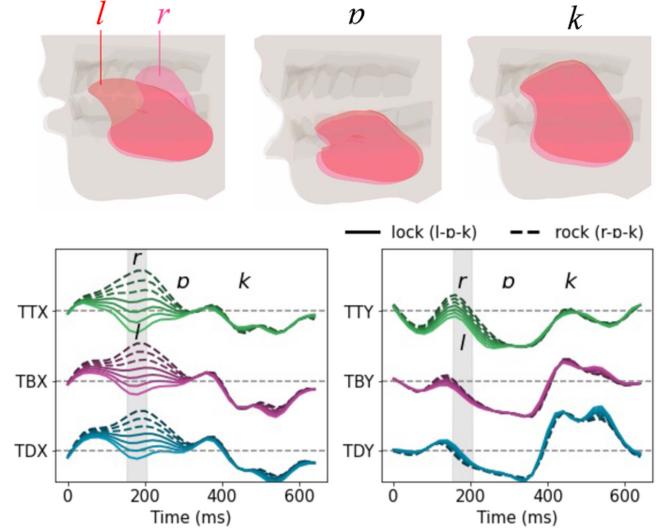


Fig. 5. Articulatory traces encoded for speaking “lock” and “rock”, denoted by the different line styles, “—” and “- -”, respectively. The bottom panel shows the midsagittal displacements of TT, TB, and TD, and the top panels show snapshots of the corresponding vocal tract anatomy. In the snapshots, the vocal tract of “lock” and “rock” are overlaid with separate colors, orange and pink, respectively. The shaded region indicates the window of “l” or “r”. The color is darkened while interpolating from “lock” to “rock”, where the line style again indicates the recognized words.

$\alpha + Art("rock") \cdot (1 - \alpha)$, where $Art(\cdot)$ means articulatory encoding. The Fig. 5 bottom panel shows the interpolated traces with $\alpha \in \{1.0, 0.8, 0.6, 0.4, 0.2, 0.0, -0.2\}$, gradually changing pronunciation from “lock” to “rock”. Then, ASR is applied to synthesized speech to determine which word is perceived, drawing a perceptual boundary at $\alpha = 0.2$, which is denoted using different line styles in Fig. 5. This demonstration shows that SPARC provides interpretable control knobs, which can be manipulated to simulate speech sounds and observe the causal interaction between articulatory control and the generated speech.

B. Case 2: Manner of Articulation – Voice Onset Time

The next example is about the manner of articulation, especially focused on the phonological phenomena called voice onset time (VOT). While the spatial arrangement of articulators forms filters for consonants, the manner of sound further modifies sounds into perceptually distinguishable consonants. In this example, we demonstrate that the VOT can be directly implemented by the timing of the rise of loudness.

Here, a speech clip of the word “bay” is encoded by the articulatory encoder (Fig. 6 left). We only show the traces of lips on the Y axis as these are the key articulatory factors of labial

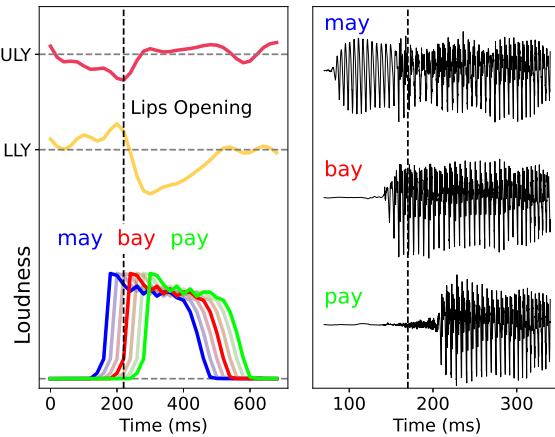


Fig. 6. Articulatory modulation example for “bay”. The left panel shows the Y axis of UL and LL with the vertical dashed line indicating the beginning of the lips opening. The loudness trajectories are depicted, which are moved back and forth, while the salient red, green, and blue colors indicate the perceived plosives. The right panel shows zoom-in wave form of synthesized audio around the lips opening, showing the different voice onset times.

consonants. We can clearly see the lips opening around 200 ms. Then, we shift the loudness trace back and forth and generate speech from the manipulated features. By shifting the loudness trace 60 ms earlier than the moment of lips opening, we can nasalize the sound from “bay” to “may” (from plosive to nasal). On the other hand, by shifting it 60 ms backward, the sound is converted to “pay” (from voiced to voiceless plosive). We can also observe the induced VOT difference along this manipulation (Fig. 6 right), which is aligned with the known VOT patterns by the nasality and voicedness of sounds.

These two proof-of-concept examples demonstrate promising utilities of SPARC as a simulator of articulatory control, and as a speech analysis tool that provides a natural and intuitive phonological interpretation and intervention.

C. Potential Applications

Based on the results, we envision several potential applications of SPARC. First, SPARC can be used as an analysis platform for investigating a phonological basis of speech without collecting high-cost articulatory data [16]. Second, the high-performance synthesizer can be used as a speech simulator, which can facilitate a control theoretic approach to speech processing or reinforcement learning of speaking agents [4], [41], [81], [82]. Third, the universal articulatory analysis can be utilized for a language learning tool or therapy, where visualizing the vocal tract can enhance the learning experience [83], [84], [85]. Lastly, a TTS system can be built upon SPARC that synthesizes articulatory control from text or higher-order structure of speech [86], [87], [88], revealing a descriptive and interpretable relationship between text and articulation. In conclusion, SPARC holds promise for various applications in both speech science and engineering.

VI. CONCLUSION

We propose a novel encoding-decoding framework of speech, SPARC, which mimics the biophysical apparatus of speech

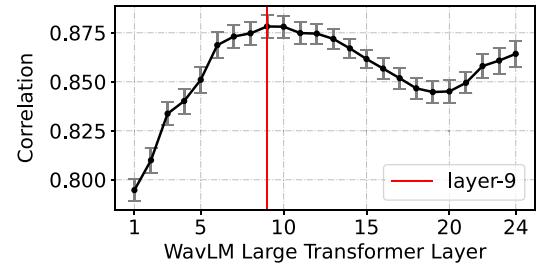


Fig. 7. Probing performance of each layer in WavLM Large on MNGU0.

production. Through large-scale training, our proposed SPARC framework achieves a high-performance articulatory inversion and synthesis that can generalize to unseen speakers, with a minimal loss of information compared to ground truth speech. To our knowledge, this is the first demonstration of universal articulatory synthesis that can scale up to an indefinite number of speakers. This universality is supported by a novel speaker identity encoding that embeds highly discriminable speaker-specific voice textures. Additionally, the encoded speaker embedding is effectively disentangled from articulatory features and allows accent-preserving voice conversion. Using SPARC, speech can be represented in a low-dimensional space, with each channel corresponding to an articulator in the vocal tract. This physical embodiment of speech allows for natural and intuitive control and interpretation.

In future work, we will scale the SPARC framework to incorporate expressive speech and singing. Also, we will improve the robustness of the system under noisy environments. These efforts will maximize the promising utilities of the articulatory coding of speech.

APPENDIX A ADDITIONAL ANALYSIS

A.1. Layer-wise AAI Performance

We selected the 9-th layer after comparing the inversion performance of each layer measured by Pearson correlation coefficients (PCC) averaged across EMA channels on 5-fold cross-validation (Fig. 7). For each fold, 100 test utterances are randomly held out and 95% confidence intervals are denoted. At the best layer (the 9th), the average correlation is 0.878 ± 0.012 . After the layer selection, all data from MNGU0 are used to fit the linear model. The parameters are estimated by ordinary least square (OLS) method. We find a minimal difference between with and without a regularization term (e.g., L2 penalty). Thus, we use OLS for a better simplicity.

A.2. Measuring Correlation Between Articulations from Different Articulatory Space.

Since the EMA sensors are not spatially aligned across participants, each EMA channel from different participants is not directly comparable. As our model is using a particular participant’s (MNGU0’s) vocal tract as a template space, we find a mapping that spatially aligns MNGU0 sensors to different

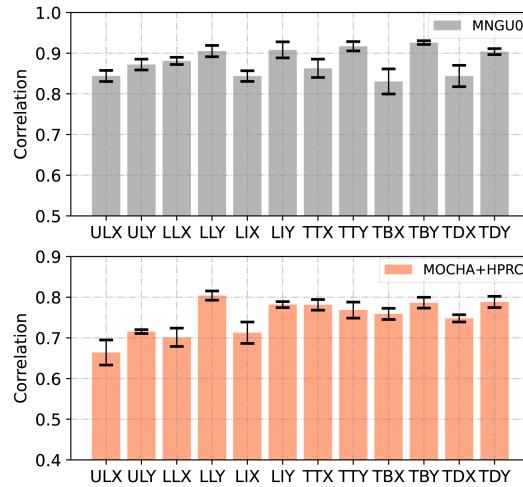


Fig. 8. Correlation of individual channels on MNGU (top) and other speakers in MOCHA and HPRC (bottom).

participants' to calculate the correlation. Following [10], we pass speech audios from a target speaker from MOCHA-TIMIT or HPRC to our inversion model, and then a linear model is trained to predict the reference EMA of the target speaker. We used a linear regression with L1 penalty (Lasso) to impose sparsity in the weights, using $\alpha = 0.01$. Scikit-learn library is used for training the model (“`sklearn.linear_model.Lasso`”) [89]. This L1 penalty reveals channel-by-channel correspondence in the affine transformation (Fig. 9; more discussed in the later section). We split the data to 5 folds where 4 folds are used for training the linear model and 1 fold is used for evaluating the model. This is done individually for each speaker in MOCHA-TIMIT (7 speakers with 27 m of speech in average) and HPRC (7 speakers with 59 m of speech in average).

A.3. AAI Performance on Individual Articulator

The performances of AAI on the individual channels are denoted in Fig. 8, measured by correlation (PCC) with 95% confidence intervals. The top panel shows performance on MNGU0 and the bottom panel shows performance after the transformations to other speakers (15 speakers from MOCHA+HPRC). In both cases, the Y-axis traces are generally better predicted than the X-axis traces.

A.4. Evidence of Geometric Isometry Between Individual Articulatory Space

As shown in [10], the weights of the linear transformations across speakers are highly concentrated within the articulators (Fig. 9; red diagonal boxes). The tongue tip is also affected by the tongue blade and dorsum, which is natural since the tongue is connected and sensor positions are not firm, but still significantly more affected by the tongue tip.

A.5. Evidence for Excluding Coronal Axes of EMA

In our articulatory encoding, we largely ignore the coronal axes of EMA data. Here, we provide evidence for such

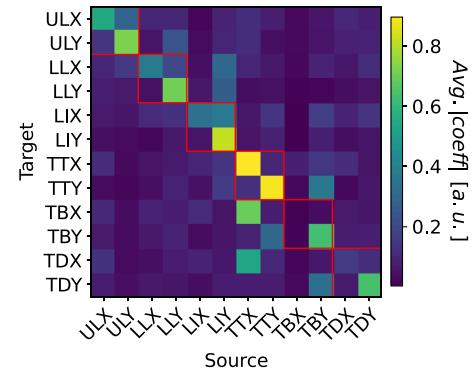


Fig. 9. Average absolute coefficients of affine transformation from MNGU0 to other speakers.

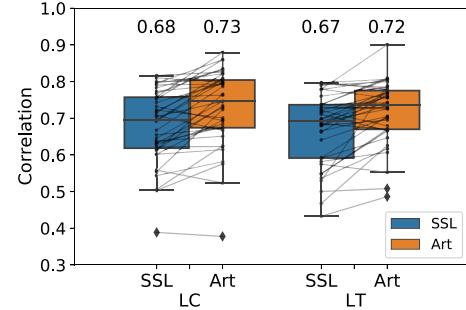


Fig. 10. Linear probing on the coronal axis of the left corner of the mouth (LC) and the lateral tongue (LT). The prediction performance by using WavLM features (SSL) and 12 other midsagittal articulatory channels (Art). Each dot is a speaker.

exclusion with a probing analysis. The same linear probing is conducted using SSL features (WavLM) and the original 12 midsagittal articulatory channels (Art). We use EMA-MAE [90] for this experiment since the dataset includes two coronal channels: the left corner of the mouth (LC) and the lateral tongue (LT). As shown in Fig. 10, both for LC and LT, the linear probing using SSL features is significantly worse than using articulatory features (for LC, SSL features: 0.68 ± 0.19 , articulatory features: 0.73 ± 0.20 ; for LT, SSL features: 0.67 ± 0.19 , articulatory features: 0.72 ± 0.17). This suggests that those coronal features can be predicted with a fair correlation (0.72-0.73) by a linear combination of the existing 12 channels. These results suggest that the coronal axes can be safely ruled out for the proposed coding framework.

APPENDIX B MODEL & TRAINING DETAIL

B.1. Pitch Tracker Configuration

For CREPE, we use a PyTorch implementation called “`torchcrepe`”.⁷ We use “full” model and the range of pitch is set as from 50 to 550. The default Viterbi decoding is used. The input waveform is resampled to 16 K Hz and z-scored within the utterance. We found that if we set the hop length as 320 to output 50 Hz (the sampling rate of SSL features), sometimes, it

⁷Retrieved from <https://github.com/maxrmorrison/torchcrepe>.

makes errors in some female speakers with high pitch levels. To mitigate that we decrease the hop length to 80 which results in 200 Hz outputs, and then, we downsample the output pitch and periodicity by a factor of 4 to match the 50 Hz sampling rate. As a side effect, we observe that the periodicity is sometimes erroneously detected with such hop length modification. Therefore, we threshold the periodicity bigger than 0.4. The periodicity thresholding is only applied for the inference time.

B.2. Loudness Implementation

The loudness is estimated as the average of the absolute magnitude of the input waveform for every 20 ms bin. This is implemented with a fixed, single-channel convolutional layer with a stride of 320 and kernel size of 320, where each weight in the kernel is set as 1/320. Likewise, the waveform is resampled to 16 K Hz and z-scored within utterance.

B.3. Speaker Identity Encoder Architecture

The CNN outputs of WavLM Large are weighted by periodicity inferred from CREPE and averaged across frames per utterance. Here, we also include the convolutional positional encoding and the projection layer applied for the CNN outputs, in the original WavLM model. That is, we use the features right before entering the Transformer encoder, which has 1024 channels. The FFN is composed of a linear layer that maps 1024 input features to 1024 output features, followed by GELU and dropout with 0.2, and then another linear layer that projects features to 64-channel speaker embedding. This FFN is the only network trained within the vocoder training and other parts are not updated.

B.4. Generator Architecture

Following the HiFi-GAN architecture, our generator is a convolutional neural network composed of transposed convolutions each followed by a multi-receptive field fusion (MRF) module [60]. Each MRF module outputs the sum of three residual block outputs [91]. For our three residual blocks, we use dilations of 1, 3, and 5 and kernel sizes of 3, 7, and 11, respectively. For our transposed convolutions, we use kernel sizes of 10, 8, 4, and 4, and each stride is half the respective kernel size. Channels are halved each layer until the final layer outputs a single channel corresponding to the acoustics.

To condition the generation on speaker identity, a FiLM layer is applied for outputs of each residual convolutional layer in the MRF module [61]. Every FiLM layer is implemented as a linear layer followed by ReLU and dropout of 0.2, and another linear layer to predict scale and center for each channel of the targeted convolution outputs. These scales and centers are then multiplied and added to the outputs of convolutional layer that the FiLM layer is attached to (feature-wise affine transform).

B.5. Discriminator Architecture

We utilize two types of discriminators as in the HiFi-GAN model: 1) a multi-period discriminator (MPD), and 2) a multi-scale discriminator (MSD) [60]. The MPD takes as input evenly-spaced input frames, and the MSD average-pools the input. Then, both models feed the processed inputs into convolutional neural networks (CNNs). We follow the same CNN architectures as the HiFi-GAN discriminators [60]. Each MPD CNN is composed of strided convolutional layers each followed by leaky rectified linear unit (ReLU) activation functions. Similarly, each MSD CNN is composed of strided and grouped convolutional layers each followed by leaky ReLUs. Like HiFi-GAN, we use five MPDs with spacings of 2, 3, 5, 7, and 11, and three MSDs that rescale inputs by 1, 2, and 4 times.

B.6. Loss Configuration

Like HiFi-GAN, our loss function is a weighted sum of the GAN loss, mel-spectrogram loss, and feature matching loss [60]. The weight for three loss terms are 1, 45, and 2, respectively. For the mel-spectrogram loss, we use the following parameters: {fs: 16000, fft_size: 1024, hop_size: 160, win_length: null, window: “hann”, num_mels: 80, fmin: 0, fmax: 8000}.

B.7. Training Detail

We use Adam optimizer with learning rate of 10^{-4} and beta of (0.5, 0.9). The model is updated for 1.5 M iterations, and the learning rate is halved by every 8 K steps which stays static after 320 K steps. For every iteration, a random 320 ms window is sampled from each clip in a batch with a size of 64.

B.8. Multilingual Fine-Tuning Detail

We fine-tune the model on the multilingual datasets for 500 K iterations. The model is initialized with a checkpoint of the model trained only on English for 1 M iterations. The training data include 555 hours of English, 1965 hours of German, 1553 hours of Dutch, 161 hours of Portuguese, 247 hours of Italian, 104 hours of Polish, 917 hours of Spanish, 1075 hours of French, 60 hours of Chinese, 22 hours of Japanese, and 11 hours of Korean. The English proportion is the train split of LibriTTS-R.

ACKNOWLEDGMENT

The authors acknowledge Dr. Abdelrahman Mohamed (Rembrand) and Prof. Alan W Black (Carnegie Mellon University) for insightful discussions at various stages of this work.

REFERENCES

- [1] C. P. Browman and L. Goldstein, “Articulatory phonology: An overview,” *Phonetica*, vol. 49, no. 3–4, pp. 155–180, 1992.
- [2] T. Chiba and M. Kajiyama, “The vowel, its nature and structure,” *Phonetic Soc. Jpn.*, 1958.

[3] G. Fant, *Acoustic Theory of Speech Production: With Calculations Based on X-Ray Studies of Russian Articulations*. Berlin, Germany: Walter de Gruyter, 1971, no. 2.

[4] S. Maeda, “Compensatory articulation during speech: Evidence from the analysis and synthesis of vocal-tract shapes using an articulatory model,” in *Speech Production and Speech Modelling*. New York, NY, USA: Springer, 1990, pp. 131–149.

[5] I. P. Association, *Handbook of the International Phonetic Association: A Guide to the Use of the International Phonetic Alphabet*. Cambridge, U.K.: Cambridge Univ. Press, 1999.

[6] J. Chartier, G. K. Anumanchipalli, K. Johnson, and E. F. Chang, “Encoding of articulatory kinematic trajectories in human speech sensorimotor cortex,” *Neuron*, vol. 98, no. 5, pp. 1042–1054, 2018.

[7] G. K. Anumanchipalli, J. Chartier, and E. F. Chang, “Speech synthesis from neural decoding of spoken sentences,” *Nat.*, vol. 568, no. 7753, pp. 493–498, 2019.

[8] C. J. Cho, E. Chang, and G. Anumanchipalli, “Neural latent aligner: Cross-trial alignment for learning representations of complex, naturalistic neural data,” in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 5661–5676.

[9] C. J. Cho, P. Wu, A. Mohamed, and G. K. Anumanchipalli, “Evidence of vocal tract articulation in self-supervised learning of speech,” in *Proc. ICASSP 2023–2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[10] C. J. Cho, A. Mohamed, A. W. Black, and G. K. Anumanchipalli, “Self-supervised models of speech infer universal articulatory kinematics,” in *Proc. ICASSP 2024–2024 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2024, pp. 12061–12065.

[11] 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 Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3451–3460, 2021.

[12] P. Birkholz, “Modeling consonant-vowel coarticulation for articulatory speech synthesis,” *PLoS One*, vol. 8, no. 4, 2013, Art. no. e60603.

[13] P. K. Krug, S. Stone, and P. Birkholz, “Intelligibility and naturalness of articulatory synthesis with vocaltractlab compared to established speech synthesis technologies,” *Proc. SSW*, 2021, vol. 11, pp. 102–107.

[14] P. Wu, S. Watanabe, L. Goldstein, A. W. Black, and G. K. Anumanchipalli, “Deep speech synthesis from articulatory representations,” *Interspeech*, pp. 779–783, 2022.

[15] M. Kim, Z. Piao, J. Lee, and H.-G. Kang, “Style modeling for multi-speaker articulation-to-speech,” in *Proc. ICASSP 2023–2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[16] Y. Gao, P. Birkholz, and Y. Li, “Articulatory copy synthesis based on the speech synthesizer VocalTractLab and convolutional recurrent neural networks,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 32, pp. 1845–1858, 2024.

[17] P. Wu et al., “Speaker-independent acoustic-to-articulatory speech inversion,” in *Proc. ICASSP 2023–2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[18] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, “Neural analysis and synthesis: Reconstructing speech from self-supervised representations,” *Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 16251–16265, 2021.

[19] H.-S. Choi, J. Yang, J. Lee, and H. Kim, “Nansy++: Unified voice synthesis with neural analysis and synthesis,” in *Proc. 11th Int. Conf. Learn. Representations*, 2023.

[20] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, “Soundstream: An end-to-end neural audio codec,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 495–507, 2021.

[21] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, “High fidelity neural audio compression,” *Trans. Mach. Learn. Res.*, 2024.

[22] Z. Ju et al., “Naturalspeech 3: Zero-shot speech synthesis with factorized codec and diffusion models,” in *Proc. 41st Int. Conf. Mach. Learn.*, 2024.

[23] T. Rebernik, J. Jacobi, R. Jonkers, A. Noiray, and M. Wieling, “A review of data collection practices using electromagnetic articulography,” *Lab. Phonol.*, vol. 12, no. 1, 2021, Art. no. 6.

[24] P. K. Ghosh and S. Narayanan, “A generalized smoothness criterion for acoustic-to-articulatory inversion,” *J. Acoustical Soc. America*, vol. 128, no. 4, pp. 2162–2172, 2010.

[25] P. K. Ghosh and S. S. Narayanan, “A subject-independent acoustic-to-articulatory inversion,” in *Proc. 2011 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2011, pp. 4624–4627.

[26] P. Liu, Q. Yu, Z. Wu, S. Kang, H. Meng, and L. Cai, “A deep recurrent approach for acoustic-to-articulatory inversion,” in *Proc. 2015 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2015, pp. 4450–4454.

[27] A. A. Attia, Y. M. Siriwardena, and C. Espy-Wilson, “Improving speech inversion through self-supervised embeddings and enhanced tract variables,” in *Proc. 32nd Eur. Signal Process. Conf.*, 2024, pp. 306–310.

[28] Y. M. Siriwardena and C. Espy-Wilson, “The secret source: Incorporating source features to improve acoustic-to-articulatory speech inversion,” in *Proc. ICASSP 2023–2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[29] M. G. Rahim, C. C. Goodyear, W. B. Kleijn, J. Schroeter, and M. M. Sondhi, “On the use of neural networks in articulatory speech synthesis,” *J. Acoustical Soc. America*, vol. 93, no. 2, pp. 1109–1121, 1993.

[30] N. Seneviratne, G. Sivaraman, and C. Y. Espy-Wilson, “Multi-corpus acoustic-to-articulatory speech inversion,” in *Proc. Interspeech*, 2019, pp. 859–863.

[31] H. Dudley, R. R. Riesz, and S. S. Watkins, “A synthetic speaker,” *J. Franklin Inst.*, vol. 227, no. 6, pp. 739–764, 1939.

[32] H. Dudley, “Remaking speech,” *J. Acoustical Soc. America*, vol. 11, no. 2, pp. 169–177, 1939.

[33] H. K. Dunn, “The calculation of vowel resonances, and an electrical vocal tract,” *J. Acoustical Soc. America*, vol. 22, no. 6, pp. 740–753, 1950.

[34] K. N. Stevens, S. Kasowski, and C. G. M. Fant, “An electrical analog of the vocal tract,” *J. Acoustical Soc. America*, vol. 25, no. 4, pp. 734–742, 1953.

[35] G. Rosen, “Dynamic analog speech synthesizer,” *J. Acoustical Soc. America*, vol. 30, no. 3, pp. 201–209, 1958.

[36] P. Mermelstein, “Articulatory model for the study of speech production,” *J. Acoustical Soc. America*, vol. 53, no. 4, pp. 1070–1082, 1973.

[37] P. Rubin, T. Baer, and P. Mermelstein, “An articulatory synthesizer for perceptual research,” *J. Acoustical Soc. America*, vol. 70, no. 2, pp. 321–328, 1981.

[38] S. Maeda, “A digital simulation method of the vocal-tract system,” *Speech Commun.*, vol. 1, no. 3–4, pp. 199–229, 1982.

[39] C. Scully, “Articulatory synthesis,” in *Speech Production and Speech Modelling*. Dordrecht, Netherlands: Springer, 1990, pp. 151–186.

[40] S. Aryal and R. Gutierrez-Osuna, “Accent conversion through cross-speaker articulatory synthesis,” in *Proc. 2014 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2014, pp. 7694–7698.

[41] P. Birkholz, L. Martin, Y. Xu, S. Scherbaum, and C. Neuschaefer-Rube, “Manipulation of the prosodic features of vocal tract length, nasality and articulatory precision using articulatory synthesis,” *Comput. Speech Lang.*, vol. 41, pp. 116–127, 2017.

[42] M.-A. Georges, L. Girin, J.-L. Schwartz, and T. Hueber, “Learning robust speech representation with an articulatory-regularized variational autoencoder,” *Interspeech*, pp. 3345–3349, 2021.

[43] S. Aryal and R. Gutierrez-Osuna, “Data driven articulatory synthesis with deep neural networks,” *Comput. Speech Lang.*, vol. 36, pp. 260–273, 2016.

[44] 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 *Proc. ISSP 2020–12th Int. Seminar Speech Prod.*, 2020.

[45] A. Van Den Oord et al., “Neural discrete representation learning,” in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 6309–6318.

[46] X. Zhang, D. Zhang, S. Li, Y. Zhou, and X. Qiu, “SpeechTokenizer: Unified speech tokenizer for speech large language models,” in *Proc. 25th Int. Conf. Learn. Representations*, 2024.

[47] A. Polyak et al., “Speech resynthesis from discrete disentangled self-supervised representations,” in *Proc. Interspeech*, 2021, pp. 3615–3619.

[48] S.-H. Lee, S.-B. Kim, J.-H. Lee, E. Song, M.-J. Hwang, and S.-W. Lee, “HierSpeech: Bridging the gap between text and speech by hierarchical variational inference using self-supervised representations for speech synthesis,” *Adv. Neural Inf. Process. Syst.*, vol. 35, pp. 16624–16636, 2022.

[49] S.-H. Lee, H.-Y. Choi, S.-B. Kim, and S.-W. Lee, “Hierspeech: Bridging the gap between semantic and acoustic representation of speech by hierarchical variational inference for zero-shot speech synthesis,” 2023, *arXiv:2311.12454*.

[50] H. Guo, C. Liu, C. T. Ishi, and H. Ishiguro, “Quickvc: Any-to-many voice conversion using inverse short-time fourier transform for faster conversion,” in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2023, pp. 1–7.

[51] S. Chen et al., “Why does self-supervised learning for speech recognition benefit speaker recognition?,” *Interspeech*, pp. 3699–3703, 2022.

[52] A. Pasad, B. Shi, and K. Livescu, “Comparative layer-wise analysis of self-supervised speech models,” in *Proc. ICASSP 2023–2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[53] M. Morrison, C. Churchwell, N. Pruyne, and B. Pardo, “Fine-grained and interpretable neural speech editing,” *Interspeech*, pp. 187–191, 2024.

[54] K. Richmond, P. Hoole, and S. King, "Announcing the electromagnetic articulography (day 1) subset of the mngu0 articulatory corpus," *Interspeech*, pp. 1505–1508, 2011.

[55] S. Chen et al., "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE J. Sel. Topics Signal Process.*, vol. 16, no. 6, pp. 1505–1518, Oct. 2022.

[56] Y. Yang et al., "StreamVC: Real-time low-latency voice conversion," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2024, pp. 11016–11020.

[57] J. W. Kim, J. Salamon, P. Li, and J. P. Bello, "Crepe: A convolutional representation for pitch estimation," in *Proc. 2018 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 161–165.

[58] R. Netzorg, B. Yu, A. Guzman, P. Wu, L. McNulty, and G. K. Anumanchipalli, "Towards an interpretable representation of speaker identity via perceptual voice qualities," in *Proc. ICASSP 2024-2024 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2024, pp. 12391–12395.

[59] Z. Fan, M. Li, S. Zhou, and B. Xu, "Exploring wav2vec 2.0 on speaker verification and language identification," *Interspeech*, 2021, pp. 1509–1513.

[60] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 17022–17033, 2020.

[61] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "Film: Visual reasoning with a general conditioning layer," in *Proc. AAAI Conf. Artif. Intell.*, 2018, vol. 32, no. 1, pp. 3942–3951.

[62] Y. Koizumi et al., "Libritts-r: A restored multi-speaker text-to-speech corpus," *Interspeech*, pp. 5496–5500, 2023.

[63] H. Zen et al., "Libritts: A corpus derived from librispeech for text-to-speech," *Interspeech*, pp. 1526–1530, 2019.

[64] C. Vieux et al., "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit," *CSTR*, 2017.

[65] R. S. McGowan, "Recovering articulatory movement from formant frequency trajectories using task dynamics and a genetic algorithm: Preliminary model tests," *Speech Commun.*, vol. 14, no. 1, pp. 19–48, 1994.

[66] A. Wrench, "Mocha: Multichannel articulatory database," 1999. [Online]. Available: <http://www.cstr.ed.ac.uk/research/project/artic/mocha.html>

[67] M. Tiede, C. Y. Espy-Wilson, D. Goldenberg, V. Mitra, H. Nam, and G. Sivaraman, "Quantifying kinematic aspects of reduction in a contrasting rate production task," *J. Acoustical Soc. America*, vol. 141, no. 5, pp. 3580–3580, 2017.

[68] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 28492–28518.

[69] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, "UTMOS: UTokyo-SaruLab system for voicemos challenge 2022," *Interspeech*, pp. 4522–4525, 2022.

[70] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "Mls: A large-scale multilingual dataset for speech research," *Interspeech*, pp. 2757–2761, 2020.

[71] K. Park, "KSS dataset: Korean single speaker speech dataset," Kaggle, 2018. [Online]. Available: <https://www.kaggle.com/datasets/bryankpark/korean-single-speaker-speech-dataset>

[72] S. Takamichi, K. Mitsui, Y. Saito, T. Koriyama, N. Tanji, and H. Saruwatari, "JVS corpus: Free Japanese multi-speaker voice corpus," 2019, *arXiv:1908.06248*.

[73] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "Aishell-3: A multi-speaker Mandarin TTS corpus and the baselines," *Interspeech*, pp. 2756–2760, 2021.

[74] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. 2018 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 5329–5333.

[75] H. Wang et al., "Wespeaker: A research and production oriented speaker embedding learning toolkit," in *Proc. ICASSP 2023-2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[76] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," *Interspeech*, pp. 2616–2620, 2017.

[77] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," *Interspeech*, pp. 1086–1090, 2018.

[78] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "VoxCeleb: Large-scale speaker verification in the wild," *Comput. Speech Lang.*, vol. 60, 2020, Art. no. 101027.

[79] J. Li, W. Tu, and L. Xiao, "FreeVC: Towards high-quality text-free one-shot voice conversion," in *Proc. ICASSP 2023-2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[80] B. H. Story, "A parametric model of the vocal tract area function for vowel and consonant simulation," *J. Acoustical Soc. America*, vol. 117, no. 5, pp. 3231–3254, 2005.

[81] J.-F. Patri, J. Diard, and P. Perrier, "Optimal speech motor control and token-to-token variability: A Bayesian modeling approach," *Biol. Cybern.*, vol. 109, pp. 611–626, 2015.

[82] T. Rakotomalala, P. Perrier, and P. Baraduc, "Trajectories predicted by optimal speech motor control using LSTM networks," in *Proc. Inter-speech 23rd Annu. Conf. Int. Speech Commun. Assoc. (ISCA)*, 2022, pp. 630–634.

[83] P. Badin, Y. Tarabalka, F. Elisei, and G. Bailly, "Can you 'read' tongue movements? evaluation of the contribution of tongue display to speech understanding," *Speech Commun.*, vol. 52, no. 6, pp. 493–503, 2010.

[84] A. Suemitsu, J. Dang, T. Ito, and M. Tiede, "A real-time articulatory visual feedback approach with target presentation for second language pronunciation learning," *J. Acoustical Soc. America*, vol. 138, no. 4, pp. EL382–EL387, 2015.

[85] B. Bernhardt, B. Gick, P. Bacsfalvi, and J. Ashdown, "Speech habilitation of hard of hearing adolescents using electropalatography and ultrasound as evaluated by trained listeners," *Clin. Linguistics Phonetics*, vol. 17, no. 3, pp. 199–216, 2003.

[86] J. Lian, A. W. Black, L. Goldstein, and G. K. Anumanchipalli, "Deep neural convolutive matrix factorization for articulatory representation decomposition," *Interspeech*, pp. 4686–4690, 2022.

[87] 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 *Proc. ICASSP 2023-2023 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[88] C. J. Cho, A. Mohamed, S.-W. Li, A. W. Black, and G. K. Anumanchipalli, "SD-HuBERT: Sentence-level self-distillation induces syllabic organization in hubert," in *Proc. ICASSP 2024-2024 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2024, pp. 12076–12080.

[89] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.

[90] A. Ji, J. J. Berry, and M. T. Johnson, "The electromagnetic articulography Mandarin accented English (EMA-MAE) corpus of acoustic and 3D articulatory kinematic data," in *Proc. 2014 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2014, pp. 7719–7723.

[91] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.



Cheol Jun Cho received the B.S. degree in computer science and engineering from Seoul National University, Seoul, South Korea. He is currently working toward the Ph.D. degree in computer science from the University of California, Berkeley, Berkeley, CA, USA, under the supervision of Dr. Gopala K. Anumanchipalli and Dr. Jack L. Gallant. His research interests include to uncover the principles of representation underlying intelligence with a focus on the speech and language domain. He also develops foundational techniques for spoken language modeling, brain-computer interfaces to restore speech for individuals with severe paralysis, and data-driven neuroscience using naturalistic paradigms.



Peter Wu received the bachelor's and Master's degree in computer science from Carnegie Mellon University, Pittsburgh, PA, USA. He is currently working toward the Ph.D. degree from the University of California, Berkeley, Berkeley, CA, USA with the Department of Electrical Engineering and Computer Sciences. His research interests include machine learning and signal processing.



Tejas S. Prabhune is a third-year undergraduate from the University of California, Berkeley, Berkeley, CA, USA. He is working toward the B.S. degree in electrical engineering and computer science. Since 2023, he is a Student Researcher with the Berkeley Speech Group, and has been a Researcher with the Quantitative Imaging and Artificial Intelligence Lab with Stanford and the Edge Computing Lab in Harvard. He is also the co-founder and CEO of Caravan, developing a dynamically-allocated cloud marketplace for GPUs. His research interests include speech-driven avatars, speech neuroprosthesis, interpretable neural speech processing, and convex optimization.



Gopala K. Anumanchipalli is the Robert E. and Beverly A. Brooks Assistant Professor in Electrical Engineering and Computer Sciences with the University of California, Berkeley, Berkeley, CA, USA, where he leads the Berkeley Speech Group. He is a Member of the Berkeley AI Research, Weill Neurosciences and Computational Precision Health. His research focuses on the science and engineering of spoken language with application in AI and healthcare. He has been recognized as a Kavli Fellow, Google Research Scholar, JP Morgan AI Research Awardee, Hellman Fellow, Noyce Innovator, among other honors.



Dhruv Agarwal is working toward the dual bachelor's degrees in electrical engineering and computer science, and business administration from the University of California, Berkeley, Berkeley, CA, USA. As a Member of the Berkeley Speech Group, he is engaged in research on speech processing, specifically analyzing self-supervised learning representations for their articulatory representations and evaluating speech generation models. He is interested in autonomous vehicles, where he leads the perception component of Berkeley's Formula Electric autonomous team. His professional experience includes an internship with the Anti-Fraud Team at SoFi Technologies and a Machine Learning internship with Typeless.