

What Do Self-Supervised Speech Models Know About Words?

Ankita Pasad Chung-Ming Chien Shane Settle Karen Livescu

Toyota Technological Institute at Chicago, USA

{ankitap, cmchien, settle.shane, klivescu}@ttic.edu

Abstract

Many self-supervised speech models (S3Ms) have been introduced over the last few years, improving performance and data efficiency on various speech tasks. However, these empirical successes alone do not give a complete picture of what is learned during pre-training. Recent work has begun analyzing how S3Ms encode certain properties, such as phonetic and speaker information, but we still lack a proper understanding of knowledge encoded at the word level and beyond. In this work, we use lightweight analysis methods to study segment-level linguistic properties—word identity, boundaries, pronunciation, syntactic features, and semantic features—encoded in S3Ms. We present a comparative study of layer-wise representations from ten S3Ms and find that (i) the frame-level representations within each word segment are not all equally informative, and (ii) the pre-training objective and model size heavily influence the accessibility and distribution of linguistic information across layers. We also find that on several tasks—word discrimination, word segmentation, and semantic sentence similarity—S3Ms trained with visual grounding outperform their speech-only counterparts. Finally, our task-based analyses demonstrate improved performance on word segmentation and acoustic word discrimination while using simpler methods than prior work.¹

1 Introduction

Self-supervised speech models (S3Ms) are effective across a variety of applications (Mohamed et al., 2022; Yang et al., 2021), including lower-level tasks such as speaker identification (Chen et al., 2022) and speech recognition (Baevski et al., 2020; Hsu et al., 2021c) as well as more linguistically complex spoken language understanding tasks (Ashihara et al., 2023; Pasad et al., 2022; Shon et al., 2022; Tsai et al., 2022;

Wu et al., 2023). However, downstream task performance alone does not reveal what knowledge is learned during pre-training and where it is encoded.

Recent work has begun studying the acoustic and phonetic content encoded in S3Ms (Abdullah et al., 2023; Hsu et al., 2021c; Ma et al., 2021; Pasad et al., 2021, 2023), and the findings have in turn helped guide model development and use (Baevski et al., 2022; Feng et al., 2022; Liu et al., 2023; van Niekerk et al., 2021; Pasad et al., 2021, 2023). However, S3Ms may encode higher-level linguistic information as well, since their model architectures (typically based on self-attention [Gulati et al., 2020; Vaswani et al., 2017]) use contextual information from the entire speech input. There has been little analysis thus far of the word-level information encoded in these models (Pasad et al., 2023; Sanabria et al., 2023). To fill this gap, our work addresses two key questions: (i) *how is word-related information distributed across frames within a word segment?* and (ii) *in which layers, and how well, do S3Ms encode segment-level pronunciation, syntactic, and semantic information?*

To investigate these questions, we use canonical correlation analysis, a standard lightweight analysis tool, along with unsupervised evaluations on tasks: acoustic word discrimination, word segmentation, and semantic sentence similarity. We present a comparative study across ten S3Ms differing in their pre-training objective, data modality, and model size. Some of the key findings include:

- The form of the pre-training objective affects which intermediate layers correlate the most with word-level properties. (Section 5.2.1)
- Word-identifying information is concentrated close to the center of each segment. (Section 5.1.2)
- Pre-trained representations from different S3Ms require varying complexities of

¹Codebase: <https://github.com/ankitapasad/layerwise-analysis>.

post-processing to use the encoded knowledge. (Section 5.1.1)

- The visually grounded S3Ms are better than speech-only S3Ms on several tasks: word discrimination (Section 5.1.1), segmentation (Section 5.1.3), and semantic similarity (Section 5.2.2).
- A task’s evaluation domain impacts the relative ranking of S3Ms, but the individual layer-wise trends remain domain-invariant. (Section 5.3)
- With a simple parameter-free word segmentation algorithm, S3M features outperform previous, more complex unsupervised approaches. (Section 5.1.3)
- With a single-layer RNN trained on a small amount of labeled data, S3M features achieve near-perfect word discrimination, outperforming prior work by a large margin. (Section 5.1.1)

2 Related Work

The research community has begun investigating how S3Ms encode a number of properties such as speaker identity (Chen et al., 2022; Fan et al., 2021; Feng et al., 2022; Liu et al., 2023; van Niekerk et al., 2021), para-linguistics (Li et al., 2023a; Shah et al., 2021), articulatory and prosodic features (Bannò and Matassoni, 2023; Ji et al., 2022; Kim et al., 2022), and phones (Abdullah et al., 2023; Hsu et al., 2021c; Ma et al., 2021; Pasad et al., 2021, 2023). Work on generative models based on S3Ms also suggests that some S3Ms learn phone-like sub-word units (Lakhotia et al., 2021; Nguyen et al., 2023).

Analysis on higher-level units, such as words in S3Ms, has been limited. Pasad et al. (2023) analyzed the extent to which different layers of S3Ms encode word identity, and Sanabria et al. (2023) found that pooled S3M representations over word segments perform competitively on the task of acoustic word discrimination. These results indicate that S3Ms encode word-identifying information. However, it is not clear from prior work how word information is distributed across frames and what aspects of words are encoded, such as their pronunciation, syntactic properties, or semantic properties; our work addresses this gap.

Task-specific probing classifiers have been a common analysis tool for speech models (Belinkov and Glass, 2019; Palaskar et al., 2019; Prasad and Jyothi, 2020) including S3Ms (Baeviski et al., 2021; Hsu et al., 2021c; Ma et al., 2021; Shah et al., 2021; Shen et al., 2023). While these probes provide an intuitive evaluation measure, the design decisions involved in training task-specific classifiers have confounding effects, making the scores hard to interpret (Belinkov, 2022; Hewitt and Liang, 2019; Ravichander et al., 2021). We use canonical correlation analysis (CCA) (Hotelling, 1936) and training-free task-based evaluation, thus bypassing the dependence on task-specific classifiers.

CCA and its more robust variants have been previously used to compare representations within and across neural networks (Raghu et al., 2017; Kornblith et al., 2019), to study text representation models (Saphra and Lopez, 2019; Tsvetkov et al., 2016; Voita et al., 2019), and more recently for the analysis of S3M representations (Li et al., 2023a; Pasad et al., 2021, 2023; Yang et al., 2023b). While classifiers require discrete labels, CCA can be used with both discrete and continuous-valued labels. CCA is also computationally inexpensive and has a closed-form solution.

Word similarity (WordSim) tasks (Faruqui and Dyer, 2014) have been commonly used for intrinsic evaluation of word vectors. However, previous work has observed some problematic aspects of these tasks (Faruqui et al., 2016), including that WordSim performance is not well-correlated with extrinsic evaluation, whereas CCA-based evaluation more consistently tracks downstream task performance (Tsvetkov et al., 2016). Our work also shares some motivation with the Zero Resource Speech Benchmark (Nguyen et al., 2020), but tasks in this benchmark require encoding isolated word segments and/or discretizing the representations. Our analyses use word segment representations extracted in context, in order to match the most common use cases of S3Ms.

The task of acoustic word discrimination (AWD) (Carlin et al., 2011) has been commonly used to evaluate segment-level acoustic word embeddings, using both unsupervised (Algayres et al., 2020; Levin et al., 2013; Peng et al., 2020; Van Staden and Kamper, 2021) and supervised models (Algayres et al., 2020; He et al., 2017; Kamper et al., 2016; Settle and Livescu, 2016). To our knowledge, only Sanabria et al. (2023) and

Van Staden and Kamper (2021) have studied the use of S3Ms for generating acoustic word embeddings for AWD. Van Staden and Kamper (2021) used ‘‘first-generation’’ S3Ms while Sanabria et al. (2023) analyzed two of the same modern S3Ms we study here. Our work provides a more comprehensive study of using multiple approaches (unsupervised mean pooling, dynamic time warping, and supervised models using S3M features), as well as a new state-of-the-art for one of the most commonly used AWD benchmarks.

Word unit discovery and segmentation are common benchmark tasks that have also been used to study speech representations (Algayres et al., 2022; Bhati et al., 2021; Cuervo et al., 2022; Dunbar et al., 2020; Nguyen et al., 2022; Sanabria et al., 2021; ten Bosch and Cranen, 2007). Previous work studying the segmentation capabilities of S3Ms includes Kamper (2022), based on first discovering phone-like units and then using them to discover word segments, and Peng and Harwath (2022b), based on thresholding the attention map of a visually grounded speech model (VG-HuBERT, which we also use here). Our study complements this prior work by comparing the layer-wise performance of a large number of S3Ms for this task and showing that a simple, training-free segmentation algorithm performs very competitively.

Textual sentence similarity is a classic task in NLP (Conneau and Kiela, 2018), but there are only a few studies investigating *spoken* utterance representations for this task (Merkx et al., 2021, 2023; Zhu et al., 2022). Some downstream tasks in the SUPERB benchmark (Yang et al., 2021) successfully use spoken utterance representations from frozen S3Ms, represented by a single mean-pooled vector. We complement our understanding of the capabilities of pooled utterance representations by performing a broad cross-model comparison.

3 Analysis Methods

We extract frame-level and span-level representations from all layers of each S3M. We use CCA to compare word segment representations with various linguistic vectors (Section 3.1) and also investigate encoded properties using training-free approaches for several tasks: (i) acoustic word discrimination (Section 3.2), (ii) word segmentation (Section 3.3), and (iii) sentence-level semantic similarity (Section 3.4).

Linguistic property	Attribute vector (dimension)
word identity	one-hot embeddings (500)
word pronunciation	acoustically grounded word embeddings (128)
part-of-speech tags	Attributes derived from PTB ¹ (45)
semantic attributes	Attributes derived from SemCor ² (41)

Table 1: *Linguistic properties that we compare to S3M representations via CCA.*

3.1 CCA-based Analysis

CCA (Hotelling, 1936) is a statistical technique that measures the relationship between two continuous-valued random vectors by evaluating the maximum correlations between their linear projections. CCA takes as input n pairs of vectors $\{(x_1, y_1), \dots, (x_n, y_n)\}$, sampled from the random vectors (or ‘‘views’’) $X \in \mathbb{R}^{d_1}, Y \in \mathbb{R}^{d_2}$, and returns *canonical correlations*, a correlation-based measure of similarity between the two views. First, we solve for the directions of maximum correlation between linear projections of X and Y : $v_1, w_1 = \arg \max_{v,w} \text{corr}(v^T X, w^T Y)$. The subsequent directions $v_i, w_i \forall i \in [2, \min(d_1, d_2)]$ maximize the same correlation subject to each new projection being uncorrelated with others in the same view. This problem has a closed-form solution requiring one singular vector decomposition.

We use projection-weighted CCA (PWCCA) (Morcos et al., 2018), a robust CCA variant commonly used in recent analysis studies (Voita et al., 2019; Pasad et al., 2021, 2023; Yang et al., 2023b). The value of a PWCCA score lies between 0 and 1. Specifically, we use PWCCA to compare word segment representations with various linguistic vectors (Table 1):

Word Identity. We measure how well S3M representations encode word identity by comparing them with word IDs. For CCA computation, we follow Pasad et al.’s (2023) approach and convert the discrete word IDs to one-hot vectors. We use this analysis to examine the location of word information within the word segment.

²<https://github.com/ytsvetko/qvec/tree/master/oracles>.

Acoustically Grounded Word Embeddings.

Acoustically grounded word embeddings (AGWEs) are written word embeddings trained jointly with acoustic word embeddings (AWE), i.e., representations of spoken word segments (He et al., 2017; Hu et al., 2020; Settle et al., 2019). A contrastive learning objective jointly optimizes AWE and AGWE models such that AWE and AGWE of the same word are closer than for pairs of different words. We use AGWEs obtained from joint AWE+AGWE training on the LibriSpeech corpus (Panayotov et al., 2015).³ We expect CCA similarity with AGWEs to measure word-level pronunciation information encoded by the S3Ms.

Syntactic Features. Tsvetkov et al. (2016) construct syntactic vectors from the Penn Treebank (PTB) (Marcus et al., 1993). For each word, an empirical probability is calculated for each of the 45 part-of-speech (POS) tags based on frequencies in the tagged corpus. This results in 45-dimensional syntactic vectors and each vector sums to 1. For example, “light” has values 0.52, 0.41, 0.05, and 0.02 for noun (*NN*), adjective (*JJ*), proper noun (*NNP*), and verb (*VB*) attributes, respectively, and zero for the rest.

Semantic Features. Tsvetkov et al. (2015) exploit word sense annotations in SemCor (Miller et al., 1993), a WordNet-annotated version of the Brown Corpus. For each word, an empirical probability is calculated for each sense attribute (26 nouns and 15 verbs), based on their frequencies in the labeled corpus. This results in 41-dimensional semantic vectors and each vector sums to 1. For instance, the vector for the word “family” has a value of 0.96 and 0.04 for *NN.GROUP* and *NN.ACT* attributes, respectively, and zero for the rest.

The resulting embedding space puts words with similar attributes closer together. For instance, the semantic vector of “family” is most similar to words with a high value for the *NN.GROUP* attribute: government, leaders, elite, platoon. This behavior differs from that of a more fine-grained distributional embedding space such as GloVe (Pennington et al., 2014) where some of the nearest neighbors for “family” are husband, father, mother, sister, and wife.

³The AGWEs used here are trained similarly to (Shi et al., 2021) and are made available by Pasad et al., (2021): <https://github.com/ankitapasad/layerwise-analysis>.

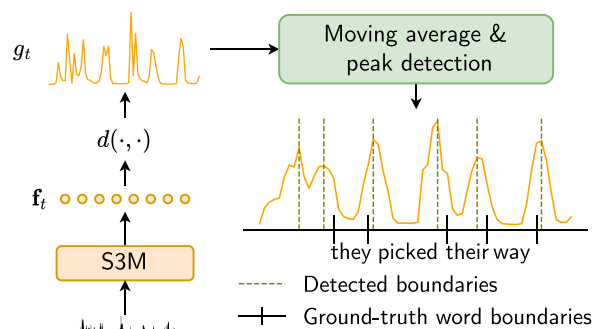


Figure 1: Our word segmentation algorithm.

We do not compare S3M representations with learned text embeddings (such as GloVe or BERT [Devlin et al., 2019] representations). Although these embeddings possibly encode richer text representations than our linguistic features, they also contain a mix of syntactic and semantic information. This would not allow us to study the syntactic and semantic features separately.

3.2 Acoustic Word Discrimination

Acoustic word discrimination (AWD) is the task of determining whether or not a pair of acoustic waveform segments ($\mathbf{X}_i, \mathbf{X}_j$) correspond to the same word (Carlin et al., 2011). A measure of dissimilarity between \mathbf{X}_i and \mathbf{X}_j is computed, and the pair is predicted to be “the same” if their dissimilarity falls below a threshold and “different” otherwise. AWD performance is reported as average precision, i.e., the area under the precision-recall curve generated by varying the threshold.

We use S3Ms for AWD in three ways. *pool-AWD* compares cosine distance after mean-pooling the frame-level features. *DTW-AWD* computes a dynamic time warping distance between segments using the cosine distance between their frame-level features. *RNN-AWD* trains a recurrent neural network on the frame-level representations, following Shi et al.’s (2021) approach but using phone sequences for supervision as in Hu et al. (2020).

3.3 Word Segmentation

We ask how well S3M representations can perform word segmentation “intrinsically”. We design a straightforward training-free algorithm to leverage the behavior of frame-level representations near word segment boundaries (see Figure 1).

Given a sentence comprising T frames, first, we extract the frame-level features \mathbf{f}_t ($1 \leq t \leq T$) from an S3M layer and perform mean and variance normalization, for each channel, across all \mathbf{f}_t 's to get the normalized $\hat{\mathbf{f}}_t$'s. Then we compute the dissimilarity between adjacent frames $d(\cdot, \cdot)$ to get $g_t = d(\hat{\mathbf{f}}_{t+1}, \hat{\mathbf{f}}_t)$, and smooth g_t with a moving average. Finally, we use a peak detection algorithm to identify adjacent frames with higher dissimilarity than the surrounding frames. While peak detection algorithms have been commonly used for phoneme (Cuervo et al., 2022; Kreuk et al., 2020; Räsänen et al., 2011) and word segmentation (Bhati et al., 2021; Cuervo et al., 2022), most prior word segmentation methods (with the exception of Peng and Harwath, 2022b) have relied on explicit training of the segmentation models while our approach does not.

The detected word boundaries are evaluated using standard metrics: precision, recall, F1-score, and R-value, using a tolerance window of 20ms following prior work (Kamper, 2022).

3.4 Sentence-level Semantic Similarity

Finally, we ask whether S3Ms encode any semantic content at the utterance level. We evaluate utterance representations from S3Ms on *spoken STS* (Merkx et al., 2021), a spoken (read) version of the popular semantic textual similarity (STS) dataset (Conneau and Kiela, 2018). STS consists of sentence pairs annotated with a semantic similarity judgment. For each utterance in a pair, we extract a sentence-level representation from an S3M layer and use cosine similarity between these representations to predict semantic similarity. We report Spearman's ρ correlation between the annotated human judgments and the predicted similarity scores.

4 Experiment Details

We present analysis for *ten S3Ms* differing in (i) pre-training objective, (ii) data modality (using either speech or image-speech pairs), and (iii) model size. The pre-trained checkpoints are obtained from publicly available sources.⁴ For

⁴*wav2vec2*: <https://github.com/facebookresearch/fairseq/tree/main/examples/wav2vec>
HuBERT: <https://github.com/facebookresearch/fairseq/tree/main/examples/hubert>
WavLM: <https://github.com/microsoft/unilm/tree/master/wavlm>
AV-HuBERT: <https://github.com/facebookresearch>

word-level analysis on LibriSpeech (Panayotov et al., 2015) (Sections 4.2 and 4.3), we use ground-truth word alignments generated by the Montreal Forced Aligner (Lugosch et al., 2019; McAuliffe et al., 2017).

4.1 Background on S3Ms

S3Ms are trained with an objective function formulated to solve a pretext task on unlabeled data. In a typical model architecture, the raw audio (or filter bank features) is first passed through convolutional layers (or a linear projection). Then, the resulting frame-level *local* features are processed through self-attention layers. The models we use have 7 convolutional (or 1 linear) and 12 or 24 transformer layers.

All the models in this work use a masking-based pretext task, using both the left and the right context to recover the masked segment (target). The target comes from either the local features (*wav2vec 2.0*⁴ [*wav2vec2*] [Baevski et al., 2020] and *FaST-VGS+* [Peng and Harwath, 2022a]) or from one of the intermediate transformer layers, represented as a discrete cluster ID (*HuBERT*⁵ [Hsu et al., 2021c], *WavLM*⁶ [Chen et al., 2022], *AV-HuBERT*⁷ [Shi et al., 2022], and *VG-HuBERT* [Peng and Harwath, 2022b]). Models of the first type are trained with a contrastive loss and the latter with a classification loss. The classification loss for *WavLM* uses cluster IDs from *HuBERT*'s intermediate layers (the same ones used in *HuBERT*'s iterative pre-training). Unlike *HuBERT*, *WavLM* augments the input data to simulate noisy/overlapped speech.

The visually grounded⁸ *FaST-VGS+* and *VG-HuBERT* models are initialized with pre-trained *wav2vec2-Base*⁹ and *HuBERT-Base*,¹⁰

av.hubert FaST-VGS+: <https://github.com/jasonppy/FaST-VGS-Family> *VG-HuBERT*: <https://github.com/jasonppy/word-discovery>.

⁵*wav2vec2* and *HuBERT Base* models are pre-trained on 960 hours LibriSpeech, and the corresponding *Large* models on 60k hours LibriLight data.

⁶*WavLM -Base* is pre-trained on 960 hours LibriSpeech and *WavLM-Large* on 94k hours consisting of LibriLight, GigaSpeech, and VoxPopuli.

⁷*AV-HuBERT* models are pre-trained on LRS3.

⁸It is arguable whether such visually grounded models are "self-supervised" since the visual signal provides a form of supervision. We include them here since they are in many ways similar to speech-only S3Ms and have similar use cases.

⁹For *FaST-VGS+*, CNN, self-attention, and cross-attention layers are added before training on SpokenCOCO.

¹⁰For *VG-HuBERT*, the top 3 layers are reinitialized before training on SpokenCOCO.

respectively, thus providing a way to isolate and analyze the effect of visual grounding. These models are trained with a cross-modal contrastive loss with an appended *CLS* token, a fixed-dimensional utterance-level representation. *AV-HuBERT* is trained on a lipreading dataset with a pre-training objective that uses multi-modal discrete units. In the case of the audio-visual models, we use only the audio branch for our analysis.

4.2 CCA Evaluation

For CCA similarity with word ID (*CCA-word*), we sample $\sim 7k$ word instances across 500 distinct words from the dev-clean subset of LibriSpeech. We represent the word segment using either a single frame, by mean-pooling across a quarter of the contiguous frames, or by mean-pooling across all the frames within the word boundaries. The single frame is sampled from one of five equidistant locations starting at the first frame. The quarter chunk of contiguous frames is extracted from one of the four quarters of the word segment.

In all other CCA experiments, we obtain word segment representations by mean-pooling across all frames within the word boundaries and compare these representations with external linguistic embedding vectors (Table 1). We sample 364k word instances across 9.9k distinct words from the LibriSpeech train-clean and train-other subsets. This sample includes the 8.6k and 4k word vocabularies from *PTB* and *SemCor* vectors, respectively.

We evaluate PWCCA with multiple data splits to avoid overfitting, using the implementation from Pasad et al. (2023).¹¹ Specifically, we run each experiment with 5 train-val-test splits. In the result figures below, we plot the mean of the 5 runs along with a shadow around the mean corresponding to the minimum and maximum values. For most result plots, the shading is not visible as there is negligible deviation among results across runs.

4.3 Acoustic Word Discrimination

We evaluate *pool-AWD* and *DTW-AWD* on “clean” and “other” partitions of the LibriSpeech development set. Our *RNN-AWD* models are trained and evaluated on Switchboard (Godfrey et al., 1992) data using the same train-dev-test

¹¹<https://github.com/ankitapasad/layerwise-analysis>.

split as prior work (Carlin et al., 2011; He et al., 2017; Jansen et al., 2013; Kamper et al., 2016) with each partition containing approximately 10k spoken word segments. We evaluate *pool-AWD* on our Switchboard dev set first to find the best layer before supervised *RNN-AWD* training. In all cases, spoken word segments are 0.5-2s in duration, and segments used for evaluation on LibriSpeech and Switchboard span 5k and 3k word vocabularies, respectively.

4.4 Word Segmentation

We consider two measures of dissimilarity between neighboring frames, Euclidean distance and cosine distance. We use a prominence-based algorithm (Virtanen et al., 2020) to detect peaks in the dissimilarity curve with a prominence value exceeding a specified threshold. For each layer in each S3M, we grid search over the choice of distance metric, prominence value threshold, and moving-average window size. We choose the best combination based on F1-scores for word boundary detection on a randomly sampled subset of the LibriSpeech dev-clean split ($\sim 2k$ utterances). We also evaluate all layers on the Buckeye (Pitt et al., 2005) validation set, and the best layer of each S3M is evaluated on the Buckeye test set.


4.5 Sentence-level Semantic Similarity

The natural speech recordings in *Spoken STS* constitute 5% (638 sentence pairs) of the original STS corpus (Merx et al., 2021). Sentences in each pair are read by four speakers, and thus, each pair has 16 speaker combinations. Each spoken sentence is represented by mean-pooling all frame-level representations from an S3M layer. For *VG-HuBERT*, we extract the utterance-level *CLS* token representation as well. As in previous work (Merx et al., 2021; Zhu et al., 2022), the predicted score for each sentence pair is the mean of the cosine similarities between their representations for all speaker combinations.


5 Findings

We present our findings in two parts. Section 5.1 investigates the spoken word knowledge learned by S3Ms, and how this information is distributed across the frames of each word segment.¹²


¹²Model trends are consistent between LibriSpeech dev-clean and dev-other results, so only dev-clean results are shown.



(a) CCA similarity between mean-pooled S3M word segment representations and word ID one-hot vectors.



(b) pool-AWD scores on LibriSpeech dev-clean.



(c) pool-AWD scores on Switchboard dev.

Figure 2: Evaluation of the word-identifying information in mean-pooled word segment representations from *Base* (left) and *Large* (right) S3Ms.

Section 5.2 looks at specific linguistic properties—pronunciation, syntactic, and semantic—of word-level and sentence-level representations.

5.1 Analysis of Frame-level Representations

We investigate the word-related information encoded by S3M layers in different frames across the word segment, specifically knowledge of word identity and word boundaries.

5.1.1 Ease of Accessing Encoded Information

Figure 2a shows layer-wise correlation scores with word ID vectors for all models.¹³ We investigate whether this word-identifying information, as evidenced by high CCA scores, is easily accessible by evaluating the word representations on AWD.

For *pool-AWD*, we see that the best and worst performing models have a large difference (Figures 2b, 2c) despite being similarly well-correlated with word ID vectors (Figure 2a).

¹³Results from all models except *VG-HuBERT* are replicated from Pasad et al. (2023).




Figure 3: *DTW-AWD* results on LibriSpeech dev-clean.

Next, we evaluate *DTW-AWD* on a subset of models (Figure 3) and find that (i) all models have a better performance than *pool-AWD* with a reduced cross-model performance gap, and (ii) the cross-model ranking is also consistent with *pool-AWD*. The cross-model performance gap is further reduced with our supervised *RNN-AWD* experiments (Table 2) and the cross-model ranking is consistent with the corresponding *pool-AWD* trends on Switchboard (Figure 2c). Our multi-view *RNN-AWD* model attains a near-perfect average

Method	AP
Multi-View RNN (He et al., 2017)	
w/ log-Mel filterbank features	0.84
w/ <i>wav2vec2-Base</i> (L8)	0.93
w/ <i>HuBERT-Base</i> (L9)	0.94
w/ <i>WavLM-Base</i> (L10)	0.95
w/ <i>WavLM-Large</i> (L20)	0.98

Table 2: *RNN-AWD* performance on the Switchboard word discrimination test set (Carlin et al., 2011); layer used specified in parentheses.

precision, significantly outperforming previous work (by >10% absolute).

These experiments suggest that some models (such as *wav2vec2*) distribute discriminative word information across frames in a way that is not easily extracted through mean-pooling and compared via cosine distance, indicating that more structured reasoning over the whole segment may be helpful, such as frame-level processing in our DTW and RNN experiments. A similar observation is made in prior work (Sanabria et al., 2023) where representing words using sub-sampled and concatenated frames instead of mean-pooling gives the most relative improvement for *wav2vec2*.

We include a detailed discussion on the layer-wise trends and effect of evaluation domain (Figures 2b, 2c) in Section 5.3.1.

5.1.2 Are All Frames Equally Informative?

Next, we analyze frame-level representations to understand how word-identifying information is distributed within word boundaries. We represent word segments either using individual frames at different locations or by pooling over frames spanning different quarters of the segment.

We measure CCA scores between word segment representations and word ID and find that frames near the center of the word segment are most informative of the word identity (Figures 4 and 5). Specifically, the single center frame and the 2nd and 3rd quarter spans are all as highly correlated with the word identity as the mean-pooled representations. These findings are consistent across all S3Ms analyzed.

We see a similar word localization trend in the AWD evaluations (Figure 5), but with a stronger bias toward the start of the segment. In particular, using the 2nd quarter span alone yields a better AP for *wav2vec2-Base* and *HuBERT-Base*. This gives




Figure 4: Correlation with word identity for *wav2vec2-Base* when using a single frame to represent a word segment.

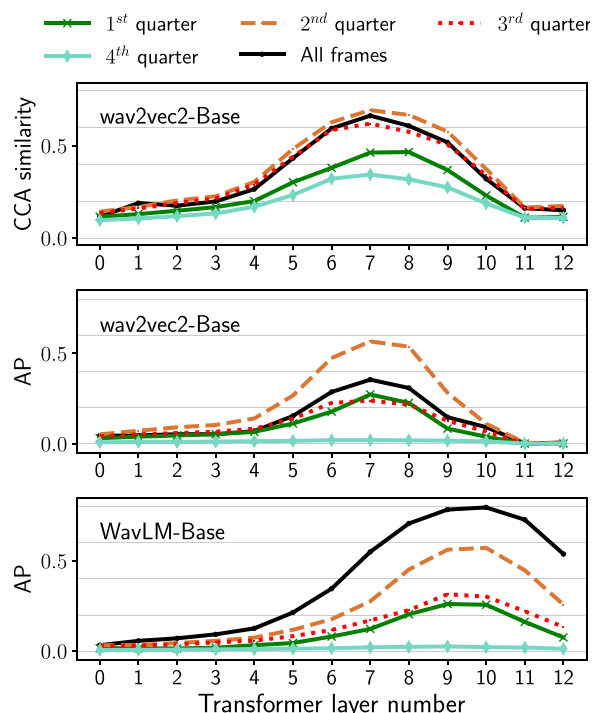


Figure 5: Correlation with word identity and AWD scores when pooling over segment quarters.

a possible explanation for their better relative performance on *DTW-AWD* and *RNN-AWD* compared to *pool-AWD* since those approaches can adjust their focus to only the most relevant frames.

5.1.3 Word Segmentation

Figure 6 shows the F1-scores of S3Ms on the word segmentation task. All of the models demonstrate non-trivial word segmentation capability.¹⁴ We observe that visually grounded models consistently outperform their speech-only

¹⁴*AV-HuBERT* is not included in this experiment as its frame rate is 40 ms, which is larger than the maximum acceptable error of 20 ms on the Buckeye word segmentation task.

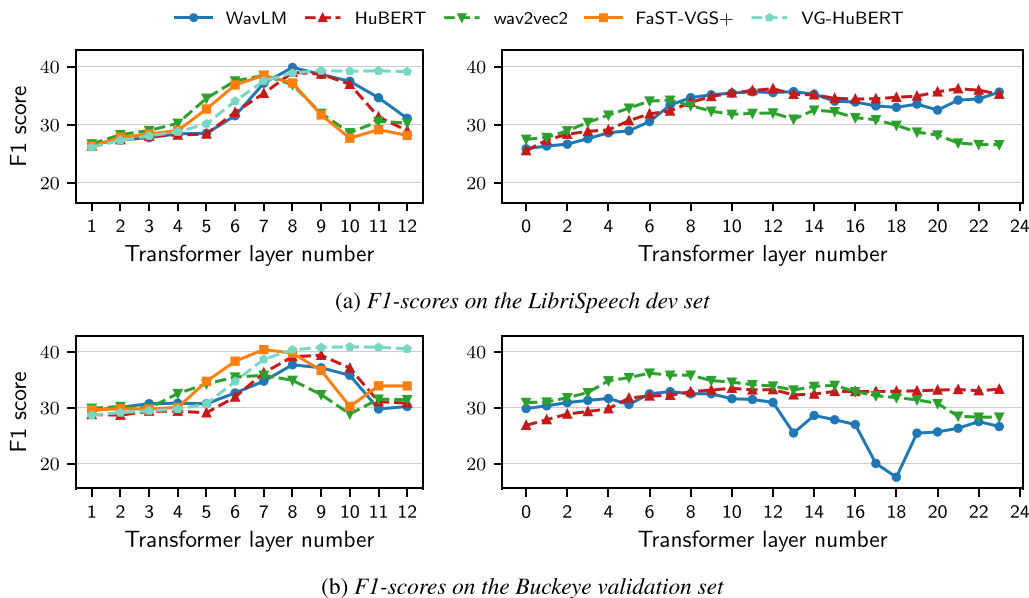


Figure 6: Unsupervised word segmentation using representations from *Base* (left) and *Large* (right) S3Ms.

counterparts, possibly because of the visual context. Further strengthening this hypothesis, we note that *VG-HuBERT* has a minimal performance drop at the final few layers, unlike other S3Ms, which can be attributed to the proximity to the cross-modal loss. *FaST-VGS+* does not see the same trend, while it is designed such that the final few layers we analyze here are only trained on the self-supervised loss but not the cross-modal loss. Similar unit discovery capabilities of visually grounded models have also been studied in prior work (Harwath and Glass, 2017; Peng and Harwath, 2022b). We include more discussion of performance comparison across S3Ms and across evaluation domain (Figures 6a, 6b) in Section 5.3.2.

In Table 3, we compare the best-performing layers in our experiments with previous word segmentation algorithms that take S3M features as inputs. With our simple training-free method, we obtain the best F1 score of 41.0% using the 10th layer of *VG-HuBERT*. This outperforms a previously published attention-based approach using *VG-HuBERT* (Peng and Harwath, 2022b) and a recent dynamic programming-based approach that also trains an autoencoder on top (Kamper, 2022). However, we note that our approach falls short in terms of R-values, implying that it tends to over-segment. This can possibly be improved by designing different criteria for hyper-parameter selection, as our criterion is solely based on the F1 score.

Method	Prec.	Rec.	F1	R-val.
<i>Prior work</i> ¹⁵				
DPDP (Kamper, 2022)	35.3	37.7	36.4	44.3
<i>VG-HuBERT</i> (Peng and Harwath, 2022b)	36.2	32.2	34.1	45.6
<i>Ours (Best Layer)</i>				
<i>WavLM-Base</i> (L8)	31.9	45.7	37.6	30.7
<i>HuBERT-Base</i> (L9)	33.8	46.6	39.2	34.9
<i>wav2vec2-Base</i> (L7)	27.0	47.2	34.3	8.9
<i>VG-HuBERT</i> (L10)	36.0	47.6	41.0	39.5

Table 3: Word segmentation performance on the Buckeye test set. Higher is better for all metrics.

5.2 Analysis of Pooled Span Representations

Next, we measure how correlated word segmentation representations are with the other linguistic properties from Table 1: pronunciation, syntactic (POS) attributes, and semantic attributes (Section 5.2.1), and we evaluate mean-pooled utterance representations on sentence similarity (Section 5.2.2). Our remaining word-level experiments consider mean-pooled word segment representations as these consistently correlate well with word ID (Figures 4 and 5).

¹⁵GradSeg (Fuchs and Hoshen, 2023) also shows impressive word segmentation results on the Buckeye dataset, but they provide results only on the validation set, making it difficult to compare.

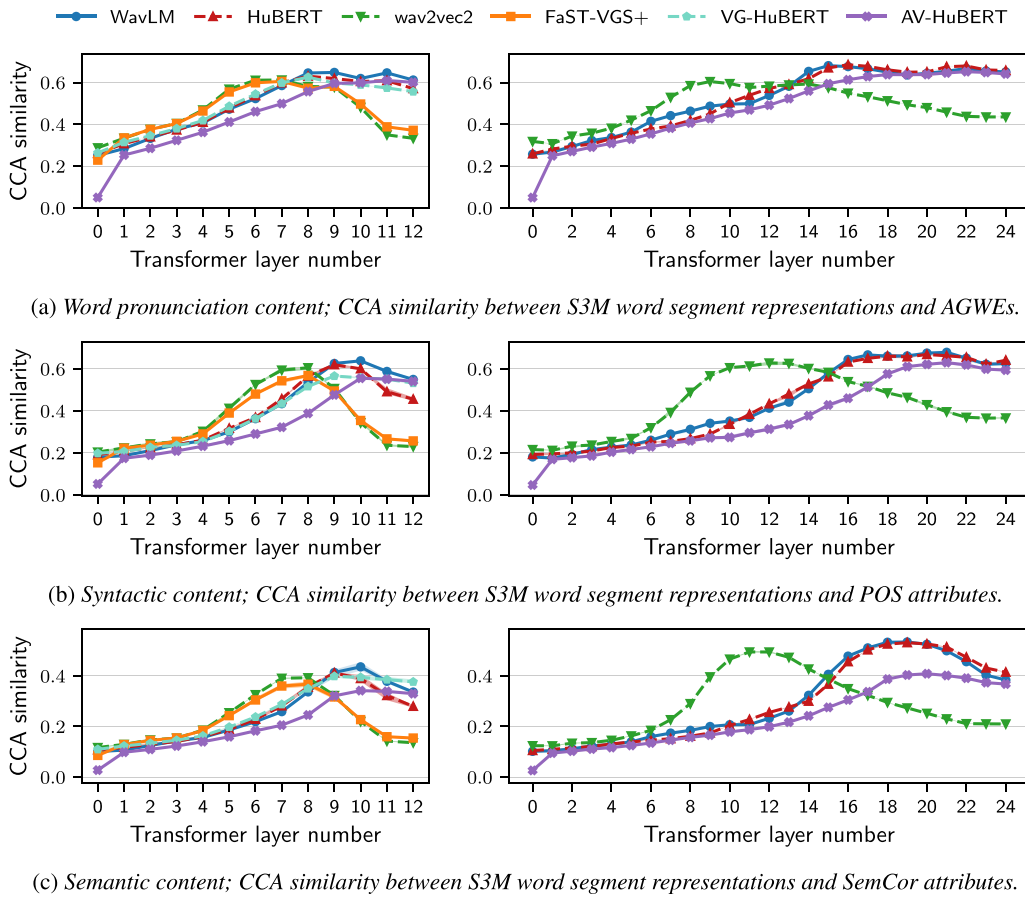


Figure 7: Measure of different linguistic properties using CCA for *Base* (left) and *Large* (right) S3Ms.

5.2.1 Similarity with Linguistic Properties

In Figure 7, we observe that models trained to recover local features (*wav2vec2* and *FaST-VGS+*) have the highest correlation at central layers, specifically layers 5–7 for *Base* models and layers 8–11 for the *Large* model. The rest of the models are trained to recover discrete units from an intermediate layer and have the highest correlation at much higher layers. This dependence on the form of pre-training objective has been observed before for lower-level acoustic and phonetic features (Pasad et al., 2023).

As seen for our other experiments (Figures 2, 6), the audio-visual models (*AV-HuBERT* and *VG-HuBERT*) see the least drop off in the final layers. These models are optimized with an audio-visual objective, suggesting that meaningful linguistic content is retained better with visual grounding.

For all S3Ms, pronunciation content (Figure 7a) is best correlated at lower layers than syntactic (Figure 7b) and semantic properties (Figure 7c). In *Base* models, the same set of intermediate

layers is best correlated with both syntactic and semantic attributes. The *Large* models, on the other hand, have a more pronounced peak for semantic than syntactic content, which in turn has a narrower plateau than the word pronunciation trends (Figure 7 right).

This differs from some observations made for BERT, a pre-trained *text* model, where different linguistic features—such as POS, constituents, dependencies, and entities—are encoded best at different layers (Tenney et al., 2019). This difference is possibly because the speech pre-training objective is mostly local with much of the model capacity (i.e., the majority of the layers) devoted to inferring local acoustic and lower-level phonetic features. Meanwhile, text models that start with higher-level segmented sub-word units have the capacity to encode fine-grained linguistic properties in different layers. BERT’s superiority in linguistic knowledge is supported by Shen et al. (2023) where BERT outperforms *wav2vec2* and *HuBERT* by 20% relative on a parsing-related probing task.




Figure 8: Performance on spoken STS task using representations from *Base* (left) and *Large* (right) S3Ms.




Figure 9: Visualization of the embedding spaces of the intermediate layers of S3Ms. Each point represents one word sample. Only the 6 most common POS tags are shown.

To qualitatively study the syntactic information encoded in S3M representations, we visualize the mean-pooled word representations from the layers with high correlation with the PTB syntactic vectors (Figure 7b). We sample $\sim 7k$ word instances across 500 distinct words and apply t-SNE to project the word representations to 2 dimensions (Figure 9). We find that, for *WavLM*, word samples with the same POS tag (especially for verbs, nouns, and adpositions) are encoded into vectors close to each other. However, the representations of *wav2vec2* are not as well-separated. These visualizations further corroborate our findings from CCA trends (Figure 7b), where *WavLM* shows a greater correlation than *wav2vec2*.

5.2.2 Sentence-level Semantics

Figure 8 shows the layer-wise performance on the spoken sentence similarity task. We include two baselines: (i) *FBank* uses mean-pooled filter-bank features as a sentence representation, and (ii) *naive* text baseline reports the fraction of word overlap in text transcripts between a pair of sentences.

Although the *naive* text baseline has a non-trivial correlation score of 0.4, the best-performing layers outperform the baselines by at least 50%. These results suggest that the mean-pooled S3M representations encode meaningful content beyond just the local acoustics and word identities.

The CLS token of *VG-HuBERT* has the best correlation score of 0.64 at layer 11, closely followed by layer 8 of *VG-HuBERT* and *FaST-VGS+*, both visually grounded models. The speech-only S3Ms we analyze outperform other S3Ms previously evaluated on this task (Merx et al., 2021; Zhu et al., 2022).¹⁶ However, they all underperform a text oracle baseline from Zhu et al. (2022) using self-supervised text embeddings (*SimCSE-unsup-RoBERTa*), which has a correlation score of 0.77.

5.3 Effect of Domain on Task-based Evaluation

Prior work evaluating S3Ms on downstream tasks has demonstrated how the relative ranking of S3Ms may be influenced by the domain of an S3M’s pre-training data as well as the evaluation methodology (Hsu et al., 2021b; Tsai et al., 2022; Yang et al., 2021; Zaiem et al., 2023b). For instance, similarly to all our task-based experiments (Figures 2b, 6, and 8), the SUPERB benchmarks (Tsai et al., 2022; Yang et al., 2021)¹⁷ and Zaiem et al. (2023b) report instances where some *Large* S3Ms under-perform their *Base* counterparts on downstream tasks.

Next, we discuss our takeaways related to the effect of (mis-)match between the domain of pre-training data and task data on some of our analysis experiments.

¹⁶The comparison with Merx et al. (2021) is based on Pearson’s correlation, not reported here.

¹⁷<https://superbbenchmark.org/leaderboard>.

5.3.1 Acoustic Word Discrimination

We evaluate *pool-AWD* on both LibriSpeech (Figure 2b), a read speech domain, and Switchboard (Figure 2c), a conversational speech domain. We observe that the relative ranking of S3Ms differs for the two settings. For instance, *AV-HuBERT* has better performance on Switchboard, outperforming all *Base* models, whereas all other S3Ms have higher scores on LibriSpeech. *WavLM-Large* outperforms *WavLM-Base* on Switchboard but the larger model underperforms on LibriSpeech. In both cases, the domain of pre-training data provides a potential explanation. Specifically, *AV-HuBERT* models are pre-trained on TED videos (Afouras et al., 2018) and *WavLM-Large* is pre-trained on a mix of data (Chen et al., 2021; Wang et al., 2021) including orated speech and spontaneous speech, whereas all other S3Ms are trained on read speech domains (Hsu et al., 2021a; Kahn et al., 2020; Panayotov et al., 2015).

We note that some cross-model rankings are consistent across evaluation domains. For instance, *HuBERT* and *WavLM*, both pre-trained to predict discrete cluster IDs from intermediate layers, outperform *wav2vec2*, which is trained to recover local features. As seen for other task-based evaluation (Sections 5.1.3, 5.2.2), the visually grounded models, *FaST-VGS+* and *VG-HuBERT*, outperform the speech-only *Base* models, *wav2vec2* and *HuBERT*, used to initialize them.

Additionally, we observe that the layer-wise trends for all S3Ms are consistent across evaluation domains and follow a similar dependence on the pre-training objective as noted by our previous results (Section 5.2.1) and some prior work (Pasad et al., 2023).

5.3.2 Word Segmentation

We evaluate word segmentation on LibriSpeech (Figure 6a) and Buckeye (Figure 6b). Similarly to previous findings, we observe that the relative ranking of S3Ms differs for the two settings. Specifically, S3Ms pre-trained solely on LibriSpeech (*wav2vec2-Base*, *HuBERT-Base*, *WavLM-Base*) take a much larger hit in performance when evaluated on Buckeye, and the visually grounded models, on the other hand, have a slightly better performance on Buckeye than on LibriSpeech. Again, the layer-wise trends for most S3Ms are invariant to the evaluation domain. *WavLM-Large* does not follow this trend

and more than half of the layers have a drastically poorer performance on Buckeye. We hypothesize that the hyperparameters (tuned on LibriSpeech, Section 4.4) transfer better for other *Large* models than for *WavLM-Large*, due to domain mismatch, as discussed above for *pool-AWD* (Section 5.3.1).

6 Conclusion

The analyses presented here further our understanding of S3Ms, specifically their representation of word-level properties. Some of our findings corroborate patterns found in earlier work; for example, the most linguistically “deep” information appears to be encoded best in a small set of intermediate layers, and pre-training objective and model size influence layer-wise trends. We contribute new findings about previously unstudied aspects of S3Ms, such as the distribution of word information within word segments and the encoding of syntactic and semantic features. Most importantly, the comparison of a large number of models using the same analyses and tasks, and the study of multiple word-level properties, enables a more complete understanding of the space of S3Ms. As an additional product of this work, we obtained strong results on multiple benchmark tasks, outperforming prior work using simple models based on frozen S3M representations.

Our work studies which S3M layers are better (or worse) at encoding certain linguistic properties. Previous studies have used similar findings to guide modeling decisions when adapting pre-trained models for downstream tasks (Pasad et al., 2023; Xie et al., 2022; Yang et al., 2023a), including the choice of which layers to drop, distill,¹⁸ or reinitialize (Chang et al., 2022; Choi et al., 2021; Hsu et al., 2021c; Hwang et al., 2022; Li et al., 2023b; Pasad et al., 2021; Zaiem et al., 2023a). We therefore expect our findings to inform design choices for both model development and their utilization for downstream tasks. For instance, for all the S3Ms we study, our analysis reveals that linguistic content is most prominent within the intermediate layers (Figure 7). Since layers encoding semantic content should be particularly beneficial for language understanding tasks, our findings suggest an exploration of alternative strategies to the common

¹⁸We use the term “distill” to encompass various modeling variants such as transfer learning and using post-processed activations as targets, in addition to model distillation.

practice of adding a prediction head to the topmost layer (Shon et al., 2022, 2023).

Our analyses have addressed several questions about S3Ms' word-level representations, thereby providing a foundation to address more challenging questions. For example, a natural next step is to ask how much (and where) phrase- and sentence-level properties, such as constituents, dependencies, and entities, are encoded. For some tasks, such as word segmentation, they are far from solving the task, although our results with S3Ms are stronger than prior work. Finally, we have noted (as have some prior studies) that larger models are not always better by all measures, raising the question of what the additional model capacity provides and whether there is a better way to train and utilize larger models.

Acknowledgments

We thank the anonymous reviewers and the action editor for their time and helpful feedback. This work is partially supported by AFOSR grant FA9550-18-1-0166.

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*. <https://doi.org/10.21437/Interspeech.2023-2131>
- Triantafyllos Afouras, Joon Son Chung, and Andrew Senior. 2018. LRS3-TED: A large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*.
- Robin Algayres, Tristan Ricoul, Julien Karadayi, Hugo Laurençon, Salah Zaiem, Abdelrahman Mohamed, Benoît Sagot, and Emmanuel Dupoux. 2022. DP-Parse: Finding word boundaries from raw speech with an instance lexicon. *Transactions of the Association for Computational Linguistics (TACL)*. https://doi.org/10.1162/tACL_a_00505
- Robin Algayres, Mohamed Zaiem, Benoît Sagot, and Emmanuel Dupoux. 2020. Evaluating the reliability of acoustic speech embeddings. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2020-2362>
- Takanori Ashihara, Takafumi Moriya, Kohei Matsuura, Tomohiro Tanaka, Yusuke Ijima, Taichi Asami, Marc Delcroix, and Yukinori Honma. 2023. SpeechGLUE: How well can self-supervised speech models capture linguistic knowledge? In *Interspeech*. <https://doi.org/10.21437/Interspeech.2023-1823>
- Alexei Baevski, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2021. Unsupervised speech recognition. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. 2022. Data2vec: A general framework for self-supervised learning in speech, vision and language. In *International Conference on Machine Learning (ICML)*.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Stefano Bannò and Marco Matassoni. 2023. Proficiency assessment of L2 spoken English using wav2vec 2.0. In *IEEE Spoken Language Technology Workshop (SLT)*. <https://doi.org/10.1109/SLT54892.2023.10023019>
- Yonatan Belinkov. 2022. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*. https://doi.org/10.1162/coli_a_00422
- Yonatan Belinkov and James Glass. 2019. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics (TACL)*. https://doi.org/10.1162/tACL_a_00254
- Saurabhchand Bhati, Jesús Villalba, Piotr Żelasko, Laureano Moro-Velazquez, and Najim Dehak. 2021. Segmental contrastive predictive coding for unsupervised word segmentation. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1874>
- Michael A. Carlin, Samuel Thomas, Aren Jansen, and Hynek Hermansky. 2011. Rapid evaluation of speech representations for spoken term discovery. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2011-304>

- Heng-Jui Chang, Shu-wen Yang, and Hung-yi Lee. 2022. DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit BERT. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP43922.2022.9747490>
- Guoguo Chen, Shuzhou Chai, Guanbo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khundampur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Yujun Wang, Zhao You, and Zhiyong Yan. 2021. GigaSpeech: An evolving, multi-domain ASR corpus with 10,000 hours of transcribed audio. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1965>
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei. 2022. WavLM: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing (JSTSP)*. <https://doi.org/10.1109/JSTSP.2022.3188113>
- Hyeong-Seok Choi, Juheon Lee, Wansoo Kim, Jie Lee, Hoon Heo, and Kyogu Lee. 2021. Neural analysis and synthesis: Reconstructing speech from self-supervised representations. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Alexis Conneau and Douwe Kiela. 2018. SentEval: An evaluation toolkit for universal sentence representations. In *International Conference on Language Resources and Evaluation (LREC)*.
- Santiago Cuervo, Maciej Grabias, Jan Chorowski, Grzegorz Ciesielski, Adrian Łańcucki, Paweł Rychlikowski, and Ricard Marxer. 2022. Contrastive prediction strategies for unsupervised segmentation and categorization of phonemes and words. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP43922.2022.9746102>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Ewan Dunbar, Julien Karadayi, Mathieu Bernard, Xuan-Nga Cao, Robin Algayres, Lucas Ondel, Laurent Besacier, Sakriani Sakti, and Emmanuel Dupoux. 2020. The zero resource speech challenge 2020: Discovering discrete subword and word units. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2020-2743>
- Zhiyun Fan, Meng Li, Shiyu Zhou, and Bo Xu. 2021. Exploring wav2vec 2.0 on speaker verification and language identification. In *Interspeech*.
- Manaal Faruqui and Chris Dyer. 2014. Community evaluation and exchange of word vectors at wordvectors.org. In *Association for Computational Linguistics (ACL): System Demonstrations*. <https://doi.org/10.3115/v1/P14-5004>
- Manaal Faruqui, Yulia Tsvetkov, Pushpendre Rastogi, and Chris Dyer. 2016. Problems with evaluation of word embeddings using word similarity tasks. In *1st Workshop on Evaluating Vector-Space Representations for NLP*. <https://doi.org/10.18653/v1/W16-2506>
- Chi-Luen Feng, Po-chun Hsu, and Hung-yi Lee. 2022. Silence is sweeter than speech: Self-supervised model using silence to store speaker information. *arXiv preprint arXiv:2205.03759*.
- Tzeviya Sylvia Fuchs and Yedid Hoshen. 2023. Unsupervised word segmentation using temporal gradient pseudo-labels. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP.1992.225858>
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong

- Zhang, Yonghui Wu, and Ruoming Pang. 2020. Conformer: Convolution-augmented transformer for speech recognition. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2020-3015>
- David Harwath and James Glass. 2017. Learning word-like units from joint audio-visual analysis. In *Association for Computational Linguistics (ACL)*. <https://doi.org/10.18653/v1/P17-1047>
- Wanjia He, Weiran Wang, and Karen Livescu. 2017. Multi-view recurrent neural acoustic word embeddings. In *International Conference on Learning Representations (ICLR)*.
- John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. In *Empirical Methods in Natural Language Processing (EMNLP)*. <https://doi.org/10.18653/v1/D19-1275>
- Harold Hotelling. 1936. Relations between two sets of variates. *Biometrika*. <https://doi.org/10.1093/biomet/28.3-4.321>
- Wei-Ning Hsu, David Harwath, Christopher Song, and James Glass. 2021a. Text-free image-to-speech synthesis using learned segmental units. In *Association for Computational Linguistics (ACL)*.
- Wei-Ning Hsu, Anuroop Sriram, Alexei Baevski, Tatiana Likhomanenko, Qiantong Xu, Vineel Pratap, Jacob Kahn, Ann Lee, Ronan Collobert, Gabriel Synnaeve, and Michael Auli. 2021b. Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training. In *Interspeech*.
- Wei-Ning Hsu, Yao-Hung Hubert Tsai, Benjamin Bolte, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021c. HuBERT: How much can a bad teacher benefit asr pre-training? In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- Yushi Hu, Shane Settle, and Karen Livescu. 2020. Multilingual jointly trained acoustic and written word embeddings. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2020-2828>
- Dongseong Hwang, Khe Chai Sim, Zhouyuan Huo, and Trevor Strohman. 2022. Pseudo label is better than human label. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2022-11034>
- Aren Jansen, Samuel Thomas, and Hynek Hermansky. 2013. Weak top-down constraints for unsupervised acoustic model training. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP.2013.6639241>
- Hang Ji, Tanvina Patel, and Odette Scharenborg. 2022. Predicting within and across language phoneme recognition performance of self-supervised learning speech pre-trained models. *arXiv preprint arXiv:2206.12489*.
- Jacob Kahn, Morgane Rivi re, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazar , Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, Tatiana Likhomanenko, Gabriel Synnaeve, Armand Joulin, Abdelrahman Mohamed, and Emmanuel Dupoux. 2020. Libri-Light: A benchmark for ASR with limited or no supervision. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP40776.2020.9052942>
- Herman Kamper. 2022. Word segmentation on discovered phone units with dynamic programming and self-supervised scoring. *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*. <https://doi.org/10.1109/TASLP.2022.3229264>
- Herman Kamper, Weiran Wang, and Karen Livescu. 2016. Deep convolutional acoustic word embeddings using word-pair side information. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP.2016.7472619>
- Eesung Kim, Jae-Jin Jeon, Hyeji Seo, and Hoon Kim. 2022. Automatic pronunciation assessment using self-supervised speech representation learning. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2022-10245>
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. Similarity of neural network representations revisited. In *International Conference on Machine Learning (ICML)*.
- Felix Kreuk, Joseph Keshet, and Yossi Adi. 2020. Self-supervised contrastive learning for

- unsupervised phoneme segmentation. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2020-2398>
- Kushal Lakhotia, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, and Emmanuel Dupoux. 2021. On generative spoken language modeling from raw audio. *Transactions of the Association for Computational Linguistics (TACL)*.
- Keith Levin, Katharine Henry, Aren Jansen, and Karen Livescu. 2013. Fixed-dimensional acoustic embeddings of variable-length segments in low-resource settings. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. <https://doi.org/10.1109/ASRU.2013.6707765>
- Yuanchao Li, Yumnah Mohamied, Peter Bell, and Catherine Lai. 2023a. Exploration of a self-supervised speech model: A study on emotional corpora. In *IEEE Spoken Language Technology Workshop (SLT)*.
- Zhengyang Li, Thomas Graave, Jing Liu, Timo Lohrenz, Siegfried Kunzmann, and Tim Fingscheidt. 2023b. Parameter-efficient cross-language transfer learning for a language-modular audiovisual speech recognition. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*.
- Oli Liu, Hao Tang, and Sharon Goldwater. 2023. Self-supervised predictive coding models encode speaker and phonetic information in orthogonal subspaces. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2023-871>
- Loren Lugosch, Mirco Ravanelli, Patrick Ignoto, Vikrant Singh Tomar, and Yoshua Bengio. 2019. Speech model pre-training for end-to-end spoken language understanding. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2019-2396>
- Danni Ma, Neville Ryant, and Mark Liberman. 2021. Probing acoustic representations for phonetic properties. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP39728.2021.9414776>
- Mitchell Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn treebank. *Computational Linguistics*. <https://doi.org/10.21236/ADA273556>
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal forced aligner: Trainable text-speech alignment using kald. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2017-1386>
- Danny Merckx, Stefan L. Frank, and Mirjam Ernestus. 2021. Semantic sentence similarity: Size does not always matter. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1464>
- Danny Merckx, Sebastiaan Scholten, Stefan L. Frank, Mirjam Ernestus, and Odette Scharenborg. 2023. Modelling human word learning and recognition using visually grounded speech. *Cognitive Computation*. <https://doi.org/10.1007/s12559-022-10059-7>
- George A. Miller, Claudia Leacock, Randee Tengi, and Ross T. Bunker. 1993. A semantic concordance. In *Human Language Technology*. <https://doi.org/10.3115/1075671.1075742>
- Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaløe, Tara N. Sainath, and Shinji Watanabe. 2022. Self-supervised speech representation learning: A review. *IEEE Journal of Selected Topics in Signal Processing (JSTSP)*. <https://doi.org/10.1109/JSTSP.2022.3207050>
- Ari Morcos, Maithra Raghu, and Samy Bengio. 2018. Insights on representational similarity in neural networks with canonical correlation. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Tu Anh Nguyen, Maureen De Seyssel, Robin Algayres, Patricia Rozé, Ewan Dunbar, and Emmanuel Dupoux. 2022. Are word boundaries useful for unsupervised language learning? *arXiv preprint arXiv:2210.02956*.
- Tu Anh Nguyen, Maureen de Seyssel, Patricia Rozé, Morgane Rivière, Evgeny Kharitonov, Alexei Baevski, Ewan Dunbar, and Emmanuel Dupoux. 2020. The Zero Resource Speech Benchmark 2021: Metrics and baselines for

- unsupervised spoken language modeling. In *NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing*.
- Tu Anh Nguyen, Eugene Kharonov, Jade Copet, Yossi Adi, Wei-Ning Hsu, Ali Elkahky, Paden Tomasello, Robin Algayres, Benoit Sagot, Abdelrahman Mohamed, and Emmanuel Dupoux. 2023. Generative spoken dialogue language modeling. *Transactions of the Association for Computational Linguistics (TACL)*. https://doi.org/10.1162/tacl_a_00545
- Shruti Palaskar, Vikas Raunak, and Florian Metze. 2019. Learned in speech recognition: Contextual acoustic word embeddings. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP.2019.8683868>
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. LibriSpeech: An ASR corpus based on public domain audio books. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP.2015.7178964>
- Ankita Pasad, Ju-Chieh Chou, and Karen Livescu. 2021. Layer-wise analysis of a self-supervised speech representation model. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. <https://doi.org/10.1109/ASRU51503.2021.9688093>
- Ankita Pasad, Bowen Shi, and Karen Livescu. 2023. Comparative layer-wise analysis of self-supervised speech models. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP49357.2023.10096149>
- Ankita Pasad, Felix Wu, Suwon Shon, Karen Livescu, and Kyu J. Han. 2022. On the use of external data for spoken named entity recognition. In *North American Chapter of the Association for Computational Linguistics (NAACL)*. <https://doi.org/10.18653/v1/2022.naacl-main.53>
- Puyuan Peng and David Harwath. 2022a. Self-supervised representation learning for speech using visual grounding and masked language modeling. In *AAAI Workshop on Self-supervised Learning for Audio and Speech Processing*. <https://doi.org/10.21437/Interspeech.2022-10652>
- Puyuan Peng and David Harwath. 2022b. Word discovery in visually grounded, self-supervised speech models. In *Interspeech*.
- Puyuan Peng, Herman Kamper, and Karen Livescu. 2020. A correspondence variational autoencoder for unsupervised acoustic word embeddings. In *NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. <https://doi.org/10.3115/v1/D14-1162>
- Mark A. Pitt, Keith Johnson, Elizabeth Hume, Scott Kiesling, and William Raymond. 2005. The Buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability. *Speech Communication*. <https://doi.org/10.1016/j.specom.2004.09.001>
- Archiki Prasad and Preethi Jyothi. 2020. How accents confound: Probing for accent information in end-to-end speech recognition systems. In *Association for Computational Linguistics (ACL)*. <https://doi.org/10.18653/v1/2020.acl-main.345>
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. SVCCA: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Neural Information Processing Systems (NIPS)*.
- Okko Räsänen, Unto K. Laine, and Toomas Altsaar. 2011. Blind segmentation of speech using non-linear filtering methods. *Speech Technologies*. <https://doi.org/10.5772/16433>
- Abhilasha Ravichander, Yonatan Belinkov, and Eduard Hovy. 2021. Probing the probing paradigm: Does probing accuracy entail task relevance? In *European Chapter of the Association for Computational Linguistics (EACL)*. <https://doi.org/10.18653/v1/2021.eacl-main.295>

- Ramon Sanabria, Hao Tang, and Sharon Goldwater. 2021. On the difficulty of segmenting words with attention. In *Second Workshop on Insights from Negative Results in NLP*. <https://doi.org/10.18653/v1/2021.insights-1.11>
- Ramon Sanabria, Hao Tang, and Sharon Goldwater. 2023. Analyzing acoustic word embeddings from pre-trained self-supervised speech models. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP49357.2023.10096099>
- Naomi Saphra and Adam Lopez. 2019. Understanding learning dynamics of language models with SVCCA. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Shane Settle, Kartik Audhkhasi, Karen Livescu, and Michael Picheny. 2019. Acoustically grounded word embeddings for improved acoustics-to-word speech recognition. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP.2019.8682903>
- Shane Settle and Karen Livescu. 2016. Discriminative acoustic word embeddings: Recurrent neural network-based approaches. In *IEEE Spoken Language Technology Workshop (SLT)*. <https://doi.org/10.1109/SLT.2016.7846310>
- Jui Shah, Yaman Kumar Singla, Changyou Chen, and Rajiv Ratn Shah. 2021. What all do audio transformer models hear? Probing acoustic representations for language delivery and its structure. In *IEEE International Conference on Data Mining Workshops (ICDMW)*.
- Gaofei Shen, Afra Alishahi, Arianna Bisazza, and Grzegorz Chrupała. 2023. Wave to syntax: Probing spoken language models for syntax. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2023-679>
- Bowen Shi, Wei-Ning Hsu, Kushal Lakhota, and Abdelrahman Mohamed. 2022. Learning audio-visual speech representation by masked multimodal cluster prediction. In *International Conference on Learning Representations (ICLR)*.
- Bowen Shi, Shane Settle, and Karen Livescu. 2021. Whole-word segmental speech recognition with acoustic word embeddings. In *IEEE Spoken Language Technology Workshop (SLT)*. <https://doi.org/10.1109/SLT48900.2021.9383578>
- Suwon Shon, Siddhant Arora, Chyi-Jiunn Lin, Ankita Pasad, Felix Wu, Roshan Sharma, Wei-Lun Wu, Hung-Yi Lee, Karen Livescu, and Shinji Watanabe. 2023. SLUE Phase-2: A benchmark suite of diverse spoken language understanding tasks. In *Association for Computational Linguistics (ACL)*. <https://doi.org/10.18653/v1/2023.acl-long.496>
- Suwon Shon, Ankita Pasad, Felix Wu, Pablo Brusco, Yoav Artzi, Karen Livescu, and Kyu J. Han. 2022. SLUE: New benchmark tasks for spoken language understanding evaluation on natural speech. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP43922.2022.9746137>
- Louis ten Bosch and Bert Cranen. 2007. A computational model for unsupervised word discovery. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2007-429>
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT rediscovers the classical NLP pipeline. In *Association for Computational Linguistics (ACL)*. <https://doi.org/10.18653/v1/P19-1452>
- Hsiang-Sheng Tsai, Heng-Jui Chang, Wen-Chin Huang, Zili Huang, Kushal Lakhota, Shu-wen Yang, Shuyan Dong, Andy T. Liu, Cheng-I Jeff Lai, Jiatong Shi, Xuankai Chang, Phil Hall, Hsuan-Jui Chen, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung-yi Lee. 2022. SUPERB-SG: Enhanced speech processing universal performance benchmark for semantic and generative capabilities. In *Association for Computational Linguistics (ACL)*.
- Yulia Tsvetkov, Manaal Faruqui, and Chris Dyer. 2016. Correlation-based intrinsic evaluation of word vector representations. In *1st Workshop on Evaluating Vector-Space Representations for NLP*. <https://doi.org/10.18653/v1/W16-2520>

- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Guillaume Lample, and Chris Dyer. 2015. Evaluation of word vector representations by subspace alignment. In *Empirical Methods in Natural Language Processing (EMNLP)*. <https://doi.org/10.18653/v1/D15-1243>
- Benjamin van Niekerk, Leanne Nortje, Matthew Baas, and Herman Kamper. 2021. Analyzing speaker information in self-supervised models to improve zero-resource speech processing. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1182>
- Lisa Van Staden and Herman Kamper. 2021. A comparison of self-supervised speech representations as input features for unsupervised acoustic word embeddings. In *IEEE Spoken Language Technology Workshop (SLT)*. <https://doi.org/10.1109/SLT48900.2021.9383625>
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems (NIPS)*.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, Ilhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*. <https://doi.org/10.1038/s41592-020-0772-5>, PubMed: 32094914
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives. In *North American Chapter of the Association for Computational Linguistics (NAACL)*. <https://doi.org/10.18653/v1/D19-1448>
- Changan 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 *Association for Computational Linguistics (ACL)*. <https://doi.org/10.18653/v1/2021.acl-long.80>
- Felix Wu, Kwangyoun Kim, Shinji Watanabe, Kyu J. Han, Ryan McDonald, Kilian Q. Weinberger, and Yoav Artzi. 2023. Wav2seq: Pre-training speech-to-text encoder-decoder models using pseudo languages. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSP49357.2023.10096988>
- Shuo Xie, Jiahao Qiu, Ankita Pasad, Li Du, Qing Qu, and Hongyuan Mei. 2022. Hidden state variability of pretrained language models can guide computation reduction for transfer learning. In *Findings of Empirical Methods in Natural Language Processing (EMNLP)*. <https://doi.org/10.18653/v1/2022.findings-emnlp.422>
- Gene-Ping Yang, Yue Gu, Qingming Tang, Dongsu Du, and Yuzong Liu. 2023a. On-device constrained self-supervised speech representation learning for keyword spotting via knowledge distillation. In *Interspeech*.
- Mu Yang, Ram CMC Shekar, Okim Kang, and John HL Hansen. 2023b. What can an accent identifier learn? Probing phonetic and prosodic information in a wav2vec2-based accent identification model. In *Interspeech*.
- Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhota, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Ko-tik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung-yi Lee. 2021. SUPERB: Speech processing universal performance benchmark. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1775>
- Salah Zaiem, Robin Algayres, Titouan Parcollet, Slim Essid, and Mirco Ravanelli. 2023a. Fine-tuning strategies for faster inference using

speech self-supervised models: A comparative study. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. <https://doi.org/10.1109/ICASSPW59220.2023.10193042>

Salah Zaiem, Youcef Kemiche, Titouan Parcollet, Slim Essid, and Mirco Ravanelli. 2023b. Speech self-supervised representations benchmarking: A case for larger probing heads. *arXiv preprint*

arXiv:2308.14456. <https://doi.org/10.2139/ssrn.4733627>

Jian Zhu, Zuoyu Tian, Yadong Liu, Cong Zhang, and Chia-wen Lo. 2022. Bootstrapping meaning through listening: Unsupervised learning of spoken sentence embeddings. In *Findings of Empirical Methods in Natural Language Processing (EMNLP)*. <https://doi.org/10.18653/v1/2022.findings-emnlp.81>