

[b] = [d] - [t] + [p]: Self-supervised Speech Models Discover Phonological Vector Arithmetic

Kwanghee Choi¹, Eunjung Yeo¹, Cheol Jun Cho², David Harwath¹, David R. Mortensen³
¹UT Austin, ²UC Berkeley, ³CMU

Abstract

Self-supervised speech models (S3Ms) are known to encode rich phonetic information, yet how this information is structured remains underexplored. We conduct a comprehensive study across 96 languages to analyze the underlying structure of S3M representations, with particular attention to phonological vectors. We first show that there exist linear *directions* within the model’s representation space that correspond to phonological features. We further demonstrate that the *scale* of these phonological vectors correlate to the degree of acoustic realization of their corresponding phonological features in a continuous manner. For example, the difference between [d] and [t] yields a voicing vector: adding this vector to [p] produces [b], while scaling it results in a continuum of voicing. Together, these findings indicate that S3Ms encode speech using phonologically interpretable and compositional vectors, demonstrating phonological vector arithmetic. All code and interactive demos are available at <https://github.com/juice500ml/phonetic-arithmetic>.

1 Introduction

The same year that word2vec (Mikolov et al., 2013a), a self-supervised method for distributional dense word representations, was introduced, Mikolov et al. (2013b) showed that these representations encode linguistically meaningful relations through vector arithmetic (fig. 1), *i.e.*,

$$\mathbf{v}_{\text{king}} - \mathbf{v}_{\text{man}} + \mathbf{v}_{\text{woman}} \simeq \mathbf{v}_{\text{queen}}, \quad (1)$$

providing a ready explanation for how word2vec represents semantics.¹ In this paper, we ask: Do self-supervised speech models (S3Ms) represent phonology in an analogous compositional manner?

¹Many have argued that the analogy test used in Mikolov et al. (2013b), which differs from ours, can be fragile in practice; see § A.1 for details.

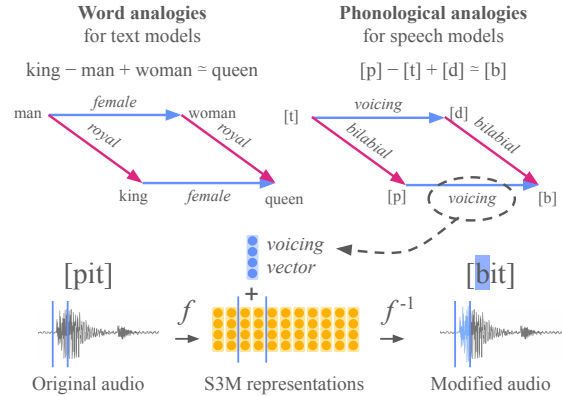


Figure 1: Comparing analogies for text and speech. Word representations enable semantic analogies², while speech representations enable phonological analogies. Such analogies (§ 3) can be used to control speech synthesis in a phonologically grounded manner (§ 4).

S3Ms, a class of neural networks trained on large quantities of unlabeled speech, have demonstrated strong performance across various downstream tasks, including speech recognition, synthesis, and spoken language understanding (Baevski et al., 2020; Hsu et al., 2021; Chen et al., 2022). Consequently, many studies have sought to understand what properties of S3M representations support this performance.

Early analyses primarily investigated **what** information S3M representations encode (Pasad et al., 2021). Empirical studies have shown that S3Ms organize speech according to relative distances that reflect acoustic similarity (Choi and Yeo, 2022), and that these representations form clusters corresponding to phonetic units (Wells et al., 2022; Choi et al., 2025). However, **how** this information is structured is still underexplored.

Building on prior observations and linguistic intuition, we aim to refine the current understanding of S3Ms. In § 3, we propose a hypothesis: *Phono-*

²Word analogies borrowed from Ethayarajh et al. (2019).

logical features are represented linearly within S3M representations, allowing phonological analogies to emerge. For example, consider the phone quadruplet [b], [p], [d], and [t]. [b] and [p] form a voiced-voiceless bilabial plosive pair, and [d] and [t] form a voiced-voiceless alveolar plosive pair.³ This yields two symmetric phonological analogies:

$$[b] : [p] = [d] : [t] \quad (\text{voicing}) \quad (2)$$

$$[b] : [d] = [p] : [t] \quad (\text{POA}) \quad (3)$$

These analogies yield the following approximations in each corresponding vector \mathbf{r} extracted from S3Ms (details in § 3.1.4):

$$\mathbf{r}_{[b]} \simeq \mathbf{r}_{[p]} + (\mathbf{r}_{[d]} - \mathbf{r}_{[t]}) \quad (\text{voicing}) \quad (4)$$

$$= \mathbf{r}_{[d]} + (\mathbf{r}_{[p]} - \mathbf{r}_{[t]}) \quad (\text{POA}) \quad (5)$$

which yields two symmetric compositional phonological vectors: a voicing vector $\mathbf{v}_{\text{voi.}} = \mathbf{r}_{[d]} - \mathbf{r}_{[t]}$ in eq. (4) and a change of POA vector⁴ $\mathbf{v}_{\text{POA}} = \mathbf{r}_{[p]} - \mathbf{r}_{[t]}$ in eq. (5).

To empirically evaluate this hypothesis, we tested 19 phonological features from PanPhon (Mortensen et al., 2016).⁵ We find that analogies based on all 19 phonological features consistently hold in S3M representations.

In § 4, we further explore the scale of each phonological vectors. In detail, we introduce a scalar $\lambda \in \mathbb{R}$ into eq. (4):

$$\mathbf{r}_{[b]} \simeq \mathbf{r}_{[p]} + \lambda \cdot (\mathbf{r}_{[d]} - \mathbf{r}_{[t]}). \quad (6)$$

We hypothesize that *the scale λ will control the acoustic characteristics associated with a phonological vector in a continuous manner*. For instance, we expect the scale λ of the voicing vector $\mathbf{v}_{\text{voi.}}$ to correspond to the degree of voicing of the segment.

To validate this hypothesis, we train a vocoder f^{-1} to approximate the inverse of S3M f :

$$\mathbf{R} = f(\mathbf{x}) \quad (7)$$

$$\mathbf{x} \simeq f^{-1}(\mathbf{R}), \quad (8)$$

³Voicing refers to the vibration of the vocal folds. Bilabial and alveolar describe place of articulation (POA), with bilabials produced by bringing both lips together and alveolars by placing the tongue tip against the alveolar ridge.

⁴We can also name it as alveolar-to-bilabial vector or negative coronality vector.

⁵Syllabic, sonorant, continuant, delayed release, lateral, nasal, strident, voice, spread glottis, anterior, coronal, distributed, labial, high, low, back, round, tense, and long.

where \mathbf{x} is the input speech signal and $\mathbf{R} \in \mathbb{R}^{T' \times F}$ its S3M representation before pooling (details in § 3.1.4). We modify the representation \mathbf{R} by adding scaled phonological vectors (eq. (6)) and resynthesize the audio through the vocoder f^{-1} (eq. (8)).

Eight phonological features that can be directly quantified through acoustic measurements on \mathbf{x} were selected: height (high and low), backness, roundness, nasality, sonority, stridency, and voicing. We observe that acoustic measurements strongly correlate with the scale λ , for both interpolation ($|\lambda| \leq 1$) and extrapolation ($|\lambda| > 1$) of phonological vectors. These results suggest that S3M representations encode phonological features not as purely binary distinctions but as a continuum through specific vector directions and scales.

In summary, our contributions to advancing the understanding of S3M representations are twofold:

- **Direction:** We show that S3M representations exhibit phonological vector arithmetic, *i.e.*, existence of compositional vectors that align with phonological features (§ 3).
- **Scale:** We show that the scale of phonological vectors corresponds to the degree of their associated phonological features, leading to interpretable control of speech synthesis (§ 4).

2 Datasets

We use two phonetically transcribed and manually segmented datasets, TIMIT and VoxAngeles, to evaluate both English-specific and cross-linguistic generalizability of our findings.

TIMIT (Garofolo et al., 1993) contains utterances and their phonetic segmentations of 630 English speakers with balanced phonetic and dialectal coverage. To avoid ambiguities in extracting phonological features, we exclude diphthongs from our analysis. In addition, plosive closures are merged with their subsequent releases.⁶

VoxAngeles (Chodroff et al., 2024) provides phonetic segmentations for audio from the UCLA Phonetics Archive from 95 languages across 21 language families, providing broader phonetic diversity. As the dataset has no predefined train/test splits, we use the full set. Because VoxAngeles excludes English, it enables evaluation of whether S3Ms generalize to phones that do not occur in English.

⁶Plosives in TIMIT are separately segmented into closure and burst, different to IPA definitions.

3 Experiment 1: Direction of Phonological Vectors

In this section, we test the first hypothesis: whether linear phonological vectors that satisfy phonological analogies exist within S3M representations, by observing their directions (eqs. (4) and (5)).

3.1 Method

3.1.1 Creating phonological analogies

Phonological features capture fundamental properties of speech sounds (Trubetzkoy, 1949; Jakobson et al., 1951), such as voicing, place of articulation (e.g., bilabial, coronal, velar), and manner of articulation (e.g., plosive, fricative, nasal). To comprehensively analyze such features, we use PanPhon (Mortensen et al., 2016) to extract discrete phonological features for each phone p : $\mathbf{h}'_p = \text{PanPhon}(p) \in \{-1, 0, 1\}^{21}$. PanPhon yields 21 phonological features with values + (present), 0 (not applicable), or - (absent). For example, /b/ can be represented as [+voice, +labial, -nasal, +anterior, 0tense, ...]. To binarize the features, we expand each value +, 0, and - to be [1, 0], [0, 0], and [0, 1], producing $\mathbf{h}_p = \text{extend}(\mathbf{h}'_p) \in \{0, 1\}^{42}$.

A quadruplet of phones $\mathbf{p} = (p_1, p_2, p_3, p_4)$ yields two symmetric analogies, e.g., eqs. (4) and (5). We construct the quadruplet set Q by filtering from every possible quadruplets drawn from the phone vocabulary $\mathbf{p} \in V$:⁷

$$Q = \{\mathbf{p} \in V \mid \mathbf{h}_{p_1} - \mathbf{h}_{p_2} = \mathbf{h}_{p_3} - \mathbf{h}_{p_4}\}, \quad (9)$$

such that each quadruplet $\mathbf{p} \in Q$ denotes two symmetric analogies $p_1 : p_2 = p_3 : p_4$ and $p_1 : p_3 = p_2 : p_4$, or, equivalently, $\mathbf{r}_{p_1} \simeq \mathbf{r}_{p_2} + \mathbf{r}_{p_3} - \mathbf{r}_{p_4}$. Note that the analogies are not required to be minimal pairs and may differ along multiple phonological features, where we did not observe any systematic influence of minimal vs. non-minimal pairs (§ B.6).

We obtain 236 and 468 quadruplets (hence double the analogies) from the TIMIT test set and the full VoxAngeles dataset, respectively. We excluded the quadruplets where there is no PanPhon mapping for any one of the phones, or if any of the phones has less than 50 occurrences within the dataset, to ensure reliable cosine similarity estimates, which reduced the number of phones from 47 to 43 for TIMIT, 567 to 57 for VoxAngeles. No

⁷ $\mathbf{p} \in V$ means that all phones p_1, p_2, p_3, p_4 belong to V .

quadruplets were obtained for the “constricted glottis” and “consonantal” feature for both datasets,⁸ resulting in $21 - 2 = 19$ phonological features being tested.

3.1.2 Measuring cosine similarities

Through preparation steps described above, we now have a vector representation for each phone: $D = \{(p_i, \mathbf{r}_i)\}$. To quantitatively measure whether each quadruplet $\mathbf{p} = (p_1, p_2, p_3, p_4)$ is consistent with the underlying representations, we calculate the average cosine similarity:

$$\overline{\text{cos}}(\mathbf{p}) = \mathbb{E}[\cos(\mathbf{r}_{p_1}, \mathbf{r}_{p_2} + \mathbf{r}_{p_3} - \mathbf{r}_{p_4})], \quad (10)$$

where the average is calculated by randomly sampling each phone representation from D .

To quantify the uncertainty of similarity estimate, we use bootstrapping, similar to Choi et al. (2024b). We randomly sample 1000 phones each with replacement, and compute the averaged cosine similarity. In our preliminary experiments, we found that each bootstrap estimate is stable, resulting in very small across-replicate variance. As such, we construct the 99% confidence interval (CI) using 10 replicates, to avoid excessive computation.

3.1.3 Comparing cosine similarities

We compare the above cosine similarities with two baselines. The **same-phone baseline** provides an upper bound by measuring the similarity between the same phone p_1 from different utterances:

$$\overline{\text{cos}}^+(\mathbf{p}) = \mathbb{E}[\cos(\mathbf{r}_{p_1}, \mathbf{r}'_{p_1})], \quad (11)$$

where \mathbf{r}'_{p_1} is another random sample of p_1 from D .

Similarly, **different-phone baseline** measures between p_1 and random phone that is not p_1 :

$$\overline{\text{cos}}^-(\mathbf{p}) = \mathbb{E}[\cos(\mathbf{r}_{p_1}, \mathbf{r}_{\text{not } p_1})]. \quad (12)$$

For a well-structured representation, we expect the similarities to satisfy the ordering:

$$\overline{\text{cos}}^-(\mathbf{p}) < \overline{\text{cos}}(\mathbf{p}) < \overline{\text{cos}}^+(\mathbf{p}). \quad (13)$$

We assess whether this ordering holds by comparing the confidence intervals (CIs) of the estimates, verifying that the upper CI of the left term is below the lower CI of the right term. Equation (13) is analogous to an ABX test (Chaabouni et al., 2017),

⁸low, constricted glottis, spread glottis, long, and consonantal in TIMIT, and constricted glottis, lateral, and consonantal in VoxAngeles were missing.

where the target representation is expected to be more similar to another instance of the same phone (eq. (11)) than to its approximation (eq. (10)), and more similar to its approximation (eq. (10)) than to representations of other phones (eq. (12)).

To summarize the behavior over all quadruplets, we define the **success rate**, which is the proportion of quadruplets whose similarity scores satisfy the ordering in eq. (13) such that phonological analogies hold:

$$S(Q) = \frac{1}{|Q|} \sum_{\mathbf{p} \in Q} \mathbb{1}[\overline{\text{cos}}^-(\mathbf{p}) < \overline{\text{cos}}(\mathbf{p}) < \overline{\text{cos}}^+(\mathbf{p})], \quad (14)$$

where $\mathbb{1}$ denotes the indicator function, returning 1 if the condition holds, and 0 otherwise. We additionally evaluate performance using an alternative metric, described in § A.1.

3.1.4 Speech representations

We use spectral representations as baselines. We employ log mel spectrograms (**MelSpec**) and mel-frequency cepstral coefficients (**MFCC**), using the default parameters of librosa (McFee et al., 2015).

We compare these baselines to three monolingual S3Ms trained on English: **wav2vec 2.0** (Baevski et al., 2020), **HuBERT** (Hsu et al., 2021), and **WavLM** (Chen et al., 2022), using the LARGE configuration. We extract 25 layerwise representations per model. We defer additional details to § A.2.

Given the speech utterance $\mathbf{x} \in \mathbb{R}^T$ of length T , speech representation is a matrix $\mathbf{R} = f(\mathbf{x}) \in \mathbb{R}^{T' \times F}$, where the temporal dimension (number of frames) is reduced to $T' \simeq T/s$ with the model’s stride size s .⁹ Our goal is to extract a single vector $\mathbf{r} \in \mathbb{R}^F$ for each phone segment, given its start and end times (t_s, t_e) . Following Pasad et al. (2021, 2023), we conduct average pooling after applying feature slicing. § B.2 conducts additional comparisons on different slicing methods.

3.2 Results

3.2.1 S3Ms vs. Spectral representations

Figure 2 compares the success rates of S3Ms and spectral representations on phonological analogies using TIMIT. The last layer of HuBERT (94%) and WavLM (92%), as well as the middle layer of wav2vec 2.0 (61%) substantially outperforms spectral features MFCC (19%) and MelSpec (0%).

⁹The evaluated S3Ms use $s = 320$ for 16kHz input audio.

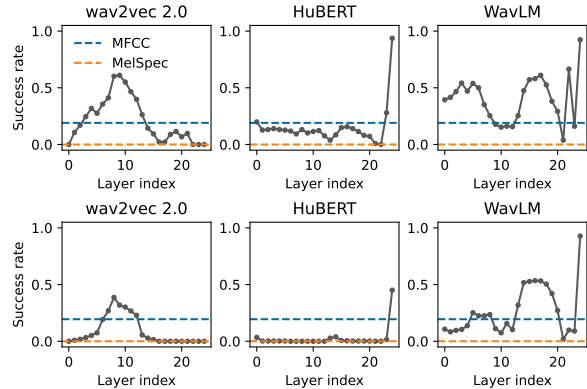


Figure 2: Comparing S3Ms with spectral representations on TIMIT (top) and VoxAngeles (bottom).

Each S3M exhibits distinct layerwise behavior: wav2vec 2.0 peaks in the middle layer, whereas HuBERT and WavLM peak in the last layer. This behavior is consistent with prior observations that measure layerwise phonetic information through probing (Pasad et al., 2023; Choi et al., 2025). Our results extend these findings by showing that S3Ms exhibit phonological compositionality.

Further, we observe that a greater number of analogies hold in the middle or final layers compared to layer index 0. We hypothesize that the need for deeper layers suggests that S3Ms benefit from increased contextualization when forming abstract phonological vectors. We explore this hypothesis further in § 3.2.3 and § B.2.

Additionally, we compare phone recognizers that are fine-tuned from S3Ms to assess the impact of the phone recognition task on the phonological analogies (§ B.4). We also find that alternative evaluation metrics can yield different layerwise trends in Figure 2 (§ B.1).

3.2.2 Do S3Ms generalize to unseen phones?

We examine whether phonological analogies hold for phones from unseen languages using VoxAngeles dataset (Figure 2). Of the 468 analogies, 316 (68%) contain at least one phone that does not exist in the English (TIMIT) phone set.

Consistent with the trends observed in § 3.2.1, S3Ms continue to achieve higher success rates than spectral representations. In particular, WavLM, HuBERT, and wav2vec 2.0 achieve success rates of 93%, 45%, and 39%, respectively, compared to 19% for MFCC and 0% for MelSpec. This indicates that English-only S3Ms capture phonological structure beyond English-specific phones.

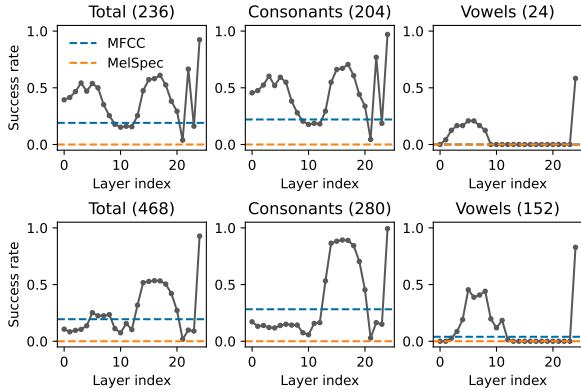


Figure 3: Comparing consonant-only and vowel-only quadruplets on TIMIT (top) and VoxAngeles (bottom) for WavLM. Number within the parenthesis denotes the number of quadruplets. We exclude cases where a quadruplet contains both consonants and vowels.

3.2.3 Vowels vs. Consonants

We analyze WavLM representations, which achieve the highest overall success rates (Figure 2). We observe three prominent peaks in success rate: (1) a first intermediate peak between layers 0–10, (2) a second intermediate peak between layers 10–20, and (3) the highest peak at the final layer.

To further investigate the origin of different peaks, we separate phonological analogies into vowels and consonants. As shown in Figure 3, vowels are associated with the first intermediate peak across both datasets. In contrast, consonants exhibit more complex behavior: in TIMIT, consonants contribute to both intermediate peaks, whereas in VoxAngeles they are associated with the second intermediate peak. Nevertheless, both vowels and consonants peak on the final layer.

We suggest that differences may stem from the distinct acoustic-temporal properties of vowels and consonants. Vowel cues tend to be temporally localized, whereas consonantal cues are often distributed across surrounding segments.¹⁰ Given that deeper layers are more likely to leverage broader contextual information, we speculate that phonological features requiring less temporal context often peak earlier in the network (*e.g.*, vowels), and vice versa (*e.g.*, consonants).

In summary, these findings indicate that the strong performance of S3Ms is closely tied to their ability to leverage contextual information. This

¹⁰For example, aspirated plosive [p^h] in the word *apply* can be inferred from multiple cues spanning a broader temporal context, including formant transitions in the preceding vowel, burst energy at release, subsequent aperiodic noise, and partial devoicing of the following [l].

conclusion is further supported by experiments that limit the temporal window size (§ B.2). Moreover, results suggest that different temporal complexities are preferentially contextualized at different layers, while the final layer gathers them into a unified representation.

We additionally tested other factors beyond the consonant-vowel distinction that may dictate such layerwise trends. However, we found that neither individual phonological features (§ B.5) nor phonological distances between analogies (§ B.6) substantially influence these layer-wise patterns, leaving more underlying causes for future investigation.

4 Experiment 2: Scale of Phonological Vectors

In this section, we test the second hypothesis: whether the scale λ of the phonological vectors (eq. (6)) correlates with acoustic measurements associated phonological features, by training a vocoder to invert S3Ms.

4.1 Method

4.1.1 Modifying representations through phonological vectors

We define each phonological vector \mathbf{v} using the PanPhon features \mathbf{h} in § 3.1.1. We define the phonological vector as the difference between the mean representations of phones with and without the feature i (not necessarily minimal pairs):

$$\mathbf{v}_i = \mathbb{E}_{\mathbf{h}[i]=+1}[\mathbf{r}] - \mathbb{E}_{\mathbf{h}[i]=-1}[\mathbf{r}]. \quad (15)$$

Motivated by the analysis of § 3.2.3, we separately compute consonants and vowels, using consonant-derived vectors for consonants, and vice versa. For example, the voicing vector is obtained by subtracting the averaged representations \mathbf{r} of all voiced consonants ($\mathbf{h}[\text{voi.}] = +1$) with all unvoiced consonants ($\mathbf{h}[\text{voi.}] = -1$). We additionally analyze the sample efficiency of eq. (15) (§ B.9) and compare with single phone pair constructions (§ B.8).

As mentioned in eq. (6), we apply vector \mathbf{v} to the frames corresponding to the target phone p . Given its start and end frame indices t'_s, t'_e (§ 3.1.4), the modified representation $\tilde{\mathbf{R}}$ is defined as:

$$\tilde{\mathbf{R}}_t = \begin{cases} \mathbf{R}_t + \lambda \mathbf{v} & (t'_s \leq t < t'_e) \\ \mathbf{R}_t & (\text{otherwise}), \end{cases} \quad (16)$$

where the scaling factor $\lambda \in \mathbb{R}$ controls the strength of the modification and \mathbf{R}_t denotes the representation at timestep t . Finally, using the vocoder model

Phonological feature	High	Low	Back	Round	Nasal	Sonorant	Strident	Voice
Acoustic measurement	F1	F1	F2	F2	F1BW	HNR	COG	COG
Expected correlation sign	–	+	–	–	–	+	+	–

Table 1: Summary of phonological features, their associated acoustic measurements, and expected correlation signs. The phonologically expected correlation sign is denoted as + (positive) or – (negative). Five types of acoustic measurements are used: first formant (F1), second formant (F2), first-formant bandwidth (F1BW), harmonic-to-noise ratio (HNR), and center of gravity (COG). Further details are provided in § A.3.

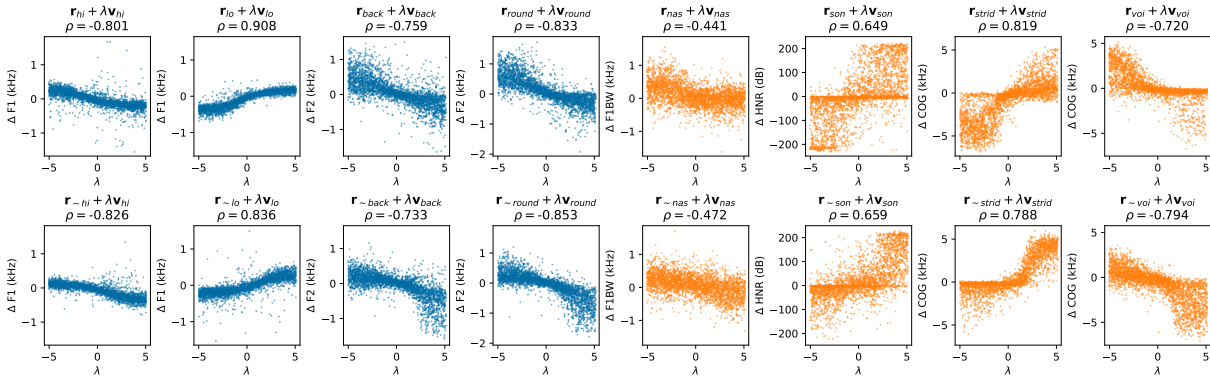


Figure 4: Comparison between the phonological vector scale λ and the acoustic measurements (§ A.3) on TIMIT. ρ denotes Spearman’s rank correlation coefficient. Blue and orange plots indicate vowels and consonants, respectively. The empirically observed correlation signs match the theoretical expectations shown in Table 1. Further, plots demonstrate the linearity of phonological vectors, resulting in monotonic (but not necessarily linear) changes in acoustic measurements.

f^{-1} , we reclaim the expected speech through resynthesis: $\tilde{x} = f^{-1}(\tilde{\mathbf{R}})$. We use the final layer representations from WavLM, as it has been shown to be effective for both TIMIT and VoxAngeles (§ 3.2).

4.1.2 Analyzing modified representations through acoustic measurements

To assess whether the scaling factor λ reflects the degree of associated phonological feature realization, we extract acoustic measurements from the modified representation $\tilde{\mathbf{R}}_t$. Specifically, we use the vocoder f^{-1} to resynthesize speech \tilde{x} from $\tilde{\mathbf{R}}$, and then compute the acoustic measurements for the targeted phonological feature. These measurements quantify how the degree of feature realization in the resynthesized speech varies with respect to λ . For example, increasing λ for the voicing vector is expected to yield a higher degree of voicing in the resynthesized target phone.

To quantify these effects, we compare acoustic measurements before and after modifying representations with scaled phonological vectors. We assess the relationship between the scale λ and the change of acoustic measurements Δ using Spearman’s rank correlation coefficient (ρ).

4.1.3 Vocoder training

We train two vocoder models based on the Vocos vocoder (Siuzdak, 2024): an English model using LibriTTS (Zen et al., 2019), and a multilingual model using FLEURS-R (Ma et al., 2024). Vocos mitigates synthesis artifacts commonly introduced by aggressive upsampling and demonstrates robustness to out-of-distribution inputs (Siuzdak, 2024), making it well suited for analyzing S3M representations. Details of vocoder training can be found in § A.4. We quantitatively evaluate resynthesis quality using acoustic measurements in § B.12.

4.1.4 Phonological vectors

We test eight phonological vectors with well-established corresponding acoustic measurements: *high*, *low*, *back*, and *round* for vowels, and *nasal*, *sonorant*, *strident*, and *voice* for consonants. We summarize the acoustic measurements on § A.3 and the expected correlation sign in Table 1. We additionally analyze the non-orthogonal relationships between these phonological vectors in § B.9. We also examine the influence of the *round* phonological vector across multiple acoustic measurements in § B.11.

To estimate the phonological vector \mathbf{v} , we use




Figure 5: Applying round vector to front vowel [i], where there is no front rounded vowel in English. Orange and blue arrows denote F2 and F3, respectively, which are all decreasing for $\lambda > 0$.




Figure 6: Applying the voicing vector to phone [b]. Orange arrows denote the onset of voicing in the subsequent vowel. For increasing values of λ , the voicing onset time is moved earlier, extending into the closure segment for large values of λ .




Figure 7: Applying the strident vector to phone [b]. Orange arrows point to the burst, where increasing λ removes the burst. Blue arrows denotes the increasing energy around 4~8kHz.




Figure 8: Applying the nasal vector to phone [b]. Orange arrows point to the burst, where increasing λ removes the burst. Blue arrows point to the low frequency murmur introduced with increasing nasalization.

TIMIT train split and a fixed randomly selected subset of languages for VoxAngeles. For each vector, we modify and resynthesize 3000 utterances drawn from the remaining data, *i.e.*, the TIMIT test split and the remaining VoxAngeles languages. For each utterance, we sample $\lambda \in U(-5, 5)$ from the uniform distribution and a phone segment (with replacement), modify its representation using $\lambda \mathbf{v}$, and resynthesize using the vocoder f^{-1} .

4.2 Results

4.2.1 Quantitative analyses

Figure 4 compares the acoustic measurements of the original and the modified-resynthesized speech on TIMIT. Across all features, the correlation signs observed in our experiments exactly matches the theoretically expected signs (Table 1). Further, we observe consistent and monotonic relationships between the vector scale λ and the resulting acoustic measurements. Our results confirm that the phonological vectors behave in accordance with their

intended interpretation as linear directions in the learned representation space, inducing monotonic but not necessarily linear changes in acoustic measurements. Figure 26 in § B.10 further shows that the same trends hold for VoxAngeles, demonstrating the generalization abilities to unseen phones during training.

Further, we found the effects of the phonological vectors are continuous rather than binary. For example, increasing the scale of the voicing vector does not abruptly toggle voicing on or off. Instead, it produces smooth shifts in COG, reflecting a graded change in the degree of voicing. This suggests that S3M representations encode phonological features not merely as categorical contrasts, but as continuous directions. This property enables fine-grained control over acoustic variation along individual phonological dimensions.

We also observe robust extrapolation well beyond the interpolation range $|\lambda| \leq 1$. $|\lambda| > 1$ still yields acoustically interpretable outputs, fur-

ther supporting the linear structure of phonological vectors within the representation space.

There were three exceptions for extrapolation. For sonorance, the curve is effectively saturated for already-sonorant segments, reflecting the fact that sonorants cannot easily be made “more sonorant.” Similarly, voiced consonants and non-stridents also show comparable saturation effects.

4.2.2 Qualitative analyses

We complement the quantitative results with a qualitative inspection of spectrograms. For each phonological feature, we resynthesize audio using $\lambda = (-5, -2, -1, 0, 1, 2, 5)$. We visualize the modified phone with 500ms of context on both sides to show both the local effect and potential coarticulation (Figures 5 to 8).

The **rounding** vector is applied to the high front unrounded vowel [i] in Figure 5. For $\lambda > 0$, all formants are being lowered together, consistent with an acoustic hallmark of lip rounding. Note that English has no front rounded vowels, demonstrating that rounding vector generalizes to unseen phones.

The **voicing** vector is applied to the voiced bilabial plosive [b] in Figure 6. For $\lambda < 0$, the onset time of voicing after the plosive is increased. For $\lambda > 0$, the model decreases the voice onset time (VOT). For large values, the voicing of the subsequent vowel is extended into the closure of the plosive, exhibiting negative VOT.

The **strident** vector is applied to [b] in Figure 7. For $\lambda > 0$, frication above 4kHz increasingly emerges, matching the spectral signature of strident fricatives. Notably, it also removes the burst characteristics of plosives when increasing stridency. This indicates that S3M representations encode not only static spectral envelopes, but also internal temporal structure (burst vs. frication), and modify them coherently along the phonological dimension.

The **nasal** vector is applied to [b] in Figure 8. Increasing λ introduces nasal acoustic cues: weakening of the burst, and introduction of a low-frequency murmur. As with stridents, these modifications affect both temporal and spectral structure associated with the underlying manner of articulation, not merely its coarse spectral profile.

Taken together, our analyses show that phonological vectors induce phonologically coherent changes in speech representations. These effects follow expected trajectories for individual phonological features, vary continuously rather than categorically, and in some cases extrapolate in lin-

guistically interpretable ways. Moreover, the vectors modulate not only spectral envelopes but also temporal cues, indicating that S3M representations encode rich internal structure. These findings further connect S3M representations to scalar and multi-valued features in phonological theory (see Gnanadesikan (1997) for a survey).

5 Related works

Word analogies. Mikolov et al. (2013b) demonstrated linear analogies in word embeddings, showing syntactic and semantic relations through vector arithmetic. Kim and de Marneffe (2013) extended to continuous semantic scales for adjectives, which is analogous to scaling phonological vectors with λ . Also, Pennington et al. (2014) use analogical tasks to compare embeddings, and Ethayarajh et al. (2019) provide a mathematical explanation for linear analogies, motivating our work for phonological analogies. Additionally, Levy and Goldberg (2014) raised concerns about word analogy evaluation, which motivated Fournier et al. (2020) to propose alternative methods based on comparing relational directions rather than representations directly.

Linear representation hypothesis (LRH) suggests that human-interpretable features are linearly represented within models’ hidden representations (Elhage et al., 2022; Park et al., 2024; Modell et al., 2025), theoretically supporting the existence of phonological vectors. LRH is also closely related to steering vectors (Subramani et al., 2022; Turner et al., 2023), which control model behavior by adding interpretable vectors to model representations, motivating our controllable speech synthesis.

Speech model interpretability. Prior work mainly focused what information is encoded, *i.e.*, spectral (Choi and Yeo, 2022), phonetic (Choi et al., 2024a), articulatory (Cho et al., 2023, 2024a), syllabic (Baade et al., 2025; Cho et al., 2025), lexical (Peng and Harwath, 2022; Pasad et al., 2024), syntactic (Shen et al., 2023), and semantic information (Pasad et al., 2024). More closely related to our work, several studies examined how information is encoded: layerwise analyses revealed linguistic hierarchy (Pasad et al., 2021, 2023), representations store information within similarities (Choi and Yeo, 2022; Choi et al., 2024b), and hierarchy among similarities (Abdullah et al., 2023; Choi et al., 2025).

Phonological analogies. For S3Ms, Gauthier

et al. (2025) showed morphological inflection induces linear geometry and Nakamura et al. (2025) identified certain phonological axes within S3Ms. Li et al. (2021) also analyzed phone recognition models and identified voicing and aspiration vectors. Chaabouni et al. (2017) used phonological analogies and ABX tasks as evaluation tools for multimodal speech representations. Others also leveraged phonological analogies in text settings: Silfverberg et al. (2018) showed that phonemic text embeddings can learn phonological vectors without supervision and Zouhar et al. (2024) used phonological analogies to evaluate them. In contrast, our work provides a large-scale cross-lingual analysis on S3Ms and further explores the scale of phonological vectors.

Controllable speech synthesis approaches often adopt interpretable features, including phonetic posteriorgrams (Zhao et al., 2019; Morrison et al., 2024), articulatory features (Cho et al., 2024b; Krug et al., 2025), or phonological features (Staib et al., 2020; Tännander et al., 2024), to enable interpretable control. While they rely on explicitly designed representations, our phonological vectors are driven from self-supervision. While we do not directly compare synthesis performance, we expect our emergent phonological vectors from S3Ms can be leveraged in future work on such applications.

6 Conclusion

We show that S3Ms, trained only on speech without phonological supervision, learn *linearly composable* and *scalable* phonological vectors. These findings advance both speech processing and linguistics by clarifying the internal structure of S3M representations and refining our understanding of phonological features. In speech processing, our findings enable intuitive interpretations of S3M representations and fine-grained control of speech synthesis along phonological dimensions. In linguistics, they provide empirical evidence that phonological features can emerge from acoustic regularities and motivate a view of phonological features as continuous rather than strictly binary.

Limitations

This work explored only a limited region of the space of possibilities. Only a small number of S3Ms were investigated. Different models behaved differently and it is not possible to isolate what the causes of these differences are on the basis of the

results reported here. Although there can be multiple possible systems of articulatory and acoustic features, our studies only investigated the feature system assumed by PanPhon. This makes it difficult to draw firm conclusions about whether what is important, in extracting phonological vectors, is identifying a coherent feature that minimally delineates a phonological natural class or simply identifying a consistent phonetic difference. Finally, synthesis results are influenced not only by the S3Ms but also by the vocoder. Since we evaluate only a single vocoder, some observed behaviors may reflect vocoder-specific characteristics rather than properties of the S3Ms alone.

Ethics statement

Our work focuses on phonological understanding of the self-supervised speech representations through the lens of vector arithmetic. All of our experiments are conducted using publicly available datasets (§ 2), which were collected and released under licenses appropriate for research use. We do not collect new data, nor do we include personally identifiable information beyond what is already present in these datasets.

While our speech synthesis experiments (§ 4) are intended solely for scientific analysis, they may have broader societal implications if misused. We do not evaluate or claim applicability to the generation of misleading or deceptive content. We release code, demo, and models for reproducibility and research purposes only, and we encourage future work to consider appropriate safeguards when applying similar techniques in downstream or user-facing systems.

Use of AI assistants

AI assistants were used in the preparation of this manuscript. Specifically, they were employed primarily for code auto-completion, minor text editing, and grammar polishing. Nevertheless, all scientific content, code implementations, results, and analyses were conceived, verified, and finalized by the authors.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 2238605. We are grateful to Calvin Chang, Christina Bjorndahl, the anonymous ACL reviewers, and the area chair for their helpful comments.

References

- Badr M Abdullah, Mohammed Maqsood Shaik, Bernd Möbius, and Dietrich Klakow. 2023. An information-theoretic analysis of self-supervised discrete representations of speech. In *Interspeech*.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. Common Voice: A massively-multilingual speech corpus. In *LREC*.
- Alan Baade, Puyuan Peng, and David Harwath. 2025. SyllableLM: Learning coarse semantic units for speech language models. In *ICLR*.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. In *NeurIPS*.
- Christina Jean Bjorndahl. 2018. *A Story of /v/: Voiced Spirants in the Obstruent-Sonorant Divides*. Cornell University.
- Paul Boersma and David Weenink. 2025. Praat: doing phonetics by computer. <https://www.fon.hum.uva.nl/praat/>. Version 6.4 [Computer program].
- Rahma Chaabouni, Ewan Dunbar, Neil Zeghidour, and Emmanuel Dupoux. 2017. Learning weakly supervised multimodal phoneme embeddings. In *Interspeech*.
- Guoguo Chen, Shuzhou Chai, Guan-Bo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khudanpur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Zhao You, and Zhiyong Yan. 2021. Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio. In *Interspeech*.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, and 1 others. 2022. WavLM: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16.
- Cheol Jun Cho, Nicholas Lee, Akshat Gupta, Dhruv Agarwal, Ethan Chen, Alan Black, and Gopala Anumanchipalli. 2025. Sylber: Syllabic embedding representation of speech from raw audio. In *ICLR*.
- Cheol Jun Cho, Abdelrahman Mohamed, Alan W Black, and Gopala K Anumanchipalli. 2024a. Self-supervised models of speech infer universal articulatory kinematics. In *ICASSP*.
- Cheol Jun Cho, Peter Wu, Abdelrahman Mohamed, and Gopala K Anumanchipalli. 2023. Evidence of vocal tract articulation in self-supervised learning of speech. In *ICASSP*.
- Cheol Jun Cho, Peter Wu, Tejas S Prabhune, Dhruv Agarwal, and Gopala K Anumanchipalli. 2024b. Coding speech through vocal tract kinematics. *IEEE Journal of Selected Topics in Signal Processing*, 18.
- Eleanor Chodroff, Blaž Pažon, An Baker, and Steven Moran. 2024. Phonetic segmentation of the UCLA phonetics lab archive. In *LREC-COLING*.
- Kwanghee Choi, Jee-weon Jung, and Shinji Watanabe. 2024a. Understanding probe behaviors through variational bounds of mutual information. In *ICASSP*.
- Kwanghee Choi, Ankita Pasad, Tomohiko Nakamura, Satoru Fukayama, Karen Livescu, and Shinji Watanabe. 2024b. Self-supervised speech representations are more phonetic than semantic. In *Interspeech*.
- Kwanghee Choi and Eun Jung Yeo. 2022. Opening the black box of wav2vec feature encoder. *arXiv preprint arXiv:2210.15386*.
- Kwanghee Choi, Eunjung Yeo, Calvin Chang, Shinji Watanabe, and David R Mortensen. 2025. Leveraging allophony in self-supervised speech models for atypical pronunciation assessment. In *NAACL*.
- Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. 2021. Un-supervised cross-lingual representation learning for speech recognition. *Interspeech*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, and 1 others. 2022. Toy models of superposition. *arXiv preprint arXiv:2209.10652*.
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. In *EMNLP*.
- Kawin Ethayarajh, David Duvenaud, and Graeme Hirst. 2019. Towards understanding linear word analogies. In *ACL*.
- Louis Fournier, Emmanuel Dupoux, and Ewan Dunbar. 2020. Analogies minus analogy test: measuring regularities in word embeddings. In *CoNLL*.
- Mark JF Gales, Kate M Knill, Anton Ragni, and Shakti P Rath. 2014. Speech recognition and keyword spotting for low-resource languages: Babel project research at cued. In *4th Workshop on Spoken Language Technologies for Under-Resourced Languages (SLTU 2014)*.
- John S. Garofolo, Lori F. Lamel, William M. Fisher, Jonathan G. Fiscus, David S. Pallett, and Nancy L. Dahlgren. 1993. Darpa timit: Acoustic-phonetic continuous speech corpus cd-rom, nist speech disc 1-1.1.
- Jon Gauthier, Canaan Breiss, Matthew K Leonard, and Edward F Chang. 2025. Emergent morpho-phonological representations in self-supervised speech models. In *EMNLP*.

- Amalia Elisabeth Gnanadesikan. 1997. *Phonology with ternary scales*. University of Massachusetts Amherst.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhota, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. HuBERT: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29.
- Yannick Jadoul, Bill Thompson, and Bart De Boer. 2018. Introducing parselmouth: A python interface to praat. *Journal of Phonetics*, 71.
- Roman Jakobson, C Gunnar Fant, and Morris Halle. 1951. Preliminaries to speech analysis: The distinctive features and their correlates.
- Jacob Kahn, Morgane Riviere, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazaré, Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, and 1 others. 2020. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP*.
- Joo-Kyung Kim and Marie-Catherine de Marneffe. 2013. Deriving adjectival scales from continuous space word representations. In *EMNLP*.
- Masahiko Komatsu, Shinichi Tokuma, Won Tokuma, and Takayuki Arai. 2002. Multi-dimensional analysis of sonority: perception, acoustics, and phonology. In *Interspeech*.
- Paul Konstantin Krug, Christoph Wagner, Peter Birkholz, and Timo Stich. 2025. Precisely controllable neural speech synthesis. In *ICASSP*.
- P. Ladefoged and S.F. Disner. 2012. *Vowels and Consonants*. Wiley.
- Peter Ladefoged. 1996. *Elements of acoustic phonetics*. University of Chicago Press.
- Omer Levy and Yoav Goldberg. 2014. Linguistic regularities in sparse and explicit word representations. In *CoNLL*.
- Xinjian Li, Juncheng Li, Florian Metze, and Alan W Black. 2021. Hierarchical phone recognition with compositional phonetics. In *Interspeech*.
- Min Ma, Yuma Koizumi, Shigeki Karita, Heiga Zen, Jason Riesa, Haruko Ishikawa, and Michiel Bacchiani. 2024. FLEURS-R: A restored multilingual speech corpus for generation tasks. In *Interspeech*.
- Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. librosa: Audio and music signal analysis in python. *SciPy*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Wen tau Yih, and Geoffrey Zweig. 2013b. Linguistic regularities in continuous space word representations. In *NAACL*.
- Alexander Modell, Patrick Rubin-Delanchy, and Nick Whiteley. 2025. The origins of representation manifolds in large language models. *arXiv preprint arXiv:2505.18235*.
- Max Morrison, Cameron Churchwell, Nathan Pruyne, and Bryan Pardo. 2024. Fine-grained and interpretable neural speech editing. In *Interspeech*.
- David R. Mortensen, Patrick Littell, Akash Bharadwaj, Kartik Goyal, Chris Dyer, and Lori Levin. 2016. PanPhon: A resource for mapping IPA segments to articulatory feature vectors. In *COLING*.
- Tomohiko Nakamura, Kwanghee Choi, Keigo Hojo, Yoshiaki Bando, Satoru Fukayama, and Shinji Watanabe. 2025. Discrete speech unit extraction via independent component analysis. In *International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*. IEEE.
- Kiho Park, Yo Joong Choe, and Victor Veitch. 2024. The linear representation hypothesis and the geometry of large language models. In *ICML*.
- Ankita Pasad, Chung-Ming Chien, Shane Settle, and Karen Livescu. 2024. What do self-supervised speech models know about words? *Transactions of the Association for Computational Linguistics*, 12.
- Ankita Pasad, Ju-Chieh Chou, and Karen Livescu. 2021. Layer-wise analysis of a self-supervised speech representation model. In *ASRU*.
- Ankita Pasad, Bowen Shi, and Karen Livescu. 2023. Comparative layer-wise analysis of self-supervised speech models. In *ICASSP*.
- Puyuan Peng and David Harwath. 2022. Word discovery in visually grounded, self-supervised speech models. In *Interspeech*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. MLS: A large-scale multilingual dataset for speech research. In *Interspeech*.
- Tarun Pruthi and Carol Y Espy-Wilson. 2007. Acoustic parameters for the automatic detection of vowel nasalization. In *Interspeech*.
- Gaofei Shen, Afra Alishahi, Arianna Bisazza, and Grzegorz Chrupala. 2023. Wave to syntax: Probing spoken language models for syntax. In *Interspeech*.
- Miikka Silfverberg, Lingshuang Jack Mao, and Mans Hulden. 2018. Sound analogies with phoneme embeddings. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2018*.

- Hubert Siuzdak. 2024. Vocos: Closing the gap between time-domain and fourier-based neural vocoders for high-quality audio synthesis. In *ICLR*.
- Marlene Staib, Tian Huey Teh, Alexandra Torresquintero, Devang S Ram Mohan, Lorenzo Foglianti, Raphael Lenain, and Jiameng Gao. 2020. Phonological features for 0-shot multilingual speech synthesis. In *Interspeech*.
- Peter Strevens. 1960. Spectra of fricative noise in human speech. *Language and speech*, 3.
- Nishant Subramani, Nivedita Suresh, and Matthew E Peters. 2022. Extracting latent steering vectors from pretrained language models. In *ACL*.
- Chihiro Taguchi, Yusuke Sakai, Parisa Haghani, and David Chiang. 2023. Universal automatic phonetic transcription into the international phonetic alphabet. In *Interspeech*.
- Christina Tännander, Shivam Mehta, Jonas Beskow, and Jens Edlund. 2024. Beyond graphemes and phonemes: continuous phonological features in neural text-to-speech synthesis. *Interspeech*.
- Nikola Trubetzkoy. 1949. *Grundzüge der Phonologie*, volume 7. Travaux di Cercle Linguistique de Prauge.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. Steering language models with activation engineering. *arXiv preprint arXiv:2308.10248*.
- Changhan Wang, Morgane Riviere, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux. 2021. VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation. In *ACL*.
- Dan Wells, Hao Tang, and Korin Richmond. 2022. Phonetic analysis of self-supervised representations of english speech. In *Interspeech*.
- Qiantong Xu, Alexei Baevski, and Michael Auli. 2022. Simple and effective zero-shot cross-lingual phoneme recognition. In *Interspeech*.
- Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. LibriTTS: A corpus derived from librispeech for text-to-speech. In *Interspeech*.
- Guanlong Zhao, Shaojin Ding, and Ricardo Gutierrez-Osuna. 2019. Foreign accent conversion by synthesizing speech from phonetic posteriorgrams. In *Interspeech*.
- Vilém Zouhar, Calvin Chang, Chenxuan Cui, Nate B Carlson, Nathaniel Romney Robinson, Mrinmaya Sachan, and David R Mortensen. 2024. Pwesuite: Phonetic word embeddings and tasks they facilitate. In *LREC-COLING*.

A Additional details

A.1 Item- and offset-based analogy tests

For the word analogy $a : b = a' : b'$, Mikolov et al. (2013b) evaluated using item-based analogy tests, *i.e.*, testing whether $b' \simeq a' - a + b$. Subsequent work has shown that offset-based tests, which assess whether $a - a' \simeq b - b'$, provide a more robust measure of linguistic regularity (Levy and Goldberg, 2014; Fournier et al., 2020). In this work, we adopt both perspectives: we primarily use item-based tests in § 3, and employ offset-based tests as a secondary analysis in § B.1.

The choice of item-based analogy test (success rates) in § 3 is motivated by the characteristics of speech, where the number of paired phones associated with a phonological feature is much smaller compared to the number of words within a word relation. As a result, offset-based tests on TIMIT and VoxAngeles discard many analogies (36 quadruplets compared to 236 in TIMIT, and 112 compared to 468 in VoxAngeles), which prevents further analyses such as comparisons between vowels and consonants (§ 3.2.3). Further, unlike in text, each phone in speech is realized through multiple utterances. Accordingly, quantities such as $a' - a + b$ do not correspond to a single point estimate but rather to a distribution, for which we report confidence intervals on success rates.

A.2 Details on S3Ms

In § 3.2.1, we compared three widely used S3Ms. We use the LARGE configuration for all models, consisting of 7 layers of 1D CNNs followed by 24 transformer blocks, a total of approximately 300M parameters. We extracted representations from each of the 24 transformer layers (index 1~24) as well as one from the CNNs (index 0), yielding 25 layerwise representations per model.

wav2vec 2.0 (Baevski et al., 2020) is trained on 60k hours of English read speech from LibriLight audiobook dataset (Kahn et al., 2020). The model is trained using a contrastive loss with on-the-fly learnable codebooks.

HuBERT (Hsu et al., 2021) is also trained on LibriLight but with a different training strategy. It uses a predictive loss, where the targets are k-means cluster assignments, either MFCC or the previous training iteration of HuBERT-base.

WavLM (Chen et al., 2022) additionally incorporates GigaSpeech (Chen et al., 2021) and VoxPopuli (Wang et al., 2021), which include both

read and spontaneous speech. Its loss is similar to HuBERT’s, but it includes an additional speech denoising objective.

A.3 Acoustic measurements

In § 4, we compare eight phonological features and their corresponding acoustic measurements. We use Parselmouth (Jadoul et al., 2018), a Python interface to Praat (Boersma and Weenink, 2025), to compute the measurements.

Formants, *i.e.*, resonant frequencies of the vocal tract, are used for measuring vowels (Ladefoged, 1996; Ladefoged and Disner, 2012). The first formant (F1) is used for the vowel height (high or low), such that high vowels show lower F1, and vice versa. Second formant (F2) is used for backness and roundness, such that back vowels and round vowels have lower F2.

Bandwidth of F1 (F1BW), *i.e.*, the frequency range around F1 with significant energy, is used for measuring nasality (Pruthi and Espy-Wilson, 2007). Nasal sounds exhibit increased damping from nasal cavities, resulting in spreading out the formant energy; a broader F1 bandwidth.

Center of gravity (COG), *i.e.*, amplitude-weighted average frequency of the spectrum, is used for measuring voicing (Bjorndahl, 2018) and stridents (Stevens, 1960). For voiced sounds, COG decreases due to the presence of the voicing bar, whereas it increases for stridents due to frication in the upper frequency range.

Harmonics-to-noise ratio (HNR), *i.e.*, the ratio of periodic (harmonic) energy to aperiodic (noise) energy, is used for measuring sonorance (Komatsu et al., 2002). Sonorant sounds tend to exhibit higher HNR due to their periodic structure.

A.4 Details on vocoder training

To train the neural vocoder f^{-1} , we follow the overall setup of Vocos (Siuzdak, 2024), with modification to the network architecture and training configuration. We slightly modify the network architecture dimensions, such that we can use either WavLM or MFCC representations as input. For faster convergence, we increase both the batch size and the learning rate by a factor of eight relative to the original configuration.

We use two speech synthesis datasets to train a vocoder. **LibriTTS** (Zen et al., 2019) is a multi-speaker English dataset with 585h of read speech. **FLEURS-R** (Ma et al., 2024) contains 1.3kh of enhanced read speech in 102 languages.

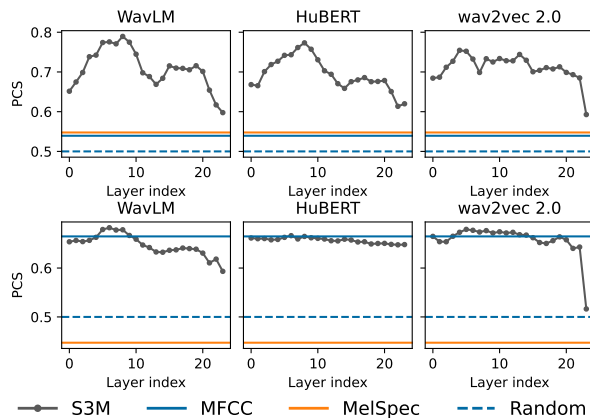


Figure 9: Comparing pairing consistency score (PCS) of S3Ms and spectral representations using TIMIT (upper) and VoxAngeles (lower).

B Additional results

B.1 Layerwise analysis via offset-based test

B.1.1 Settings

For the offset-based test, we use the pairing consistency score (PCS) (Fournier et al., 2020). PCS measures the separability between offsets drawn from the same relation and offsets drawn from mismatched relations using a binary classification criterion.

We construct PCS categories by grouping phone pairs that share the same phonological feature differences. For each category, we compute representation offsets between phone pairs belonging to the same category (correct offsets). These are contrasted with mismatched offsets, which are formed by shuffling the second phone in each pair. We use averaged representations per each phone to compute their offsets.

PCS is computed by treating the offset similarity as a binary classification problem, where correct offsets are labeled as positive instances and mismatched offsets as negative instances. We report the area under the receiver operating characteristic (ROC) curve as the evaluation metric. A random baseline with no relational consistency yields a PCS of 0.5.

B.1.2 Results

Figure 9 presents the layerwise PCS results for S3Ms and spectral representations, and we compare these findings with the success rate results in § 3. Consistent with § 3.2.1, the PCS results confirm the superiority of S3Ms over spectral representations across layers. Moreover, PCS scores generally increase in later layers relative

to layer 0, suggesting increased contextualization in deeper representations. In line with § 3.2.2, we also observe non-negligible performance on unseen phones, although performance remains lower than that achieved on seen phones.

Despite these overall consistencies, the detailed layerwise trends differ from those observed using success rates. Unlike the patterns reported in § 3.2.1, PCS exhibits relatively similar behavior across layers, with peak performance typically occurring in intermediate layers rather than at the final layer. This suggests that the most informative representations for offset-based relational consistency may reside in middle layers rather than the last layer with the sudden peak of success rates. We leave a deeper investigation of the discrepancies between different evaluation metrics for future work.

B.2 Feature vs. Audio slicing

B.2.1 Settings

We compare two common pooling methods for obtaining a phone vector \mathbf{r} from the representation matrix $\mathbf{R} = f(\mathbf{x})$.

Feature slicing (Pasad et al., 2021, 2023) performs temporal average pooling on the slice of the representation matrix:

$$\mathbf{r} = \text{avgpool}(f(\mathbf{x})[t'_s : t'_e]), \quad (17)$$

where $t'_s = \lfloor t_s/s \rfloor$ and $t'_e = \lceil t_e/s \rceil$ are the corresponding indices after temporal downsampling. As transformer layers have a global receptive field, it yields a contextualized representation, even though prior work shows that it tends to be locally dominated (Baevski et al., 2020; Hsu et al., 2021).

Audio slicing (Choi et al., 2024a,b), in contrast, slices the waveform directly:

$$\mathbf{r} = \text{avgpool}(f(\mathbf{x}[t_s : t_e])). \quad (18)$$

This method restricts the temporal receptive field strictly to the phone segment, removing the neighboring context. Empirically, this has been shown to lead to clearer separation in cosine similarity comparisons (Choi et al., 2024b). If the model requires minimum window size $w > (t_e - t_s)$, we ensure the sliced audio is at least length w by adding equal margins around the segment.

B.2.2 Results

In Figures 10 and 11, we observe that feature slicing is more effective than audio slicing for S3Ms,

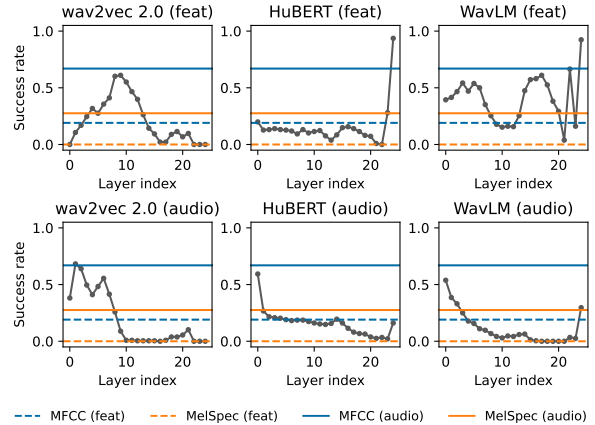


Figure 10: Comparing success rates for S3Ms with spectral representations using TIMIT. We denote feature and audio slicing as feat and audio, respectively.




Figure 11: Comparing success rates for feature sliced S3Ms with spectral representations using VoxAngeles. We denote feature and audio slicing as feat and audio, respectively.

whereas the opposite holds for spectral representations. Overall, feature-sliced WavLM last-layer representations shows strongest performance overall, where we primarily use for all the experiments.

Comparing feature and audio slicing for S3Ms further supports our conclusion in § 3.2.3, where audio slicing removes contextualization by limiting the temporal receptive field size.

In contrast, audio slicing is beneficial for spectral representations. Especially, MFCC reaches a success rate of 67% and 50% for TIMIT and VoxAngeles, respectively, which is surprisingly high. We suspect the difference comes from spectrum magnitude normalization. For feature slicing, spectral representations’ magnitude is normalized per utterance, where audio slicing directly normalizes the magnitude for each phone segment, likely leading to performance improvement.

However, as we show in § B.3, MFCC representations’ cosine similarities tend to collapse toward 1.0, making them substantially harder to utilize, leading to worse synthesis performance in § B.13.

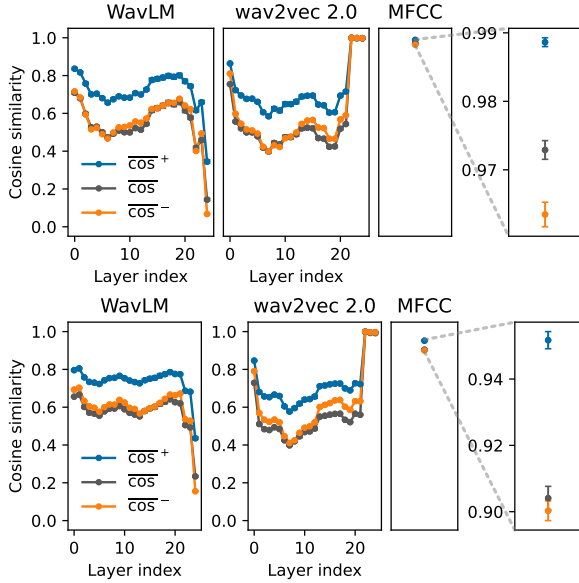


Figure 12: Comparing averaged similarities \bar{C} for feature sliced S3Ms with audio sliced MFCC on TIMIT (upper) and VoxAngeles (lower). We calculate the 99% CI considering quadruplet-wise cosine averages.

B.3 S3Ms avoid anisotropic collapse

B.3.1 Settings

To observe the absolute similarity values, we define the **averaged similarity**:

$$\bar{C}(A) = \frac{1}{|A|} \sum_{\mathbf{p} \in A} \overline{\text{cos}}(\mathbf{p}), \quad (19)$$

and estimate a CI over the set of quadruplet-wise similarities.

B.3.2 Results

Figure 12 shows the averaged cosine similarities for WavLM, wav2vec 2.0, and MFCC. For later layers of wav2vec 2.0, all similarities approach 1.0, reflecting anisotropic collapse (Ethayarajh, 2019), resulting in reduced success rates in Figure 2. MFCCs also behave similarly, with much smaller margins compared to the last layer of WavLM. On the other hand, WavLM does not show collapsing behavior, likely leading to more easily usable phonological vectors for synthesis in § 4.

B.4 Impact of fine-tuning with phone recognition task

B.4.1 Settings

We compare XLSR-53, a multilingual S3M, and their fine-tuned variants for phone recognition, Wav2vec2Phoneme and MultiIPA. To compare




Figure 13: Comparing pre-trained S3M (XLSR-53) and fine-tuned phone recognition models (Wav2vec2Phoneme and MultiIPA) on TIMIT.

these three models, we use the same settings of § 3.2.1.

XLSR-53 (Conneau et al., 2021) is trained with the wav2vec 2.0 contrastive objective on multilingual datasets spanning 53 languages from CommonVoice (Ardila et al., 2020), Multilingual LibriSpeech (Pratap et al., 2020), and BABEL (Gales et al., 2014), total of 56k hours.

Wav2vec2Phoneme (Xu et al., 2022) fine-tunes XLSR-53 for phone recognition using a CTC loss. The model is trained on the same datasets, using automatically generated phonemic transcriptions from multilingual G2P systems.

MultiIPA (Taguchi et al., 2023) in a similar manner to Wav2vec2Phoneme, but uses a smaller subset consisting of seven languages.

B.4.2 Results

Figure 13 compares XLSR-53 with their fine-tuned phone recognizer counterparts. Both fine-tuned models improve success rates relative to XLSR-53 across nearly all layers, indicating that fine-tuning for phone recognition generally strengthens phonological structure. However, two models differs in their layerwise behavior. Because the final layer precedes the CTC head, it is optimized to produce representations that linearly separate phones. Wav2vec2Phoneme shows a sharp increase in success rate in the deeper layers, suggesting that it develops more abstract phonological structure, whereas MultiIPA’s improvements are concentrated in earlier layers. This contrast may be related to differences in language coverage. As Wav2vec2Phoneme is trained on substantially more languages, it may encourage the model to learn more abstract, general phonological structure, while MultiIPA’s smaller language set may reduce the pressure for such abstraction. An interesting avenue for future work is to explore phone recognition as post-training strategies for S3Ms, promoting the emergence of abstract phonological vectors.

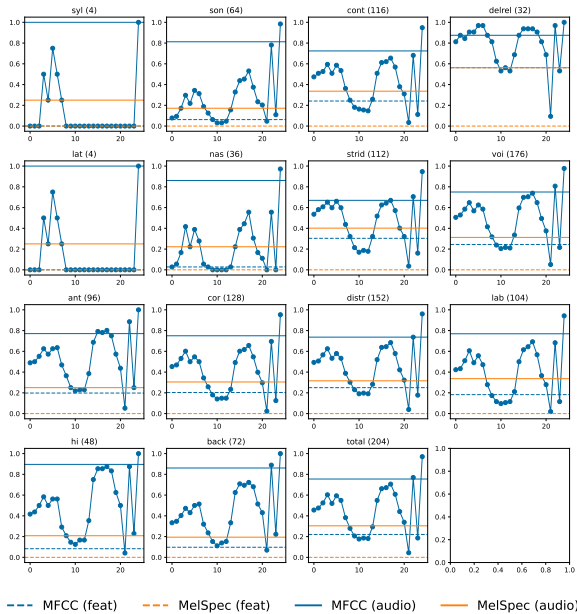


Figure 14: Phonological feature-wise success rates for consonants on TIMIT using WavLM representations.

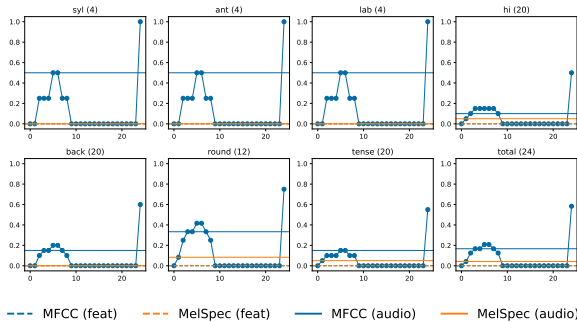


Figure 15: Phonological feature-wise success rates for vowels on TIMIT using WavLM representations.

B.5 Different phonological features does not lead to different layerwise trends

B.5.1 Settings

We observe the success rates of individual phonological features on both TIMIT and VoxAdeles. For each feature, we consider only the quadruplets for which the phonological vector in eq. (9) is nonzero, *i.e.*, where either $\mathbf{h}_{p_1}[i] \neq \mathbf{h}_{p_2}[i]$ or $\mathbf{h}_{p_1}[i] \neq \mathbf{h}_{p_3}[i]$ for the feature i under consideration. Following § 3.2.3, we evaluate consonants and vowels separately when computing feature-wise success rates, where quadruplets that contain a mixture of consonants and vowels are excluded from the analysis.

B.5.2 Results

For both TIMIT (Figures 14 and 15) and VoxAdeles (Figures 16 and 17), the layerwise trends are largely consistent across phonological features,

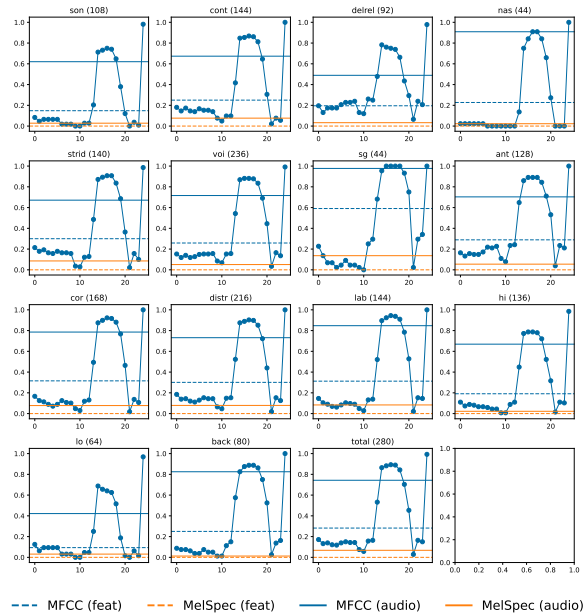


Figure 16: Phonological feature-wise success rates for consonants on VoxAdeles using WavLM representations.

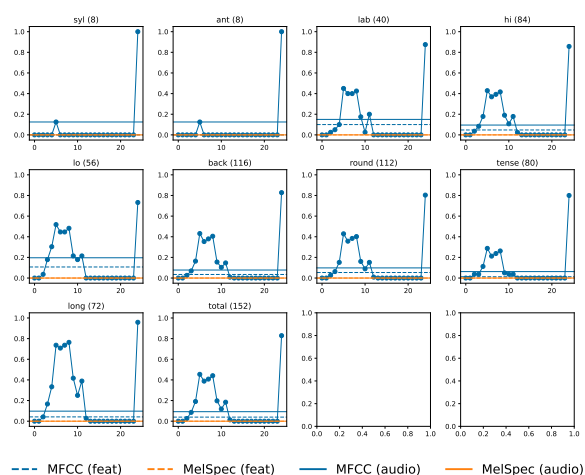


Figure 17: Phonological feature-wise success rates for vowels on VoxAdeles using WavLM representations.

with no clear feature-specific differences. One exception is the syllabic (syl) and lateral (lat) in Figure 14. However, the four samples are the permutations of the quadruplet of [l], [n], [l], and [n], where the syllabic consonants behave similar to vowels in terms of their reduced temporal variation.

B.6 Phonological distances does not lead to different layerwise trends

B.6.1 Settings

We observe the success rates on both TIMIT and VoxAdeles per phonological distance. Phonological distance between two phones is defined as the number of differing phonological features. For

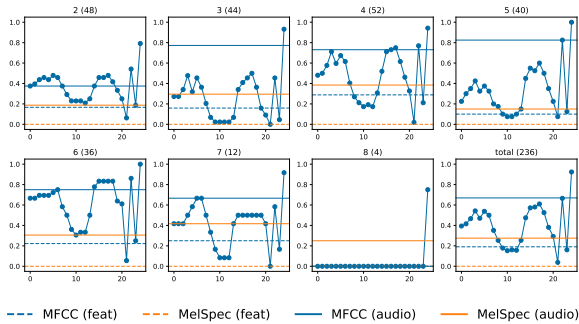


Figure 18: PanPhon feature distance-wise success rates on TIMIT using WavLM representations.

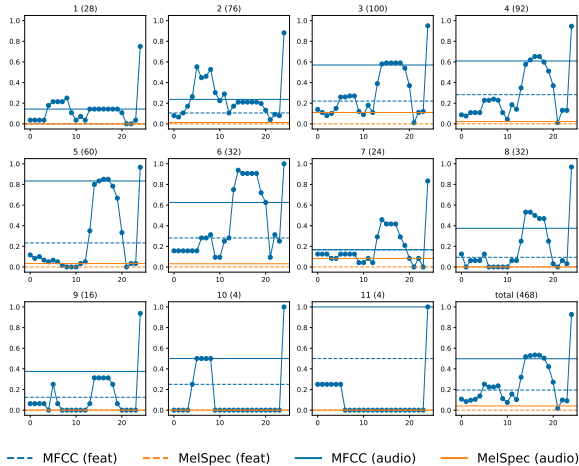


Figure 19: PanPhon feature distance-wise success rates on VoxAngeles using WavLM representations.

each quadruplet, we compute the distances between (p_2, p_3) and (p_2, p_4) , and assign the quadruplet to a distance bin based on the maximum of these two values. It ensures that each quadruplet is grouped according to its most phonologically divergent contrast.

B.6.2 Results

For both TIMIT (Figure 18) and VoxAngeles (Figure 19), the layerwise trends remain largely consistent across different phonological distances, with no variation observed as distance increases. One exception occurs in distance bins containing only four quadruplets. However, these cases arise from permutations of a single phone set, similar to § B.5.

B.7 Measuring sample efficiency for phonological vector extraction

B.7.1 Settings

In this experiment, we empirically evaluate how many phone representations are needed to extract accurate phonological vectors. To measure the accuracy of extracted phonological vectors, we treat

the vector extracted from the full training dataset in § 4.1.1 as the ground truth. We then approximate this vector using smaller subsets of the training data. We quantify approximation quality using cosine similarity.

We evaluate eight phonological vectors in § 4.1.4. For each positive and negative averaged representation in Equation (15), we randomly sample N representations with replacement. We repeat the approximation 1000 times with random samples and visualize the resulting cosine similarities using histograms. We evaluate $N \in \{1, 4, 16, 64, 256\}$ sampled representations.

B.7.2 Results

For both TIMIT and VoxAngeles (Figure 20), cosine similarity consistently improves as N increases across all phonological vectors. When $N = 256$, cosine similarity approaches 1.0, suggesting that a few hundred samples suffice for accurate approximation.

B.8 Using a single phone pair for phonological vectors

B.8.1 Settings

We compare phonological vectors extracted using the full dataset with those obtained using only a single phone pair for the positive and negative averaged representations in § 4.1.1. For example, for the voicing phonological vector, we evaluate whether using only the $([b], [p])$ pair leads to a good approximation.

We use the results from § B.7 as baselines for comparison. Similarly, the vector obtained from the full training dataset is treated as the ground truth. For each phonological vector, we evaluate all possible phone pairs and report the resulting distribution of cosine similarities with the ground truth via histograms.

B.8.2 Results

Figure 21 shows the distribution of cosine similarities, using § B.7 as the baseline. In TIMIT, each phone has few hundred to more than a thousand samples. However, the cosine similarity remains around 0.5, similar to using only four samples in the subsampling experiment. These results indicate that using a fixed phone pair produces different vectors from those obtained using diverse phones, despite maintaining positive cosine similarity with the ground truth. This is likely because the two

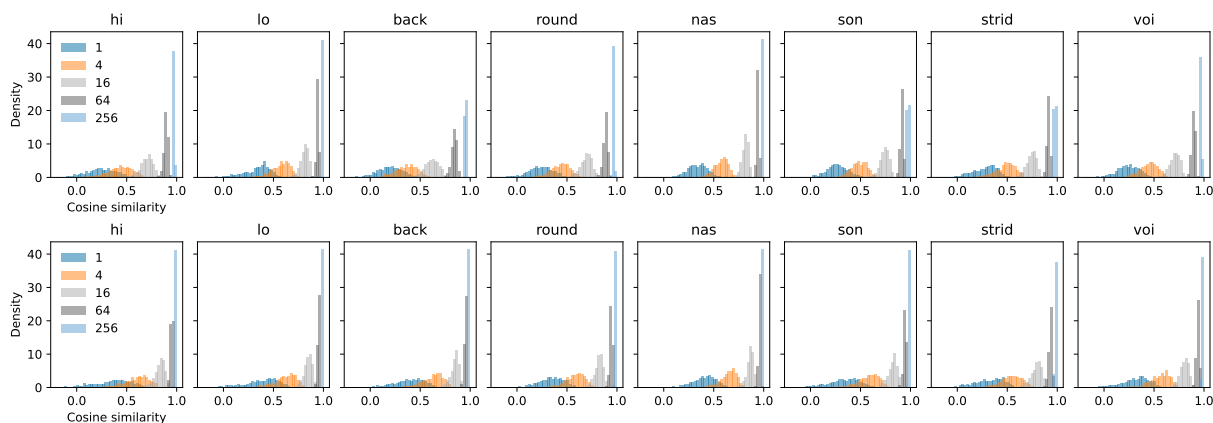


Figure 20: Approximating various phonological vectors on TIMIT (upper) and VoxAngeles (lower) with different sample size $N = \{1, 4, 16, 64, 256\}$. Each histogram depicts the distribution of cosine similarities.

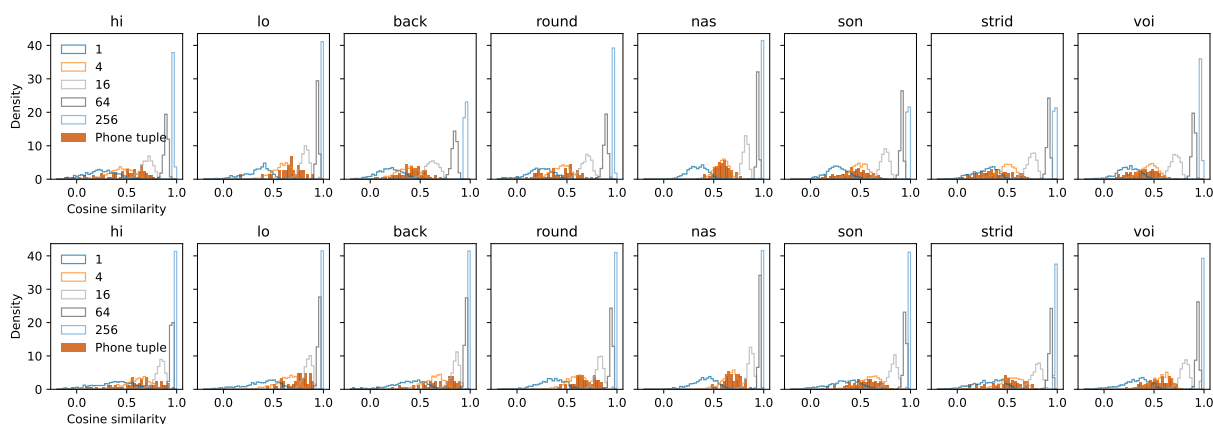


Figure 21: Approximating various phonological vectors on TIMIT (upper) and VoxAngeles (lower) using a fixed phone pair. Each histogram depicts the distribution of cosine similarities. Figure 20 is overlaid for comparison.

phones in a pair often differ in multiple phonological features, causing the resulting vector to capture mixed contrasts.

B.9 Comparing similarities between phonological vectors

B.9.1 Settings

We investigate the relationships among different phonological vectors by measuring their pairwise cosine similarities. Specifically, we compute cosine similarities between the eight phonological vectors defined in § 4.1.4.

B.9.2 Results

Figures 22 and 23 visualize the pairwise cosine similarities between phonological vectors for TIMIT and VoxAngeles. We mainly discuss cases with absolute cosine similarity values that are greater than 0.5. Across both datasets, we observe several phonologically interpretable patterns. First of all, vowel-related and consonant-related phonological

vectors exhibit near-orthogonal similarities.

For vowels, the high and low vectors show strong negative cosine similarity, reflecting their opposing acoustic properties. Additionally, for VoxAngeles, round vector show positive similarity with high and back, and negative with low. It is acoustically consistent with the rounding, which lowers both formants, F1 and F2. High vowels are characterized by low F1, low vowels by higher F1, and back vowels by lower F2.

For consonants, nasal, sonorant, and voice vectors exhibit positive similarity, where nasals are always sonorants, and sonorants are almost always voiced. Also, strident and sonorant vectors show negative similarity, where stridents are not sonorant and vice versa.

Overall, these results demonstrate that cosine similarities between extracted phonological vectors capture meaningful phonological relationships.




Figure 22: Cosine similarities between different phonological vectors drawn from TIMIT.

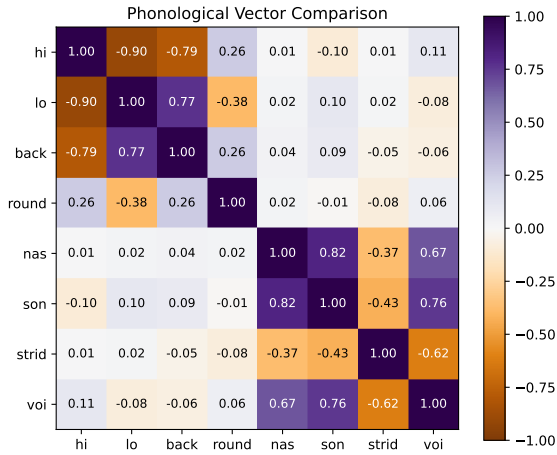


Figure 23: Cosine similarities between different phonological vectors drawn from VoxAngeles.

B.10 Resynthesis works for unseen languages

B.10.1 Settings

Following § 4.2.1, we use the same settings of § 4.1 by using representations from the final layer of WavLM. As mentioned in § 4.1.2, we construct train-test splits by randomly selecting languages for evaluation. For resynthesis, as described in § 4.1.3, we use the vocoder trained on FLEURS-R, a multilingual speech dataset.

B.10.2 Results

A comparison between Figure 26 and Figure 4 shows that VoxAngeles with FLEURS-R-trained vocoder exhibits trends that are nearly identical to those observed on TIMIT with LibriTTS-trained vocoder. These include monotonic relationships between acoustic measurements and the phonological vector scale with the smooth interpolation and ex-

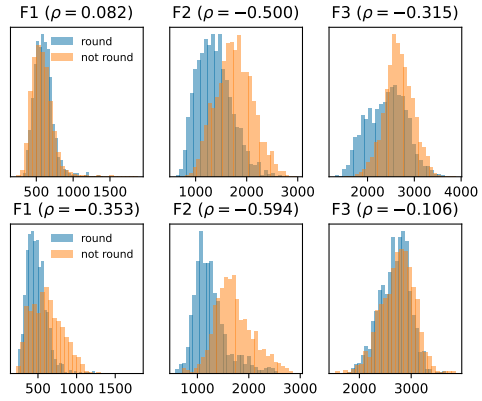


Figure 24: Acoustic measurements of F1, F2, and F3 for rounded (blue) and unrounded (orange) vowels on TIMIT (upper) and VoxAngeles (lower). ρ indicates Spearman's rank correlation coefficient between the roundness and the acoustic measurements.

trapolation behavior. The close correspondence between seen and unseen language settings indicates that WavLM representations generalize effectively to unseen languages through phonological vectors.

B.11 Case study: vowel rounding

B.11.1 Settings

To examine how a phonological vector influences multiple acoustic measurements, we analyze the first three formants (F1, F2, F3) for the rounding vector. Lip rounding lengthens the front cavity of the vocal tract, which generally lowers all formants, while the specific behavior of each formant varies by language (Ladefoged, 1996; Ladefoged and Disner, 2012).

To compare the behavior of the phonological vector with actual acoustic measurements, we first measure the formant distributions of rounded and unrounded vowels in the datasets. We then examine whether the strength of the phonological vector λ reflects these dataset-specific patterns.

B.11.2 Results

We first show the histograms of formant measurements of the vowels from TIMIT and VoxAngeles in Figure 24. For TIMIT, F2 and F3 differences are more prominent than F1, whereas for VoxAngeles the differences are more pronounced in F1 and F2 than in F3.

We then examine how formants vary with the strength of the phonological vector λ in Figure 25. The same dataset-specific patterns emerge: the weakest correlation appears in F1 for TIMIT and in F3 for VoxAngeles. This suggests that the round-




Figure 25: Comparing the phonological vector weight λ with acoustic measurements on TIMIT (upper) VoxAngeles (lower) using WavLM. We observe three measurements, F1, F2, and F3, for the round vector. ρ indicates Spearman’s rank correlation coefficient.

ing phonological vector captures dataset-specific phonetic realizations in a data-driven manner.

B.12 Acoustic measurements remain stable after resynthesis

B.12.1 Settings

To assess the stability of the trained vocoder and the acoustic measurements, we compare the acoustic measurements computed from the original waveform \mathbf{x} and from the resynthesized speech $\tilde{\mathbf{x}} = f^{-1}(f(\mathbf{x}))$. We set $\lambda = 0$, corresponding to identity resynthesis without modifying the WavLM representations.

B.12.2 Results

We visualize the differences in acoustic measurements using density plots in Figure 27. Across all tested consonants and vowels, the distributions are

tightly concentrated around zero, indicating that the resynthesis and acoustic measurement pipeline leads to reliable analyses.

B.13 Phonological vectors from MFCCs are ineffective for resynthesis

B.13.1 Settings

Following § 4.2.1 and § B.10, we adopt the same settings of § 4.1. We evaluate both audio-sliced and feature-sliced MFCC representations, as § B.2 demonstrated that audio slicing can improve the effectiveness of MFCCs.

B.13.2 Results

Figures 28 and 29 show the results for audio-sliced and feature-sliced MFCCs, respectively. Across both slicing strategies, the acoustic measurements exhibit little to no correlation with the scale of the corresponding phonological vectors. The only exceptions are weak correlations observed for back vowels by the audio-sliced MFCCs and stridents by the feature-sliced MFCCs. However, these effects are inconsistent, particularly given that no corresponding correlations are observed on TIMIT. Moreover, observed correlations are substantially weaker than that of WavLM, as shown in Figures 4 and 26. Taken together, these results indicate that MFCC-derived representations are ineffective as phonological vectors for synthesis experiments.

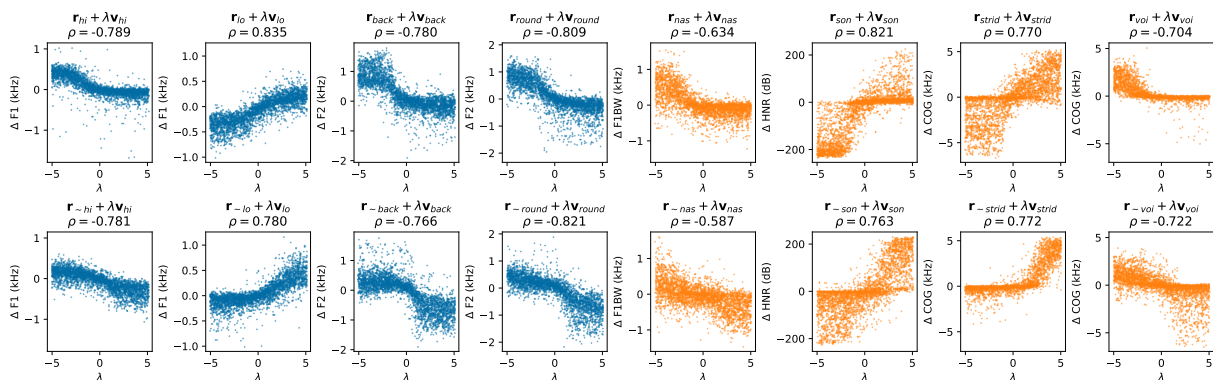


Figure 26: Comparing the phonological vector weight λ with acoustic measurements on VoxAneles using WavLM. ρ indicates Spearman’s rank correlation coefficient. Blue and orange plots indicate vowels and consonants, respectively.

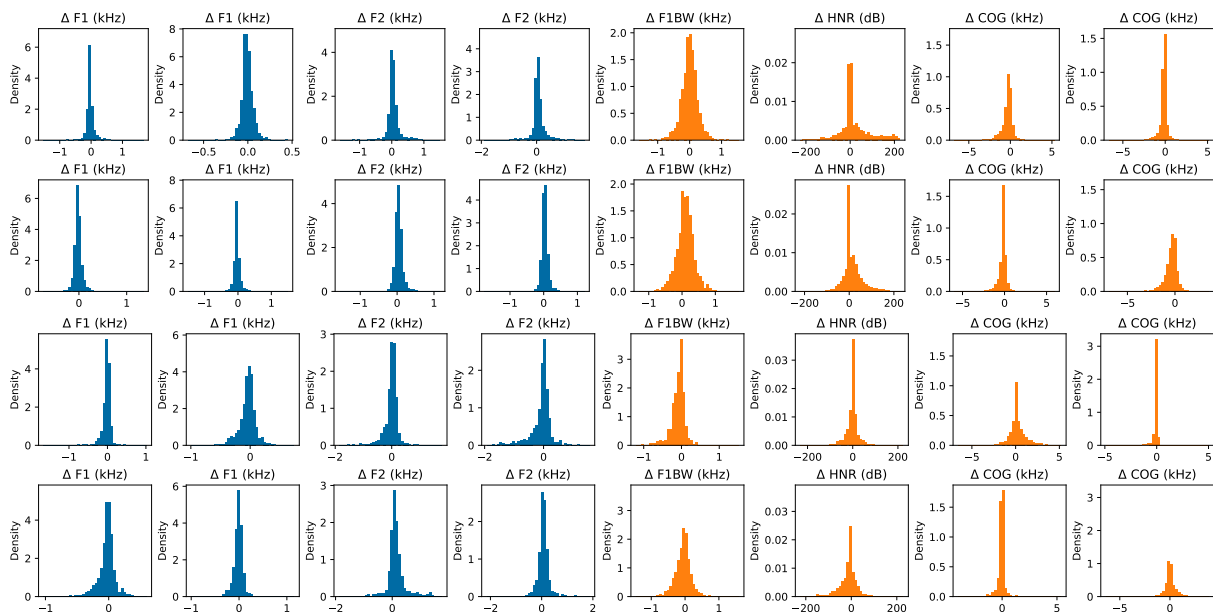


Figure 27: Comparing the acoustic measurements of original and synthesized speech ($\lambda = 0$) on TIMIT (upper two rows) and VoxAneles (lower two rows). We use the same range of x-axis from y-axis in Figures 4 and 26. We observe that the differences are highly centralized to zero, ensuring the stability of resynthesis through the vocoder.

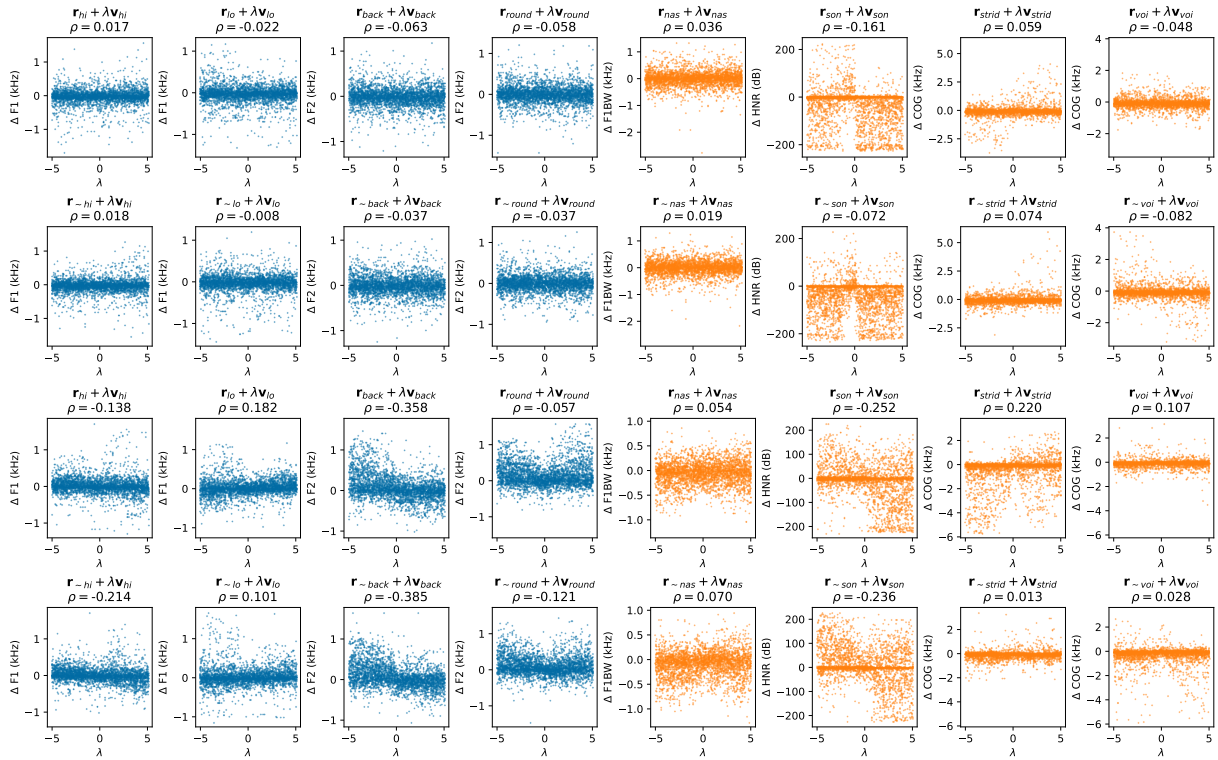


Figure 28: Comparing the phonological vector weight λ with acoustic measurements on TIMIT (upper two rows) and VoxAdeles (lower two rows) using phonological vectors from audio-sliced MFCC. ρ indicates Spearman's rank correlation coefficient. Blue and orange plots indicate vowels and consonants, respectively. There is little to no controllability, with the exception of weak correlation on back vowel on VoxAdeles.

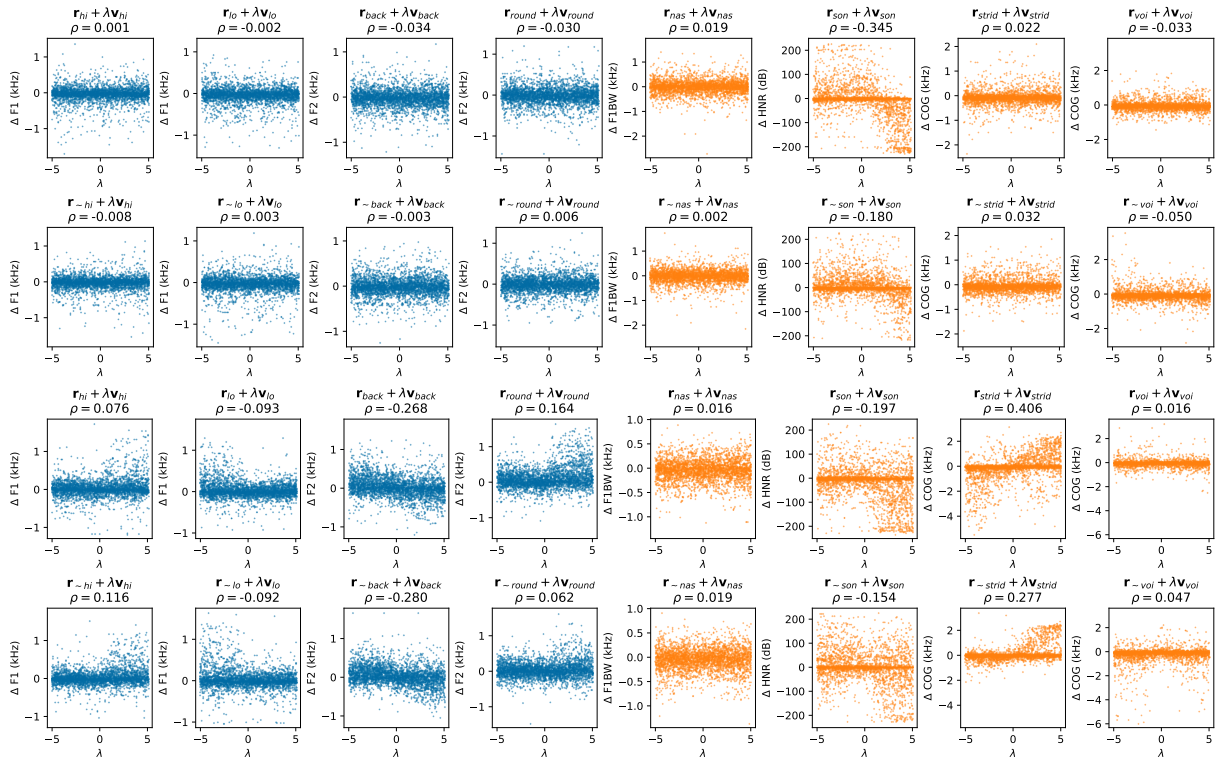


Figure 29: Comparing the phonological vector weight λ with acoustic measurements on TIMIT (upper two rows) and VoxAdeles (lower two rows) using phonological vectors from feature-sliced MFCC. ρ indicates Spearman's rank correlation coefficient. Blue and orange plots indicate vowels and consonants, respectively. Similar to Figure 28, there is little to no controllability.