

Identifying Speaker Information in Feed-Forward Layers of Self-Supervised Speech Transformers

Tzu-Quan Lin* and Hsi-Chun Cheng* and Hung-yi Lee* and Hao Tang†

* Graduate Institute of Communication Engineering, National Taiwan University, Taiwan

E-mail: tzuquanlin@gmail.com, leo5470@gmail.com, tlkagkb93901106@gmail.com

† Centre of Speech Technology Research, University of Edinburgh, United Kingdom

E-mail: hao.tang@ed.ac.uk

Abstract—In recent years, the impact of self-supervised speech Transformers has extended to speaker-related applications. However, little research has explored how these models encode speaker information. In this work, we address this gap by identifying neurons in the feed-forward layers that are correlated with speaker information. Specifically, we analyze neurons associated with k-means clusters of self-supervised features and i-vectors. Our analysis reveals that these clusters correspond to broad phonetic and gender classes, making them suitable for identifying neurons that represent speakers. By protecting these neurons during pruning, we can significantly preserve performance on speaker-related task, demonstrating their crucial role in encoding speaker information.

I. INTRODUCTION

Self-supervised Transformers have enjoyed a great success in learning speech representations that can be applied to a wide range of tasks [9, 21, 23, 16, 2, 5, 22]. Their main application has largely been on automatic speech recognition, but the impact of self-supervised Transformers has extended to speaker-related tasks [11, 30, 6], rivaling the state of the art [27]. Despite the successes, few (except [8] and [28]) have analyzed how speaker information is encoded in these models. Even in [8] and [28], the analyses have been limited to probing classifiers. There is generally a lack of studies and tools for analyzing speaker information in self-supervised speech Transformers.

A recent study [19] has shown that compressing the feed-forward layers (also known as FFNs) in Transformers is particularly detrimental for speaker tasks, hinting that speaker information is more prominent at feed-forward layers compared to other parts of the Transformer. The finding provides us an approach to studying speaker information through pruning. In addition, the interpretation of feed-forward layers as key-value memory [14] has proven useful for analyzing the feed-forward layers in Transformers. In particular, given a property of interest (such as phones, gender, or pitch), one can identify a set of neurons in FFNs that is responsible for representing a property [20]. The combination of pruning and identifying property neurons provides a novel approach to analyzing feed-forward layers in Transformers.

In this work, we study speaker information in self-supervised Transformers by identifying property neurons related to speakers. We take the aforementioned pruning methodology as our evaluation: a set of neurons are deemed important

if pruning them leads to severe degradation in performance. In other words, the performance should be largely preserved if the set of important neurons are *protected* during pruning. If we are able to identify a set of neurons that does not hurt a speaker task much when protected during pruning, then that set of neurons are the ones where the speaker information is prominent.

There are a few challenges when adopting this methodology. The first challenge is that as the number of classes within a property increases, the number of identified neurons decreases. As we will see later in the experiments, this seems to be an inherent drawback of the approach proposed in [20]. Given the large number of speakers in a typical speech dataset, using speaker IDs directly would result in very few property neurons being identified. A subsequent challenge is to decide what alternative property to use for identifying speaker-related neurons, especially under the constraint that speaker labels are unavailable.

To address this, we opt for the classes discovered by k-means (with a small k) on self-supervised features and i-vectors [10]. While x-vectors [29] offer stronger performance, they require supervised training. In contrast, i-vectors are label-free, aligning better with our unsupervised pipeline. Inspired by prior works [25, 24] showing that phonetic information can aid the learning of speaker embeddings, we hypothesize that identifying speaker-related neurons requires both phonetic and speaker information. Accordingly, we use both SSL and i-vector-based clusters, which we expect to correspond to broad phonetic and speaker classes, respectively. This unsupervised design allows us to identify speaker-relevant neurons without relying on any labels. An added benefit of this approach is that no labels are needed when identifying neurons with the k-means clusters.

Our experimental results show that by protecting the neurons discovered using clusters on self-supervised features and i-vectors, we can significantly preserve speaker information when pruning FFNs and improve pruning effectiveness. Our approach, without requiring labels, rivals the approach proposed in [20] which does use phone and gender labels. Our approach works especially well in the one-shot pruning setting. We are able to prune about 70% of the total parameters in FFN in a single step without causing significant performance degradation.




Fig. 1. An example showing the set relationships of SSL cluster neurons (left), i-vector cluster neurons (middle), and one of the possible alternatives for identifying i-vector cluster neurons (right).

II. METHODOLOGY

In this section, we review the approach proposed in [20] and discuss the extension made in this work. At the high level, to find a set of neurons, we need to decide what it means for a neuron to be activated and how it correlates with a property.

Whether a neuron is activated can be determined by looking at individual frames. For a frame $v \in \mathbb{R}^d$, a neuron is considered activated when its value is among the top $\lambda\%$ in the frame. Formally, a frame is binarized into

$$b(v) = [\mathbb{1}_{v_1 > \text{top } \lambda\%(v)} \cdots \mathbb{1}_{v_D > \text{top } \lambda\%(v)}]^\top, \quad (1)$$

where λ is a hyperparameter. This binary vector is called the activation pattern in [20]. Once the activation is determined, we can correlate the activations with the properties of interest.

A. SSL Cluster Neurons

As argued in the introduction, we run k-means on the self-supervised features to obtain clusters and correlate the activation patterns with the clusters. Specifically, for a cluster c , we define G_c^{ssl} to be the set of neurons

$$\{d : p([b(v)]_d = 1 \mid \text{frame } v \text{ is in cluster } c) > \rho\%\}, \quad (2)$$

where ρ is a thresholding hyperparameter. The probability is computed with simple counts of co-occurrences. The classes used in [20] are phones and the discovered neurons form clusters of vowels, consonants, and semi-vowels. We will later in the experiments show that clusters of self-supervised features discovered by k-means already form a good separation of these broad categories.

Similar to the findings in [20], we find that certain neurons are correlated with speech, and get activated with all the clusters. We follow [20], selecting the non-overlapping set among different clusters

$$P_{ssl} = \bigcup_{c=1}^{k_{ssl}} G_c^{ssl} - \bigcup_{1 \leq c < j \leq k_{ssl}} G_c^{ssl} \cap G_j^{ssl}, \quad (3)$$

where k_{ssl} is the number of SSL clusters. We refer to the set of neurons P_{ssl} as SSL cluster neurons.

B. i-vector Cluster Neurons

The approach of how neurons are discovered is general and can be easily extended by conditioning them on other properties of speech. Since our focus is on analyzing the speaker information, we further cluster the i-vectors of utterances and label all the frames in an utterance with their i-vector cluster ID. This leads to further conditioning of the neurons, and we define the set of neurons $G_{g,c}^{ive}$ for an i-vector cluster g and SSL cluster c as

$$\left\{ d : p([b(v)]_d = 1 \mid \text{frame } v \text{ is labeled i-vector cluster } g, \text{ and SSL cluster } c) > \rho\% \right\}. \quad (4)$$

To obtain the neurons conditioned just on i-vectors, the approach taken by [20] is taking the intersection of all SSL clusters

$$G_g^{ive} = \bigcap_{c=1}^{k_{ssl}} G_{g,c}^{ive}. \quad (5)$$

Given the formal definitions, we can now see that as k_{ssl} increases, the discovered neurons naturally decreases.

It is not entirely clear why the neurons need to be conditioned on SSL clusters and then later taking the intersection. There are a few alternatives that can achieve a similar need.

Alternative A We can take the union

$$A_g^{ive} = \bigcup_{c=1}^{k_{ssl}} G_{g,c}^{ive}. \quad (6)$$

instead of intersection in Eq.5.

Alternative B We can drop the conditioning on SSL clusters altogether and define

$$B_g^{ive} = \left\{ d : p([b(v)]_d = 1 \mid \text{frame } v \text{ is labeled i-vector cluster } g) > \rho\% \right\}. \quad (7)$$

Alternative C We can follow Alternative B but further intersect SSL cluster neurons outside the probability

$$C_g^{\text{ive}} = B_g^{\text{ive}} \cap \bigcap_{c=1}^{k_{\text{ssl}}} G_c^{\text{ssl}}. \quad (8)$$

All of the above alternatives provide a set of neurons related to i-vector clusters. Regardless of how the neurons related to i-vectors are discovered, the last step is still to find the non-overlapping ones

$$P_{\text{ive}} = \bigcup_{g=1}^{k_{\text{ive}}} G_g^{\text{ive}} - \bigcup_{1 \leq g < j \leq k_{\text{ive}}} G_g^{\text{ive}} \cap G_j^{\text{ive}}, \quad (9)$$

where k_{ive} is the number of i-vector clusters. Here, G_g^{ive} can be replaced with any of the three alternatives mentioned above. It is not immediately clear why we would choose one or the other, and we will look into these alternatives in the experiments. Figure 1 illustrates how the neurons are identified and their set relations.

C. Protecting Neurons while Pruning

There are two feed-forward layers in a Transformer block. We prune the 3072 hidden dimensions after the first feed-forward layer by removing rows from the first layer and the corresponding columns in the second layer. Following [15], we prune the dimensions (i.e., neurons) that have small ℓ_1 -norm of the weights. Given the set of neurons we discovered in previous sections, we can decide to keep (or protect) them even when they should have been pruned.

We choose two pruning approaches as our evaluation: one-shot pruning [17] and iterative pruning [15, 12]. For both pruning approaches, we have a target sparsity to achieve. One-shot pruning prunes the neurons that have low ℓ_1 of the weights to the target sparsity in one step. Iterative pruning, however, iterates between pruning and fine-tuning, using a schedule to gradually prune until the sparsity reaches the target.

III. EXPERIMENTAL SETTING

In the experiments, we will use the methodology outlined in the introduction, identifying neurons that are potentially important for encoding speakers and later protecting them during pruning. We use the SUPERB benchmark [32], specifically speaker identification, to probe the speaker information of self-supervised speech Transformers. In addition to speaker identification (SID), we will also look at phoneme recognition (PR), emotion recognition (ER), and intent classification (IC). The self-supervised models we study are MelHuBERT [18] and wav2vec 2.0 [3], where MelHuBERT takes mel spectrograms as input and wav2vec 2.0 takes raw waveform samples. Our method is also applicable to other models such as HuBERT [16] and WavLM [5], but their pre-training targets are not publicly released, making continued pre-training and pruning difficult to evaluate.

For identifying the neurons, we set $\lambda = 1\%$ as the threshold to compute the activation patterns and $\rho = 1\%$ for correlating activation patterns with clusters. We run k-means on the

9th layer of MelHuBERT to obtain the SSL clusters. Recall that we hypothesize in the introduction that speaker-related neurons may rely on both phonetic and speaker information. We therefore intentionally select the 9th layer, as it produces phonetically prominent clusters. We use the same frame-wise cluster assignments for both MelHuBERT and wav2vec 2.0 features.

For iterative pruning, we prune 128 dimensions every 25,000 steps and continue pretraining the model after each pruning step using a learning rate of 1e-5. We use a batch size of 4 for MelHuBERT and 12 for wav2vec 2.0. Pruning continues until 512 hidden dimensions remain in the FFNs. For one-shot pruning, we first identify the SSL cluster neurons and i-vector cluster neurons in the unpruned model, then prune all other neurons in a single step. Since the number of neurons identified vary across layers, different layers may retain different numbers of dimensions after one-shot pruning. To ensure a fair comparison, we prune the model to the same average number of dimensions across pruning methods. Specifically, the model is pruned to an average of 867 dimensions per layer for MelHuBERT and 780 dimensions per layer for wav2vec 2.0.

IV. RESULTS

Below we will present three sets of results. We will first study the k-means clusters and how they capture broad phone classes. Next, we will present the comparison between protecting the neurons we discovered and regular pruning. Additionally, we also compare to the approach in [20], as our approach is largely derived from theirs. Finally, we will study the impact of the hyperparameters, in particular, how the number of clusters and the different alternatives of discovering neurons impact pruning.

A. Interpreting SSL and i-vector Clusters

For the SSL clusters, given their rich phonetic information [32, 31, 26, 7] and their main applications in automatic speech recognition [3, 16, 1, 13], we hypothesize the clusters are likely to be correlated with broad phone classes. Figure 2 shows the result of running k-means on MelHuBERT features with $k_{\text{ssl}} = 3$. We find that the first cluster predominantly corresponds to vowels, the second includes vowels and voiced consonants, and the third mainly contains voiced and unvoiced consonants. This result is very much aligned with the clustering results shown in [20], where vowels and consonants are separated with semi-vowels inbetween the two. For the i-vector clusters, when $k_{\text{ive}} = 2$, they closely match the gender classes. The first cluster has a 97.23% purity of being male, while the other cluster has a purity of 89.82% of being female. This again aligns well with the gender results in [20].

The decision to use a small number of clusters has been discussed in previous sections. The main issue is that the number of discovered neurons decreases significantly when using a large number of clusters, as is shown in Eq. 5. We will study this problem empirically later, and this seems to be an inherent problem of the approach in [20].

TABLE I

THE RESULTS FOR ITERATIVE AND ONE-SHOT PRUNING ON MELHUBERT AND WAV2VEC 2.0. WE USE “\ CLUSTER NEURONS” TO DENOTE PRUNING WHILE PROTECTING THE CLUSTER NEURONS IN THIS WORK, AND “\ PROPERTY NEURONS” TO DENOTE PRUNING WHILE PROTECTING PROPERTY NEURONS [20]. ITERATIVE AND ONE-SHOT PRUNING RESULTS ARE NOT DIRECTLY COMPARABLE DUE TO DIFFERENCES IN MODEL SPARSITY. NOTE THAT LIN ET AL. [20] USE LABELS FOR IDENTIFYING PROPERTY NEURONS, WHILE OURS IS PURELY UNSUPERVISED.

Model	Pruning Method	SID (ACC%) \uparrow		PR (PER%) \downarrow		ER (ACC%) \uparrow		IC (ACC%) \uparrow	
		Iterative	One-Shot	Iterative	One-Shot	Iterative	One-Shot	Iterative	One-Shot
MelHuBERT	unpruned	63.96		8.17		57.52		72.24	
	regular pruning	51.04	54.08	12.28	9.11	54.96	55.92	59.32	67.65
	\ cluster neurons	52.66	62.96	12.22	9.11	57.52	58.67	68.81	82.64
	\ property neurons [20]	54.10	58.27	10.8	9.00	57.11	54.82	71.13	68.78
wav2vec 2.0	unpruned	75.18		5.74		63.43		92.35	
	regular pruning	57.41	63.42	14.54	11.17	59.34	60.77	89.45	91.03
	\ cluster neurons	58.34	70.67	14.34	11.26	59.43	62.89	94.06	95.65
	\ property neurons [20]	57.35	66.05	14.11	9.70	61.39	61.28	91.22	91.88




Fig. 2. Amounts of broad phone classes within the clusters produced by running k-means on self-supervised features with $k_{ssl} = 3$.

B. Protecting Neurons during Pruning

We present our main results in Table I. While all tasks suffer after pruning, speaker identification (SID) suffers the most by as much as 23.6% relative for wav2vec 2.0 after iterative pruning. Protecting both cluster neurons and property neurons help recover the performance loss due to pruning. However, protecting the cluster neurons is on par with protecting property neurons, while being significantly better in one-shot pruning, recovering most of the performance loss.

Despite not being our goal, protecting cluster neurons also brings certain benefits for emotion recognition (ER) and intent classification (IC). In one-shot pruning, our approach can even get achieve a better performance than the unpruned results. We suspect that the unpruned results suffer a mild overfitting, and pruning can serve as a regularization and brings improvement [4]. Protecting cluster neurons has little effect on phone recognition (PR).

Overall, we are able to prune about 70% of the parameters with one-shot pruning while maintaining most of the performance for speaker identification.

C. Analysis and Ablation Study

To gain further understanding of how the hyperparameters affect the identified neurons and the pruning performance, in this section, we conduct analysis and ablation studies for the proposed approach. For simplicity and consistency, all experiments here are done on MelHuBERT with iterative

pruning.

Effect of Different Numbers of Clusters. We first analyze the impact of the number of clusters k_{ssl} used in k-means. The results are shown in the first and second rows of Table II. When k_{ssl} is set to 39 or 100, the performance across the four downstream tasks degrades significantly. As we have alluded to in the introduction, we attribute this to the problem of not being able to identify enough neurons to protect. Figure 3 shows how the number of identified neurons changes when we use different cluster sizes, and we can see that when the number of clusters is large, few neurons are identified except for the first and the last layers.

Effect of Removing SSL or i-vector Cluster Neurons. Next, we evaluate the impact of not protecting either SSL cluster neurons or i-vector cluster neurons. The results are presented in the third and fourth rows of Table II. We observe that both types of neurons are essential for maintaining performance, as the removal of either negatively affects the results. Interestingly, SSL cluster neurons appear to be more critical than i-vector cluster neurons for SID. This finding aligns with previous research suggesting that phonetic information is useful for the learning of speaker embeddings [25, 24].

Alternatives to Identifying i-vector Cluster Neurons. Finally, we explore the alternatives of identifying i-vector cluster neurons described in Section II-B. The results, presented in the fifth, sixth, and seventh rows of Table II, show that conditioning and then intersecting outperforms the other alternatives in three out of four tasks, with IC as the only exception. Overall, conditioning and then intersecting shows better performance than other alternatives. We again attribute this to the number of neurons identified. As shown in the bottom part of Figure 3, the i-vector cluster neurons identified by conditioning and intersecting are fewer compared to those from Alternative A and Alternative B, allowing for more efficient pruning that avoids protecting too many neurons. In contrast, Alternative C fails to identify enough number of i-vector cluster neurons in the middle layers.

TABLE II

ANALYSIS AND ABLATION STUDY FOR THE PROPOSED APPROACH. THE SUBSCRIPT DENOTES THE DIFFERENCES WHEN CHANGING ONE OF THE DESIGN DECISIONS COMPARING TO THE PROPOSED APPROACH.

	SID (ACC%) \uparrow	PR (PER%) \downarrow	ER (ACC%) \uparrow	IC (ACC%) \uparrow
$k_{ssl} = 100$	49.19 _{-3.47}	14.20 _{+1.98}	55.36 _{-2.16}	60.61 _{-8.20}
$k_{ssl} = 39$	49.98 _{-2.68}	14.26 _{+2.04}	54.94 _{-2.58}	60.79 _{-8.02}
without SSL clusters	48.84 _{-3.82}	14.15 _{+1.93}	55.29 _{-2.23}	64.62 _{-4.19}
without i-vector clusters	49.86 _{-2.80}	13.89 _{+1.67}	57.23 _{-0.29}	69.62 _{+0.81}
Alternative A	51.84 _{-0.82}	13.78 _{+1.56}	56.17 _{-1.35}	69.89 _{+1.08}
Alternative B	51.21 _{-1.45}	13.48 _{+1.26}	56.63 _{-0.89}	71.55 _{+2.74}
Alternative C	50.44 _{-2.22}	13.79 _{+1.57}	55.27 _{-2.25}	62.53 _{-6.28}
Proposed Method	52.66	12.22	57.52	68.81




Fig. 3. The number of i-vector cluster neurons in different layers for different number of SSL clusters k_{ssl} (top) and different alternatives to identify i-vector cluster neurons (bottom).

V. CONCLUSION

In this work, we introduce the concept of SSL cluster neurons and i-vector cluster neurons. Our approach better retains performance when these neurons are protected during pruning. Notably, it is particularly effective for one-shot pruning, even surpassing the unpruned results in some cases. Additionally, we conduct analysis and ablation studies on the neuron identification process, addressing gaps in prior research. Overall, these cluster neurons play a significant role in encoding speaker information.

REFERENCES

- [1] Alexei Baevski and Abdelrahman Mohamed. “Effectiveness of Self-Supervised Pre-Training for ASR”. In: *ICASSP*. IEEE, 2020, pp. 7694–7698.
- [2] Alexei Baevski et al. “data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language”. In: *ICML*. 2022.
- [3] Alexei Baevski et al. “wav2vec 2.0: A framework for self-supervised learning of speech representations”. In: *Advances in Neural Information Processing Systems*. Vol. 33. Curran Associates, Inc., 2020, pp. 12449–12460.
- [4] Yuheng Bu et al. “Information-Theoretic Understanding of Population Risk Improvement with Model Compression”. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 34.4 (2020), pp. 3300–3307.
- [5] Sanyuan Chen et al. “WavLM: Large-scale self-supervised pre-training for full stack speech processing”. In: *IEEE Journal of Selected Topics in Signal Processing* 16.6 (2022), pp. 1505–1518.
- [6] Zhengyang Chen et al. “Large-scale self-supervised speech representation learning for automatic speaker verification”. In: *ICASSP*. IEEE, 2022, pp. 6147–6151.
- [7] Kwanghee Choi et al. “Self-Supervised Speech Representations are More Phonetic than Semantic”. In: *Interspeech*. 2024, pp. 4578–4582.
- [8] Shammur Absar Chowdhury, Nadir Durrani, and Ahmed Ali. “What do end-to-end speech models learn about speaker, language and channel information? A layer-wise and neuron-level analysis”. In: *Computer Speech & Language* 83 (2024), p. 101539.
- [9] Yu-An Chung et al. “An unsupervised autoregressive model for speech representation learning”. In: *Interspeech*. 2019, pp. 146–150.
- [10] Najim Dehak et al. “Front-End Factor Analysis for Speaker Verification”. In: *IEEE Transactions on Audio, Speech, and Language Processing* 19.4 (2011), pp. 788–798.
- [11] Zhiyun Fan et al. “Exploring wav2vec 2.0 on speaker verification and language identification”. In: *Interspeech*. 2021, pp. 1509–1513.
- [12] Jonathan Frankle and Michael Carbin. “The lottery ticket hypothesis: Finding sparse, trainable neural networks”. In: *ICLR*. 2018.
- [13] Changfeng Gao et al. “Self-supervised pre-training for attention-based encoder-decoder asr model”. In: *TASLP* 30 (2022), pp. 1763–1774.
- [14] Mor Geva et al. “Transformer feed-forward layers are key-value memories”. In: *Proceedings of the 2021 Con-*

- ference on *Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2021, pp. 5484–5495.
- [15] Song Han et al. “Learning both Weights and Connections for Efficient Neural Network”. In: *Advances in Neural Information Processing Systems*. Vol. 28. Curran Associates, Inc., 2015.
- [16] Wei-Ning Hsu et al. “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units”. In: *IEEE/ACM transactions on audio, speech, and language processing* 29 (2021), pp. 3451–3460.
- [17] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. “SNIP: Single-shot Network Pruning based on Connection Sensitivity”. In: *ICLR*. 2019.
- [18] Tzu-Quan Lin, Hung-yi Lee, and Hao Tang. “MelHuBERT: A simplified HuBERT on Mel spectrograms”. In: *ASRU*. 2023, pp. 1–8.
- [19] Tzu-Quan Lin et al. “Compressing transformer-based self-supervised models for speech processing”. In: *arXiv preprint arXiv:2211.09949* (2022).
- [20] Tzu-Quan Lin et al. “Property Neurons in Self-Supervised Speech Transformers”. In: *SLT*. 2024, pp. 401–408.
- [21] Shaoshi Ling and Yuzong Liu. “DeCoAR 2.0: Deep Contextualized Acoustic Representations with Vector Quantization”. In: *arXiv preprint arXiv:2012.06659* (2020).
- [22] Alexander H Liu et al. “DinoSR: Self-Distillation and Online Clustering for Self-supervised Speech Representation Learning”. In: *Advances in Neural Information Processing Systems*. Vol. 36. Curran Associates, Inc., 2024.
- [23] Andy T Liu, Shang-Wen Li, and Hung-yi Lee. “TERA: Self-Supervised Learning of Transformer Encoder Representation for Speech”. In: *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021), pp. 2351–2366.
- [24] Yi Liu et al. “Introducing phonetic information to speaker embedding for speaker verification”. In: *EURASIP Journal on Audio, Speech, and Music Processing* (2019).
- [25] Yi Liu et al. “Speaker embedding extraction with phonetic information”. In: *Interspeech* (2018), pp. 2247–2251.
- [26] Kinan Martin et al. “Probing self-supervised speech models for phonetic and phonemic information: a case study in aspiration”. In: *Interspeech*. 2023, pp. 251–255.
- [27] Victor Miara, Theo Lepage, and Reda Dehak. “Towards Supervised Performance on Speaker Verification with Self-Supervised Learning by Leveraging Large-Scale ASR Models”. In: *Interspeech*. 2024, pp. 2660–2664.
- [28] Mukhtar Mohamed et al. “Orthogonality and isotropy of speaker and phonetic information in self-supervised speech representations”. In: *Interspeech*. 2024, pp. 3625–3629.
- [29] David Snyder et al. “X-vectors: Robust dnn embeddings for speaker recognition”. In: *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE. 2018, pp. 5329–5333.
- [30] Yingzhi Wang, Abdelmoumene Boumadane, and Abdelwahab Heba. “A fine-tuned wav2vec 2.0/HuBERT benchmark for speech emotion recognition, speaker verification and spoken language understanding”. In: *arXiv:2111.02735* (2022).
- [31] Dan Wells, Hao Tang, and Korin Richmond. “Phonetic analysis of self-supervised representations of english speech”. In: *Interspeech*. 2022, pp. 3583–3587.
- [32] Shu-wen Yang et al. “SUPERB: Speech processing Universal PERFORMANCE Benchmark”. In: *Interspeech*. 2021, pp. 1194–1198.