

# Normalization through Fine-tuning: Understanding Wav2vec2.0 Embeddings for Phonetic Analysis

Yiming Wang and Jiahong Yuan

University of Science and Technology of China, China

E-mail: wangyiming@mail.ustc.edu.cn, jiahongyuan@ustc.edu.cn Tel: +86-551-63606825

**Abstract**—Phonetic normalization is essential in speech recognition and analysis, ensuring the comparability of features extracted from raw audio data. In the current paradigm of fine-tuning pre-trained large transformer models, however, it is not treated as a necessary preprocessing step; rather, it is implicitly carried out within the models themselves. This study examines how normalization emerges in transformer-based speech models, with a focus on Wav2vec2.0. Understanding this process is a crucial step toward effectively leveraging model embeddings as phonetic features for speech analysis. Our results show that fine-tuning Wav2vec2.0 can achieve phonetic normalization by selectively suppressing irrelevant information, such as speaker sex, while enhancing task-relevant information, such as tones and finals. In multi-task learning, however, the embeddings of fine-tuned models retain information for each task without compromising performance, suggesting that suppressing task-irrelevant information is not strictly required for effective classification. These results demonstrate that phonetic normalization can be flexibly achieved in speech models, providing insights into potential parallels in human speech perception.

## I. INTRODUCTION

Transformer-based speech models, such as Wav2vec2.0, have shown remarkable success in phonetics-related tasks, including phoneme recognition, tone recognition, and speaker identification. To leverage Wav2vec2.0 embeddings for phonetic analysis, it is essential to understand how these embeddings encode information, particularly with regard to phonetic normalization. For example, when analyzing the acoustic differences among lexical tones to uncover inherent tonal features in a language, speaker-specific variations, e.g., females have a higher pitch than males, must be normalized to allow for meaningful comparisons of tonal features across speakers.

Probing experiments have demonstrated that the embeddings of pre-trained Wav2vec2.0 models encode various types of information, with different layers capturing different aspects. Our study takes this a step further by investigating whether and how fine-tuning Wav2vec2.0 models for a specific task can suppress information irrelevant to that task. Moreover, we examine how the embeddings differ when the model is fine-tuned for multiple tasks compared to a single task.

Our main findings are summarized as follows:

- Phonetic normalization is achieved through the fine-tuning of Wav2vec2.0. The process, from the initial to the final layers of a fine-tuned Transformer, varies depending

on the type of information: for tones and finals, the information is initially enhanced and subsequently suppressed, while for speaker sex, the information is progressively suppressed through the layers. The observed differences can likely be attributed to the way information is encoded across the various layers of the pretrained Wav2vec2.0 model.

- The embeddings from a model fine-tuned for multiple tasks, such as tone and sex classification, effectively retain information of both tone and sex. The effectiveness of these embeddings is comparable to those fine-tuned for each task separately.
- For models fine-tuned for a single task, all task-irrelevant information is suppressed. For instance, the embeddings of a model fine-tuned for tone recognition suppresses both sex and final information. This is not an obvious outcome from a machine learning perspective, as embeddings containing both tone and sex information perform equally well for optimizing the objective of tone classification.
- The combination of results 2 and 3 have an important implication for phonetic analysis. Phonetic normalization does not necessarily require the suppression of task-irrelevant information. For example, sex or final information can be either suppressed or preserved in features that effectively distinguish tones. This finding provides new insights into how phonetic normalization can be flexibly achieved in speech models and how it is realized in human speech perception.

## II. RELATED WORK

### A. Normalization

Normalization in human speech perception has long been studied [1]. [2] investigated the effect of F0 on the perception of vowels using synthetic sounds, showing that the perceptual category boundary between vowels shifted when the F0 was doubled. [3] found that listeners can determine where the pitch of an isolated vowel, produced by a stranger, falls within that stranger's pitch range, without prior familiarity or context. [4] used high-density electrocorticography to record neural population activity while participants listened to sentences that varied in intonational pitch contour, phonetic content, and speaker. They found that the cortical representation of intonation relies on relative-pitch encoding, rather than absolute-pitch encoding.

---

This work was supported in part by USTC grant KJRW-202302.

In phonetic research, numerous vowel normalization methods have been proposed in the literature [5], [6]. Regarding pitch and tone normalization, commonly utilized techniques include z-scores, semitones, and the 5-point scale [7], [8]. In speech technology, techniques such as speaker adaptive training (SAT), vocal tract length normalization (VTLN), and Maximum Likelihood Linear Regression (MLLR) are extensively employed for speaker normalization and adaptation [9].

In deep learning, a variety of normalization techniques have been developed to speed up training and mitigate issues related to vanishing and exploding gradients. These include batch normalization, layer normalization, weight normalization, and spectral normalization, among others. It's important to note that these methods do not incorporate explicit speaker or phonetic normalization as discussed above.

### B. Wav2vec2.0

Wav2vec2.0 is a self-supervised learning framework based on Transformers that can learn speech representations from raw audio data [10]. The framework processes the speech signal with a multilayer convolutional network to extract latent features, which are then fed into vector quantization and transformer networks. It is pre-trained on a noise contrastive binary classification task, which enables the model to capture a rich amount of information about speech, as demonstrated by probing experiments that show their effectiveness on a wide range of tasks [11]–[13].

The pre-trained models of Wav2vec2.0 can be fine-tuned for speech recognition using labeled data [10]. [14] fine-tuned Wav2vec2.0 with a CTC loss to recognize suprasegmentals such as syllables, tones, and pitch accents. Their results demonstrated a 50% error reduction in Mandarin tone recognition compared to previous studies. [15] proposed fine-tuning Wav2vec2.0 with a cross-entropy loss to classify tones in an utterance on a frame-by-frame basis, which not only improves tone classification accuracy but also generates frame-level representations suitable for tonal analysis.

Using features extracted from a pre-trained Wav2vec2.0 model for phoneme recognition, [16] discovered that the features from the first 10 layers as well as the last few layers provide very poor performance, while layers 15-19 achieve the best phoneme recognition accuracy. In this study, we will investigate the performance of features extracted from various layers of fine-tuned Wav2vec2.0 models across multiple tasks.

## III. DATA AND METHOD

### A. Data

Our experiments were conducted on the Aishell-1 dataset [17], which is widely used for Mandarin ASR. It consists of 165 hours of read speech and word transcriptions in Mandarin Chinese from 400 speakers. The dataset is divided into train, dev, and test sets with 150, 10, and 5 hours of speech, respectively.

### B. Forced Alignment

We trained an HMM-GMM based forced aligner on the dataset using the HTK toolkit<sup>1</sup> and the pronouncing dictionary included in the dataset. The dictionary transcribes words into initials and tonal finals in Pinyin, a Roman alphabet system of phonetic transcription for Mandarin Chinese.

From forced alignment, we assigned a label to each frame in an utterance. The central frame within the boundary of a tone or phone received the corresponding label whereas all other frames were labeled as 'O', a special category that will be ignored in computing cross-entropy loss in fine-tuning. The framerate, set at 20ms, matches the output framerate of Wav2vec2.0.

### C. Fine-tuning Wav2vec2.0 for framewise classification

Pairs of audio and frame-aligned labels in the training set were used to fine-tune a Wav2vec2.0 large model pre-trained on 960 hours of Librispeech audio (libri960\_big.pt) for framewise classification. In this task, every frame in an utterance is classified into a label. In a conventional approach to framewise classification (e.g. [18]), the features of a frame are locally extracted and the frames are largely independent of each other. In our approach, features of all frames in an utterance are interdependent due to the self-attention mechanism in the transformer model.

During training, a randomly initialized linear projection is added on top of the contextual representations of Wav2vec2.0 to map the representations into label tokens. The labels belong to a set of  $C$  categories, excluding a specific category 'O', which is ignored in the loss computation. The network's predictions  $\hat{y}_{i,c}$  represent the probability that token  $i$  belongs to category  $c$ . The true labels are represented by a one-hot encoded vector  $y_{i,c}$ , where  $y_{i,c} = 1$  indicates that token  $i$  belongs to category  $c$ , and  $y_{i,c} = 0$  otherwise.

The entire network is then optimized by minimizing a cross-entropy loss as described in Equation 1, with the category of 'O' being excluded during the loss calculation. Specifically, for each token  $i$  and each category  $c \in \{1, 2, \dots, C\}, c \neq O$ , the cross-entropy loss is computed as:

$$\text{Cross-Entropy}(y, \hat{y}) = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1, c \neq O}^C y_{i,c} \log(\hat{y}_{i,c}) \quad (1)$$

where  $N$  is the number of tokens in the dataset,  $y_{i,c}$  is the true label for token  $i$ , and  $\hat{y}_{i,c}$  is the model's predicted probability for token  $i$  being in category  $c$ . The experiment was carried out on the Fairseq platform<sup>2</sup>.

Initially, only the output classifier is trained for the first 10k updates, after which the Transformer is updated as well. We set the max tokens to 1 million, which is equivalent to 62.5 seconds of audio with a sampling rate of 16 kHz, and the learning rate to  $1e-5$ . The total number of updates is set to 100k.

<sup>1</sup><https://htk.eng.cam.ac.uk/>

<sup>2</sup><https://github.com/facebookresearch/fairseq>

During inference, each audio utterance was fed forward through the fine-tuned model. The model output at the central frame of every phone or tone in the test set was compared with the reference label to calculate the classification accuracy. The features from each of the 24 Transformer layers were extracted at the central frame of every phone or tone in the test set for analysis and visualization. These features have a dimension of 1024.

It's worth noting that while during training only one frame of each tone or phone has a target label and all other frames are labeled as category 'O', which is ignored in computing loss, during inference no frames are classified as 'O', and in most cases, all frames within the boundary of a phone or tone get the same classification. Examples of training and prediction labels are listed below:

- *Training labels:* O O O O O T1 O O O O O T5 O ...
- *Predictions:* T1 T1 T1 T1 T1 T1 T1 T1 T1 T5 T5 T5 T5 ...

#### D. Multitask Training for Simultaneous Encoding of Sex, Tone, and Final

In order to test whether the model can simultaneously encode sex, tone, and final, we employed a multitask training approach. During training, the model was designed to perform two classification tasks, denoted as task1 and task2. Each task contributes a loss term, with  $L_1$  and  $L_2$  representing the losses for task1 and task2, respectively. The final loss function used for optimization is given by:

$$L = L_1 + L_2$$

This combined loss function allows the model to optimize both tasks simultaneously. By employing this multitask learning strategy, we observed the performance of the model across both classification tasks to evaluate its ability to encode the multiple attributes of sex, tone, and final in a single unified representation.

## IV. RESULTS

### A. Classification accuracy

TABLE I  
ACCURACY FOR MONO-TASK AND MUTI-TASK TRAINED MODEL

	t	st	f	sf	tf
Tone accuracy	95.95%	95.42%	-	-	95.15%
Final accuracy	-	-	98.37%	98.42%	97.00%

Using the method described above, we fine-tuned Wav2vec2.0 for tone, phone, and sex classification, respectively. To test whether different tasks can be encoded simultaneously, we also trained the combinations of the tasks which are sex&tone(st), sex&final(sf), and tone&final(tf). There are five tones in total, labeled T1 to T5 (with T5 representing the neutral tone), and 39 finals. The classification accuracy on the test set is reported in Tabel I, demonstrating

the excellent performance of the fine-tuned models on these tasks. A plausible explanation is that, although only a single frame of each tone or phone is utilized to train classifiers, its features encapsulate rich contextual information learned through the self-attention mechanism in Transformer. Meanwhile, the similar performance between the multi-task model and mono-task model shows that sex and tone/final can be simultaneously encoded within a shared feature space.

### B. Normalization

To observe how normalization is performed in the model fine-tuned for different tasks, we extracted features from all Transformer layers of the finetuned and pre-trained models, compared these features. We reduced the features to 100 dimensions with PCA and computed correlation through Singular Vector Canonical Correlation Analysis for Deep Learning Dynamics and Interpretability (SVCCA)[19], [20], and utilized UMAP for visualization. The features were extracted at the central frame of all tones/finals in the test set.

In the SVCCA analysis, we computed the correlation between the extracted features and tones, sex, and finals, respectively. A high correlation indicates a strong similarity between the features and the corresponding labels.

Figure 1 shows the SVCCA correlation of features extracted from the pre-trained model at every layer of the Transformer. A couple of interesting observations can be made. Firstly, the features from the pre-trained model exhibit superior performance in sex correlation compared to tone and phone. The sex identification was conducted at the segment level rather than the utterance level. This means that the model must classify each segment of tone/final into a sex, which is much more complex than using a whole utterance to compute. Secondly, features from different layers demonstrate varying performance. Specifically, for sex correlation, the lower and upper layers are higher than the middle layers, whereas for tone and final correlation, the middle layers are higher.

Figure 2 compares the normalization processes under different training objectives. For the task targeted during fine-tuning, the SVCCA correlations exhibit an upward trend from the middle layers to the final layers, indicating that this part of the network is primarily influenced by the fine-tuning task. In contrast, in the earlier layers, sex correlation remains high while tone and final correlations are lower. This pattern reflects the behavior of the initial pre-trained model, where sex characteristics dominate before the model is adjusted by fine-tuning. These observations highlight how different training tasks affect the model's layer-wise representations and the varying degrees of normalization that occur throughout the network. Additionally, Figure 3 illustrates how the SVCCA correlations for non-target tasks change across the model layers. We observe that these correlations are gradually normalized after the middle layers, indicating that the model increasingly focuses on the target task as fine-tuning progresses, while non-target task features are progressively suppressed.

Features undergo transformations across different layers, resulting in new representations. This transformation process




Fig. 1. SVCCA correlations on features extracted from different layers of pre-trained model across three different tasks




Fig. 2. SVCCA correlations on features extracted from models optimized for different target tasks




Fig. 3. SVCCA correlations on features extracted from models optimized for non-target tasks

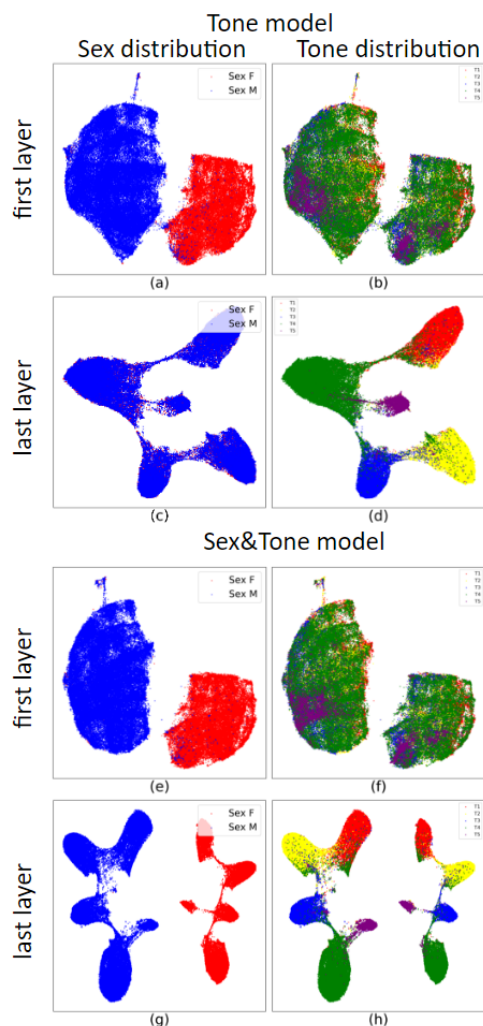


Fig. 4. UMAP visualization of embeddings extracted from the first and last layers of the tone and tone&sex models.

is observed as a decrease in the correlation of SVCCA when features are modified in ways that do not benefit downstream tasks. Conversely, if the transformations render the features useful, they are re-extracted in later layers, indicated by an increase in SVCCA correlations. Features that are irrelevant to the task at hand are normalized and discarded.

The sex normalization in the model fine-tuned for tone classification is also evident in the UMAP<sup>3</sup> visualization of the features. Figure 4 display the features extracted from the first and last Transformer layers of two fine-tuned models: one trained on the tone classification task (a,b,c,d) and the other on the sex&tone classification task (e,f,g,h). In the visualizations on the left, features are colored based on sex, while on the right, they are colored based on tone. In the first layer, both models divide the features into two groups, corresponding to male and female. However, in the last layer, an interesting difference emerges: the tone-only model mixes the features by sex, while the tone and sex model continues to separate them. Despite this, both models perform well on the tone classification task. This observation suggests that there are

multiple ways for a speech model to encode different tasks effectively.

## V. CONCLUSION

Transformer-based speech models have achieved remarkable success in automatic speech recognition. However, their potential for speech analysis — particularly how to effectively leverage their embeddings — remains underexplored. This study investigates how model embeddings encode information in the speech signal, with a focus on phonetic normalization. Understanding this process is a critical step toward effectively utilizing these embeddings for speech analysis.

Compared with the pretrained Wav2vec2.0 model, fine-tuned models encode more task-specific information in their learned representations. Our results show that fine-tuning Wav2vec2.0 can effectively achieve phonetic normalization by selectively suppressing irrelevant information, such as speaker sex, while enhancing task-relevant information, such as tones and finals. In multi-task learning, however, the embeddings of

<sup>3</sup><https://umap-learn.readthedocs.io/en/latest/>

fine-tuned models preserve information for each task without compromising performance, suggesting that suppressing task-irrelevant information is not strictly necessary for effective classification.

These results provide new insights into the flexible mechanisms by which phonetic normalization can be achieved in speech models and how similar processes may operate in human speech perception.

#### REFERENCES

- [1] K. Johnson and M. J. Sjerps, "Speaker normalization in speech perception," *The Handbook of Speech Perception, Second Edition*, pp. 145–176, 2021.
- [2] R. L. Miller, "Auditory tests with synthetic vowels," *The Journal of the Acoustical Society of America*, vol. 25, pp. 114–121, 1953.
- [3] D. N. Honorof and D. H. Whalen, "Perception of pitch location within a speaker's f0 range.," *The Journal of the Acoustical Society of America*, vol. 117, pp. 2193–2200, 2005.
- [4] C. Tang, L. S. Hamilton, and E. F. Chang, "Intonational speech prosody encoding in the human auditory cortex," *Science*, vol. 357, pp. 797–801, 2017.
- [5] P. Adank, R. Smits, and R. Hout, "A comparison of vowel normalization procedures for language variation research," *The Journal of the Acoustical Society of America*, vol. 116, pp. 3099–3107, 2004.
- [6] C. Voeten, W. Heeringa, and H. Velde, "Normalization of nonlinearly time-dynamic vowels," *The Journal of the Acoustical Society of America*, vol. 152, pp. 2692–2710, 2022.
- [7] P. Rose, "Considerations in the normalisation of the fundamental frequency of linguistic tone," *Speech Communication*, pp. 343–352, 1987.
- [8] J. Zhang, "A comparison of tone normalization methods for language variation research," in *32nd Pacific Asia Conference on Language, Information and Computation*, Hong Kong, Dec. 2018, pp. 823–831.
- [9] D. Giuliani, M. Gerosa, and F. Brugnara, "Improved automatic speech recognition through speaker normalization," *Computer Speech & Language*, pp. 107–123, 2006.
- [10] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "Wav2vec 2.0: A framework for self-supervised learning of speech representations," *arXiv preprint arXiv:2006.11477*, 2020.
- [11] J. Shah, Y. K. Singla, C. Chen, and R. R. Shah, "What all do audio transformer models hear? probing acoustic representations for language delivery and its structure," *ArXiv:2101.00387*, 2021.
- [12] D. Ma, N. Ryant, and M. Liberman, "Probing acoustic representations for phonetic properties," *Proceedings of ICASSP 2021*, 2021.
- [13] M. Bartelds, W. de Vries, F. Sanal, C. Richter, M. Liberman, and M. Wieling, "Neural representations for modeling variation in speech," *Journal of Phonetics*, vol. 92, 101137: 1–17, 2022.
- [14] J. Yuan, N. Ryant, X. Cai, K. Church, and M. Liberman, "Automatic recognition of suprasegmentals in speech," *ArXiv:2108.01122*, 2021.
- [15] J. Yuan, X. Cai, and K. Church, "Improved contextualized speech representations for tonal analysis," in *Proceedings of Interspeech 2023*, Aug. 2023, pp. 4513–4517.
- [16] A. Baevski, W.-N. Hsu, A. Conneau, and M. Auli, "Unsupervised speech recognition," in *NeurIPS 2021*, Dec. 2021, pp. 27 826–27 839.
- [17] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *Proceedings of O-COCOSDA 2017*, 2017.
- [18] N. Ryant, M. Slaney, M. Liberman, E. Shriberg, and J. Yuan, "Ighly accurate mandarin tone classification in the absence of pitch information," in *Proceedings of Speech Prosody 2014*, 2014, pp. 673–677.
- [19] M. Raghu, J. Gilmer, J. Yosinski, and J. Sohl-Dickstein, "Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability," *Advances in Neural Information Processing Systems*, vol. 30, pp. 6076–6085, 2017. [Online]. Available: <http://papers.nips.cc/paper/7188-svcca-singular-vector-canonical-correlation-analysis-for-deep-learning-dynamics-and-interpretability.pdf>.
- [20] A. Morcos, M. Raghu, and S. Bengio, "Insights on representational similarity in neural networks with canonical correlation," *Advances in Neural Information Processing Systems*, vol. 31, pp. 5732–5741, 2018. [Online]. Available: <http://papers.nips.cc/paper/7815-insights-on-representational-similarity-in-neural-networks-with-canonical-correlation.pdf>.